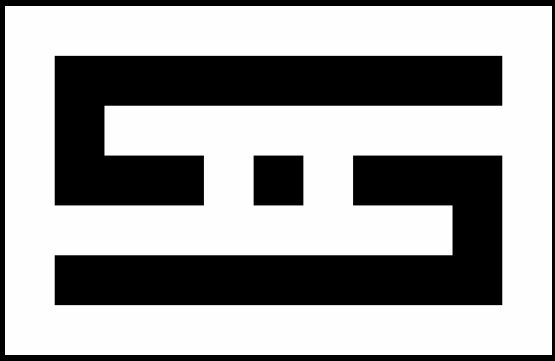

SCHIT HAPPENED

The Story of the World's Most Improbable Start-Up

by Jason Stoddard

DARK VERSION

Typesetting by Lucas van den Bosch
lucas.bosch@tum.de
lblb on Head-Fi.org

Created with LaTeX and KOMA-Script
Text in Bitstream Charter
Headings in Lato

To Mike, for going along with this crazy idea.

To Rina, for support, naming the company,
and tolerating the whole mess.

To Eddie, for helping us really get started
when we needed it most.

To Tony, for testing dang near everything we've made.

To Alex, for taking over, making things right,
and running the show.

To Dave, for doing the hardest stuff.

Hey all,

I'm Jason Stoddard, the co-founder (with Mike Moffat) of Schiit Audio. You may have seen some of my posts on Head-fi. But you haven't heard the full story of Schiit … why we started the company, how we've grown from June 2010 to today, the good stuff … and the bad stuff. So, in a burst of insanity, I thought, "Why not turn this into a book?"

Why? Well, in addition to Schiit, I've written plenty of stuff, including about 30 published stories and 3 books. Jude and the team at Head-fi were kind enough to provide this space for me to serialize this story. The plan is to post a chapter a week until it's done, which should take us through the summer.

Now, you may be wondering, "Why the heck should I read some crazy book about Schiit?" Well, there's nobody forcing you to. But if you're interested in the inner workings of a true by-the-bootstraps startup (did you know we started in a garage, with no external investment?) or if you've been thinking about starting your own business (beyond the usual platitudes of "hey, you should incorporate, and there's this stuff called 'cash flow' and stuff"), or if you're just curious

about how a modern audio company works (from who's actually engineering the gear, to how it gets made, to the ups and downs of the day to day working of the company), you may enjoy it, or find it helpful and informative if you want to start a company yourself. I have no illusions that this will be the next business best-seller—if it ever makes it into print—but I'm having fun writing it, and I'm hoping you'll like reading it.

Disclaimer: There'll be plenty of technical stuff in here, because, after all, I am an engineer. You won't have to understand it all, but you may be more entertained if you have a techy side. Later on, there'll be some early photos from Schiit, as well as a video or two, that have never been seen before.

Thanks for reading, and I hope you enjoy it!

All the best,
Jason Stoddard

2014

Forward

Christmas Presents Until the End of Time?

So, do you think it'll go? Do you think they'll sell?" Mike Moffat asked, looking at the first assembled Asgard on the engineering bench in my garage. He was being Mike-fidgety, rocking from heel to heel in the small, chilly space.

"Well, on paper it looks good," I told him. "But you know how that works. They'll either sell, or we'll have Christmas presents until the end of time."

Mike laughed, a little nervously. Because he knew how it goes. You can plan and study, do endless market research and cost studies, run focus groups and get tons of input from key prospects and do all the little things that companies do to procrastinate and dither and second-guess before putting out another "gotta have" product … and things can still go sideways.

But this isn't a story about stuff like that. This is

a story about gut feelings, good guesses, and not following the herd. And succeeding.

This is the story of Schiit Audio, the world's most improbable start-up.

Yes. Schiit. Let's start with that as an improbability factor. What company in its right mind would name itself that? I mean, if you were a marketing agency and proposed that name to a client, how would they react? You'd be picking your butt up off the pavement outside their headquarters, post-haste.

But that isn't all that made us a crap candidate for succeeding. Consider:

- We started this with no outside funds, no VC, no crowdfunding
- We'd both been out of audio for about 15+ years—more on that later
- We went with direct sales, even though that had only really worked for one other company—Emotiva
- We started with no staff, in my garage
- We decided to make everything in the USA, even though the prevailing wisdom of the time was "China's the World's Manufacturing Floor, why even try to compete?"

- And, in a complete burst of insanity, we decided to start with inexpensive products

Ah, and it's now probably past time I introduce myself. I'm Jason Stoddard, Co-Founder of Schiit Audio. Mike Moffat's my business partner. Our official titles are "Head" and "Number 2" respectively. Hey, Mike asked for it. No, we don't take ourselves too seriously here. I won't bore you with our full CVs (that's fancy-speak for *wut we dun*), but you may have heard of Mike Moffat. He was the founder of Theta (the first one, the analog one), in the late 1970s. You can blame him, at least in part, for resurrecting tube audio. He was the first person to use 6DJ8s in audio. He installed Philip K. Dick's stereo systems. He sold amps to L. Ron Hubbard (no, you can't make this schiit up). Then, in the 1980s, he became the Father of the DAC with Theta Digital. His DSPre was the first standalone DAC on the market, and it was a showstopper—its own digital filter algorithms running on Motorola DSPs so powerful they couldn't be exported into the Soviet Union, for a start. Theta mopped up in the DAC world for several years, then Mike founded Angstrom, the maker of the world's first upgradable surround processor. From there, Mike moved into entertainment, creating complex systems for digital movie distribution. At least until I tempted him

away with Schiit.

I'm … well, I'm confused. I'm a published, award-winning science fiction author (strange-andhappy.com), a summa cum laude BS Engineering analog geek (schiit.com) and 20-year veteran of the marketing wars at another company I founded (centric.com). I've done stuff as strange as lecture Harvard professors on virtual world marketing, and as driven as earning my way to Vice-President, Engineering at Sumo at age 25, which nominally made me Ed Miller's boss—he was the founder of Souncraftsmen, Sherwood and Great American Sound, and head of engineering for SAE, to drop some names. Not that Ed cared, he just did his own thing. He was cool.

I'm the one writing this book. You can blame it all on me. I have no illusions of this being a best-seller, or of it changing the world. But I think we have an interesting story—one that others can learn from, both in and out of audio.

"Oh yeah?" you ask, leaning back and crossing your arms. "Well, I ain't gonna read no sixty thousand words about some small-time company just to get few phrases that belong on Sucksessories posters."

Cool. Gotcha. So I'll cut to the chase. If you're only interested in business intelligence, you won't have to read any further than the next 7 bullet points:

1. Shooting to be the next billion-dollar mass-market company is insane—you might as well buy lottery tickets.
2. Niche is where it's at—specifically a niche where people can get in fistfights over the color of a knob.
3. Pick a niche you know and love and do something nobody else can do—"me-too" never works.
4. Be memorable—this isn't about getting *everyone* to *like* you, this is about getting *some* people to *love* you.
5. Go direct—distribution is a poisonous remnant of 19th century economics in a disintermediated world.
6. Run from both conventional marketing wisdom and the social media mavens—both of them are geared towards the mass market with eight-digit ad budgets and multiple decades to build a brand.
7. Don't think this'll be easy—this is hard work, but you'll also be having a whole lot of fun if you're doing it right!

Okay, now you're skeptical. You're thinking: But

I just read a book from (insert the name of some multibillion-dollar-valuation corporate CEO here), and he said it's easy to reach the masses and change the world, and it seems like anyone can do it, why would I shoot for less than that?

That's cool. That is, if you're lucky enough to come up with something different enough to merit venture funding, if you get through all the rounds with the team and product intact, if something better doesn't come out of nowhere, if the public whims don't change, if you don't get ousted before the real money starts, if you're cool with 100 hour weeks and lots of travel and losing touch with the real fun of creation and becoming a new salesman with his dog-and-pony show for the money guys in Silly Valley or Singapore or wherever the money is in this moment, more power to you. Go ahead and create the next Amazon, Apple, Facebook, or Google. This book isn't for you, and you can stop reading now.

But I have a lot of friends who have gone down that route. Brilliant people. Hard workers. They don't have any problem with all of the above. They go even farther, begging and scraping to keep the team stapled together when the money gets thin, mortgaging everything they have on

the One Big Idea …

… and doing it again
when the first one doesn't get past angel funding.
… and doing it again
when the second one doesn't get its second round.
… and doing it again and again and again,
as many times as it takes.

Bottom line, there are plenty of billion-dollar ideas out there. Making one into a real company that succeeds isn't just a lot of work. It's about money, luck, connections, money, luck, money, and luck. And more luck.

This story is for people who don't have a lot of the above. For people who are shooting to create a company that might do a million a year, or ten million, or maybe a hundred million, eventually, way out in the future.

So, if you'd like to know more about where we came from, how we got started, why the crazy name, where the fixation on Norse mythology came from, our first successes, our first failures, what we screwed up later (hell, if you buy anything from us after reading how much we mess things up, it'll be a minor miracle), how we develop products and market and go to shows and work with suppliers and do everything in

the USA except the Magni wall-wart, and about ten thousand other things, read on.

If you're looking for a story that will make you an instant millionaire, cure cancer, repel an alien invasion, or thwart the plans of an evil CEO to turn the world into a dystopic corpocrat future, you've come to the wrong place.

Oh, and about those Asgard Christmas presents?

To date, we haven't had to give a single one away. In fact, staying in stock through the holidays is one of the hardest things to do around here. Yes. Even 42 months later.

And with that, let's flash back 20+ years …

Chapter 1

The Line is Down. Here's an Undocumented Test Rig. Fix It.

My first day at Sumo in 1989 was maybe the most bizarre first-day-on-the-job ever. I literally just walked in the door to find Ed Miller pacing the hallway, eyes darting from left to right. He'd clearly been waiting for me to show up.

"Oh good, you're here," he said, motioning for me to follow him. I trailed him past a series of nondescript offices, out onto the factory floor, and into a cluttered, messy little room with a giant glass window that looked out onto the PCB assembly line.

"We have a problem," Ed said. "I hope you can fix it. I took a look, but I don't have the notebooks. But if you can't fix it, maybe you can whip something up." I wasn't really listening to Ed. I was staring at the battered blue-and-white engineering workbench. Or trying to. Because the entire top surface of the workbench was

littered with about fifty pounds of discarded resistors, capacitors, output transistors, screws, wires, solder blobs, candy wrappers, scrawled post-its, and other unidentifiable electronic and non-electronic trash. And this was what was on top of a hunk of low-nap industrial carpet. Below it was about another inch of the same kind of junk. Below that, there was another piece of carpet, considerably more frayed, and another inch or two of junk.

Yes, this was what the previous engineer did. Instead of cleaning up, he just threw another rug over it and piled more junk on top.

"So, want to have a look?" Ed said, pointing at a battered aluminum Bud box in the middle of the carpet-desk. It had ten red LEDs on it, a socket for a power transistor, and a pushbutton.

"What is this?" I asked, trying to come back to reality.

"It's a MOSFET matcher. But it's not working. And the line is down," Ed said, pointing out the window to the PCB assembly team. Outside, ten women looked back at me, arms crossed, clearly waiting for something. "They can't get back to work until the MOSFET matcher is fixed."

Ah, crap. It finally sank in. They were waiting for this test rig. Everything was at a dead stop without it.

And the job of fixing it had just been dumped on me. I said nothing, suddenly realizing just how much I was in over my head. I was seven months out of college. I didn't even know MOSFETs needed matching. I didn't know what they matched. And I certainly didn't expect to see some homebuilt device at my great new engineering job. Hell, I'd just come out of working on spread-spectrum communications at Magnavox APS's "black hole" lab, where $100k gate array prototypes were all over the place.

"Do you have a schematic?" I asked, feeling a little ill.

"No. That's the problem."

Great. Just great.

At that moment, I thought about asking for my old job at Magnavox back. They told me I had an open door any time. And it was easy. Real easy. Hell, they got pissed when I worked too fast, because they couldn't bill out their entire project to the government. They'd never throw me into something like this. Hell, they'd given me two

weeks just to play with the layout software when I started at Magnavox.

Maybe I didn't want to be in audio after all.

But no. I wanted to be in audio. I loved audio. I had a speaker company on the side, one I'd started in college. My senior project in engineering had been a switched-capacitor adaptive noise reduction system. I had years in audio. I always dreamed about going into audio. I'd almost fallen off the chair when I'd seen the Help Wanted ad in the paper, and its enticing offer to "join the Sumo engineering team, and advance the future of audio." I'd been beside myself when I was asked to come in and interview.

And that was when I made the decision to stay. I couldn't let it all go now. *I'd figure it out. Somehow.*

"I'll take a look," I told Ed, and shooed him off.

And then I sat there, crapping my pants. The women were still staring at me. Nobody was working. Everything was in my hands.

No pressure, no pressure at all.

Business lesson 1: say you can do it. Then deliver—at all costs.

I can make up a story about how I brilliantly fixed the tester, but I really just got lucky. The problem was just a bad connection. I reattached it, and it worked. I was a hero.

But I still knew absolutely nothing about what I'd gotten into. Looking at the schematics of the Sumo products was like looking at a Olde English codex—massively confusing and completely incomprehensible. Why did they use so many parts? What did they all do?

Don't laugh. I was a green engineer. Very green. Sumo was only the second "real" job I'd had since graduating. And it was leagues away from the regimented, spotless, cutting-edge environment of Magnavox APS, where I'd mainly done software and PCB layout. It was messy, old-school, a union shop, and part of a larger company (Califone) that owned it, at least for part of the time I was there. I didn't believe companies like Sumo existed in the brave new world of nearly-1990. When I interviewed, Ed showed me their layout room, where they taped up all the PC board artwork, and their blueprint machine, where they ran copies of hand-drawn schematics. Those things

had been done with software and pen plotters at Magnavox.

"You don't actually use those things, do you?" I asked him, looking around for the CAD workstations. But there were none in sight.

Ed just laughed, a little nervously, and shooed me off to the rest of the interview. At the time, I thought his titter meant, "Of course not." But it turned out it meant, "How else would we do it?"

And yes, the first boards I did for Sumo were taped up. It was what had to be done. And I learned a lot from doing it. In 6 months, I understood what made a good amplifier (or so I thought) and was designing new Sumo products as Chief Engineer. In two years, I was Vice-President, Engineering. And by the time I left, I'd gotten us into new markets with new products, and made the company much, much more efficient in terms of production and parts commonality.

But it wasn't all rosy.

This second anecdote is should be a poster for *What Not To Do, Ever,* at any company.

My second job at Sumo after the MOSFET tester heroics was debugging their shiny new Andromeda II amplifier. It was selling briskly, be-

cause it was one of the few amps that could drive the insane load of the Infinity Kappa 9, a popular speaker at the time. The problem was, Andromeda IIs were also coming back, blown up, in scary numbers. I dove in, went back to the books, and quickly noticed a big problem with the amp: the P-channel MOSFETs were only rated for 75% of the current as the N-channel MOSFETs, which meant they were much weaker than Sumo's previous engineer might have expected. Combined with a slow breaker-based protection system, that could end with blown-up parts in a hurry. I found an alternate part and dropped them in the amp, then had to deal with some instability problems. Eventually, it worked and didn't oscillate.

But the revised amps, in my testing, still smoked. And by smoked, I mean, parts would literally catch on fire, and flames would come out of the vent holes. What's more, it did it after the protection system did its job and blew a breaker!

It was a mystery. What made it even more mysterious was the fact that the amp didn't always smoke. After several rounds of testing, an epiphany hit: think like a customer. What would they do if their amp stopped making music when the breaker blew? They'd probably turn up the

volume on the preamp to see what it would do.

And, you know what? Every time I turned the volume up after the breaker blew, the amp smoked.

Digging deeper, I found that the new power supply didn't completely shut down when the breakers opened. Which meant the front end, and the drivers, could completely drive themselves to death (fire) when the breakers blew. I rigged up a circuit to shut it down, and voila—the revised Andromeda II behaved as it should. No blown MOSFETs. No smoke. Just a relay click and nothing.

So, I did what any idealistic young engineer would do—I told everyone in the management meeting that we had to stop shipping Andromeda 2s immediately, revise the boards for the new outputs and the more advanced protection system, and then ship new, perfect, safe amplifiers.

So what did they say? Come on, you know what's coming, don't you?

They said, "Are you crazy? We have to make numbers this month. Ship them anyway."

And what did they do? They shipped them, of course. Lots and lots of them. I think about 70 % of them eventually came back, and

contributed to an extreme service load that never went away in my 5-year tenure.

Business Lesson 2: don't ship stuff that blows up. Ever.

Yeah, I know, it sounds like common sense. But it's amazing how common sense can go by the wayside in companies that live and die by receivables financing. When you hear "we have to make numbers for the end of the month," be scared. Very scared.

Sumo was, by and large, the company that taught me what not to do.

Exhibits of What Not To Do:

1. The example above—shipping stuff that you know will break. Come on, this isn't rocket science.
2. Don't ever tell anyone, "I don't care that you say it can't be done at that price. I'm VP of Marketing, I make four times what you do, and I've already sold a bunch to our dealers. That's the price, make it work."
3. Re the above: never sell anything you haven't made yet. Period. Ever. Nor tell someone about stuff that's coming up, or show products

that aren't yet products. If they aren't on shelves ready to ship, they don't exist. (It took us a while to remember this one when we started Schiit—oops.)

4. Don't lose customer returns. Or use them to fix other customer returns. Yes, I know, more common-sense stuff.
5. Don't try to go too broad. In Sumo's case, this meant getting into speakers. The Sumo Aria was an amazing planar speaker. Also, it was an amazing pain, because most of them broke, and they were made on a contracting basis for us by an outside company.
6. Re the above: speakers seem to be easy. Anyone can put a driver in a box and consider themself a speaker designer. That's why there are so many speakers out there. What you really want, when you're creating your niche company, is something with high barriers to entry. Electronics are a lot tougher than speakers, with lots of components and safety standards and FCC. There will be less competition, which is a good thing.
7. Don't be cheap, especially to the point of having the checks be late. Especially paychecks.

But Sumo had one thing going for it. Its heart was in the right place. We were trying to make inexpensive components that could compete with

the “best of the best.” We didn’t do insanely overwrought chassis for megabuck amps. And this taught me to be efficient, and work with what we had. It’s one of the reasons we’re good at production engineering now.

Business lesson 3: don’t dwell on the negatives—learn from them.

And Sumo was where I was converted from a hardcore objectivist to subjective-objectivism.

“Ah hell,” some of you are saying now. “I don’t want to hear this hoo-ha about how all amps sound different. Properly engineered amps run within their limits all sound the same, anything else is placebo/misdirection/insanity!”

Amen, brother. Or that’s what I would have said in 1989. By 1990, I wasn’t so sure.

When I started at Sumo, I already had better amps than they made. At least in my mind. Two Carver M-1.5T monoblocks, 350 W per channel, sporting Bob Carver’s latest and greatest Magnetic Field Power Supply, or something like that. They were lightweight and ran cool. Light-years ahead of Sumo’s giant, hot, heavy Andromeda II, which was only rated at 200 W per channel.

And I knew those amps well. I had a small company manufacturing speakers, called Odeon. (See the note on speakers above.) We used the Carvers for testing, development, and demos. We had them up to wall-shaking levels. They were pretty damn good.

So, when Sumo's president said, "Hey, take home an Andromeda 2, and let me know what you think," I wasn't too excited. I put it off. And I put it off again. I didn't want to tell him the Carvers were better, or lie and say the Sumo was better.

Eventually, I took one home. I told the other guys in the speaker company, "Hey, look at this dinosaur, there's no way it will beat the Carvers." It was all a joke.

Until we turned it on.

Holy crap. Not only did Sumo's "underpowered" amp wipe the ground with the Carvers in terms of higher output, it also sounded better. Way better. The rest of the team came out of the other room and just stood there, dumbfounded. I think Eddie spoke first. "So, when do we get one of these to keep?" he asked.

And that's what started me down the road to subjective-objectivism. Not pure subjectivism, of

course—measurements are still very important. But it led me to dig into the reasons why the Sumo amp sounded different than the Carvers. It started with simple things, like current capability and rated power.

It continued into gain stage structure and out of band performance, and led to me dramatically changing Sumo's amplifiers—to the point of doing a zero-feedback, single-gain-stage preamp shortly before I left (Artemis—very, very rare) and several no-overall-feedback amplifiers (The Ten, The Five, Andromeda III) that measured as well as the full-loop-feedback amps that preceded them, but sounded much better.

And yeah, yeah, I know: *You're crazy. All amps sound the same!*

Business lesson 4: don't discount personal experience.

There are a thousand other stories about Sumo, but let's cut to the part where I met Mike Moffat of Theta Digital.

At the time, Sumo just happened to be in the same business park as Theta. I knew who the company was, of course. Everyone in audio knew

Theta. They were really tearing it up in the DAC market. And Sumo, like everyone else at the time, was working on our own DAC. In many ways it presaged what we're doing today—it was modular and upgradable, and the DAC card itself could be added to a Sumo preamp.

But I didn't want to meet Mike. No. It was too intimidating. Plus, they made expensive stuff, and we made cheap stuff. Plus, he'd probably be a golf-playing blowhard who was too full of himself.

But Sumo and Theta shared a components sales rep, and she kept insisting I meet with them for drinks. After a while, I relented. And I forgot all my rationalizations for avoiding the meeting.

Mike was, and is, a character. Instead of being uptight and high-and-mighty, he was very casual and approachable. Hell, after he got to know us at Sumo, he'd sneak out of Theta to come over to use our bathrooms to change into a suit when he was going to the opera. He didn't want his employees to see him in a suit, but the A opera was the only place he could hear unamplified music.

At Theta, I found the company I wanted Sumo to be.

Mike's Theta ran on incentives. Employees were paid bonuses based on the number of units shipped, and on their individual performance, and on stepping up to do tedious things, like upgrades. Some of his techs made several times their salary in bonuses. The office was casual to the point of not even having part numbers for their parts—instead of a 05-1225, a 1 kΩ 1/2 W 5 % resistor was called a 1 kΩ 1/2 W 5 % resistor.

Yes, I know, not very exciting. But consider the results: in 1/2 the space of Sumo, Theta was selling 10× the dollar volume of products. Their net profit was easily 8× that of Sumo. I told the president of Sumo this. His response: "That's stupid, paying people bonuses. Then you have variable salary cost, you can't predict it." I wanted to say, "Well, it seems a whole lot more stupid to run an inefficient business like this," but for once, I said nothing.

It was becoming very clear that nothing I could do would change the way the company ran.

Business lesson 5: be open to meeting new people, and transformative ideas.

And that's why I started moonlighting for Theta. Mike Moffat was really intrigued by the idea of doing an inexpensive DAC. After more dinners and more drinks, we finally hatched the idea of Cobalt. The Cobalt 307 was my design, with input from Theta—a true hybrid of Sumo's ideas and Theta's ideas.

Cobalt blew up the DAC market, selling 1000 per month for some time—which happened to be about 2.5× the total market size according to one industry pundit. The combination of solid name and inexpensive price really set the high-end world on fire.

And—I think it's important to note here—"inexpensive" would seem pretty pricey today. The Cobalt 307 was $599. In 1993. That's about $970 today. If we'd been able to sell Cobalt direct to the customer, like we do Schiit, it would have been $349.

Yes, that's how much the dealer takes. More on that later.

But in 1993, selling direct wasn't feasible. We would have had to take out full-page ads in all

the magazines to the tune of $20 000 or so a month, and we would have had to have multiple employees in a full-time call center to take orders. There was no Amazon Marketplace. No Shopify. No pay-per-click advertising. Hell, there was no viable internet. It was a different world.

At Theta, I also designed the discrete, current-feedback output stage of the top-end Theta Gen V, mainly on a bet. Mike and Dave—Mike's lead engineer at the time—were convinced that op-amps were the way to go, but I'd learned enough about discrete design to know they were wrong.

"I can design a stage that will work better than any op-amp," I told them.

"Even on measurements?" Dave asked.

"Even on measurements."

They took the bet, and I came up with a design that was an exercise in insanity. 260 parts on a 4 inch×6 inch Teflon circuit board, with two PCM63 DACs in balanced configuration.

But it beat the op-amp stage, both in measurements and in the listening room. And that's how the Theta Gen 5 was the first discrete output DAC that Theta made.

Business lesson 6: take a chance, do crazy things … a lot of times, it's worth it.

Even as Theta was kicking ass, audio was getting, well, weird. Theta stuff cost a lot to make, so it was priced very high. And we had the dealer vig to pay, of course.

But Theta's products weren't priced high because they were lookers. It was all about the technology inside.

Theta's competitors took a different tact: make it pretty and even more expensive. Theta's balanced Gen V was $ 5500. Mark Levinson "outdid" Theta with a $ 16 000 DAC. Krell upped the ante with $ 32 000 amplifiers. The magazines ate it up. The race towards "gold-plated Bentley" audiophilia was on.

Mike and I didn't get it. He didn't want to put a $ 250 board in a $ 2500 chassis. He wanted to make game-changing stuff. But it seemed the magazines were only interested in the megadollar price tags. Eventually, that led Mike to start Angstrom and get into the field of surround sound.

Me? I went evil. I went into marketing …

Chapter 2

15 Years On the Marketing Front Lines

"Marketing?" I know some of you are asking. "What does that have to do with engineering?"

Well, not much. But, like I said, I'm confused. In addition to an engineering major, I also took enough English classes to be an English minor—and my GPA in English was higher than for my summa cum laude engineering degree. So that's a bit of foreshadowing right there.

Also, back in the Sumo days, when I did the first brochure for our speaker company, Odeon—which is an odyssey in itself, from getting kicked out of Vasquez Rocks for the photo shoot (we didn't know about things like "permits" and "insurance" back then) to the Cretaceous-era desktop publishing software—the first dealer said, "Well if you can't make it in speakers, you definitely have a future in advertising."

At the time, I brushed off the comment. I wanted to make audio stuffs, not brochures!

But I kept coming back to marketing. At Sumo, when the VP exited, I ended up doing the copy, layout, photo art direction, etc for their brochures and print ads. Theta counted on me to do the brochure for the Cobalt 307 (which I'll post if I can find one, because the copy there really is the first expression of the balls-out attitude that is a hallmark of Schiit—we even made fun of the gold-plated audiophilia that was starting to take over at the time, proudly saying the Cobalt 307 was cheap because of mass production, rather than being "handcrafted by happy elves in Wichita.")

Fun Fact: the Cobalt 307 had blue LEDs for one reason only: to thumb our nose at Krell. Until the Cobalt 307, blue LEDs were astoundingly expensive (about $10) but we were able to get some of the first inexpensive ones around (less than $1.)

And now I should add a disclaimer: we're still not talking about Schiit for a while. This is still the run-up to the company. If you want to read about Schiit and Only Schiit, you'll have to wait to the next chapter.

This is the tale of Centric (centric.com), a company I am still involved with. Centric is a company that does marketing for tech companies, food companies, and many other kinds of organizations, including some high-end audio firms. Centric just passed its 20th anniversary a few weeks ago.

But First, Let's Talk About the Marketing Industry

If you're expecting this to read like the boozy exploits you see on Mad Men, prepare to be disappointed. The top-tier, multi-billion-dollar ad agencies are working with clients who are the corporate equivalent of the Rockefellers, Rothschilds, and the Sultan of Brunei—companies with so much money they could buy your town as a joke.

We didn't work in that rarefied realm. Consider this: the cost of an average 30 second Superbowl ad is $10 million. This counts $4 million for the airtime, and $6 million for production, logistics, pre- and post-distribution, social media, online media, etc, etc.

Consider just two facts:

- $10 million is 6× to 10× higher than a typical annual marketing budget for an $50 million to $100 million tech company—and this includes marketing salaries
- $10 million is larger than the annual revenue of all but a handful (literal handful) of audio companies

So yes. Rothschilds. Rockefellers. Sultans. That's what we're talking, when we're talking Superbowl ads. And more. Did you know that Toyota spends $100 million to $150 million in advertising to launch a new car—and this is not a full year of advertising, this is during the launch months? Did you know that a single brand at P&G, such as Tide, can have a $50 million to $100 million annual advertising budget, every year, for decades? That's why when your agency starts trotting out "branding examples" from the big names, and suggest you emulate them, you run. Fast. They have nothing to do with the reality of a by-the-bootstraps company.

But I'm getting ahead of myself. Let's disambiguate this whole "marketing biz" thing a bit. I'm sure some of you are sitting there, wondering:

a) What the heck is the difference between marketing agencies, advertising agencies, PR, social media, etc?

b) Why the heck companies think they need marketing—shouldn't the best product win?

Okay, so let's break down the difference between agencies, circa 2013:

Advertising Agency. Primarily develops and places advertising. Many agencies identify themselves as a "creative agency" (thems the dudes that drink a lot) or "media agency" (thems the guys who buy the ads and make a profit on them), or both. But today, it's more complicated. Are they primarily broadcast? Primarily online? Both? Do they do print? Outdoor? In-store? Native? Social? Ad agencies can do all of that.

Interactive Agency. An ad agency, but subtract the print, broadcast, and outdoor, and add web and mobile development.

Social Agency. An ad agency, but one that annoys your friends where they hang out online, like a crazed cybernetic door-to-door salesman.

Design Agency. An ad agency, but subtract the focus on persuasion and turn up the emphasis on great art and visual communication.

PR Agency. In the past, this was your conduit to the press. They knew the editors and could help you get placements. Today, that's evolving

rapidly as conventional media (like magazines, newspapers, etc) crumble and online media/ blogging/social rises.

Marketing Agency. Like all the agencies above, with different strengths in different areas. Usually focused on one or more niches. May drink less. May drink more. Centric is a marketing agency.

"Well, hell," you're saying. "Do I need all those agencies to succeed?"

No. Not at all. You may not need a single one of them. I'll get to that. But for the moment, let's cut back to the birth of Centric, and why marketing?

The Centric Rationale

My rationale for starting a marketing company was something like this:

1. Hey, I did this for Sumo and Theta and my own company, so I have some experience.
2. It's not a manufacturing company, where you have inventory, overhead, labor, distribution, etc—it's a lot easier to get started.
3. The products usually don't ever catch anything on fire.

In retrospect, not the best reasoning. But hey, I was 28. Leaving a VP of Engineering job to start a business in a field I knew nothing about seemed perfectly sensible at the time. So, as audio went into exponential price expansion, I jumped ship and started Centric in January of 1994.

Now, if you're a A resident, you might be thinking, "Hmm, wait, isn't that when the 1994 Northridge Earthquake hit?"

Right. I started the company exactly one week before the earthquake, and I was living about 10 miles away from Northridge when it hit. I was renting a house on a hill above the San Fernando Valley, and I clearly remember waking up in the early morning, sitting up in bed to look out the window at the valley, and watching them shut down the whole starry mess of it, grid by grid, as shelves toppled inside the house and the fires began outside.

And I remember thinking, *Holy crap, if this just hit downtown LA, kiss this business goodbye. Because there ain't no downtown no more.*

In earthquakes, what matters is how far away you are from the epicenter. If the epicenter was Sylmar, where I was living, OK, that's bad, but not the end of the world. If the epicenter was

downtown, and it was strong enough to knock over nearly everything in a house 40 miles away, that was, like, The Big One. End story. Full stop.

And again, remember—no internet. And I didn't have TV. Didn't believe in it. I was into audio, remember? So I had no idea where the epicenter was. I sprinted over the crap on the floor and went out to the car to listen to the radio. And the first thing they said: *LA. It was in LA.*

Yep. Done. Pack it up.

And that's what I did. I got in the car and went up to Valencia to see if some friends were all right. It was like driving in a zombie apocalyptic horror movie, with trailer parks burning on one side, toppled phone poles and smoldering transformers, and nobody, nobody on the road.

Later, we found out the epicenter was Northridge, and that A itself was still intact. While tragic for Northridge, it turned out not to be the end of the world. It just made driving to the service bureau to get film output monumentally sucky, due to the downed freeways. Remember. No internet. No FTP. Big file mean sneakernet, man.

And two days after the quake, I was driving to see a new prospect for Centric, XLO Audio, in

Rancho Cucamonga.

Proving that it's never the end of the world.

And, over the next few years, we added companies like Threshold, Infinity, 3D Systems, Pioneer, Veeco, Compaq Capital, HP, and a whole bunch of other tech, industrial, and consumer electronics companies to the list.

And for a while, everything was glorious. We rode the wave of the first internet boom, doing some of the earliest web development work, earliest e-commerce, earliest web marketing … all built on the basis of personal incentives, like I'd seen at Theta.

We were even smart enough to avoid the worst of the Web 1.0 downturn, though it did hurt when we went from 7 optical networking startup clients to 0 in a single year.

But, in the end, life was good. Marketing was fun. I got to see all sorts of crazy new cutting-edge technology, and the clients loved me because I could talk to the engineers and scientists and not be dismissed as "the agency freak." I even had a hell of a science fiction moment at one client, when they were showing off their new Pico-Force measuring system based on atomic force

microscopy, where they could actually unfold individual protein strands and manipulate them at the molecular level.

"That's like the nanomanipulators in Neal Stephenson's Diamond Age," I said.

At that point, two of the scientists turned around to look at me, eyes open and jaws slack in shock.

"You read that too?" one of them asked.

"Where did you think we got the idea?" the other said.

And we also got to do a lot of cool, cutting-edge stuff in marketing. In addition to some of the earliest web development and online marketing, we were able to do some of the earliest social work for Warner Brothers, and built HP's presence in the virtual world of Second Life, as well as "the largest virtual experience ever" in the words of MIT Tech Review, on the David Rumsey Maps project. We've constantly experimented with what's new in marketing.

And … paradoxically, that's why we're more conservative today. We haven't seen the results from social marketing, unless it's for an entertainment company. The big bang in virtual never happened.

Mobile is very, very important, but who knows if that will extend to augmented reality?

Condensed Marketing Stuff Follows

"So what does this mean to someone who wants to start their own company?" you ask. "Or to someone who's just on the outside, thinking about it?"

Well, to summarize what we learned in the past two decades, and give you the "key takeaways" (sorry, lapsed into corp-speak there):

1. **Most companies are too terrified to be effective at marketing.**
   Show them something amazing, something catchy, something incredibly effective, and the first reaction (at most clients) is, "Wow, this is wonderful, let's do it!" Then, two days later, an email appears. It usually goes like this: "Our CEO/lawyers/accountant/design intern/marketing director's daughter/fish/dog looked at it and we're concerned that it may be too 'out there …" Yep. Done. *Key takeaway: Don't be scared to stand out.*
2. **This terror can affect everything they do, so they may not be effective at *anything*.**
   The second-guessing of great ideas doesn't

stop at marketing. It usually extends all across the organization, to product development and customer service. That's why you get so many me-to products and crap customer service. "But our competition is doing it," whines the product manager. "But the competition doesn't provide any better support," says the director of customer service.

Key takeaway: A race to the bottom helps nobody. Don't benchmark yourself into mediocrity.

3. **Most companies have no idea what to do in marketing.**

"Let's do social, I heard it's cheap and easy," or, "I'm tired of the website, let's change it," or, "Well, all of our competitors are going to that show, so we need to be there," or, "I know the magazines are getting less and less effective every year, but I think we need to be in the books," is the rule of the day.

Key takeaway: marketing should be a portfolio strategy, with the most money going to the most effective and measurable tactics, with detailed analytics on what is working and what isn't, with a small percentage reserved for experimentation on "new" or "interesting" ideas. If marketing doesn't make money, it shouldn't be done. Period.

Why is it important:

- **Fear is the mind-killer.** Fear is the little death … no, wait, that's from Dune. But it's true. Second-guessing your first reaction to something you love usually doesn't result in great things. Trust your gut. And remember, even if you've seen it 50 times, most of your prospects are only seeing it for the first time. If the marketing creative work stops you, it works. Do it.
- **Kill the fear before it spreads.** If you start a company, and instantly start worrying about if your product has every little feature that your competitors do, you might as well name it the RX-4001i-Re A and hope that someone mistakes you for Epson, or some other company that's been around since the earth cooled and can get away with crap like that.
  Apple's products never have the best specs, most features, biggest shiniest displays, etc—and yet, even now, they're the highest-value brand and company in the world.
- **Marketing is important, but don't do it blindly.**
  Today, you can measure any aspect of anything you do online, down to which ad drove which specific sales of which product. Get the reports. Sit down with the agency and torture them until they bring out the one dude

who really understands them, and have him explain it to you.
Do more of the stuff that works, and less of the stuff that doesn't.

- **Corollary 1: don't believe nearly everything an agency tells you.**
  They're going to trot out these ancient case studies about how branding is done, how P&G has built the Tide brand, or how Toyota built its brand, etc, and imply that those are the right models for you. 100 % total bullschiit. These brands had hundreds of millions to billions of dollars to spend on a single product or model, over decades. They're Rockefellers and sultans, completely disconnected from reality. You aren't. *Create your own can't-be-ignored product and personality.*
- **Corollary 2: see above, ×10 000 if it's "something new."**
  Agencies love "something new." It's usually confusing and not measurable, and they are more likely to win an award for it. So they'll trot out a case study about how someone got like 12 billion views on YouTube or 1 million Facebook likes or 3 million Twitter followers, but they'll leave out the convenient fact that (a) the company also had a $ 150 million ad campaign running at the time, or (b) they're

a celebrity, or (c) they were just damn lucky. *Forget chasing new/easy/cheap. Marketing is none of the above.*

- **Corollary 3: Mass advertising is unmeasurable, and almost never works for smaller budgets.** This is why agencies love it. Well, at least the first part. Smaller budgets are defined as $10 million or less. They'll try to dazzle you with reach and frequency and such, but bottom line, you're not going to track a magazine ad or TV spot back to a specific purchase.
  Stay online. Measure. Refine. Do better.
- **Corollary 4: Mass social almost never works, unless you're an entertainment company**. Entertainment properties have fans. They're natural for social. Almost every other company isn't. People are there to talk to their friends, not BUY NOW. You're entering their living room, their pub, and their coffee house. They don't like it. Social produces 10× the results of conventional advertising for entertainment, and 1/10 the result of conventional advertising for everyone else, in our experience. *Forget big social—it's a distraction that can eat your company.*
- **Corollary 5: on the other hand, micro-social almost always works, unless you're**

a dick. Finding the small, specific, passionate communities that are interested in your products, whether they are barbecues, espresso machines, audio gear, or high-end bicycle accessories, is almost always worth it. Going out, joining these communities, answering questions that come up, and not selling at all is a wonderful way to get the word out. But don't think you're King Salesman of the Universe out to convert the masses, or start attacking other brands, moderators or forum members. One problem: most agencies are too lazy to do this hard work. And it is hard work.

Pay lots of attention to micro-social, and be prepared to post, respond, meet new friends, piss some people off, delight some others, and become part of your specific niche.

So, Do I Need Marketing?

Bottom line, this is important because marketing is one of the most important things you'll do. It will be critically important to the success of your company. You may not need to have a single agency to do it, but you will need to get the word out—in a memorable, compelling way.

"But that's not true!" bleats one member of the audience. "I hear that Gen Yers are so cynical and jaded to marketing that it doesn't work anymore. I hear we're moving into a post-advertising future."

LOLOLOROFLCOPTER. No. Sorry. In 500 years, when we've all enhanced ourselves to be perfect physical examples of the human species, immortal and all-knowing, or uploaded ourselves to the grid, or devolved into a dystopic hunter-gatherer existence that can only communicate in leet-speak, there will be marketing. There will be ads. You can bet on it. And the successful companies and organizations will know how to use it effectively.

It's true we're moving into a different ad regime, though. Gen Y doesn't like screamy, shouty, "This is the biggest bestest most amazing product in the universe, it will transform your world, and happy bunnies will follow you wherever you go." Because words like "best" and "amazing," and "super," have been overused.

Gen Y, in general, wants to know more about the nuts and bolts. Spare the superlatives. Give them the facts.

But, you know what? Whether it's an AMA on Reddit, a post on smokingmeatforum.com, or a

banner ad on Gizmodo … it's still advertising—and still marketing. If you, as a company principal, can do some of the marketing basics, it might be enough to save you from having to hire an agency.

If you can't, shop very carefully, ask a lot of questions, measure everything they do, remember that you're not a Rockefeller, and remember this short advice:

1. The most important thing is your website and e-commerce system.
2. The second most important thing is how they work on mobile devices.
3. The third most important thing is press, and by press we mean mentions and articles both online and off, in and out of the niche press.
4. Online ads are probably next, but make sure you can track all the way to a sale. You're shooting for a cost per sale that's less than the profit on the sale. Don't let them tell you anything else.
5. Everything else comes after: shows, brochures, t-shirts, lifesize figurines of your founder, skywriting, heat-activated urinal billboards (which are actually a thing), sponsoring your own events, laser-blasting your logo on the surface of the moon, etc …

Okay, one more anecdote, and then we'll move on to the founding of Schiit. Which wasn't called that at first. Actually, it had no name. But I'm getting ahead of myself.

During my time at Centric, I built it up to a big, successful, multi-million dollar business. We did very well. Not bragging, just facts. And, one day, I sat back and wondered, What's next?

What's next turned out to be Lisa, AKA Rina, my wife of 13 years. Would I have started Schiit 9 years earlier if I hadn't met her? Probably not. But she challenged me enough to keep pushing, keep expanding what I could do, that she certainly got me into the right head-space to start something new.

It started when Lisa and her writing buddy, Jen, announced they were going to write a book and get it published. This was 2002.

Now, I'd done some writing in the past, and I had even sold a couple of things. I knew how hard it was. I had written dozens of stories, and never really gotten anywhere with them. So I muttered something vague and wished them good luck, and figured that would be the end of that.

Eight months later, they had a book contract. I couldn't be outdone, so I pulled out the computer and started writing again. The end result is my own three novels and about 30 published stories, as well as a 1st place win in the Writers of the Future contest, being a finalist for a Theodore Sturgeon award, and twice a finalist for a Sidewise Award.

The point is: I had this capability all along. But I didn't do anything about it until someone (figuratively) kicked me in the butt.

Who's going to kick you in the can? When will you do your writing, or company-building, or adventuring, or whatever you want to do?

Chapter 3

From Death, Rebirth: Armageddon 2009

All great things come to an end. And in 2009, I thought Centric might come undone.

We'd weathered the web development downturn, and we'd ridden through two business hiccups that were either our fault or just the changing winds of the marketing times, but I'd never seen anything like the complete and utter disaster that was two-double-ought-nine.

Clients slashed budgets. New management jettisoned us. Proposals sat forever, or were teleported onto the world filled with single unmatched socks and pen caps. And, to top it off, one of our biggest projects ever, a near-$500k development of a kid's virtual world, went slowly and painfully—then finally turned into a major debacle when the initial traffic brought the site to its knees.

Sue, my business partner at Centric, summed it up at the end-of-the-year Centric party. "The only

good thing we can say about this year is that it's over. Slam the door, nail it shut, and never look back."

But out of that disaster, we got Schiit.

Why?

First, let's start with the ass-kicking factor. Like Lisa, that year kicked us in the ass. And it made us think. For Centric, it led to an entirely new office (moving from the Academy of Television Arts and Sciences building in North Hollywood, with an office overlooking the Emmy statue, to an old wine shop in old town Newhall), an entirely new way of working with our staff, with more flexibility, more freedom, less overwork … and by mid-2010, Centric was back on track, and doing better work than ever.

Second, the audio factor. I'd toyed with the idea of starting another audio company from time to time, but I'd always been distracted by the "real" work of marketing, and by memories of how hard it was to work through distribution. 2009 gave me more time to think about it.

And finally, writing—and a fortuitous gift. In 2009, I was deep in writing mode, working on two of my own novels, planning more, submitting

stories, attending writing groups. And writing takes a lot of time. And, for me, writing also means time without distractions. I'm not one of those coffee-shop wordcrafters who can work with screaming children running around their chair and baristas barking names at 110 dB. Hell, I can't write if there's a TV on in the other room.

The solution? Use headphones. At first, just the Apple earbuds that came with my iPhone. Yes. Don't barf. We all have to start somewhere, right? They let me drown out the distractions and write. It was all good.

Except, in the back of my mind, a little voice kept whispering: *this could be better.* A friend gave me a pair of V-Moda earbuds. And they were better than the Apple earbuds. Which made the voice in the back of my mind louder. But it was still OK. I was focused on my writing. I could think about audio later.

Then, the fortuitous gift. A friend sent me a Chinese tube headphone amp, simply because he traveled to China a lot, and knew I used to be into audio, and … and I sat there looking at this intricate thing, thinking, How the hell can this be only $ 300? No wonder manufacturing is dead in the USA.

Of course, it didn't work so great with the earbuds, being too noisy for them. My wife bought me a pair of AKG 701s, mainly because they seemed to be highly regarded and relatively inexpensive. They worked pretty well with the tube amp. Good enough that I began to understand what some of that "tube magic" was.

A sidenote: I'd never really been into tubes until Schiit. Sumo was all solid state. So was Theta. I knew Mike did something with tubes, way back when, but that was it.

I used that combo for a time, but I kept looking at it, and wondering, Can this be even better?

On a whim, I tried the headphones with an old Sumo prototype that never made it to production—the Sumo Antares integrated amp. Oh, Sumo never made an integrated, you say? You're almost right. We only made one of them.

But it was a speaker amp. Would it light up the headphones? What would it sound like? Was I totally insane?

Then I hooked it up and listened. And sat there listening for hours. This was it. This was what those headphones needed. So much more detail, control, and—and, well, kinda etched top end,

and, well, it was kinda noisy, but you can't have it all, can you? And it was really, really good, this ancient, 60 W PC speaker amp.

That really set my mind going. Headphones were efficient. They didn't need 60 W. Which meant the power supplies could be regulated to kill the noise. And you could easily do Class A. And you could play with super-simple topologies that simply wouldn't work in the speaker realm. It would allow me to do things that simply weren't practical before—and that could be a lot of fun!

Quick Notes: Speaker Amps and Headphone Amps

Headphone amps and speaker amps don't really have to be different in execution or topology. That's why you see people using speaker amps with some headphones. But the devil is in the details:

Speaker amps are usually all about maximizing efficiency in order to deliver high watts at low distortion and moderate noise into a known load (4 Ω to 8 Ω). Because of this, they tend to converge around some common, well-known

topologies that meet this need—the most common of which is the Lin topology—differential input with some voltage gain, VAS with more voltage gain, and current gain stages afterwards, with overall feedback.

Headphone amps don't need to be super-efficient or super-powerful, but they have to be very, very quiet—20× to 200× lower noise than a typical speaker amp, if you expect to run anything but planars. And they also have to be ready for loads from 600 Ω to 16 Ω—a much broader range than speaker amps. The result is that you now have the freedom to design around many different topologies, including single-stage and overall-feedback-free designs, as well as the standard Lin variants as used in speaker amps.

And yeah, I know there are non-Lin speaker amps, there are current-feedback speaker amps, and circlotron-style speaker amps, and transformer-coupled speaker amps, and Class D speaker amps, but those are outliers. The bottom line is that your common speaker amp is most likely Lin topology, two voltage gain stages (counting the front end) and two or three current gain stages afterwards. A headphone amp can be anything from a single op-amp to a Class A follower to tube OTL to Lin.

Sidenote 2: The original Asgard was supposed to be an ultra-high-power amplifier, delivering a full watt of power. Yes, I know how silly this seems today.

Sidenote 3: The only Lin amp we do is Magni—so, ironically, Magni is closest to a speaker amp in the Schiit family.

On to the DAC

Up until this time, I'd been listening from a computer source. From the analog outputs.

Yes, I know, I've committed every headphone sin known to mankind, I should be purged from this planet, I'm a cloth-eared idiot. But all that thinking about amps got me wondering about DACs. I had an old Cobalt 307, and I found that my MacBook had optical outputs, and that Monoprice made funky cables that went from 1/8 inch Toslink to regular Toslink.

Soon, I had the Cobalt running into the Antares, and again—what a revelation! This antique DAC and geriatric amp were doing some amazing things. I didn't want to write. I wanted to sit and listen to music.

But—they had to be doing a lot more interesting stuff with DACs and such these days, right (don't laugh, I'd been out of the game a long, long time.) I started to spend a lot of time online, researching what was out there. I discovered Head-Fi. I read about ten thousand reviews.

And I sat there, stunned. All the energy that got sucked out of two-channel audio when it started going down the road to ever-bigger price tags was back, and bigger than ever. I showed it to Mike.

"It's like high-end about 1980," Mike said. "Just getting started. Before we went insane."

Of course, Mike didn't know that I was going to start a company and drag him into it. I still didn't know for certain myself.

But thoughts kept piling on each other: What if we could do something here? What would we do? Where would we make it? How would we sell it? Dealers again? How would that ever work in a world where Chinese manufacturers were selling direct on eBay? And direct? The only company I knew selling direct was Emotiva, and I had no idea how they were doing.

Fun fact: Centric actually subleased office space from Dan Laufman in 1995 to 6, when he was running a PCB assembly and contract manufacturing business. Yes, the Dan Laufman that would go on to found Emotiva, after getting tired of doing OEM work (that means, in English, *makin stuff fo other peeps.*) By 2009, though, we'd fallen out of touch.

But, hmm, direct. Direct changed everything. Because it cut out the reps, the distributors/warehousers, and the dealers.

A Quick Primer on High End Economics

Let's pause for a quick look at how pricing works in the high-end world. Cue everyone in traditional high-end audio hating me now. If I die of mysterious circumstances, you know why.

Here's how it works with a traditional distribution chain:

- Reps take 8 % to 10 % (thems the guys who go out and sell your stuff to dealers)
- Warehousing can take 5 % (at the dealer or third-party, if you need it)

- Dealers take 40 % to 50 %[1] (thems the guys who take the order)

That means that 48 % to 65 % of the cost of a product can be in its distribution. So, the chain that sells, shelves, and stores the product take 1/2 to 2/3.

That means the manufacturer—that is, the company that engineers, designs, certifies, tests, packages, ships, markets supports, warrants, and repairs the product gets 1/3 to 1/2 of the retail cost.

Go back and read that again. *The guys who put it on a shelf get as much, or more than, the company that creates and supports the product.*

Yes, I know. Insane.

This is why, when I was last into audio, manufacturers would set MSRP at 4× to 6× of their fully burdened production cost. Your $ 499 amp? Under the old rules, they paid $ 80 to $ 120 to make it, including labor and overhead. But the

1 This is highly variable depending on the product, and, in some cases, is changing for the better these days—and it's different for mass consumer products, which operate on much lower margins. Best Buy doesn't make 40 % on computers.

manufacturer might only see $200 of that $500, with the rest going to distribution.

Now, don't get me wrong. Dealers provide a service to customers by letting them compare a whole lot of different products. This is definitely worth something. And we are losing that as they go away.

But is it worth what they're charging in today's world? Especially when you can offer in-home trials and easy returns, and when (in headphone audio), local meets allow people to compare all sorts of gear? And when you can set up shop on Amazon and have them be your warehouse?

In the past, it really was a different world. Audio companies were completely dependent on getting the connected, aggressive reps that would get them into the right dealers. If the dealer required local warehousing or co-op money for advertising or spiffs for the salespeople (AKA, the mob boss visiting you for his protection money), you did it. Because there wasn't any other choice. If you didn't do it, the dealers would sell the competing products that did.

But in 2009, we didn't have those constraints. And, looking around, I saw the roster of dealers had already shrunk considerably in the last

15 years. It seemed the pendulum was already swinging away from old-style distribution.

And we knew how easy it was do set up an e-commerce site.

And we were a marketing company, after all.

At that moment, I stopped wondering. And started thinking: *Yes. Let's do something with this.*

Chapter 4
“You Always Say You Have Schiit to Do, Why Don’t You Just Call It That?”

It’s a long road from thinking to doing, though, especially when you’ve been out of the game for so long. I was rusty, incredibly rusty. I’d forgotten a lot of what I knew about engineering, simply because I hadn’t used it in so long. And a lot of stuff had changed.

Change. As a single example, let’s consider a conversation I had at Semicon, the semiconductor manufacturer’s trade show, about 2001 or so. I was talking to an engineer about a new product, and mentioned that I’d been an engineer and had worked with VLSI gate arrays back at Magnavox—cutting-edge prototypes that cost $100 000 each.

“How many gates?” the other engineer asked me.

“About 100k,” I told him.

He laughed. “Not the price, the number of gates.”

“Right, a hundred thousand.”

The other engineer laughed even harder and waved a hand, as if dismissing a servant. “100k? We put stuff like that in toasters today.”

Yep. 12 years took $100k prototypes to $1 commodities. And now, 12 years later, we’re all carrying smartphones. In another 12 years, we’ll be wondering how we got along without Google Ambient and pervasive intelligent packaging.

But back to Schiit. There were a billion questions before we got started:

- Was everything SMD (surface mount electronics, very hard to work with)?
- Where had component costs gone?
- What kind of components were people using for audio these days—was discrete design even feasible any more?
- Where should you get them? In the old days, the reps would come visit us and give us data books and quotes (yes, as in ink printed on paper), but how did it work now?
- Same for PC boards—what were the costs now?
- What PC board layout software did people use these days?
- How about schematic capture?
- How about CAD for drawing up chassis—in the old days, we just used Illustrator?

- What did it cost to make a chassis?
- Where would you get it done?
- Who'd do the transformers?
- What were transformer costs like?
- How would we assemble the products? Sumo had done in-house assembly of PC boards and products, Theta had gone to board houses (that is, people who put parts in PC boards) and only done final assembly, and lots of people were going to complete contract assembly (typically in China) and only marketing and supporting the products.
- What safety approvals did we need?
- What did a modern R&D setup look like? We were familiar with Audio Precision analyzers, but we'd never used one, and didn't know what they made and what they cost. Mike still had his ancient Stanford spectrum analyzer, which is what we used to use in "the old days", but was that good enough? I had an old HP 339 analyzer—same questions applied there.
- Where would we do it all?

But you can obsess over all those questions and turn them into a giant wall that keeps you from ever doing anything (see Chapter 2 on big-company paralysis by analysis), or you can do some digging. I dug. And quickly found out

that a lot of the information I needed was online. Component availability and cost, PCB cost, transformer cost—that was all there. Most surprisingly, all of them were much cheaper than we'd paid twenty years ago. Thank globalization, or the downturn, or whatever, but even parts that were made in the USA were far, far less expensive than we expected. And, despite the bleating of the apocalyptic crowd about how "there ain't no more good audio parts out there", there were actually plenty—and even more new options when you got on the surface-mount side.

But then there was software. Mike used (and still uses) Altium for doing schematics and PCBs. Altium is 100 % old-school big-ticket software. Hint: if they don't list the price, you don't want to ask. Altium starts at about $ 7200.

Not a good price, especially since I was still in the "foolin around" phase. The company had no name. Centric was deep in the tank. I didn't want to pay for anything I didn't really need. In fact, as I thought about it, I decided on a goal: Can we start this for $ 10k, including the costs of the first run of products?

$ 10k we could gamble with. If it went nowhere, well, hell, as Mike said, Christmas presents. We'd

lost more money at Centric on bankrupt companies and deadbeat buyers. $10k was doable. It wouldn't thrill me to flush it down the toilet, but it wouldn't kill me, either.

But $7200 suddenly meant that I would be looking at something more like $20k. I wasn't ready to commit to that. So I looked a bit more, and discovered that software, like everything else, had changed. In addition to Altium, there were other, lower-cost alternatives like Eagle. And there was a funky little open-source program called KiCad, built by some crazy student as a senior project. It was free.

Yes, free.

Interlude: Let's pause for a moment to salute open-source software. Who would have ever expected that Microsoft would have been routed by open source? Who would have predicted that open-source is what powers most of the internet? Who would have known that there would be opportunity created all over the internet by software like Linux, Wordpress, Drupal, Joomla, and a hundred others?

Just consider Google. Google gives away things like Android, Gmail, Google Apps for Business,

Google Drive, and plays with crazy stuff like self-driving cars and longevity enhancement. But Google is really an ad company. Yep. That's how they make their money. They sell ads. Their ads show up everywhere—unobtrusive text ads in search and on websites, or banners served up on millions of sites across the desktop and mobile world. And those ads drive the free software that billions of people use. Pretty amazing business model.

But, back to the free software? Sure, why not give it a try.

Most free stuff isn't very well worked out, and crashes a lot, but Kicad was pretty full-featured and stable. It was missing some very standard features (like cut and paste! And undo!) Still, I'd worked with inconvenient software before, some of it very expensive (I'm looking at you, Photoshop 2.5.) I decided that Kicad was good enough to use for the first layouts.

And the gamble paid off. Not only did Kicad work well enough for those first layouts, I still use it today. Now, Kicad is much more full-featured, and has a solid roadmap of updates and an active developer community. It even has undo (wow.)

CAD? Yeah, we looked at a number of open-source options, and eventually ended with Alibre (now Geomagic.) How'd we arrive at that decision?

1. I remembered their name from the days when we were doing marketing work for MSC Software.
2. They had a $99 special.

Yeah. We're cheap. What can I say?

But the CAD story doesn't end there. In fact, only today (2013) are we really getting serious about 3D CAD. The learning curve for parametric modeling was just too steep. To start, we went back to the old days—2D drafting in Illustrator.

The 01 Cometh

Okay, so how did we end up with Schiit's simple, minimalistically-elegant chassis? It's almost entirely a story of economics. To make something that would compete with Chinese prices, we needed a cheap box. Period.

That immediately threw out a couple of things:

- Lots of little pieces and fancy cosmetics, like Chinese tube amps

- Anything machined out of a solid piece of aluminum—you don't want to know what that costs when you aren't making Apple-esque quantities in China

In the past, we'd do a steel clamshell and a thin aluminum front panel, like we did on Cobalt. Three pieces. Two steel. That's good, because steel is far less expensive than aluminum—and if it's damaged in handling, it's easy to send it back to be refinished.

But we had another problem: getting rid of heat. Amps need heatsinking, especially Class A amps like we were thinking about making. Traditionally, you use board-mounted heatsinks or hang a slice of heatsink extrusion out the back of the amp. The problem with those approaches was that they were pricey for the amount of heat we had to get rid of. In the case of the heatsink extrusion, it also meant another cosmetic part, and another chassis component, to deal with.

That's why we soon decided to use the chassis itself as a heatsink—economics. The problem with that was that it killed the old "steel clamshell and front panel" design. Steel doesn't work very well as a heatsink. The chassis would have to be aluminum. And it would have to be fairly thick

aluminum, too, so it could effectively spread the heat.

At first, I thought about extrusions.

> **Note:** extrusions are where you take metal and squeeze it through a form, like toothpaste. Except way hotter.

In fact, originally, the 01 (our extremely imaginative name for our first product from our as-yet-unnamed company) was supposed to be a 7 inch×7 inch×2 inch sharp-cornered square, with the volume pot set exactly in the middle and slot vents on the top. The 7 inch×7 inch square would be a custom extrusion, and the top and bottoms would be flat aluminum panels.

One catch: what do you extrude?

- **The outside perimeter, like we originally planned?** Wow, that's a BIG extrusion. Sure, you can do it. Bring $ 30 000 or so. No.
- **An L-shape for top and front?** Nope, not enough surface area to get rid of the heat. And it's still a big extrusion. No deal.
- **A U-shape for front, top, and bottom?** Now you have the area you need, but now you have to deal with extrusion tolerances. The open

end of the “U” might be too far open—or too closed. Notgonnahappen.com.

But a U-shape … that could be bent from sheet aluminum. And it could be grained before bending. And you could precisely control the tolerances. And, combined with a steel inner “sled”, you had a simple two-piece chassis.

Of course, the first drawings were still 7 inch×7 inch square. In fact, the first prototype of what would become an Asgard was designed for a 7 inch×7 inch square. At least until reality intruded, in the form of the transformer.

The transformer. It was great, because it was small, cheap, and efficient, and mounted right on the PC board. It was a flaming hunk of crap because it would peg an EMF meter across the room. This meant that any PC board traces running close to it would automatically pick up hum from its magnetic field. In a square chassis with the transformer in the center, that meant nearly every trace. It hummed so bad it was unlistenable.

And nothing would kill it completely—none of the off-the-shelf transformers I had, nor a custom one I had made at MSI Transformer, with fancy

interleaved windings and a copper ring to cut down the field.

But if I simply moved the fancy MSI transformer away from the board, the hum disappeared—whisper quiet.

So that design went in the trash, and I drew up a new chassis. That's how the 01 got its 6 inch×9 inch form factor with an offset volume pot—to move the transformer away from the input traces and circuitry. Seems really simple, in retrospect. Like I said, I was rusty.

But that was really just the first step—figuring out what the chassis would be. Next was the big question: Would it be inexpensive enough to make us competitive? For that, we needed quotes.

Let's pause here and talk about manufacturing. When you need something made, you have two choices:

1. **Buy the machinery to make it yourself.**
    a) In the case of aluminum and steel chassis, this means $100k CNC mills, punch presses, laser cutters, a precision brake, a timesaver, an anodizing tank, powdercoating equipment and an oven, silkscreening gear or laser engraving.

   b) Plus people to run all this equipment.
   c) Plus stuff you probably just can't do—try to get a new anodizing shop approved in California. Have fun with that.
2. **Contract with someone who can supply finished parts.**
   a) Very simple!
   b) Until they screw up.
   c) More on (b) later. As Mike says, "There is nothing more certain than death or taxes than your metal supplier will screw up eventually."

Okay, so let's say you're sane and go with (2). Now you need to find a manufacturer who understands:

1. What "consumer level" finishing is
2. That your deadlines actually mean something
3. You expect them to hold close to the pricing they quoted after the first run
4. You are not an aerospace company or government contractor (translation: you are not made out of money)

This really isn't as bad as it sounds. Clear communication with any outside supplier is absolutely key. Most metal suppliers in the USA are not doing consumer products. They're making instrument panels for submarines, or screws for

aircraft, or heavy frames for industrial equipment, or precision-machined stuff for scientific gear.

This means that if you expect to get consumer-level products (that is, nicely finished with a very low rate of cosmetic imperfections), you need to go in and show them. Clearly explain what has to be perfect, and what isn't cosmetic. And pay for a "first article",—this means, sample—to see how close they can get. If they can't get it in 1, run.

Anyway, back to MFG.com and local suppliers.

In Southern California, we're lucky to have the remnants of a manufacturing base. It's largely left over from the aerospace heyday, but the ones who survived have learned how to do consumer and industrial products. So we had a few to choose from.

But I also wanted to get a bigger perspective, so I also looked at MFG.com. If you're looking to have something made, it's a great sanity check. You can select from manufacturers around the world, or limit it to the USA, and you can have virtually anything quoted, from machining to plastic to metal injection molding. If you limit the search to the USA, don't expect a flood of quotes, but you will definitely end up with some options.

Now, if all of the above sounds like a lot of work, it was. Between doing schematics and laying out boards and researching suppliers and screwing up the first design and having to do it all over again—while at the same time buying a new scope and other assorted test equipment, getting prototype parts, doing research online, etc, it always seemed like I was running out to the garage (where the workbench was).

"I've got schiit to do," I'd tell Lisa, and disappear.

She's endlessly patient, but one day, she'd finally had enough. "Why don't you just call it Schiit?" she shot back, crossing her arms.

"Call what schiit?"

"The new company. You're always saying you've got schiit to do. Why not just call it Schiit?"

At first, I laughed. *A company called Schiit? No sane company would do that.* If we proposed that name to any Centric client, I imagined what they'd say. *Way too out there. Can't believe you'd propose that. Piss off too many people. What a crazy idea.* Then they'd fire us.

But I'd had 15 years of marketing playing it safe, second-guessing everything we did, and watering down every great idea until it was meaningless.

Maybe you can blame my decision on that history. Maybe it was nothing more than that.

And this company wasn't about playing it safe. Hell, we were trying to reach Chinese prices here in the USA. And do it without a million-dollar investment. That was about as crazy as it got.

"Nobody would ever forget it", I replied, finally.

"It would cut down your marketing costs", Lisa agreed.

"And we could say we make some really good Schiit."

Lisa laughed. "Why not? Go ape Schiit."

"And Schiit happens", I agreed.

"If you don't have our stuff, you're up Schiit creek", Lisa added. I nodded and sat back. Suddenly it didn't seem so crazy. Hell, the word was meaningless for, what, 80% of the world that didn't speak English? And if you spelled it funny, it could sound vaguely German.

Hell, the Teutonic connection opened up all sorts of stuff, including all the old Norse mythology. They named, like everything, from gods to spears to crows and forks. An endless source of non-alphanumeric names.

“Norse and German aren’t the same”, you say? Well, it doesn’t matter. We have comic books and movies to learn from. We’re dumb Americans. We mix stuff up. What’s more, we’re from California, where Thai-Mexican fusion food sounds like a good idea (and really is.)

And a name like Schiit would be unforgettable. Nobody could ignore it.

And, fact is, great marketing polarizes. Some people hate it. Some love it. An ad that hits the middle ground of “nice” is pure crap. Which is what most companies shoot for. Might as well cash out the whole marketing budget, roll logs of $100 bills, and have a big bonfire.

But we weren’t here to hit a nice middle ground. We were here to be unforgettable. (And hey, we didn’t have that many hundred dollar bills, either.)

And in that moment, everything gelled. *We would be Schiit.*

Chapter 5

$800 In Screws?

It's funny what sets off your doubtometer.

The name "Schiit" never did. From that first conversation, it stuck. Hundreds of hours building and testing prototypes, the first of which were massive failures, never fazed me. Getting into tubes for the first time and working with 200V rails didn't scare me, once I figured out how durable, simple, and and fuss-free tubes are. Running around to a bunch of different metal vendors to get prototypes was no problem. Placing big orders for electronics components didn't even register.

But when it came time to order screws, that $800 almost brought the whole mess down.

This was about the same time as the conversation with Mike Moffat in the Foreward. The design work was done, and we had working prototypes that sounded good and worked well. But the huge work of getting a working inventory and

production system in place wasn't done. We didn't have all the parts we needed, we didn't have a place to make the products, nor people to do the work.

And—to be clear—this is huge work. If you dismiss it as "only purchasing", or "only operations", your business is going to be headed for trouble. And, once you've grown up and have someone to handle your purchasing and ops, you'll still be in trouble if there's no oversight.

Aside: There may be one other thing more sure than death and taxes (besides "your metal vendor will screw up eventually.") That is: "If you have one critical part that you can only get from one company, an unsupervised purchasing manager will order the wrong one—or will neglect to tell you that it'll be out of stock for 22 weeks until the production run."

By the time I got serious about purchasing parts, it was February or March of 2010. This was only a short time after that conversation I had with Mike Moffat in the foreward. It was also about the same time that Centric starting showing some serious signs of life. And with marketing picking up, things were actually looking better on the "ain't going out of business" front.

So, I had less incentive to work on Schiit. I remember sitting there in front of the computer and thinking, "$ 800 in screws? What will we do with $ 800 in screws if this doesn't go somewhere? We have no use for $ 800 in screws."

And I sat there for a long time, wondering just what the hell we were doing. How did we expect to just start up a new company in a field we we'd been out of for so long, and expect that it would simply work?

In that moment, the whole thing could have come undone. I'd told Mike Moffat that we should do a new company, and we'd laughed about the name, but there were no formal documents. He might even decide just to keep working in entertainment.

The doubts piled on: Would people really buy these things online? Would they think it was all a joke? Would our hacky-ass ecommerce system even work (more on that later.)

It actually took a few minutes to press that button. For $ 800 in screws. $ 800 in screws that might have undone the company. Remember, at this point it was still a money-drain and a time-sump. It would have been a lot easier to give into the doubts.

And that brings up an important point. If you start a business, there will be doubts. Lots and lots of doubts. There will be days when you'll take $5 for the whole mess. There will be days you want to quit. These doubts, and these dark times, will be far larger than anything you can imagine if you're working for someone else, even if the company is muttering about downsizing and layoffs. Because the whole mess is on you. There's nobody else to fall back on. There's nobody else to blame.

Production, Garage Style

"When I was back at Theta, we used to have a great board house," Mike said. We were past the $800 screw order, and were talking on the phone about how we'd actually get everything made. "They're in Simi Valley."

"I was thinking we'd just hand-solder everything to start", I told him. "The stuff is simple enough."

"Ooohh-kayy", Mike said, doubtfully. "But who's going to do it?"

"I'll do it", Rina chimed in, before I could answer. She'd been listening to my side of the conversation and had figured what we were talking about. At

the time, she was starting to get her own business off the ground, and was looking for extra cash anywhere she could find it.

"Rina says she'll do it", I told Mike.

"Ohhhh-kayy", he said, even more doubtfully.

"Mike's skeptical", I told her.

"Has he forgotten that I solder better than him?" she shot back. "How many of those Theta amps did he make himself?" I said nothing. It made a lot of sense. Rina knew electronics, and electrical, and had more experience soldering than any of us.

"I think we let her have a shot at it", I told Mike, thinking, We can always switch to someone else if it doesn't work, or go to a boardhouse if we ever do a second run of these things.

"Your call", Mike said. "So who's going to put them together?"

"Me, for now."

Silence from Mike. This was his even more skeptical mode.

"We'll get help when we need it", I told him.

“Better start looking now.” Mike said. Mike’s always more of the planner. He thinks ahead. “And where are you going to make it?”

“In the garage.”

“The garage?” Skeptically.

“Hey, if it’s good enough for HP and Apple, it’s good enough for us”, I told Mike. “Remember, Christmas presents.”

More silence.

Schiit was a fundamentally different company than anything Mike had done before. In the old days, you brought out your biggest bestest baddest product first, then moved down to less expensive gear. This works well in a market where the press can be the spokesperson and the dealer can be the psychologist for a customer investing thousands of dollars. But nobody is going to throw down their credit card on a multi-thousand-dollar piece of gear from an unknown company. Schiit was always intended to work the opposite way—start with the inexpensive products, build a following, and move up.

“I’ll get you the phone number to the board house”, Mike told me.

Aside: Rina (sometimes helped by me or Jean, Centric's bookkeeper and ex-Tektronix module assembler) ended up soldering about 1000 boards for Schiit, before we caved under the pressure and went to Jaxx Electronics, the assembly house we still use. The first 600 or so Asgards and Valhallas were assembled by me, before Eddie came on the scene—more on that later.

> **Aside aside:** And, actually, we didn't build them all in the garage. We stuffed boards at the kitchen table (and, in Rina's case, on a 1966 Corvette hood), bent resistors while watching movies, built in the garage in a 20 foot by 3 foot wide space (no kidding), and burned in and did listening testing in the living room.

Duct Tape and Baling Websites

In 2010, websites and e-commerce, for Centric, was easy. Add money, and we could build whatever custom site and e-commerce system you want, with whatever features and workflow you needed.

Of course, Schiit didn't have money. Or at least none I wanted to spend. But we still needed a

website and an e-commerce system secure enough to handle a reasonable number of orders. And this led us to an entirely new approach, one we use to this day for start-ups with limited budgets.

First, we did a Wordpress template to our own custom design.
Getting your own Wordpress template only costs a few hundred dollars, especially when you're doing the design (you can blame me for the Schiit brand, design, aesthetics, and copy, by the way.) Wordpress gives you a very versatile content management system, so you can easily maintain the site by yourself.

Second, we hooked up to an easy-to-integrate payment processor. In those days, it was Google Checkout. Rina did the hack-and-paste code that allowed people to put products in a cart and check out.

Third, we figured out a rough shipping cost for domestic and international shipments, and used manual shipping calculation. Of course, this was inaccurate, cumbersome, and painful. I wouldn't do it again.

But, in the end, we had what we needed: a working e-commerce website for a few hundred dollars—and a few dozen hours of time.

Update: Today, the general principles above hold true, but you have more options. You could easily use a platform like Shopify to get up and selling fast, without any code, and with much more robust shipping and payment options. You could build the whole thing in Squarespace. You can buy a commerce-friendly responsive Wordpress theme from ThemeForest for $ 40 or so, which will have tons more features than we ever imagined. Hell, you can even just set up an Amazon store and have them fulfill (take orders and ship products) for you. Bottom line: selling online is easier and less costly than ever. Don't let it stop you from having your own business.

Okay, Let's Talk Business

Before I go farther, let's talk business. Real business. As in, business plans, business structure, all that good stuff.

No, don't roll your eyes. I know, you can read about this in pretty much any "Start UR Own Biz!" book, but let's apply the sharp point of experience and turn up all the key points to 11.

First, business plans are, in general, an incredible waste of time. I know someone is

going to blow up about this point, but I'll stand by it. I have written about a dozen serious business plans since leaving college, and every single one of them was seriously researched, complete, and sounded compelling.

Not one of those businesses ever got off the ground.

Why? At least in part because business plans are big and intimidating. A standard business plan template has dozens of sections and subsections, asks for broad knowledge across a wide range of disciplines, demands decent writing skills, and requires some serious number crunching. It's a lot easier to sit and stare at the thing, thinking, "I ain't never gonna finish this," or "I have no idea what these bozos are talking about," than to finish it.

Because of this, business plans promote paralysis by analysis. If you really want to fill in all the blanks on a business plan, you'll:

1. Eat up an incredible amount of time that could be used for getting started
2. Stir up a thousand doubts that can keep you from ever starting
3. Be so amazingly exhausted that you might not want to do a business at all

"But business plans are what you need to get capital," someone at the back says.

Uh-huh. Right. Trust me, if you don't have a working product that's making money, you're not getting capital even if your business plan was written by Hemingway. Period. And no, I don't care if you're friends with one of the board members. All VCs know that business plans are fundamentally BS.

"But it helps you keep your eye on the big picture," someone else says.

Um. No. The big picture changes every day. This is not the slow, distribution-centric world of thirty years ago. Today, a new competitor can pop up on Amazon overnight—from literally anywhere in the world. Online pundits can make or break a new company with a single post. And even traditional companies are moving faster, and getting into new, unexpected product segments.

What business plans promote isn't big-picture thinking, they promote "railroad syndrome." As in, the business plan is the rails, and you're a train. It's easy to continue driving down the same wrong path until it's too late, because:

1. It's what's in the plan, so it must be true

2. You spent so much time researching/writing, it really has to be true
3. If it's not true, you don't want to spend all that time again to figure out what is now true.

So, throw away that business plan. Forget it. Pay attention to your market. Learn your market. And keep learning. Because it changes every day.

"But I don't want to just wing it," yet another audience member says. "I want some structure in my business. What can I do besides a business plan?"

Okay, fine. Let's try something new. I'll call it a Business Brief. It can be no more than a page long. It's not for getting capital. It's not for answering every question. It's about having some answers to the most important questions. To create a Business Brief, answer these questions:

1. What will this company do that no other can do?
2. If others can do this, or are doing this, how are you significantly better?
3. Why would someone pay money for it?
4. How will they find out about it?
5. How much money do you need to start it?

The goal isn't a dissertation—single sentence answers are ideal. Let's do this for Schiit.

1. What will this company do that no other can do?
Make amazing-sounding, amazing-looking high-end audio products in the USA for prices similar to Chinese manufacturing.

2. If others can do this, or are doing this, how are you significantly better?
Nobody else truly manufacturing in the USA can beat our prices; we also have unique aesthetics and compelling features.

3. Why would someone pay money for it?
Because it's a helluva deal, and they laughed their butt off when they heard our name.

4. How will they find out about it?
By people with no sense of humor carping about the name to their friends on forums. (No, seriously: through an unforgettable brand and direct engagement in micro-social activities.)

5. How much money do you need to start it?
$ 10 000, and 2 years of no salary.

See? Easy. And very easy to change when the game changes. A business brief makes you answer the two key questions of what you do and why it matters.

Second, you incorporate. Full stop.

Don't even think about silly stuff like partnerships or sole ownership. If you are making things that plug into a wall, even with CE and FCC certifications, you need to be a corporation. Period. Yes, it's expensive ($1000 or so in California), and yes, it's a pain in the ass (as in keeping your personal and company assets completely separate, corporate minutes, resolutions, etc), but here's why you incorporate:

Let's say someone wants to listen to your great new tube amp. While in the bathtub. What's more, they love it so much they give it a big hug in the warm, watery depths. They die. Their family does not understand that stupidity does not give someone carte blanche to free money and sues your company.

- If you are a corporation, the corporation either pays for a successful defense, reaches a suitable bribe—er, settlement—to make the family go away, or loses and pays or goes bankrupt. It will be a terrible time for you, but they can't touch your own personal investments, house, cars, etc.
- If you are anything but a corporation, they can go after everything you have, whether or

not it was yours before starting the company. And by everything, this means everything.

Third, you truly understand “cash flow.”

They call it cash flow for a reason. For about two years, you get to watch the cash flow from your customers, through your hands, and back out to your vendors. And that’s about it. A fast-growing company eats cash like mad. You’ll be reinvesting everything you make in growth. And there won’t be any left over for you.

- Yes, that’s right. Expect no salary for a couple of years.
- Yes, I know, that’s unrealistic if you don’t have the savings or an alternate form of income.
- Yes, I know that’s not fair because you can’t find anyone to give you free money, and it’s holding you down, you could take on Musk and The Resurrected Jobs with one hand tied behind your back.

It’s not fair, but it’s the way things work. If you can’t afford to put in some money up front and have no salary, you’ll need to start a company that requires little or no capital, and can be done in your off-time from your real job.

Wow, this is starting to sound like a business book. And it's taking far too long. So let's cut to one more aside, and then close it up for now.

The Schiit Ass Guard?

Believe it or not, we never connected "Asgard" to "Ass Guard" until people started to comment on it after launch. So no, "Schiit Asgard" isn't an inside joke for "Schiit Ass Guard." Or maybe the joke's on us.

Chapter 6

The First Order Is ... For Something We're Not Selling

Launching a product isn't like live theater in one respect: at the theater, you've got a play date. The show's gonna go on, whether you're ready or not. It doesn't matter if all the costumes were lost because a drunk truck driver drove them down a ravine, or if the lead actor is sick, or if you really don't have the whole performance gelled. You need to get on stage and do *something*.

So, with a product launch, you're lucky in at least one respect: you can pick the date. And you can move it if things aren't ready. And, if you're not stupid, and don't talk about the product until it launches, then nobody will be the wiser. You'll look like a company that profoundly has its Schiit together.

> Oh, how I wish I could jump in a time machine, go back 4 years, and yell, "Never talk about Ragnarok and Yggdrasil until they are DAMN

GOOD AND READY to launch."

But launching a new product, especially when you're also launching a new company, is like theater in at least one respect: you're baring yourself to the ruthless examination of the public. *What will they say? What did you mess up? Is it gonna be "meh" or "omg?" What could have been better about it? What if everyone laughs you out of the game? What competitor did you miss? What if you, well, just screwed up?*

Because I gotta believe that even Steve Jobs, when he got up on stage with the first iPod, had no idea how it would go. And some of the first press commentary was pretty scathing. "Too expensive, from a niche company nobody pays attention to, why would you want to put all your music on one device?" But we all know how that played out.

Look, I work with creative people every day. And not one of them can sail blithely into a client review, thinking, "They're gonna love it, no question." Because they might not love it. They might think it's stupid. They may even make some very pointed, personal remarks about how the creative director is an unoriginal hack. I've seen it happen.

And that's why creative people get so cynical. "The client won't get it. Give them something easy and obvious. I'll save my best work for myself."

Except you can't. Because then you really are a hack.

Do you think engineers are any less creative than artists? Do you think they're hurt any less by savage commentary that questions their competence?

Do you think this might have something to do with how so many audio companies act like they're living in an ivory tower, dispensing wisdom from on high? Or something to do with the fact they'd rather not talk to customers, and relax behind the walls of distribution?

Almost Competent

Anyway, enough with the emo stuff. When we launched Asgard and Valhalla, we had a chance to look supremely competent—and had to settle for "almost."

It was June 15 of 2010. We had about 20 Asgards built and ready to ship. It was time to make the

website live, send out the press releases, and see what the public would say.

There was one little catch, though. We had no Valhallas. As in, we had exactly one working prototype board without a chassis. As in, the prototype wasn't even fully worked out yet. I knew it kinda worked, but I wanted Mike's expertise on the tube side to get it fine-tuned.

So, yeah, we launched with 1/2 of a product line. Like I said, almost competent. We could have shut the hell up about the Valhalla and surprised everyone a couple of months later. But no, we had to go and show that we were going to have a full line.

This is what you call "ego talking."

This gets you in trouble. Shut up. Perfect the product. Then launch it. Anything else isn't "product launch." It's "product escape."

And yeah, I know, everyone likes to talk about what's coming up. A lot of companies do it. But that doesn't mean it's right. Product escape blunts the impact of the launch. By the time you've gotten it out, everyone might be tired of hearing about it. (I'm hoping that isn't the case with Ragnarok and Yggdrasil, but we'll see.)

And, wouldn't you know it … the first order we got was for a Valhalla.

Of course.

How to Launch (Not an ICBM)

Okay, let's take a little tour of that marketing niche known as PR, or public relations.

Nobody really knows why it's called this, because it would be more aptly known as *press relations*. Because them's the guys who have the relationships with the editors, writers, opinion leaders, market makers, dudes with a blog and a million unique visitors a month, guys with 40 000 forum posts, etc. And because of those relationships, your PR guys can get you "free advertising" in the form of mentions and reviews.

Please note: the scare quotes are not there just for show. To a good PR company, "free advertising" equates to "pay us $ 4000 to $ 20 000 per month for the chance of coverage in the WSJ." Not exactly free.

The reality is, a brick can get free coverage. Probably not in the Wall Street Journal, though.

That is, if the brick can write, send emails, and follow a few simple rules.

Here's how you do it:

- **Find your press contacts and get their emails.**
  These are usually on the site, under Contact. You're usually shooting for the Editor/Managing Editor/First Name on the Editor List. Don't shoot low, unless you're trying to get into the mainstream—in which case, research who is writing about audio, on, say, Gizmodo, and send it to them. And the top Editor, too.
- **Write a real press release and put it on your site, with photos (at least.)** A real press release doesn't read like marketing. If you go on and on about how your product makes music sound so real that you'll feel like you dropped acid and traveled back to 1968 to tour with The Doors personally, or how every other product is complete crap, it's going to go in the trash bin. If it's more than 400 words, it's not going to get read. If it doesn't follow the inverted-pyramid journalistic style, it's getting canned. Here's a basic formula that works:
  a) **Headline:** what you are introducing, in a few words

b) **Subhead:** why it's important
c) **First paragraph:** everything they need to know about what the product is and why it's important, in about 60 words max.
d) **Second paragraph:** a quote from an important person in the company—showing personality here is fine
e) **Third, 4th, 5th paragraphs:** product details
f) **Final paragraph:** pricing and availability
g) **About the Company** block—keep this short and nonhyped

- **Write an email addressed to each editor personally,**
  tell them the most important thing about the product that will get their attention, and link to the press release on your site. This is where you can have some more fun and show some personality, but remember what the real goal is here: finding something they'll consider interesting enough to write about. What's the most standout thing about the product? Start with that.

This is exactly what we did when we launched Asgard and Valhalla, and it resulted in coverage on virtually every audio site, as well as breakout coverage on Wired, Engadget, Gizmodo, and TechCrunch. YMMV.

And yeah, I know, you thought there were payoffs and backroom deals going on here. Sorry to disappoint you.

The Deluge

Before we launched, I worried that we were gonna fall flat on our face. Within 2 days after launch, I was terrified we weren't going to be able to keep up.

It was madness. In less than two hours after the press release went out, the first online articles showed up. Then, a thread, *Cool Looking Schiit,* was posted on Head-Fi by Roscoeiii.

The emails started pouring in. They were a mix of disbelief and delight. Disbelief at the name, and delight at the looks and the price of the products. We got emails from prospective buyers, engineers, Mike's old friends, my old friends, other manufacturers (including Audeze—one of our first emails), writers, bloggers, audio press, mainstream press.

Then that first Valhalla order came in. Rina called to let me know. She knew we were nowhere near to shipping any Valhallas. She wasn't thrilled. I looked up the order online, and thought I

recognized the name. I Googled it, and crazily enough, it was a reviewer—Vade Forrester, who wrote for SoundStage.

Ah, hell. The first order was for a reviewer. For a product we wouldn't be selling for two months.

"So what do we do now?" she asked me.

"Contact him," I told her. "Make sure he saw that it was a pre-order only. And offer him an Asgard to try in the meantime."

Sidenote: don't offer pre-order. Ever.

"You're the marketing guy. You do it," she told me. "I gotta go stuff some boards."

"But I've got like a million emails!"

"And whose idea was it to do those pre-order Valhallas?" she shot back.

So contacting Vade fell on me. And good thing I did—he took the Asgard loaner, liked it, and wrote a nice review on it for Soundstage. Unfortunately, he wasn't able to get the Valhalla in there, though he ended up liking it even better.

And all the other emails fell on me. It was overwhelming. How overwhelming? I actually went home from my marketing company, claiming

illness. And that wasn't really far off the mark. My guts were churning as I realized, *Holy schiit, we may actually have something here. Now what?* I called Mike.

"Hey, uh, Mike, I think we might have a winner here with Schiit," I told him.

"Yeah? Cool." Mike replied, sounding unconcerned.

"No, I mean really. People are going crazy. I have like a hundred emails to answer."

"That's a good thing, isn't it?" Mike asked.

"But, you know, we never really did the business details," I reminded him. "You still want to be part of this, right? You still want to help?"

And that's the truth. We had no formal agreement in place when we started up. Just a couple of old engineers, playing with gear. But when it gets real, you have to get real. And Mike, I knew, was doing Hollywood work. What if he couldn't break away from that? What would I do?

What had I done, going and launching a new company?

"Of course I'll help," Mike told me.

“But this might get big.”

“We’ll make it work,” Mike told me. “One way or another, we’ll make it work.”

“We’re going to need that DAC now,” I reminded him.

“Ah,” Mike said, pausing for a long time. “Does it *have* to have USB?”

“Yes.”

Mike groaned. “It might be good if you want to print your music.”

“Mike!”

Mike grumbled a bit, but promised he’d start thinking about it.

As soon as I hung up, the phone rang again. It was Jude from Head-Fi. Yes, that Jude. The founder. No, I didn’t slip him a Krispy Kreme box full of Franklins. He called us. On the first day.

Holy schiit, again. I knew who Jude was, of course, from the press research we did. But we never thought we’d hit the biggest audio forum on the planet so hard, so fast.

Jude had a lot of questions—many of which seemed to boil down to, “Are you guys insane?”

How can we set the prices so low? Did it sound any good—what were we comparing it to? Were we making enough margin to be a sustainable business? What plans did we have for the future? Did the stuff look as good as the pictures?

The answers, in order, should have been, "We're good production engineers and crappy CFOs, we certainly hope it does or we won't be around for long, hell if I know, answering about a billion emails, and yes." I don't remember all I really said, but I must have sounded confident enough to convince Jude to buy an Asgard and try it out for himself.

Which was both exciting and terrifying.

Exciting, because in a couple of days the founder of the biggest headphone site was going to be listening to our Schiit, and terrifying, because if he didn't like it, we wouldn't have to worry about having a company for long.

The New Normal

Luckily for us, Jude liked the Asgard. A lot of early owners liked it, too, and added their impressions. I jumped on Head-Fi with the truly awful

username of SchiitHead and began answering some questions.

And, in the evening, we built. Because the orders were coming in. Lisa stuffed and soldered boards, I tested and assembled them, and she shipped them the next day after burning in overnight. We could, full-out, assemble about 10 Asgards a day. Of course, most days weren't full-out.

The orders kept coming in, and we kept shipping. It became the "new normal." In the little time I had, I finished up the Valhalla tweaks, got Mike's blessing on the sound, and got the PCB artwork and metal drawings out for production.

For a while, things became almost sustainable. Or at least tolerable. I began running numbers in my head, and decided that this could end up being a decent hobby business. Maybe I could put in an outbuilding behind the house so we'd have enough space to run it out of, and not have to spend money for an office.

Yes, I know, don't laugh. But hindsight is always 20/20. I wonder what I'll be thinking, when I look back on 2014.

Then the Valhalla metal came in, and our world imploded.

Chapter 7

Metal Debacle, Valhalla Style

"Hey Mike, the new Valhalla metal just came in," I said, holding one of the outer aluminum chassis in my hand.

"Great! I want one of those—" Mike began.

"They're junk," I said, cutting him off."

Silence on the other end of the line.

"The second shipment of Asgards came in. They're junk, too."

More silence. Then: "How many?" I looked around the garage, which was now piled high with crumpled sheets of foam that had been protecting the outer chassis. Dozens of U-shaped pieces of aluminum covered every horizontal surface. I'd already gone through every box. And every chassis I pulled out was complete crap in one way or another. All of them looked like they'd been refinished after the metal had been bent. What had been a smooth curve on the top

and bottom of the front panel was wonky and uneven, where someone had manually tried to re-grain the parts. Some still had deep cracks at the bends, indicating why they'd tried to refinish the parts. They'd bent the metal, but this time it had cracked. And they'd tried to fix it.

"All of them," I told Mike.

"How bad is it?" Mike asked.

"They're unsellable. They used the wrong temper. Stuff cracked. They tried to fix it."

"Fudge," Mike didn't say. He said something much more descriptive than that.

"And we're out of Asgards. And it's a week before we said we'd ship Valhallas. People are already asking when they're gonna ship."

Mike sighed. "And you've called the metal shop."

"Yeah. They said that they can't guarantee the alloy and heat-treatment they get from the mill anymore. And they think it's cosmetically acceptable."

"Bullschiit! Time for a new metal house." Mike was pissed. "I hate being right … there's nothing more certain than … "

" … your metal vendor will screw up eventually," I finished for him.

Of course, I'm compressing this story a bit, kinda like a 128 kbit/s MP3. I'd already been down to the metal vendor by the time I called Mike. I had them try with a new lot of metal. I'd already had them try a different way of finishing it. And we didn't have any answers. And when your metal guys start telling *you* what's cosmetically acceptable (and you don't agree), run. Fast.

But it wasn't a joke. It was real. After a few weeks of shipping products, answering emails, and getting into a rhythm, I was crushed. I didn't have an alternate metal supplier. And these guys didn't want to help. They did mainly industrial control panels. They thought of us as the picky, pain-in-the-butt small client. And to them, we were.

And for the second time, I wondered if I really wanted to get into manufacturing again. It was clear that Centric would have a good year. Mike wasn't working full-time on Schiit yet. Neither of us were going to see any money for a very long time. Maybe it was time to pack it up and go home.

But that's just fear. Fear is normal. It's OK to be scared a bit. It keeps you on your toes. You think about doing stupid things like abandoning the company, and then you come back to your senses.

So, what did we do?

We started looking for a new metal shop, of course. At the same time, we got to use the "Backordered," notice on the site, and pushed out the Valhalla release date by a month. Little did I know how used to being in backorder we'd get. Nor did I know how often new product release dates would slip. To date, we've only been on time once.

To find a new metal supplier, we used both MFG.com and through personal contact to local suppliers. Most could be eliminated from a first round of quotes—4× to 7× higher than what we were paying. What this meant was that they were an aerospace supplier, usually. Not a good fit. Nor was it a good fit if they were only making machine tool front panels and industrial controls—they weren't able to show any examples of "consumer finish."

In the end, Mike found the metal guys we use to this day. They were local—only about 20 minutes away—so Mike took the initiative to go down

and meet with them. They'd already made up an unanodized sample from our print—and it was beautiful, with consistent, perfect grain and nicely finished edges. Literally a hundred times better than we ever got from (name redacted.) For the first time, I saw what our stuff could look like—and it was very nice indeed.

The problem, of course, was the wait. No metal vendor is fast, unless you bring wheelbarrows full of cash and park them outside their offices. And even then, maybe not. When your metal is bad, it's 4 to 8 weeks of delay to get it fixed. Period. And that's assuming you don't have to go out and find a new vendor.

What's worse about the wait is the nail-biting part. Wondering, *Will it look like the sample, or will they screw it up, too?* Because that could easily happen. They could buy the wrong alloy and temper, they could try to fix it too, they could mess up the anodizing, a hundred things can happen. And you won't know until those boxes show up at your garage (er, I mean, "loading dock.")

Metal and Manufacturing, a Triptych

Comment 1: there are many ways to finish metal. There's no right way or wrong way. Graining, bead-blasting, etching, etc—as long as it produces a consistent, consumer-level finish, it's fine. But the way we do ours is somewhat unique. We grain the aluminum first as a flat sheet, then bend, anodize, and screen it. This requires unique tools that won't mar the grain, as well as a specific alloy and temper so the aluminum doesn't crack when it's bent. This method is a very inexpensive way to produce good-looking chassis—with one catch. If it's scratched, nicked, dented, or marred in any way, it goes in the recycle bin. You can't refinish it once it's been bent.

Comment 2: the importance of an inexpensive chassis. Let me cover this now, because I'm sure I'll be asked. Why do we use a process that results in chassis that can't be refinished if they're damaged in production? Because it's inexpensive, and it allows simple, two-piece chassis designs. And an inexpensive chassis is key to a high-value product. At the higher-end of high-end, it's not uncommon for the chassis to cost 3× to 10× more than the parts that go in it. And that's fine, if what you're looking for is audio art. But if you're looking for value, you have to drive the chassis

cost down to a level below the rest of the components—you know, the stuff that *actually makes the product work*. This is why our chassis cost a *lot* less than what goes in them, across the board, at all levels.

Comment 3: this is the reality of manufacturing. If you're looking for a get-rich-quick-work-2-hours-a-week-from-home-with-auto-reproducing-spambot-software deal, making things ain't for you. Stuff *will* go wrong. You *will* have to deal with it. Oh, you say you're going to make it yourself on your own machines for full control? Yeah, let us know how that goes when the machinist quits/when you get the wrong metal/when the machine breaks/when you start chewing up parts for no reason. As Mike says, "Bringing a product to market is like screwing a gorilla. You aren't done until the gorilla's done."

The New New Normal

"The new metal's here," I told Mike, about 5 weeks later.

"And?"

"And I don't want to open it," I admitted.

"*****." I said nothing. We were in deep backorder, and well past the intro date for the Valhalla. People were screaming. If the metal was junk, we might not recover from it.

"Open it," Mike said. I did … and it was perfect. It looked just like the sample. The anodizing was great, and the screens were even better than the old suppliers. We were back in business!

Rina and I went back to work. Soon, there was another "new normal," with two different amps on the line. We were working late into the night, almost every night. The Valhalla got some very good reviews. And I was finally happy about the quality of the metal we were shipping.

In the midst of that euphoria, I got a second call from Jude at Head-fi.

"You know, Can-Jam is coming up," he told me. "It's at RMAF in Colorado."

"I'd love to go, but I don't know if we're ready for shows," I told him.

"But a lot of people are asking about you," Jude told me. "Maybe you could share a space with Sennheiser. They were asking about amps."

Wait. *Did he say what I thought he just said?*

"With Sennheiser?" My voice cracked a bit.

"Yes, Sennheiser." I was still in shock. "*The* Sennheiser?"

Jude laughed. "There's only one, as far as I know."

For a long time, I couldn't say a thing. Sennheiser … and Schiit amps? Would they laugh us off the table when they first heard the name?

But I couldn't let the opportunity pass by. "Let's do it," I told Jude.

"Cool. I'll have you ship out one of each of your amps to Sennheiser at the hotel, and it'll be great to meet you there."

And that's how we got roped into our first show—where we screwed up Sennheiser's plans, and insulted at least one industry bigwig …

Chapter 8

We Screw Up Sennheiser and Insult Some Big Guys

Okay. Have any of you guys been to a trade show?

If you have, you're probably groaning and nodding your head right now. You also may be sheepishly recalling some boozy 4 AM nights out, when you knew you had to be in the booth the next morning at 9 AM sharp. If that's the case, skip the next few paragraphs, unless you want a particularly snarky take on what trade shows are actually about.

On Trade Shows

Trade shows are where people come together a meet, face-to-face and in-person, to demonstrate products that are usually targeted at a specific niche. Yes. As in the companies actually fly people from all over the world to get together, swap flu strains, go out to expensive and uncomfortable company dinners, embarrass themselves by drinking too much in front of current and prospective

customers, chase union labor trying to find their products and their booth, bribe union labor to make sure they get their stuff first, work like dogs to set up and tear down the exhibits, stand on your feet all day and try not to look miserable as people talk crap about your company as they walk by, clean up after the one guy whose hangover got a little out of hand, lose the briefcase of the new CFO in your booth storage, get new pants when you realize your stuff doesn't actually work (or blows up in front of your biggest prospect), and be derided by all your co-workers for being chosen to go on such a wonderful vacation on company expense.

Sounds like paradise, doesn't it? Well, for all the different kinds of shows they have, you wouldn't think it was so bad. In addition to the big shows that everyone knows about, like CES, ComicCon and the various auto shows, some trade shows include:

- The Natural Food Expo
- The National Work Truck Association
- Semicon
- Conspiracy Con
- The National Coffin Exhibition
- The World Toilet Summit

Now, you might be thinking in this internet-driven,

mobile-aware, Amazon-grocery-delivery day and age, trade shows are seeming like, well, a buggy whip shop in the automobile era. But they keep happening, again and again, despite advances in communications, TSA-mediated travel, and economic downturns.

Why?

Part of this is the "well, we can't not be there" theory. As in, "Well, if we don't show up, people might not think we're doing so well, and all the competition is going to be there, and we might miss out on something important." Hint: people know exactly how you're doing, whether you're there or not.

Part of this is the "I get to see all my old friends in the business" theory. Yeah, and if your company just flew them all out once a year, it would probably be cheaper than going through all the logistics of a show.

And part of this is the "Hey, I wanna close some new biz," theory.
And this is still a pretty good theory if you're on the distribution side of things. Stores and distributors do come to shows, and you may get a chance to meet with them there. You may even

close a deal. But if you're selling direct, that's not a good bet.

So, let's get this out of the way: if you're selling direct and you're at a trade show, you'd better (a) have something you want the press to see, or (b) really, really like trade shows.

So Why The Hell Did You Go?

So why get into this screed about trade shows? Because Can-Jam is part of the Rocky Mountain Audio Fest, which is a trade show. It's a lot less stressful and insane than being on the main floor at CES, but it's still a trade show. We knew that going in. Centric has helped tons of companies produce, market, and exhibit at trade shows, from Semicon to SEMA, with budgets ranging into the quarter-million dollar range.

Yes. Read that again. *A quarter million dollars.* For one show. Audio, we got it easy. Now you see why I say it'd be cheaper to just fly your colleagues out.

And, during the course of my career, I'd been to about a dozen CESes. I'd set up complete systems while epically hung-over. I'd been the idiot wondering what I'd do when I had to get

up the next day and talk to customers in a semi-coherent fashion. I'd done all the stupid. And then some.

So why did we go to exhibit with Sennheiser at RMAF Can-Jam that year? Simple:

1. It was Sennheiser, like duh.
2. We knew it was a smaller show, so how bad could it be?
3. I thought I was smarter than those heavy-drinking days past.
4. We were still only 4 months old as a company, and thought we might learn a thing or two.

There was one little snag. I was still working full-bore at Centric, and there was a client meeting I couldn't get out of on Friday, the first day of Can-Jam. I could be there for the weekend, but not on the first, opening day.

No problem, I figured. Rina and I would fly in Saturday morning, but before that, I'd send two boxes—an Asgard and a Valhalla—to the show hotel, attention Sennheiser. They could grab them and set them up Friday, and we'd join them on Saturday.

It's funny, because looking back on it, I can only shake my head at all the things we missed—like "What sources would Sennheiser be us-

ing? Should we bring one?" and "How about signage?" and "What about literature?" And, and, and … If I'd had our trade show specialist at Centric running the show, she would have strung me up. Hell, she probably still would to this day.

But hey, it was our first show. Jude was going to be there. The Sennheiser guys had a big outfit behind them. What could possibly go wrong?

Waiting While Rome Burns

As it usually is with such things, our flight was later than expected. Which meant we touched down in Denver International at about 11 AM. Getting a rental car and going to the hotel ate another hour. So, all in all, it was about noon when we arrived. We grabbed our badges and headed for the show floor—but we hadn't even walked into the Can-Jam ballroom before Jude came shooting out of the room, blinking in recognition.

"Hey, are you Jason?" he asked. "From Schiit?" I barely had time to nod before Jude added, "Hey, I thought you were bringing some amps for Sennheiser."

My stomach flipped over. *What did he just say?*

"I shipped them," I croaked out.

"Where?"

"To this hotel, to Sennheiser's attention."

"Hmm, they didn't find them." Jude didn't seem really upset, but my mind was still in full panic mode. *The amps weren't there? The Sennheiser amps? The ones they needed for the show? That* SENNHEISER *needed? Needed before half the show was over?*

"I'll go check at the desk," Rina offered, and took off looking for the amps.

"Well, let's go to the booth," Jude said. "We got Sennheiser set up with some loaners, so it's not the end of the world. I'm glad you guys could come … "

Maybe this would be alright, I thought, half-listening as I followed him into the room. Can-Jam was being held in a giant hotel ballroom. That year, it was set up as a series of tables along all the outer walls, with a few outrigger table clusters. There was also what looked like a band setting up in the large open area.

Jude saw my look. "Oh, don't worry about that," he told me. "It's a set of instruments that play through headphones, so you can play live and not disturb anyone."

And he was right. Over the next day and a half, people would beat on the drum pads and produce no noise other than an anemic thwack of a stick on a hunk of plastic. But they looked like they were having fun.

The Sennheiser booth was just another single table that year, as were most of the exhibits. The headphone revolution had only really started, and even Sennheiser seemed a little surprised to be there.

Sidenote: It's really amazing how much the industry has grown up since then. Now, professional banners, backwalls, table graphics, custom tablecloths, and a much more carefully orchestrated presence are the order of the day.

They didn't have anything other than the show-provided, block-printed SENNHEISER sign up on the black drape behind the table. Two guys stood there, hands behind their backs, in the classic I'm-bored-at-a-tradeshow pose. Another guy was hunched over in one of the two chairs that

fronted the table, listening intently to the Sennheiser HD 800s. On the table were some acrylic headphone stands holding a set of HD 600s, the then-new HD 598s, and a pair of wireless headphones—maybe the RS 180s, I think. They were being driven by a small amp I didn't recognize, connected to a massive CD player.

Jude made the introductions, while my mind raced on, full of doom-laden scenarios where the amps had gotten lost in transit, they wouldn't be at the show, we'd lose even more face in front of Sennheiser, etc. I recall him saying something about how they'd borrowed another amp and the CD player to get them up and running. Disaster, total disaster.

But even then, Jude didn't seem to think so. He took me over to the Head-fi booth, where an early Schiit fan was demoing an Asgard. That was cool, but all I really wanted was to deliver on what we promised to Sennheiser—a couple of amps. Now that we had the new metal from a different supplier, they were finally looking the way I wanted them to be, and I wanted to show them off, dangit!

That's when Rina arrived—thankfully carrying a couple of familiar boxes.

“Got em,” she said.

“Great! Let’s get them to Sennheiser!” I double-timed it back over to the Sennheiser booth, where we started the process of swapping out the amps. Luckily, there were no hitches at all—the Asgard worked perfectly, and the Valhalla was soon happily glowing and powering the HD 800s. The Sennheiser guys took a listen, nodded and said some nice words, and we were set.

Rina Runs the Company

“So, do you work for Sch … ah … I mean … ah … how do you pronounce it?” the lead Sennheiser guy asked Rina.

“Schiit,” she said. “Schiit Audio.”

Senn guy grinned, a little unsure of how to take it from there. Rina rescued him. “Yep, I make the products,” she told him.

“Make?” he asked, even more off-kilter.

“Yeah, I stuff the boards and solder them,” she said.

“Sometimes with some help,” I said, not wanting to look too small.

“And I print the orders, and do the shipping,” she added.

The other Sennheiser guy laughed. “So you run the company, while he—” pointing at me—“plays with designs?”

“Pretty much,” she agreed.

They got a good laugh out of that. From there, we lapsed into comfortable show-smalltalk: the traffic seemed slow for a Saturday, it was busier yesterday, where was everyone, etc. A show can be jammed like Comic-Con on opening morning, and show staff will still complain it was slow. I slowly relaxed. This was more like it. I could do this.

And, to be honest, we had our share of interesting visitors. One was John Broskie, of tubecad.com fame. I was thrilled to meet him, since I’d used his software for some early Valhalla calculations (which Mike dismissed, then checked and pronounced them good—he has a very big case of “not invented here” syndrome, which, given his history in audio, is probably warranted.) Broskie also provided some of the clues that led us to Lyr’s Dynamically Adaptive output stage, but that’s a story for another chapter.

Another aside: want to get into audio? Start hanging out in places like DIYAudio.com and reading sites like tubecad.com and Nelson Pass' DIY site. The leading edge of audio is really at places like these—usually not fully worked out, sometimes completely unrealizable, buried in tons of other cruft and bitching—but it is there. Then, start building stuff. You'll quickly learn what works and what doesn't, at least in a seat-of-the-pants manner. It'll also be immensely helpful to understand the basics of analog (and digital) design, focusing on control theory. Then get yourself a QuantAsylum QA400 or some other inexpensive analyzer and start seeing how your designs actually do on the measurement side. Then try to break them, loan them to friends and see what they say, and start figuring out what separates a "consumer-friendly" product from a hobby product, if you want to produce it. But again, I'm getting ahead of myself …

Broskie seemed astounded by our products—in that they were quite inexpensive, and made in the USA. He wrote about them in his blog, and comes back every year to see what we've come up with. I certainly hope he's amused by our growing show presence, relative to that first sketchy year.

And, of course, we met the press. Some seemed skeptical, some seemed impressed. By the next day (no late nights for me these days, thank you very much—we are insanely boring at shows) I had the "inexpensive, made in USA" spiel down pretty good. And it looked like it was going down pretty well. Back then, we were really the vanguard of inexpensive, made-in-USA product, so it was really surprising for a lot of people to hear. Another Chinese-made inexpensive tube amp? Meh. Made here? Hmm, maybe there's something to this.

The First Ragnarok … Was Lyr

That Can-Jam is also where we met Audeze for the first time. It was funny, really, because they were the first company that wasn't really impressed by the staggering 1 W power output of the Asgard. "Four watts is more like it," they told me. "And our driver will take 15 W."

That was the eye-opener that led to Lyr, just a few months later. I'd been playing with higher-power designs, truly insane stuff by headphone standards (you know, like 6 W to 8 W), but I hadn't really planned on selling them, except as a stunt. Like "this thing has so much power, you have to

take off the protective sticker with the disclaimer that you might blow up your headphones if you use it."

So—there you go—the truth is, Lyr was originally going to be our Ragnarok. An insanely powerful amp that people would buy simply because it was nuts.

And then Audeze happened, and changed our plans. That's why we accelerated the development of Lyr—because of the orthodynamic revolution.

The Anonymous Guy

And then there was the one incident, with the company CEO that shall remain nameless. Like I said, the "inexpensive, made in USA" spiel was going very well. Most everyone who heard it seemed thrilled that we were trying to bring back affordable, high-end products. So it kinda threw me for a loop when someone didn't seem so pleased about it.

Late in the day on Sunday, a guy came up to the Sennheiser booth. His nametag was flipped around, so I didn't know who he was (note, this probably wasn't a deliberate thing—nametags

have a habit of doing that.) If I was less green, I probably would have recognized him from another show, or from his company's press materials. But that day, he was just another anonymous dude.

Anonymous dude picks up the Asgard roughly and squints into the vent-holes in the top, as if trying to read tea leaves. And from his expression, he didn't like the fortune he saw. He turned it over and over, ran his hand along the grain, and twiddled the volume pot, all the while his expression getting more and more grim.

"How can you make this for this price in the USA?" he barked out, finally putting the Asgard down, then moving on to inspect the Valhalla.

And—it's funny—nobody had asked me that yet. So I took this as a chance to show off, and be a little snippy.

"I think it's because most manufacturers are lazy," I said. "They don't even try to make things here anymore. It's easier to just throw up your hands and say, 'well, just make it in China, because everyone else is doing it,' than to actually do the research, find the vendors here that are doing inexpensive quality work, and make it yourself."

This didn't improve Anonymous Guy's mood. His brows furrowed even more deeply as he scowled at the Valhalla.

"That's it, huh?" he asked, as if in challenge.

"And," I added, throwing gasoline happily on the bonfire, A lot of companies are really bad at production engineering—it takes a lot of work to make something simple and inexpensive, but if you go to China, you can simply throw parts at it until it works."

Anonymous guy glared at me. His jaw worked, as if he wanted to say something, but couldn't get it out. Finally he just shook his head and walked away.

"Who was that?" Rina asked.

"Hell if I know," I told her. Not really caring. There are always some angry guys around. Who could he possibly be?

It wasn't until a few weeks later that I saw Anonymous Guy's picture, and found out what company he worked for. And the light came on. Because not only was that company manufacturing product in China, they *were also selling their own expertise in helping other companies move their own manufacturing to China.* So, it was like I'd

peed in his Cheerios and then kicked him in the nuts for good measure. No wonder he'd looked less than happy.

To this day, he doesn't speak to us.

Chapter 9

Powering Up: Lyr

Okay. So it's time for a new amp. But before we get into this, let's talk about product roadmaps. Yeah, more boring business stuff. But this kind of stuff is important—that is, if you're interested in building a few amps on a hobby basis.

What's A Product Roadmap?

In brief, it's a plan for what products you'll have, when you'll introduce them, and when you'll obsolete them. Yes. Products have lifecyles, and you need to plan for how long you expect them to be in the market. Now, this doesn't have to be some elaborate stack of Gantt charts or backed by tens of thousands of dollars in market research. But the reality is, you need to at least have an idea of:

What products you intend to sell. If you did two headphone amps and then, say, decided to make a deep-fryer, this may not be the best

strategy. If you did two headphone amps and then decided to extend the line with another amp or a DAC that works with them, this makes sense.

Where they fit in the line. Are the new products upmarket? Downmarket? Why would someone buy the new product? What need does it fill that the others don't. Note: "I wanted to try this crazy new topology" isn't a recipe for logical product line. And you do want to have a product line that makes sense—one where you don't have products that overlap each other and cause confusion. Having a 1 W Class A amp, and a 1.1 W Class AB amp, and a 0.9 W Class S amp probably doesn't make a lot of sense.

How many years they'll be around before you refresh them.
Planning on selling something "until it doesn't," isn't a recipe for success. How long do you think your products will be competitive? In mass consumer industries, you see major updates every year. In some cases, this makes sense, since the market is changing so rapidly (smartphones, tablets). In others, it makes little sense because there's no giant change in the market (receivers, dishwashers.) In niche audio, every year is too fast. Every 8 years is a little too slow.

So, coming back from RMAF, we knew we had to have a new product—one set up from the start for the high power needs of orthodynamics. And we knew where it fit in the line: above Valhalla. We didn't know how long it would be around, because, let's face it, we were only about 6 months old as a company. Nothing was obsolete yet, nor would be for a long time.

The question was: what the heck should we do?

Lyr Challenges

Deciding to do a new product isn't really worth much. You also need to have a set of target goals for it. For Lyr, our early notes were as follows:

- At least 4 W RMS into 32 Ω power output
- Tube hybrid design with rollability—this was in response to early Valhallas being non-rollable, which many customers saw as a negative
- Retain as many of the key features of our other designs as possible: no overall feedback, non-Class AB output stage, etc
- A more direct signal path than either Valhalla or Asgard, which were capacitor-coupled designs
- Same chassis size as the other products

- Sexier cosmetics, without costing a fortune

The first point (4 W into 32 Ω minimum power), coupled with the "same size chassis" spec was the biggest sticking point. Asgard was about at the limit of heat generation into that size of chassis, running 40 °C to 45 °C using the case as a heatsink. And Asgard was only 1 W output. Multiplying output by 4× would result in a small hot-plate or grille—not something that could be safely used.

Of course, Asgard was Class A, which means it runs full out all the time. When you run a real Class A amp hard, it actually runs cooler.

Aside: "Class A" is easily the most abused term in all of audiodom. "Class A" is used variously to describe:

- Real Class A amps like Asgard and Asgard 2, which run full bias all the time and cannot come out of Class A mode, ever
- High-Bias Class AB amps that run part of the time in Class A, but transition out of Class A for higher output
- Low-Bias Class AB amps that have some kind of "sliding bias" arrangement to try to keep the bias in Class A all the time
- Other sliding bias amps with different kinds

of output stages
- Preamp Class AB stages that never come out of Class A operating ranges, but are technically Class AB
- Op-amps that are internally Class B, but have the output stage biased into Class A
- Op-amps that are internally Class B, but do not have the output stage biased into Class A
- Anything else that might possibly have some aspect of Class A operation in any one part of the circuit
- Anything they think they can get away with calling Class A

Here's a hint on how to spot real Class A amps: they are big, hot and heavy. Period. Anything else, and "Class A" is probably just a slogan.

Another note on Class A: It is in vogue with some audiophiles today to dismiss "Class A" as an inefficient affectation of doddering old designers who might be touched in the head, and therefore incapable of comprehending the amazing efficiency and performance of today's Class D designs. That's cool. But there's nothing wrong with using "Class A" to accurately describe a real Class A circuit, no more than there's anything wrong with using "Class H with

switched output rails for higher power output at greater efficiency" to describe, well, a Class H amp.

Now, Where The Hell Was I?

Oh yes. Lyr. 4× the output power and not enough heatsinking for Class A operation. Plus the heat of tubes. The logical answer would be to change over to a Class AB output stage and set the quiescent current at a level where the amp wouldn't become a George Foreman grille.

Of course, being a little (a) stubborn, (b) slow, (c) petulant, (d) affected by Not Invented Here syndrome—choose one or all of the above—we decided *not* to go with the logical answer.

Instead, I wanted to create something new. Something that kept most of the characteristics of single-ended Class A operation, but seamlessly transitioned to Class A push-pull, then finally into Class AB, as power needs increased. And I didn't want it just to be a high-bias Class AB output stage.

Why? Several reasons:

- Class AB output stages, by nature, use complementary transistors (BJTs, MOSFETs, or, in some cases, more exotic devices)

- Complementary means two different kinds of devices: NPN and PNP, or N-channel and P-channel
- Complementary devices are never truly complementary—or, in other words, the "inverse" equivalent isn't just an inverted version of the other
- Since they are not truly complementary, they introduce nonlinearities as you switch from one to another in a Class AB amp—and, even though these nonlinearities can be dealt with via feedback or error correction, they are an inherent part of a complementary output stage

This is why you'll see us employing noncomplementary output stages where possible. In Asgard 2, we use only N-channel MOSFETs, one as a current source. In Mjolnir, we use Circlotron-style topology to use only N-channels as well. Same with Ragnarok. Of course, this doesn't work all the time, so Magni has a conventional Class AB output stage, with complementary devices.

So, what I wanted with Lyr was an all-N-channel output stage that would be able to "slide" out of Class A when necessary, to deliver additional power.

That's easier said than done. I investigated various sliding-bias systems and dual-mode amplifiers, building and measuring about 20 different breadboard prototypes. Most of them worked to some extent, but all of them had some significant limitation—they couldn't make it out of Class A, or distortion was too high, or they required 4 different trimpot tweaks per channel to make them work.

The problem was that we were working "off the roadmap." Class AB amplifiers are well-understood. Class A amplifiers, ditto. Something in-between doesn't have a lot of references in literature. Especially when you're talking about such low power output. And, to make it more complicated, some of the best sliding-bias arrangements are tied up in Nelson Pass' patents, which of course we can't infringe on.

The Lyr project dragged on until late December with no listenable prototype. It was getting to the point where I was considering just throwing in the towel and using a Class AB output stage, because nothing was working well enough.

The Critical Stage

Sometimes the place to look for inspiration is in implementations from earlier eras in audio (sometimes not—there are genuinely better ways to do things now.) Because, in the old days, NPN and PNP components weren't just mismatched—they were sometimes not even in the same zip code. A lot of early work avoided complementary output stages entirely. Some of these found neat ways to improve the power output of a Class A circuit. The problem was that none of them really worked the way we wanted them to.

This is where we went back to the closest match, and started tweaking. Part of it was actual S-domain control analysis, and part of it was building and testing additional prototypes, to see how close we could get to the MOSFET control characteristics we needed.

> If this was a movie, insert montage of boring engineering work with fancy camera angles, fast cuts, and a driving, heroic soundtrack.

> Re above: ha. Engineering is a lot of heads-down work. There's not a lot of heroics or drama. You know, like everything in real life.

And, late one cold night in the garage, I finally had it—a stage that would do 43.2 V into a 32 Ω load. Sharp-eyed readers will do the $P = V^2/R$ calcs and say, "Hey, that ain't 4 W, that's 7.4 W RMS!"

Right. But then when you have two channels driven, the output falls due to power supply sag, so 6 W is a nice round number.

And that's how Lyr ended up being 6 W, and not 4 W—it overperformed.

On Power Ratings. Okay, please let me vent about this one thing. **Power. Ratings. Are. Done. In. RMS. Per. Channel.** You do not (a) rate at peak power, (b) add power output together for both channels, (c) use an artificially low load impedance to make the output look higher.

Or, well, you shouldn't, but some do. Maybe I should re-rate our stuff at 16 Ω—then you can brag to your friends about your 2.5 W Magni or 10 W Lyr. Bottom line: it's still the same amp.

The Path to Production

Back in those early days, the path to production was pretty streamlined. As soon as I had a fully

working PC board, we drew up some renderings of what Lyr would look like and announced it.

Yeah. Before the metal was in-house. Before we had boards in-house. Before we actually made a single Lyr.

And yeah, we're idiots.

But, in this case, it actually worked out. We sent out press releases on December 27th, and the renderings were sexy enough to get picked up by Wired, Engadget, and Gizmodo. We promised delivery by March 1, and actually started shipping in late February. It was the only product we ever pre-announced that met its delivery date—and the first product that sold out the first run before we started shipping.

And, it was the last product that had a run made by hand. By the time we'd started shipping, it was becoming completely clear that we couldn't do this by ourselves anymore.

Chapter 10

Our First Employee, Our First Boardhouse

In business, there are a lot of invisible lines that, once you cross them, it's hard to go back. I already covered one of those: getting incorporated. Incorporation comes with additional fees, costs, administration, etc—but it's invaluable if you want to keep your business and personal assets separate.

Now, I'm going to cover what's arguably the biggest invisible line: having employees. A business can be quite successful without employees—there are some single-person consultancies and specialized job shops doing excellent work and making good money while doing it. And there are plenty of advantages to working that way. Not least of which is that you absolutely know who's doing the work (you), what their capabilities are (yours) and who's responsible for delivering on-time and on-target (again, you.) You don't have the additional burden of a payroll, or the

additional administration of managing payroll taxes, withholding, etc, or paying a service to do so. It's simple. It's easy.

But it's also very limiting. What if you want to go on vacation? What if you are laid up? What if you get a once-in-a-lifetime opportunity that's simply too big for you to take on? That's when the single-person consultancy or job shop model breaks down.

So, you get employees. Sounds easy, doesn't it?

Nope. With employees comes a lot of responsibility. The baseline is that you have to meet payroll—and, with a growing business, you're probably running most of your profits back into production to expand. Can you afford it? Will it affect your ability to keep expanding? You need to run some numbers before you hire anyone, and plan for putting in money if things get thin.

And now you have more administration. I mentioned payroll taxes, withholding, etc, which can be handled by a payroll service at minimal cost. But you'd also better keep records of who you interviewed and who you ended up hiring. You'd better have an offer letter that spells out duties and expectations. You'd better make sure the language in that letter doesn't end up binding you

in an implied employment contract. You'd better have a clear probationary period, and make sure the employee knows it. You'd better provide clear sick leave and vacation policies. You'd better provide health care. And you'd better plan on keeping a record of any feedback or disciplinary actions you take, in case things go south. And you'd better have an employee manual, so that everyone knows the rules of the game.

Sounds terrible, I know. I made the whole thing sound like a war between employer and employee. In reality, it usually isn't. Most of the time, people are fundamentally decent, and you usually don't need to worry about armor-plating your ass in five inches of legal armor.

But if things go bad, you want it there. In the 20 years at Centric, we've been threatened with legal action twice by employees. Nothing ever came of either, but I'm very, very glad we had policies and procedures in place, just in case.

So, did Schiit have all that in place when we brought on our first employee?

No. Of course not.

Before Lyr, Schiit could have existed comfortably as a no-employees "hobby business." In fact, we planned for it. We modified our house and added storage in the attic to house some of our stock. We also seriously looked into building a shop in back of our house in order to be Schiit's permanent home. This 400 square foot box sounded palacial in comparison to the 50 or so square feet we were eking out between the cars in the garage at the time.

Aside: It's pretty hilarious to look back on it today, as we start reaching the limits of 5300 square feet—we just racked everything three levels high and leased a forklift to manage our space better.

But after the Lyr introduction, it quickly became clear that Schiit needed help. Rina couldn't build boards fast enough, even with Jean's help. At that time, Rina was also shipping everything we made, usually after burning them in overnight on the day after we made them. And I was getting tired of coming home every night from Centric and building ten or twenty amps. And I needed to spend more time with Mike, getting our first DAC hammered out (more on that later.) And

the next run of Lyr boards had just come in, and we needed to do runs of all three amps. And the order rate continued to accelerate.

Bottom line: we needed help, and we needed it fast.

As luck would have it, a very old friend of mine had been watching our progress (not coolly and dispassionately, like Wells' Martians, but actively helping along the way, making some custom tools that made assembling the early Asgards much easier.) He was a frequent visitor at the house, coming up for barbecue or wine or just to hang out. This old friend was Eddie.

When I say, "old friend," I mean, "old friend … " I've known Eddie since 7th grade, and he was involved in my first business, Odeon Loudspeakers …

Time Machine: Set Dials for 1989

Okay. Imagine you're just out of college. You've been into audio for a few years. You have big Carver amps and you build speakers. You've even sold a few of them. Some of them even sound pretty good. You have no money at all, and no experience with running a business.

So what do you do? You start a company, of course.

You try to run it while working another engineering job full-time. In the pre-internet, pre-direct-sale world of 1989. You know, no email, no internet, cellphones the size of bricks that cost $1200 and $45 per month if you didn't use them at all, and $0.45 per minute when you did, and had like 2 hours of battery life.

What's more, you decide to build them all yourself. First on the patio of your parents' house. Then in a 300 square foot unpermitted, unheated, uncooled, unpowered shed you built with $1000 of lumber from Home Depot in the back of a friend's house, and finally in a run-down, 1000 square foot industrial space in Sylmar, next to a meat packer and a body shop. And by build, I mean build. As in, sheets of MDF and gallons of paint would come in, and speakers would come out. All made with a Frankenstein arrangement of pin router jigs and templates, coupled with the world's most hot-rodded and dangerous table saw (no shields, no guards, 5× the power it was designed for, blade usually sticking out at least 4 inch above the table surface, and an 8-foot extension built for the guide.)

This was Odeon Loudspeakers, my first company.

Eddie worked at Odeon.

Odeon had no money. Almost literally. We were so strapped, we cut our own Styrofoam for packing material using the table saw. Yes, the modded table saw. No, nobody ever lost a hand. We should have. Of course, the place should have blown up any amount of times when the air was full of sawdust (we had no dust collection system) and the kerosene heaters were going full blast.

Odeon is why I always win the "we once did this stupid thing at CES" stories. If you've been to a CES dinner with a bunch of other audio industry guys, you know what I mean. "Well, there was this one time when nothing got delivered for the booth, we had to make do with rental plants and couches," or "Well, there was this one time when the prototype wasn't ready, so we had to assemble it the night before in the hotel room." Things like that.

Our Odeon/CES story goes like this:

In the old days, there were two CESes per year. One in Vegas, one in Chicago. Vegas was pretty easy. Throw stuff in the back of a van and drive there. An easy 4-hour trip.

Chicago? Not so much. Odeon couldn't afford airfare, much less freight for over a thousand pounds of speakers (yes, we made some big stuff.) But Odeon couldn't afford not to go to CES, either. We lived on orders made by distributors and dealers, and the only place we had contact with them was at shows.

> Pre-internet, remember? Distribution held all the cards.

So we had to be at Chicago. Which meant, in the infinite wisdom of less than a quarter of a century on this planet, meant: we pack up the van and drive. From California to Chicago.

Or, more precisely, Eddie and Jose drove. (Jose, my other business partner at the time, now runs a very successful specialty costume shop … you've probably seen their work in tiny little movies like Thor, the Avengers, Tron, Spider-Man, etc.)

Why didn't I go with them? Because I was able to fly out with Sumo. Sumo had money. Sumo shipped things and flew places. Odeon didn't.

Now, the sheer insanity of driving from A to Chicago in an overloaded 1970 Dodge van that was literally held together with Liquid Nails would

be funny enough, but what wins the Stupid CES Stories Folly is what happened once I flew in.

After dropping stuff at the Sumo hotel room, I headed down to the show sub-level, which was where the high-end stuff was being shown that year. I found the Odeon room, and two very tired-looking co-workers.

"Dude, Zagnut bars are real!" Eddie said, proudly whipping up a table skirt to show me what looked like ten gross of Zagnut candy bars, most still bundled into factory display packs.

"What?" I asked, completely confused.

Eddie pulled out one of the candy bars and dangled it in front of my eyes. "Zagnut! Like in Beetlejuice! They're real!"

"Uh … "

"So we picked up a bunch of them," Eddie said, gesturing at the boxes and boxes of candy bars. I didn't know what to say. They'd found a candy bar … that they saw in a movie … It didn't make any sense.

Jose came to the rescue. "So we need some money."

"Money?" I asked.

“Dude, you gotta try one, they’re good!” Eddie cut in.

“We ran out of money,” Jose said, waving him off. “We’re staying at a friend’s house, but we really need to get a motel or something, and it would be good to have some real food for a change … ”

“You ran out of money?” I echoed.

Jose nodded and pointed at the giant pile of Zagnut bars.

Suddenly it clicked. “Wait. You spent all the trip money on candy bars?”

Jose nodded.

“But it’s worth it!” Eddie said. “Everybody’s gonna trip when we get home. These are real!”

“And you didn’t have money for a motel.”

“No.”

“And you’ve been eating nothing but candy bars … ” I said, trailing off.

“For a day and a half,” Jose said. “Since we got here.”

Right. These were my business partners. Now you see why I win the Stupid CES Story Award,

every time. And why that business didn't last long.

So Why'd You Hire Eddie?

Because 1989 was a long time ago. People change. And, most importantly, Eddie was:

a) There
b) Willing to work
c) OK with piecework

"Piecework?" you're probably saying. "What's that?"

Piecework is where you tell someone, "Hey, I'll pay you $ X for each product you finish." It works great in cases where you're confident your employee isn't going to sacrifice quality to make numbers. And Eddie was, if anything, an insane stickler for quality. So we didn't have any worries there.

Now, to do piecework legally, you still need to either pay someone at least minimum wage (with piecework on top), or you need to have them not as an employee, but as a contractor. Which is what we did to start: Eddie worked for us as a contractor. Which had a lot of benefits in itself. Since he wasn't an employee, we didn't have to worry about withholding, health insurance, etc—

just pay him and give him a 1099 at the end of the year.

"Well, that's great!" you say. "I can avoid all the headaches with employees by using contractors."

Not so fast. There's a pretty specific legal definition of what a contractor is, and it may vary by state to state. If you're trying to skate by and call employees "contractors" to save cash, and the Powers That Be decide they're not contractors, but actually employees, you're in for a world of hurt.

Contractors must typically, among other things:

1. Be able to set their own hours
2. Use their own tools
3. Not have to work in a specific facility

With Eddie, we were pretty much in compliance on all 3, though he never actually took products home to work on them. He could have, though, and we wouldn't have cared.

But the fact was, we were still a small business. Eddie was working in our garage. It wasn't such a big deal—he was happy for the work, and we were happy for the help. The first few hundred amps he actually made standing up, between the 1966 Corvette and the garage shelf where we

burned-in and shipped the amps. His total work area was probably about ten square feet. I sat at the bench, testing, and Rina took his space during the day to ship orders.

And we slipped into that pattern for a while—Eddie coming in every evening, throwing something on the grille, then going out into the garage to work on Schiit. It wasn't a bad setup. And Eddie was very helpful in pointing out ways to make things easier, stuff we could change to make assembly go more quickly.

So, yeah, Eddie. He's the kind of friend who'll show up at 3 AM to fix your busted car, or hop out in the middle of an intersection to pick up a pipe wrench someone dropped, or tell you everything the body shop did wrong to your car, or will put together amazing things on a weekend just because he can, or hook you up with machining or bead-blasting, or make Schiit. He's also been the go-to guy for ultimate finish work on specialty costume, like on Thor, and he was the reviewer in Centric's experiment with internet video back in 2006 or so, called "Wineass." A quick YouTube search will pull up a few of the 140 episodes we shot. He's a bit of a character—and a great choice for our first, well, contractor.

He's also still our lead assembly guy to this day—he's probably put together 35 000 to 40 000 products by now.

Completing the Story: The Boardhouse

With Eddie on the team, we were now able to keep up with production, and even get ahead. The bottleneck was now in boards. Rina and Jean were overloaded—in fact, I probably ended up stuffing about 20 Lyrs out of the first run.

So we had the choice of either adding more staff, or going to a PC board assembly house. Mike Moffat was always in favor of the latter, and on the second Lyr run, I finally took his counsel.

"I used to use these guys—Robert, he's still in Simi Valley, I think. We could have them do it," Mike said.

"But how much will it cost?" I asked. "Do they even do through-hole stuff? How fast can we get it done."

"Dunno, dunno, dunno," Mike said. "But you can pick up that antique communication device that you loathe—the phone—and ask them."

"Why don't you do it? They know you're legit."

"You're just being lazy. You can't email for everything."

"Right. And when was the last time you helped put stuff together? And how about that DAC we have to do?" I shot back.

"I have ideas for the DAC," Mike grumbled. "But I hear you. I'll call them."

"Don't call them. Go ahead and take them the Lyr kit."

"Without a quote?"

"If we want to ship, without a quote."

And here's the funny thing. The next day, Mike went down to the board house and dropped the Lyr kit with them. A week later, we still didn't have a quote, but we had a full run of boards. Beautiful boards. Better than we ever did. I was sold. We'd never make boards by ourselves again.

> And when the bill came in a few weeks later, it was insanely inexpensive. Lesson learned: there are some things that it's best not to do yourself.

That's how, in the process of a few weeks, we went from a hobby business to something much

more real. We were still a tiny diversion in the board house's big runs, and Eddie was still working primarily with Jose, and Schiit still wasn't producing enough money to pay Mike or I a salary, but it was starting to feel like something that was, well, going to go somewhere.

A final aside: today, we're the board house's #2 customer … a fact I find pretty hilarious—and appropriate.

Chapter 11

USB Sucks!
Or, Mike Joins the 21st Century

Okay. Time for me to take a step down. Until now, most of the book has been about my designs, but now, it's time to talk about Mike Moffat and DACs.

But first, a scorecard. At this point in time, we're early in 2011. Say, 9 months old. We've introduced:

- Asgard
- Valhalla
- Lyr

And I'd talked about upcoming DACs and other fantasy products with 6Moons, further deepening my "don't talk about it" dilemma. Again, if I could go back and punch myself in the face, I would. But I really wanted everyone to know that we weren't just going to be about value products.

Anyway, on to Mike.

If you know a little bit about the history of digital audio, you know that Mike was one of the first guys to take digital seriously, and the first to introduce a standalone DAC. Yes, you can thank Mike for all the separate DAC vs CD player, DAC vs sound card, DAC vs the D/A in your phone/computer/Blu-Ray/oven/lawnmower arguments we have today.

Now, some others use technicalities to claim the "first DAC" prize, but the fact is: Mike and his company, Theta Digital, were first. The technicality is that their first DAC was actually a DAC and preamp, the Theta DS Pre. Back then, the idea of a DAC and preamp combined was, well, more than slightly strange. So that led to the introduction of the Theta DS Pro shortly afterwards, so confused audiophiles had a component to use with their uber-expensive preamps of the time.

This same Mike Moffat also designs all the Schiit DACs. So, it's not like we just decided to get into the DAC market—we have the guy *who started it all.*

It's funny. Shortly after we introduced the Bifrost DAC, I got an email from a prospective customer that went something like this:

"Hey, this looks interesting, but I'm wondering what your credentials are in digital design … let me know, please?"

To which I replied something like this:

"Well, other than having the "father of the DAC," Mike Moffat, on our team, we can recount all the stuff he brought to the table in terms of DAC design. This includes:

- Much experimentation with early 2-chassis player/DAC designs with shared clocks before the SPDIF standard was approved
- The first standalone DAC
- The first to use custom digital filters in a DAC
- The first to use DSP to run those digital filters
- First true time domain optimized digital filters, based on math perfected with a U of Iowa Professor Emeritus of Mathematics and a RAND Corp mathematician
- The first to identify jitter as a cause of audio degradation—before Mike, "bits is bits"—you can thank him for literally every jitter argument we have today
- The first to work to minimize the causes of jitter
- First to use a Stanford interval counter for jitter analysis, before "convenient" measurement via JTest

- Presented an AES paper on jitter in digital audio
- The first to introduce the AT&T ST-optical interface to address jitter issues
- The first with upgradable DACs at Theta (from the beginning in 1986)
- The first to make an upgradable surround sound processor at Angstrom
- The first DTS-capable surround processor at Angstrom
- Multiple firsts in entertainment digital media distribution engineering at Digi-Flix

But even before Mike was doing digital, he was doing audio, at Theta:

- One of the first in the tube revival in the 1970s
- The first non-12AX7 implementation for audio (revolutionary in its day)
- The first "no overall feedback" tube stage in the 1970s
- One of the first passive preamps
- Worked in the Chilean jungle looking for oil using the first 8 bit A-D converters prior to Theta
- In 1934, many years before he was born, God appeared to Mike and revealed to him the formula for amazing digital audio, which he has inscribed on 12 lead tablets …

The prospective customer's response:

"Oh, then you're going to annihilate pretty much everything, then?"

Ha. If it was only so easy.

When Mike Says USB Sucks, You Listen

Okay. Enough of the bragging. The fact is, I'm excited to have Mike Moffat as a partner, and I'm proud of his resume. He's contributed quite a bit to the digital audio realm—just look at Sony's original stance of "perfect sound forever," and their current frank discussion of jitter-reducing measures in their audio products. Quite a turnaround.

But to get back on topic, let's talk about DACs. Not Mike's original idea for a DAC, the one he came up with shortly after we started the company, but Bifrost.

Aside: Mike's original idea for a DAC is what eventually turned into Yggdrasil. Yes, we've been talking about it for that long. Yes, we're really late. But it has grown and morphed over time. And, in the early part of 2011, we weren't ready for a DAC that cost 8× as much as an Asgard.

Once we were focused on the real goal—an inexpensive DAC in the same size chassis as Asgard, Valhalla, and Lyr—that's when the arguments started.

"And it'll have optical, coaxial, and USB inputs," I opined to Mike.

"USB?" Mike gagged, miming sticking a finger down his throat to induce vomiting. "USB is for children and fools. Why would you want to use USB for audio?"

"Because it's really popular," I said. "Everyone has it—"

"If everyone was dressing up in tutus, would you?" Mike shot back.

"No, but … "

"No, but you'd say you did, just to be popular."

"We *have* to have USB," I pressed on. "There are a ton of people who only have laptops for sources, and most of those only have USB."

Mike grumbled something under his breath, then spat out: "But USB sucks. It just sucks. It was never meant for audio. It's an all-purpose, packet-based grab-bag that might be fine for printers or hard drives, but it's just crap for streaming. You

can recover the clock from the packet clock, barf, or you can have the computer and DAC do some negotiating and guess at the clock, barf, or you can turn the whole car around and drive it from the back seat with the computer providing the clock, barf."

Note: the above, for the more technical, is Mike's take on isosynchronus, adaptive, and asynchronous USB implementations.

Second note: also, remember, this is early 2011 we're talking about. Adaptive USB 1.1 was kinda the de facto "good solution", with most audio components using the truly terrifying I USB input receiver/DAC/headphone amplifier/car washer chips that weren't even adaptive. Some guys were fooling around with USB 2.0, but implementations were thin and software was iffy. We know. We tried all of them.

"We *have to have USB*," I told Mike, firmly.

Mike grumbled again.

"Remember Angstrom," I said, bringing out the big guns.

Angstrom was Mike's surround processor company. It would probably still be around today,

except for various life issues that aren't mine to talk about. But it's also the company where we discovered that going against the grain might not be the best idea, even if it ends up being the *right* idea.

Angstrom brought out one of the first inexpensive Dolby Digital decoders on the market, the Angstrom 100. It was a great product. But Mike didn't want to do video switching. Video switching, he said, had no place in a no-compromise home theater audio product. What's more, the new HDMI standard was imminent, and that was a whole new ballgame. The whole video switching deal was going to be changing, and fast. So why put it in a product when it was going to be obsolete in less than a year?

Yes, it made sense. To us. Unfortunately, to customers, it made a lot less sense, especially when Angstrom never brought out the promised separate video switcher.

"But Angstrom *was* right," Mike said. "Look at all the processors today, you throw them away when HDMI changes. Now we have the same thing with USB. You know there's gonna be a better way to do it in a few months."

“I know. But USB is a must-have. We’re dead in the water without it.”

Mike sat for a long time, saying nothing. Then his eyes lit up. “Make it upgradable,” he said. I shook my head in confusion. “Make what upgradable?”

“The USB input. Put it on a separate card. So you can change it when the technology changes.” Now, Mike was excited. He jumped up. “No. Not just USB. Make the whole thing upgradable, so you can swap the DAC as well! That’s the way you do it! Not just a throwaway product, something you can keep for as long as you want.”

“The whole thing? Upgradable?” I was skeptical. “Mike, you know we’re talking about a product that only costs a few hundred dollars.”

“Right. And that’s the brilliant part. You can buy a throwaway DAC, or you can buy ours.”

And that’s how Bifrost became upgradable. A great ending to the “Not Invented Here” syndrome—which I’ll cover some more in the business bits part of the chapter.

“And,” Mike said, pointing a finger in the air, “We’ll also make sure it’s bitperfect to the DAC, instead of using an asynchronous sample rate converter.”

"Is that good?" I asked. (Remember, I'm the analog guy.)

"Is it good? Is it good, he asks?" Mike said, recoiling from me as if I'd just asked if a Michelin-star restaurant was better than McDonald's. "It's *an absolute necessity* if you don't want to throw away all the original music data, and create some mathematical-abortion-mishmash of interpolated crap, especially if you're at a non-binary multiple of the original sample rate, like going from 16/44.1 to 24/192. No, wait. Let me guess. You're gonna tell me, in your infinite marketing wisdom, we have to crap everything up to 24/192 so we can have a number on a datasheet. Oh, boy."

"No, I'm not gonna tell you that," I said. "But what are we talking about here? How hard is it to keep everything bitperfect?"

"It's a pain in the ass," Mike said. "We'll need a microprocessor to switch the clocks, we have to reset the DAC when sample rates change, we'll need a hard relay mute, stuff like that."

"Expensive?"

"More than just throwing in an ASRC chip and being done with it," Mike said. "And we'll need firmware." I sat quietly for a bit. Firmware meant

Dave, the unsung code/digital hero of Schiit Audio. He's the third guy on the engineering team. But bringing him on meant even more expense, and we were still struggling to simply keep shipping on time.

But if Mike wants something, he gets it. So we got Dave. And we started working on the first prototype of what would become Bifrost. This was a bigger project than anything we'd taken on before, involving digital design, firmware, system integration, and analog design. I did only one thing on Bifrost—the discrete analog output stage. And that was plenty. Analog electronics in a digital box, being fed by the relatively noisy output of a D/A chip, is a whole different ballgame than a preamp or power amp. In that kind of environment, analog electronics like to oscillate. So, many iterations of compensation and filtering were in order. That took a good piece of my time.

While I was working on that, Mike and Dave were doing the rest. This included:

- Evaluation of the various USB input solutions
- Integration of the SPDIF inputs
- Microprocessor-based clock management
- DAC evaluation (one of the good things about a modular DAC is that it lets us try all the leading candidates for D/A ICs out there—we

settled on AKM because it's what sounded the best, plus it's one of the best-measuring DACs out there)

Fun fact: the first layout of Bifrost had the input selector switch backwards, pointing inside the chassis. I joked with Mike that we could put a cantilever behind the front panel button to activate it, and I think for a few moments, he actually believed me.

Of all "the rest," what took the longest was by far the USB input. Back then, you had the option of licensing code for USB 1.1, using a standard USB 1.1 input chip (the one I mentioned that has all the other stuff tacked onto it), or using one of three different USB 2.0 input solutions. The problem with the USB 2.0 inputs was that they were all kinda beta-ish in one way or another. One needed drivers for Mac and PC, despite Macs supporting the USB Audio 2.0 Standard natively. One was really ambiguous about their drivers and licensing.

And then there was C-Media. C-Media was an obscure Taiwanese company that has just introduced their CM6631 USB 2.0 input receiver. And by "just introduced," I mean, "the datasheet was

labeled Version 0.9, and the USB firmware programmer crashed after successfully programming one device—every time."

"Version 0.9 on a datasheet," Mike said, grimly, when we got the first docs. "It's a beta. Run away. Run far, far away."

But C-Media was helpful, providing the support we needed to get their device up and running. And when it was running, it sounded pretty good. And by "pretty good," I mean "as good or better than anything else we tried."

Mike was less impressed. "It's not complete crap," he said.

Later, we'd make it the USB input better with some tweaks, to the point where we thought we had one of the better-sounding implementations out there. That was where we launched with the original Bifrost—but that happens much later in the story. (And today, with the new CM6631A, I'm finally totally happy with USB. In fact, I use it most of the time at home.)

A USB Mode/USB Audio Standards Primer

Okay. Let's go to the "useful data" side of things. A lot of people are monumentally confused about

USB audio input. So here's a guide to the whys and wherefores—something I should probably put on our site.

USB Mode Versus USB Audio Standard. A ton of people are confused when I say things like, "Modi uses USB Audio 1.0 Standard over USB 2.0." They think I mean that it uses USB 1.0 as a transmission protocol. Actually, it doesn't. It runs USB 2.0 at 480 Mbit/s, but transmits audio using the USB Audio 1.0 standard.

Still confused? Okay, let's break it down.

USB Modes.

This is all about data rate. This has nothing to do with audio.

1.0: The earliest standard. So slow I forgot what it was. Not used for audio.

1.1: Transmits data up to 12 Mbit/s. Can be used to transmit audio up to 24/96.

2.0: Transmits data up to 480 Mbit/s. Can be used to transmit audio up to insane torture-the-cats-and-hard-drives rates like 32/768 and such.

3.0: Transmits data up to 5 Gbit/s (ha.) No USB Audio 3.0 standard. No USB Audio 3.0 receivers. However, USB 3.0 ports are backwards-compatible with USB 2.0, so they can be used with anything using USB Audio 2.0 standard.*

3.1: New reversible fantasy USB spec created out of Apple envy and support for people too dumb to insert a cable the right way. No products out yet. Looking forward to all the confusion coming our way because of this.

USB Audio Standards.

These are *standards* (hello, Microsoft) created to enable the transfer of audio data over USB. They are not USB modes.

USB Audio Standard 1.0. Supported by everyone. Plug a DAC using USB Audio Standard 1.0, like Modi, into any Mac or PC, many Linux systems, some phones, etc. and it will be recognized and play audio with no drivers required, up to 24/96.*

USB Audio Standard 2.0. Hello, Microsoft. There is this new standard called USB Audio Standard 2.0. And you really should support it. Because you look really dumb when anyone can

plug in a DAC using USB Audio 2.0 Standard to any Mac, many Linux machines, phones, etc and expect it to work without drivers,* while you still rely on kludgy workarounds that require drivers, ASIO setup or WASAPI setup, etc. Please please please please include this in Windows 8.2 or Windows 9 or Windows Apology for Metro, or whatever you're calling the next version.

*** So what is it with all these asterisks?** Well, it's because, in Microsoft and Apple's infinite wisdom, they've decided to save us from the threat of extreme power dissipation through USB ports with a new innovation called "port power management." What this means is that the USB port, rather than delivering the full 500 mA, or 1 A of power, as required by the USB standard, can be throttled down to use less power. Which plays merry hell with some DACs. Yes, including ours. Which means we get to educate everyone about how to turn off port power management, or, in some cases, ask them to go out and buy an externally powered USB hub to completely mitigate it. Maybe they should label the ports as "really full power/real USB spec," and "battery-lifetime-promoting, save-the-planet USB port that doesn't really provide full power." No, wait, that wouldn't fit. Never mind.

And, you know what? After writing that, I agree with Mike. USB sucks.[1]

USB 2.0, 24/192, and Beyond (and the Silliness of it All)

For Bifrost, though, locking it down to 24/96 wasn't an option. 24/96 was becoming the de facto entry level for digital audio via USB. I wanted to start with 24/192 capability from Day 1. And I got that.

But USB still had some oddities. The earlier C-Media USB receiver chip didn't work at 24/176.4, even though it did 24/192. Why? No idea. But at the time it wasn't a huge consideration.

Fun fact: the CM663 A can easily do up to 32/384. Why don't we enable it? Two reasons:

a) There is no 32 bit music, and never will be.*
b) There is no 384 kHz PCM audio for sale that we know of. Sure, there's DSD 2×

1 About 80 % of our customer support is helping resolve Windows USB issues, from driver installation to port power management, so I'm biased. But now you know why Modi is locked down to driverless USB Audio 1.0 Standard operation.

at 352.8 kHz, but that's a whole 'nother discussion.

* There's a famous napkin-scribble by a famous analog designer floating around out there on the internet somewhere, regarding the noise and precision of analog circuitry necessary for different digital resolutions. I can't find it at the moment, but it went something like this:

14 bit *to* 15 bit*: standard parts and layout*
16 bit *to* 17 bit*: attention to power supply noise, premium parts, careful layout*
18 bit *to* 19 bit*: extreme measures taken with low-noise parts, multi-layer boards, and exceptionally fine layout*
20 bit *to* 21 bit*: God's domain*

Fact is, 24 bit is 144 dB dynamic range, which is about the limit of our Stanford analyzers. The best DACs, to date, manage 19.5 to 20 Equivalent Number of Bits (ENOB), even if they are "24 bit" or "32 bit" spec'd. 32 bit is 192 dB dynamic range, which ain't gonna happen, no way, no how, not even in temperature-controlled circuits sitting within 2 feet of solid lead shielding. Consider that a stun grenade is 170 dB to 180 dB, and you'll see how crazy this is.

Perhaps it's a matter of capability. With SPDIF, we had some finite, and rather low, limits to amount of data we could transmit reliably in the past, especially if you were talking Toslink optical. That's why Theta went to AT&T glass-fiber optical

to get more bandwidth. Now, Toslink is better, but it's a rare Toslink that can do 24/192 reliably.

But with USB 2.0, and even more so, 3.0, we have no such restrictions. How big a data rate do you want? How many bits? No problem. We can make up silly numbers all day. But don't think it'll be meaningful in musical terms, if, say, we can transmit 64/1.544 Mbit/s bit depths and sample rates.

But, you know what? If you have a Bifrost, you don't have to worry. If aliens from the planet Zebtron land on our world tomorrow, bringing physics-defying technology that enables 64/1.544 Mbit/s audio transmission over USB, we'll have a USB Gen X card soon enough to handle it.

But I wouldn't hold my breath.

The Problem with Not Invented Here

Finally, let's talk business. Thanks to Mike's insistence on being different, we ended up with a truly unique DAC. And sometimes it takes that stubborn insistence, that rejection of everything "not invented here," to make something great.

But “not invented here” can bite you in the butt, too.

It’s something to watch for, if you’re going to start your own business. Too much “not invented here” hubris can delay products, reduce efficiency, and interrupt operations. Sometimes the right answer is only an internet search away. Or a great idea might be just a small tweak to a similar product.

“But wait,” you say. “Are you saying … steal from other companies? Plagiarize?”

No, not at all. It’s a balancing act. You should be aware of what your competition is doing, how other people have solved problems like yours, and what the basic industry benchmarks are. At the same time, you should have your own ideas.

And—here’s the hard part—you have to have a feel for whether or not your ideas are better than the prevailing wisdom, and also if your ideas are realistic to implement.

“So how the heck do you do that?” you ask.

Believe me, I wish I had a formula. Some companies will spend tons of time benchmarking against their competitors and running focus groups to try to determine if they’re going to be successful, but I believe this is more likely to result in mediocrity

rather than brilliance. You can't assume your competition has all the right ideas, and you can't assume a focus group is a microcosm of your entire prospect base. What's more, you can't assume that a truly great idea will make it through a focus group, because they're more likely to be confused about something that's truly unique, and has no point of reference.

Case in point: until the original iPhone was announced, everyone was wondering what kind of keypad and stylus it would use. Nobody guessed it would have neither one. It was simply insanity to consider it, at the time. Love or hate Apple, they changed the game. I think the best way to decide on when to stick to your guns on new ideas—to be stubborn, and pound the table, and insist on "if it's not invented here, it's not for us," comes down to weighing the risks and rewards.

Here's an example, using Bifrost.

Rewards of doing what everyone else was doing:

- **Faster product introduction:** If we'd done what everyone else was doing—using USB 1.1 and upsampling everything with ASRC, in

a non-upgradable platform—we would have had a product out much sooner.

- **Easier development:** we wouldn't need code, integration, multiple boards, custom connectors, etc.
- **Easier support:** no explanations about why the DAC has to be reset, resulting in the famous "clicking")

Risks of doing it like everyone else:

- **Prospect disappointment:** people were already expecting great/different/innovative products from us, not a rehash of what everyone else was doing, so Bifrost might have flopped if it had been like a lot of others.
- **Undershooting the bar:** if someone else introduces a game-changer right before your me-too product, you're gonna be in a world of hurt.
- **No story:** going with the crowd means you have no ideas of your own. No position. If you have no ideas of your own, what is your value? Sometimes you have to stand up and say, "This is what we believe in."

In the end, we decided to stand up and say, "We believe in this, and we'll take the pain to make it right."

Which is good. Because it was a pretty painful path getting to the Bifrost introduction, which was several months late. But the anecdotal results are that we did something right. Have you noticed that most of the better DACs these days are avoiding sample rate conversion, or allowing you to turn it off? And how about all those inexpensive upgradable DACs that assure your new purchase isn't going into the trash can after a few months … no, wait, that hasn't happened yet …

And now, on to the next chapter. Our most humbling experience. And the closest we came to throwing in the towel.

Chapter 12

Schiit Goes Evil?

No story about Schiit would be complete without talking about NwAvGuy and the infamous Asgard Incident. Although the latter sounds like a neat title for a modern dark fantasy/conspiracy movie, it wasn't a joking matter at the time.

If you want to look at this from a business perspective, this is about how you handle adversity. I've said it before, but it bears repeating: even if a Bernanke helicopter drops $100 million in your lap with a note saying, "startup capital for your company," business is a series of challenges. Some of these are gonna be company-changing. This one was.

But first, the scorecard: we're about 13 months into Schiit at this time. July of 2011. We still have the same three products:

- Asgard
- Valhalla
- Lyr
- And a prototype of the Bifrost (no chassis)

In terms of human years, we're still a baby company. Still working in the garage, still putting everything back in to make more products and develop new ones, still with the 1.2 person crew (if you count 20 % of my time, 30 % of Eddie's time, and about 70 % of Rina's time.)

So, we're cruising along, with Eddie building stuff, me testing, and Rina shipping, and Mike and I working on tweaking Bifrost so we could get it shipping (originally planned for July, but hey, you know, we're late on nearly everything), when we get an email.

This email has a video pasted into it, showing an AKG K701 driver flexing. With the email, there's a short message, asking if this is normal behavior when you turn an Asgard on and off.

Also, although I didn't know it at the time, the emailer also posted the same video on Head-Fi, asking "what causes this?"[1]

It's an insane day, so I send off a quick response: *Yep, it's normal.* Thinking at the time, *Yep, of course it's normal, many amps have transients when you turn them on and off.*

---

1 The whole thread is still available for your browsing: head-fi.org/t/562736/what-causes-this-amp-related

Anyway, back to the email …

The buyer wasn't thrilled by my response, so I offered to let the buyer return the amp for a refund, even though the sale was of a B-stock product, which we don't normally offer a 15-day return on.

Eventually, he took me up on the offer—but not before the Schiit hit the fan on Head-Fi.

Enter NwAvGuy

On the weekend after our buyer posted the video on Head-Fi, NwAvGuy appeared, pronouncing Schiit all a bunch of knuckledragging morons building dangerous amplifiers. Of course, I'm being a bit hyperbolic, but here's the actual first paragraph:

> My professional engineering opinion, having done both amplifier and speaker development, is this is potentially harmful and a serious flaw with the Asgard. The Schiit amps are clearly not designed for good objective performance. And when designers have other goals in mind there are often some serious side effects. NuForce doesn't seem to care the uDac-

2 has serious channel balance problems for example. They claim that bad channel balance was necessary for the best sound at the price. The Asgard may have been designed with similarly misguided, or sloppy, priorities.

Of course, Head-Fi went nuts. I was lucky enough to see this post shortly after it happened, so I jumped in, citing the turn-on transient as a relatively minor one (which is what I remembered from the engineering notebook—in retrospect, I should have just grabbed an Asgard and tested it—but it was a weekend, I was tired, blah blah), defending our design against his assumptions about us putting all the cost in the chassis, and bringing up some reasons why we didn't want to use relays (most of all, contact degradation from outgassing, which NwAvGuy dismissed out of hand, although it is clearly present on the data sheet.)

The war escalated as a Head-Fi moderator stepped in to defend us (he owned an Asgard), and others weighed in. Mike Moffat even posted (as baldr) an emotional tract, calling out NwAvGuy as a coward and a bully, hiding behind anonymity with unknown agendas. Perhaps a bit over-the-top, but I think it's important to say that engineers

are no less passionate about their products than an artist or performer. Attack the product, attack the person.

On Monday, more data came to light. The Head-Fi moderator tested his Asgard and found that the DC offset was more like 1 V to 2 V, not 150 mV. I confirmed this measurement. *Crap*. What now?

The notes on the prototype didn't jive. So I tested a few more units, then went back and tested the prototype. All the same—an order of magnitude higher than I remembered.

The explanation for the discrepancy between what was in our notebooks? I don't know. Maybe I wrote the numbers down wrong. Maybe I was measuring the wrong thing. The bottom line: never assume the old data is right. I shouldn't have. And I shouldn't have dismissed the question out of hand.

But, then again, I had no reason to suspect we had a problem. We'd tested about 1000 Asgards by this time, all through the full on-off cycle, through a single pair of Sennheiser HD 650s. The 650s were fine. No problems. Bottom line on this: don't trust anecdotal results. And don't assume everyone has HD 650s.

Aside: to this day, nobody has ever blown up a headphone with Asgard, at least to our knowledge. Lyr, yes—which is why we added the relay mute to it before the whole Asgard Incident.

After this, I went back on Head-Fi, apologized for my misrecollection, and said we'd:

a) Warn people about the need to unplug and re-plug headphones in the owner's manual
b) Investigate adding a relay mute to Asgard
c) Refund anyone who had ever bought an Asgard, from the start, if they were disturbed about its performance

The third item—an unconditional refund on every Asgard—would not have killed the company, but it would have come very close.

So, why did I offer it? Simple. I was done. I was making no money from Schiit, Centric was going great guns, and along comes an anonymous blogger and calls into question our competence and reason for existing. What other crazies would be coming out of the woodwork, I wondered? When would the next attack come?

Maybe this whole audio thing wasn't such a hot idea, I thought. *Maybe better to pull the plug completely.*

But, I sat down with Mike over dinner at a local Korean BBQ I still remember going there, but I don't remember the food at all, or even if I ate—I was really, really upset). And Mike talked some sense into me.

"This isn't the end of the world. We have Bifrost coming," he said.

"If we ever finish it," I moaned.

"It'll be done. And when it's done, we'll laugh about ever considering quitting." I sighed. "That's easy for you to say."

"No. It isn't." Mike sat back. "Theta had its share of attacks over the years. We were reported for FCC noncompliance by a competitor, and the FCC came in and shut us down. Hell, our offices were broken into, and our engineering computers were stolen. Someone was trying to get our digital filter code. The thing is, these days, the hounds are invisible. You can't touch them. They can drop in out of nowhere, say whatever they like, and have no repercussions."

"So what do we do?" I asked.

"We be better than them," Mike said. "Add that relay. Kill the current run. Don't sell another

Asgard without it. And offer to update everyone's current Asgard."

"The whole current run?" I asked. That was a pretty big investment.

"Yes. The whole current run of headphone-killing amps."

"But they aren't headphone-killing—" I began

"It doesn't matter that they are or aren't. What matters is that everyone thinks they are. Or at least enough people to matter. So, we go above and beyond. And make it good." I nodded. Mike was right.

The next day, I called the boardhouse and had them scrap the current run of Asgards. Then, in a 16-hour fit of engineering, prototyped a relay mute, added it on to the Asgard PCB, and ordered new boards, rush, from the boardhouse.

While I did this, the thread on Head-Fi grew and grew. Supporters, detractors, people bringing up other amplifiers that needed headphones unplugged, replugged, etc, etc.

Finally I was done. That was when I posted this to Head-Fi:

That said, we now understand that the precautions common to ultra-high-end (where it's well-known that turning on your multi-kilobuck amp before you turn on your multi-kilobuck preamp may involve having a very bad day) simply won't fly with inexpensive gear, so we're making the changes necessary to have our stuff be as user-friendly as possible.

Asgards will now ship with the same relay mute as Lyr, when we are back in stock. We will also offer a retrofit relay mute for Asgard and Lyr, for customers who want the convenience, and install it on any current owner's amp for free.

Aside: This offer continues to be in force for all Asgards and Lyrs. Because, even though we've alerted the owners in our database, there are (a) second- and third-hand products out there with owners we don't know about, (b) many owners who didn't want the relay, and (c) our database is never 100 % accurate—people move, etc.

But the biggest thing that happened in this thread was that NwAvGuy got banned.

Now, the reason for his ban wasn't his criticism of our products, but that isn't what mattered. Because the ban happened during the Asgard Incident, and because of what NwAvGuy posted on his blog afterwards, he'll be forever connected to us, rather than the other manufacturers and DIYers whose products he's criticized.

Oh, I don't believe you, you might be thinking. *Head-Fi protects its sponsors, that's why he got banned, end of story.*

Nope. If Head-Fi protected its sponsors, that thread simply wouldn't exist—and it certainly wouldn't exist now, almost three years later. If Head-Fi protected its sponsors, lots of negative stuff simply wouldn't happen. Think back on all the controversy you've seen here. Are the threads still there? Yes.

And—consider this—Jude contacted me, bought our products, and talked about them before we were ever sponsors. This is something Head-Fi does very well—uncovering new, interesting products and getting them out in the world, even if the companies that make them are not sponsors.

Coda: What We Did Wrong ... and Right

What we did wrong:

- Relying on memory of a measurement that was incorrect
- Minimizing a valid question because we're busy/tired/etc
- Not immediately checking the performance of the amplifier

What we did right:

- Throwing away the current run
- Redesigning and shipping ones with relays
- Recalling and retrofitting Asgards and Lyrs, free

NwAvGuy: The Good and the Bad

"So, you guys really, really hate NwAvGuy," you might be saying.

The reality is more complex. NwAvGuy did a lot of good for the industry, including raising awareness about output impedance matching and the importance of measurements. And, as with the Asgard Incident, he helped make Schiit Audio a stronger company. For those things, I'm thankful. And, if he ever surfaces again, I'd buy him a beer.

But, he also brought a lot of absolutism to the fore, like the oft-stated idea that an amp (or DAC) with good measurements is audibly transparent, and cannot be improved upon. The "us vs them" mentality of objective and subjective, audiophile and engineer—the objective-subjective divide widened considerably during, and after, NwAvGuy.

He also assumed a lot, without confirmation. Like our "expensive" chassis. Yes, they look expensive, but they are not. Like our "design by ear" philosophy. In actuality, it's more "confirm by ear." The speculation that we don't have, or know how to use, test equipment. The reality is that we have better equipment than the vaunted DScope.

And the bigger reality is: NwAvGuy, by his own admission, never touched a Schiit product.

Also on the good side, he helped kick the inexpensive headphone/DAC world into high gear with the open-source O2 and ODAC. Although neither of these designs are like anything we'd make, they're very popular.

And that's why you made Magni and Modi, you're thinking.

Actually, no. Modi is why we made Magni, and Modi actually appeared before the ODAC, at least in prototype form (see the upcoming chapter, DAC in a Toilet Paper Roll.) I doubt if we would have done Magni much differently if the O2 wasn't around (more on that in another future chapter.) I guess the biggest difference is in one of philosophy. We have a "live and let live" attitude at Schiit. We don't think we know it all, and we don't believe that our answers are always the best ones. We know how much work it takes to bring something to market, and we salute every company out there. It was only after NwAvGuy, though, that we enshrined our basic principles here: schiit.com/about/principles

So, if someone else can come along, kick us in the pants, and help us make things better, let me say in advance: thank you!

Where Did He Go?

Despite the joking about "Schiit had him offed," the reality is we have no idea who he is, or why he disappeared. But let's have a little fun and speculate, because some of these ideas I haven't seen in other places:

1. **Muzzled by a retainer.** As an audio consulting engineer, NwAvGuy's retainers would have heart palpitations if they saw him attacking other companies. If they found out about it, they may have said, "please stop doing that," without the please, and with a threat of contract termination.
2. **Hidden in plain sight.** Maybe he's now working in the industry, and knows the ramifications of attacking other companies (in short, lawyers, lawsuits, expenses, bad stuff all around, see Apple and Samsung.)
3. **Gone to the subjective darkside.** Maybe he heard some gear that was incredibly magical, but measured like crap. Maybe that rocked his world enough that he'll next be reviewing at The Absolute Sound or 6Moons.

Coda 2: Business Lessons

So, what did we learn from all of this? I can joke and say, "Well, we learned to put relays in everything", but that's not entirely true.

We did learn a lot on the engineering side, and it did push us to put one of the most advanced, analog-computer-style protection system in Mjolnir, and perhaps the most advanced protection

system ever in Ragnarok, fully microprocessor-controlled for all operational and fault states, running proprietary algorithms that both manage thermal runaway and conditions that may cause de-biasing of the output stage (such as when playing loud, compressed music.)

But most of what we learned was on the business side. It reminded us that we don't have all the answers, that we can and do make mistakes, and that we have to stand up, admit them, and make it right when we do.

Because no business, no matter how great the engineers, no matter how skilled the production team is, no matter how solid the logistics guys are, no matter how enlightened the management is, is infallible.

You screw up. Bad things happen. And you make them good.

Chapter 13

“Isn’t the Symbol for USB the Long Flat Rectangle?”

Strange title for a chapter, right? It’ll become much more clear—and much funnier—later on.

This chapter is really about three things:

1. Transitioning from a “headphone amplifier company” to a “DAC/amp company,” and, eventually, into an “Audio Products” company.
2. The difference between a hardware company and a software company, and some of the decisions you have to make if you’re going to be both. Plus a little primer on firmware/drivers/software.
3. Products that really, really, really, really don’t want to ship on time. Or so it seems. This ties into the “development time is proportional to the square of complexity.” Or incompetence. Or both.

I’ll skip the now-usual summary, because we still have the same three amps, and we’re the same basic company: a fast-growing headphone

amp manufacturer which produces (relatively) low-cost, high-value products in the USA.

I'd announced the Bifrost DAC on June 30, 2011, with a delivery date projected to be in August. Because we'd been running the prototypes for a while, I was really confident we could ship the Bifrost early, and was looking forward to a July launch. I wanted to make sure we shipped everything early from now on, like Lyr.

Now, at this point in time, we were still taking pre-orders. So, starting June 30th, 2011, prospective customers were able to order a Bifrost. This immediately caused several problems, as exemplified by these questions:

- **Hey, I just bought a Bifrost and an Asgard, do you hold them and ship them together, or can I get the Asgard early?**
  Problem: this broke our shipping quote system, which was still in the archaic realm of "guess a range, and hope you don't lose too much money on it, or overcharge so much that nobody buys."
- **Hey, if I get that Asgard earlier than my Bifrost, when does the 15-day money-back start?**
  Problem: Oops. Kinda broke that model, didn't we? We didn't have a great answer

for this, and ended up making all sorts of extended accommodation deals. Which gets fun if you're trying to refund money, say, after 60 days.

- **Hey, do you have an exact date when you're going to ship, I need to plan … ?**
  Problem: No, of course not. We didn't even have metal. So what did we do? We said we'd probably ship early, in July. Very dumb. Do not do this.
- **Hey, do you charge my card before you ship?**
  Problem: No, we never charge before we ship, but the problem is—the authorization on the card expires. Go back and try to capture transactions that were authorized a couple of months ago. Guess what? Sometimes they don't re-authorize (and, if you're using PayPal, you can't re-authorize them after 29 days at all, game over, contact the customer, send them an invoice.) Customers don't like these complexities.
- **Does Bifrost work with the GigaComplex Server X-2000 running a modified Amiga operating system with Rubycon Black Gate capacitors and opto-coupled USB 2.2 outputs?**
  Problem: Yeah, I'm having a little fun with this,

but you get the drift—lots of people asking questions about Linux, phones, etc. We'd only tested on Macs and PCs using the three most popular PC OSes at the time (XP, Vista, 7). And that was a hell of a lot of testing. So we really couldn't say. It wasn't until later, when Tony arrived, that we had a chance to start looking at Linux and phones and Chromebooks and stuff.

Which is a great segue into the difference between hardware companies, software companies, and what happens when you have to do both. So let's talk about that for a bit.

Hardware, Software, and Restaurants

Restaurants? Yes, restaurants. As in, places you go where they make food for you. Although I'm a foodie, I have exactly zero desire to ever open a restaurant (even though I've had some neat ideas. Why? Because restaurants combine the problems of manufacturing with the problems of service with an extra problem of *the stock actually goes bad.* No thanks. So, a tip of the hat to anyone who can make it as a restaurateur. That's a helluva business.

What does that have to do with hardware and

software? Well, the problems with hardware are largely a matter of manufacturing (buying/holding stock, assembly, margins), while the problems of software are largely a matter of service (finding and keeping a team happy building great code, aftersales questions, dealing with a constantly changing OS and software regime.) Of course, I'm oversimplifying, but you get the general picture.

Hardware Business Problems. When you're making hardware, like analog headphone amplifiers, your problem really comes down to this: can you make enough margin on what you're making to run the company and keep it growing.

Because, when you're talking hardware, you're talking stock. You have to buy chassis, transformers, boards, transistors, ICs, and a myriad of other parts before you can ship anything. That means there's a lot of money up front in hard bits. Then you have to assemble it (with either in-house or contract assemblers), test every product, ship it, and support it. If you're doing it right, your repair support will be minimal—your failure rate should be much less than 0.5 % (ideally 10× lower than this.) And, if the products are simple, your tech support should also be minimal.

So, for parts, you could be looking at 33 % to 50 % of the cost of your product. So if you're making a run of a hundred $ 1000 amps, you'll be out $ 33k to $ 50k at the start. Then you have the added costs of personnel, facilities, support, and shipping on top of that.

Note: never discount the importance of shipping—this is not a trivial task, and it can get very complex. Use "free shipping"—actually, "cost-inclusive shipping," to be more accurate—at your own risk.

So, in hardware, before you've shipped anything, you're out a sizable investment. Moreso when you find out the parts you spec'd aren't available, or you ordered the wrong parts, or something's wrong with them, etc. But once it's shipped, you should be able to make some money, if you're doing it right. And you can keep doing it, as long as there is market demand—but of course every product you make has a hard cost.

Software Business Problems. Software's a bit different. Your "hard costs" are actually salaries for the programming staff, your facilities costs, your admin costs, etc. These aren't tied to the amount of software you sell. If you spend $ 50k developing software, it doesn't matter if you

sell 100 copies or 1 000 000 copies at $ 99 each, because you're going to be distributing it via download or via inexpensive media.

But—and here's the big but—you have to support it. Unlike simple hardware, you have to assume that it won't install on some systems, it will conflict with other software, many customers may need some hand-holding in using it. And, as an added bonus, the OSes and other software changes from year to year. This means that what worked yesterday may not work today, and vice-versa.

What this all adds up to is that software is an on-going business—you need to keep programmers constantly employed and engaged, you need to test your product against changes in OSes and software, you need to issue updates to deal with incompatibilities, and you need to have a significant service staff to provide technical support. So, with software, even if your costs are not directly related to how much you sell, you have to sell enough to cover your costs and make a profit.

Note: of course, there are different models (open-source, SaaS, etc) that I'm not covering here, and admittedly I've never had a software company, but the principles are somewhat simi-

lar on the agency side (heavy service, overhead unrelated to amount of revenue.)

So, back to resturants. They have hard costs (ingredients), service costs (chefs, servers, etc) … and as added bonuses, their stock goes bad over time and they have additional regulation (liquor licenses, inspections) and accidents (sick people) to deal with. So, no restaurants in our future. But I'm certainly hoping others are up to the task.

What the Heck Does This Have to Do With Anything?

Yeah, I figured you'd ask that. What it has to do with Schiit is that, with the launch of Bifrost, we became a software company by default. A very, very lightweight software company, yes, but we had new complexities of firmware and drivers to deal with.

"Oh, big deal, you get the Giant Baby of the Year Award," say the hardcore software developers now.

Yeah, and that may be true. But by using firmware to run the Bifrost, and planning to release drivers for Windows, we were crossing another invisible

line in business: from a pure hardware company, to a hardware/software company.

And yes, a driver isn't the biggest deal in the world. But with this tiny piece of software, two things should be noted:

1. 80 % of our tech support is for Windows driver issues.
2. The have been updated no less than 6 times since we (finally) launched Bifrost in October 2011. The biggest change, of course, came when Windows 8 was released.

So, perhaps not a big deal in software-land, but a big deal for a hardware company. Luckily, we knew what we were getting into, and had long discussions about leaving Bifrost 24/96 only, which would eliminate the need for Windows drivers. The discussions went something like this, usually in a summer-hot garage:

"Mike, we need to offer 24/192 support. There's 24/192 music available on HDTracks," I told him.

"What, seven tracks of it?" Mike sneered.

"It's limited, yes, but people are asking for 24/192 support. And it would be a good differentiator for Bifrost, now that we've nixed the balanced outputs."

“Balanced should only be hardware-balanced,” Mike pontificated. “Two DACs. Summed single-ended. That’s what we did at Theta.”

“I know, but don’t change the subject. 24/192.”

Mike doesn’t change course fast, though. “If we’re going to do balanced, we’re going to do it right. Hell, Bifrost would sound better if we used two DACs per channel.”

“Yeah, and it would cost $700.”

“But it would sound good!” Mike insisted.

“Yeah, and you can do that one later. For now, Bifrost. 24/192.”

“You don’t have any balanced inputs on our amps anyway,” Mike said. I frowned. He had a point there. Balanced inputs—done right—required a 4-gang potentiometer, which we didn’t have back then. We had grand plans, sure, but no balanced inputs yet. “Mike. 24/192. Must have it.”

“24/192 uses a half-rate master clock to the DAC,” Mike said. “24/96 probably sounds better.”

“Is this a limitation of the AKM?” I asked.

“No, it’s a limitation of most delta-sigma DACs. The master clock can only be so fast.”

“I still want 24/192. Period.”

Mike sighed. “Who’s going to do the tech support?”

“Me, for now.”

“You’re going to want to shoot yourself,” Mike predicted.

“Maybe. But we need 24/192. We’ve tested 24/192. It works. Let’s support it.”

“Ohhh … kay,” Mike said.

And that’s how Bifrost got 24/192 support. It seems funny today, with 32/384 or even higher sampling rates. Not that there’s any PCM music available there, but hey, it’s like megapixels. Meaningless numbers to use in marketing. Buzzword compliance.

Note to self: we should do a 32/384 DAC that has a switch for “easy mode,” supporting 24/96 without drivers, and “expert mode,” where you’ll need Windows drivers. Except unlike everyone else, we’ll tell everyone why 32/384 is meaningless. See a couple of chapters back.

Hardware, Firmware, Software

With Bifrost, though, we didn't just have software. We had hardware (the DAC—chassis, motherboard, USB daughterboard, DAC/analog daughterboard), firmware (for the motherboard and USB daughterboard), and software (Windows drivers.) This was a level of complexity higher than anything we'd ever done before. Still not insanely complex, but complexity was one of the reasons we were late.

Yes, I'm getting back to the late part. But first, some of you are probably saying, "Software? Firmware? What the hell is the difference, anyway, and what do they do?"

Firmware is embedded code you don't expect to change very much, if at all. Examples of this include the code that runs the LCD display on your refrigerator, or makes the buttons on your sprinkler controller work. Firmware runs behind the scenes, usually never needing updating.

Software is installed code that can change quite a bit. This is everything from your nVidia driver to your copy of Microsoft Office. You update them from time to time, or even switch to different versions.

In the case of Bifrost, we had to develop firmware to run the motherboard, and modify firmware to run the USB input.

The motherboard firmware runs the front panel button, the bitperfect clock management system, and the hard mute protection system. Not very exciting—but necessary if you are going to have clock management like Bifrost. It's programmed via an Ethernet jack inside the Bifrost (intentionally inconveniently placed, because we don't expect to ever update it … unless there are major changes to the DAC/analog card or something like that.)

The USB firmware runs the CM6631A USB receiver, generating the specific clocks and formats that we need for our D/A implementation. It's programmed via the USB port.

"Aha!" some of you are saying. "I can update the USB firmware on the Bifrost via the USB port? What kinds of cool firmware do you have?"

Not so fast. Yes, you could, but there's really only been two versions of firmware—one for the original USB input, one for the Gen 2 USB input. There are a ton of other complexities as well. It's easy to brick a USB input by playing around with the firmware. So, just don't.

Note: We'll know if you change the firmware, because the programmer reports the current version of firmware on the card.

And firmware is where we get into the genesis of the title of this chapter. But let's talk about problems some more, first …

Products That Don't Want To Ship, Like Bifrost

Okay, back to Bifrost problems. Announced in June, stated to be shipping in August. When did it ship? Late October. Doesn't sound like much, but when you're already ordered into backorder before shipping, and you have hundreds of angry customers yelling at you—some of which waited almost 4 months to get theirs—it's a very big deal indeed.

Aside: Haven't I said something before about opening your mouth and pre-ordering? Well, it took us past Mjolnir and Gungnir to finally learn that lesson. Mainly because, "Nothing could be as bad as Bifrost." Yeah, right.

So, what went wrong? In this case, mainly mechanical problems, but there were electronic ones as well.

First, it was the DAC/Analog board. My first attempt at a hardware summer (from the DAC's differential output) was simple, elegant, good-sounding—and completely useless, because it would have fixed Bifrost at 1 V out, rather than the industry-standard 2 V RMS. So I had to go back to the drawing board, for a completely different topology.

Then, the way we were going to attach the daughterboards didn't work out. We'd expected to use press-in connectors, like I did back in the Sumo days. Those didn't work so well, because the Bifrost daughterboards stood a lot taller than the Sumo ones (they had to clear parts underneath.) We played around with various plastic options before saying, "the hell with it," and using metal standoffs and screws instead, to ensure the boards would stay in place during shipping.

Next, the board interconnections themselves. Now, there are plenty of header options out there, but not at the length we wanted to use. Which meant Bifrost headers were custom. Which meant an 8 week lead time. Which we hadn't counted on.

Note to other guys who want to start hardware companies: check the lead times. Then

add a few weeks. Some parts are 14 to 16 weeks. And sometimes they don't come on time.

Then the metal came. Unlike Valhalla, it was beautiful … but it didn't fit. I'd gotten used to getting the metal "done in one," and we hadn't done a first article for fit. So, now we were sitting on hundreds of chassis that didn't fit … and we had boards already out for assembly.

This was by far the biggest problem. Because when you're in a situation like this, you either have to redesign the board (and throw away the assembled ones) or redesign the chassis (and scrap the chassis).

In the end, we did a modification to the inner chassis that allowed it to work, then made a longer-term change—the only running change we've ever done to metal.

Another note to other guys: don't change the metal if you can help it. Having two (or more) versions of the same chassis plays hell with production—as in, "Hey, we found a box of chassis, but the boards don't fit."

In-between the non-fitting metal and the stuff that worked, we wasted a ton of time trying to

figure out how to make it work. Which is dumb. Metal, especially cosmetic parts, is fixed. You can't just "oval out" a hole in a cosmetic part, or cover up things with plates, or any other of a dozen stupid ideas we came up with. Throw it away and start over—or fix the board to fit the metal. Period.

And, when the modified metal came, we found that we'd inverted the left and right designators on the output jacks. Argh! Back for rescreening.

In short, Bifrost really didn't want to ship. It was a product from hell. And it's what we get for trying to take shortcuts, like no first articles, combined with the most complex product we'd ever made.

But, by RMAF 2011, we had a working Bifrost, in a cosmetically perfect chassis. So, what did we do? We took it to the show, of course!

This time, we weren't exhibiting with Sennheiser. We'd grown up enough to have our own booth! But we hadn't grown up enough to do anything other than take the products there. We had no signs, no literature, nothing. Lisa clipped a couple of 8.5 inch×11 inch printouts of the logo to the curtains behind us, but other than that,

there was no indication of who we were. Great marketing.

That was the last show we did like that. I knew we looked bad, and I vowed to change it for the next time.

It was also the last show that we took all the products in a single backpack. Literally. Asgard, Valhalla, Lyr, Bifrost, a pair of LCD-2s and a couple other cans—all in Rina's backpack.

That was her idea. It would save on shipping, and there was no chance they'd get lost, she said.

"So what happens when Homeland Security takes one look at the X-Ray machine and stops you for having 40 pounds of electronics in your bag?"

"Then I smile nicely at them and play dumb." I sighed. "Then I'll tell them I don't know you."

"It'll be fine," she insisted.

And—frighteningly enough—she was right. They didn't even blink at LAX. Didn't stop us. Didn't ask about what all the electronics were. Not a single question.

Nor did they stop her, or ask anything, at Denver International, when we were coming back. Makes you feel really good about flying. *3.5 ounces of*

shampoo? Bad boy, go to the little room. 40 pounds of aluminum, steel, copper, transformers, and wiring? No problem, move along.

Isn't the Symbol for USB ...

Now, since we had a cosmetically perfect Bifrost, it meant we could ship other cosmetically perfect Bifrosts, too. That day we put together the first ten Bifrosts, I felt about 1000 pounds lighter—and about 10 years younger.

At least until the next day, when we had probably 20 more.

Rina was listening to the Bifrosts, prior to cleaning, bagging and packing. She knew what an ordeal it had been getting the product to market. She knew all of our little (and big) frustrations. She knew what a sore point it was.

So, it was perhaps with a bit of fear and trepidation that she said, "You're always going on about how Mike did the first DAC, right?"

"Yeah, yeah," I said. "So?"

"But this is his first DAC in a while, right?" Rina continued.

“Yeah,” I said, wondering what she was going on about. “So what?”

Lisa sighed. “Isn’t the symbol for USB the long flat rectangle?” She pointed at the icon on the front of the Bifrost to clarify.

“Of course it is!” I said, thinking, *That’s a dumb question.*

“It’s not the round circle?”

“No, that’s the coaxial input. Why?”

“Because this Bifrost is playing USB when it’s on the round circle.”

What? I stomped over to take a look. She was right. It was happily playing music from the USB input, with the front panel LED indicating that it was on the coaxial input.

My heart skipped a beat. “Are they all like this?” I asked.

“I think so,” she said, and grabbed another.

In a few minutes, we confirmed: yes, they were all like that. Including the show Bifrost. Which we’d never noticed.

Great, just great.

Now, I knew what was happening. The front panel LEDs were controlled by the microprocessor. It was simply lighting up the wrong LED. It was nothing that changing the firmware couldn't fix.

But … changing the firmware took Dave. I had to alert him and get a new copy of the firmware. Which only took hours. But then we had to take apart every Bifrost so we could get to the Ethernet connector and re-program them. Again, not the end of the world. But Bifrost had to get that one last shot in, before it would happily ship.

And … *what about all the other Bifrosts we shipped yesterday?*

Yep. You got it. They all had the wrong firmware. We offered to replace it, but many customers didn't want it changed … as if it was a Bifrost Special Edition or something. By my estimate, there's 6 to 7 Bifrosts out there still with reversed LEDs. As well as my personal Bifrost, the original show unit. I've simply never bothered to change it.

So, if you buy a used Bifrost, and see that the LEDs are wrong, let us know … we'll change it. Or not. You have a piece of history.

Chapter 14

Technical Help Via Time Warner, and The World's Most Irritating Failure Mode

Sharp-eyed readers will note that I swapped the position of this chapter in the book, moving it up before "DAC in a Toilet Paper Roll." This isn't because I'm a terrible person and want to keep you in suspense for another week. It's because this chapter really comes before the next one, chronologically.

"Well, ya coulda made the outline right," quips someone.

Yeah, and I was probably drinking when I did the outline, so there you go.

In any case, this is the story of Tony, our second employee and first technician. It's also a story of an amazingly hard-to-diagnose production problem that, to this day, is not fully explained. You'll see why soon, but first, the wrap-up:

- At this point in time, it's late October 2011
- We have 4 products: Asgard, Valhalla, Lyr, and Bifrost
- Mike has yet to bring me his DAC in a Toilet Paper Roll and change the course of the company
- But, we are both thinking about what's next … the bigger badder as-yet-unnamed balanced amp and DAC
- We're still in the garage, though Eddie has been banned from smoking by now
- We've re-organized the garage to give Eddie more space, and a real table to work on—which meant moving the Corvette to Mike's extra garage space for a while

And, even before we brought out Bifrost, it was clear I couldn't do all the testing and repair. At least not do it, and remain sane. I was running the marketing company full-time—and, with it being busy, sometimes more than full time. I was also contracted by Penguin to write a couple of crazy books about giant robots.

Yes, I am an idiot. But hey, might be fun, right?

Yoda and Rain Man

In any case, this is how we got Tony: me freaking out from insane time-pressure, and Mike stepping in to help with a suggestion. Now, for all of you who think Mike is a Yoda-like sage who comes in with words of wisdom and a perfect plan, consider this conversation …

"Mike, I can't do this all," I told him. "Between Centric, the books, and the orders ramping up, I'm doing 16-hour days."

"Well, we should look for a tech," Mike told me.

"To work on a garage?" I said, doubtfully.

"Eddie does," Mike shot back.

"Eddie is special," I said. And he was. Eddie had long worked in the entertainment biz, so he was used to weird hours and spotty schedules. And he knew us. And he liked the free food.

Mike's expression brightened. "I think I know the perfect tech for us."

"Who?"

"Tony," Mike said. "He used to work for Time Warner Cable. He did all the tough installs that nobody else wanted to do." I shook my head

doubtfully. “Mike, a Time Warner tech is a little different—”

“He’s a little like Rain Man,” Mike said. “But once you show him something, he never forgets it.”

Rain Man? I thought. The image of a Time Warner-uniformed, confused Dustin Hoffman flickered through my mind. It wasn’t a reassuring thought. “Um, Mike—”

“He’s a good guy, I’ll call him,” Mike said, whipping out his phone.

“But he’s still gotta work in a garage,” I protested, as Mike put the phone to his ear.

“No problem,” Mike said. “He’s my stepson.”

Oh, great. Better and better. “But Mike, we can’t guarantee hours, and we don’t have healthcare and stuff—hell, we don’t even have real employees—isn’t he better at Time Warner?”

“Doesn’t matter,” Mike said. “He’s been laid off.” Then, into the phone: “Hey Tony, want to come down and work on some Schiit?”

“Mike!” I snapped

Mike held up a hand. “Yeah. Sure. Now is good. I just spent half an hour listening to my girly

business partner talk about how he's overworked. Yeah. Here's the address … "

When Mike hung up, I grumbled, "Unemployed Rain Man stepson … does he even know how to solder?"

"Probably," Mike said. "And if not, you only have to show him something once."

"What if he doesn't work out?" I asked.

"Then he doesn't work out," Mike said. "Look, I'm not trying to stuff him down your throat, but you're about as bad as me at admin stuff. I figure, Tony's available, let's give him a shot."

"But … Rain Man?"

Mike grinned. "Yep!"

The Tequila Bottle Theory of Hiring

Okay, you don't need to be a brilliant CEO to know that what we just did there isn't how you hire people. Not if you want to be successful. Or at least not by the rules of any manual on *Getting The Best And Brightest,* ghost-written for a CEO dropped into an established $10B company by his friends from Harvard.

Because, of course, you have to have a rock-solid vetting process that includes questions like "How many marbles fit in an average toilet?" and psychological profiling to determine if the candidate is a right fit. You have to set Goals and Responsibilities and identify Key Result Areas. And you must carefully analyze the results in order to determine the best possible candidates, then offer them a can't-refuse package that includes all the latest perks, from a fitness club to a concierge.

Or you can go play some basketball with the founder and get a job offer, or have gone to school at Stanford with the CIO, or have compromising pictures of the CMO with a chicken. The real world doesn't always work by the numbers.

In my businesses, I've hired probably a couple hundred people over the last 20 years. In the early days, I went by the formula and the checklist. I agonized over who to hire. And, a lot of the times, the formula won over my gut.

And every time the formula won over my gut, I screwed the company.

Because people can't be distilled down to a 2-page resume and a 1-hour interview. There are a ton of candidates skilled in the art of looking good

on paper. There are plenty who can be friendly, affable, and make all the right responses to the standard interview questions.

And yet they can still fall on their face. Because it's easy to recite a formula, but a lot harder to deal with the unpredictable Real World. People who interview well usually fall down at one of these things:

- **Only Doing it By the Book,** or "This is how we've always done it in my previous jobs." Yeah, that's fine, but the world is changing. Open your eyes. If you're incapable of adapting, then it's going to be a bad day.
- **Stunning Lack of Initiative,** or "Well, you didn't tell me to do that syndrome." Look. You're not computer software. You're a person. Small businesses don't have time to program your every move. You know what needs to be done. Take some initiative.
- **Prima Donna-itis,** or "You're so lucky to have me." Hey, I know you're a rockstar, but that doesn't mean you need to redecorate your entire division before you do a single ounce of work. And no, we will not sort your M&Ms. You're here to work. Get over yourself.

Your standard interview—or even the creative Google-esque interview—isn't going to identify

those kinds of people. They're not even particularly good at identifying the go-getters, unless you get off the script and ask them *why* and *how* they did something, rather than just focusing on *what* they did.

For example, let's look at this scenario:

- **Laid off from last job.** Big red flag, in traditional hiring. *Do not interview.*
- **Laid off from last job**—for recognizing the inefficiency in the ordering system and building a custom database to automate it, thereby eliminating their own job. *Hire this person. Immediately.*

Today, when I'm hiring, I don't ask any of the traditional questions, or the stupid trick questions, or give people tests. I just sit and talk to them, usually about what they're most interested in—which may or may not be work-related. Because you won't get to really know someone until they're comfortable talking with you.

> I joke that my ideal interview would be sitting down with the interviewee and a bottle of tequila. But of course we can't do that.

The results? Today, we have a lot less churn in the business, and a lot more long-term employees.

Selecting people with potential and ambition beats experience every day.

Tony and the Popping Lyrs

So where does that put us with Tony, you ask?

Well, Mike gave me a hint at his potential with his comment about "he used to do all the difficult installs." That's not the behavior of someone who just wants to collect a paycheck.

And Tony was excited to work with us. He'd been hearing about Schiit from Mike for a while. And even before he was laid off from Time Warner, he was done. He wanted to be part of something where he could make a difference.

And Schiit, yes, he could make a difference. Mike was right. Tony didn't take long to train at all. It didn't hurt that most of the stuff we were making was pretty easy to qualify, and it didn't hurt that our boardhouse's quality is very good. By putting Tony in the mix, all I had to do was look at a few broken products from time to time. My schedule was back to a manageable level, and everything seemed to be going well.

And, when we introduced the Bifrost, Tony loved it—because it involved programming and com-

puters. Tony is our resident Android/Linux guy, though he also has PCs and Macs as well. He's the guy who qualified our DACs on Linux, and he's the guy who knows the most about how to deal with, say, failed Windows driver installations.

So everything was going great, until we get the next run of Lyrs in from the boardhouse.

Tony came into the house, jiggling a finger in his ear. "The Lyrs are popping," he said.

A small pop is normal as the relay engages," I told him.

"No," Tony said. "Popping. Blow out your ear popping."

"It's probably a bad servo. I'll look at it later."

"They're all doing it," Tony said.

"All of them? And they've been through the pre-test?"

Tony nodded. "Yep. I pre-tested them, pre-biased, then Eddie put them in the chassis, and I set the bias again, just like you showed me." I frowned. If they passed the pre-test and biased OK, they pretty much had to work. We hadn't changed anything since the last run of Lyrs, and those had worked flawlessly. I went out to the garage to

have a look. I grabbed a Lyr—one labeled with a post-it note saying, "popper," took the back chassis off so I could probe around and see what was happening, and put it on the bench.

It tested perfectly. Right gain, right bias, right voltages, THD looked fine, no problems anywhere.

"Are you sure this is a popper?" I asked Tony.

"Yep!"

"Check another one, then give it to me."

Tony ran another Lyr through the sound check. As the relay engaged, a loud POP! came from the open-backed HD 650s—loud enough to echo in the garage.

Tony winced, taking the headphones off his head.

"You know you can test it without having them on your head," I told him. "Or, better yet, run it on the scope."

"Okay," Tony said, handing me the Lyr. I took the back off and tested the "confirmed popper."

Like the first one, it measured perfectly. *What the hell? I'd just heard it blast a headphone so loud it was amazing the thing still worked.* I turned off

the Lyr, then put the scope in one-shot mode to see what kind of DC spike was happening when the relay engaged. A few millivolts. Nothing to worry about.

“What is it?” Tony asked.

“I don’t know. It’s fine. No problem.”

“You just heard what it did,” Tony said. I nodded. Yep, I’d heard. I power cycled it a couple more times, and the Lyr continued to behave. No problems at all.

“I don’t know,” I told Tony.

“Maybe it fixed itself.”

“Maybe,” I said. Knowing things never fix themselves. I buttoned the Lyr back up and handed it back to Tony for sound check.

BANG! Another explosive pop reverberated through the garage.

“Is that the same one I just gave you?” I asked Tony.

“Yep!”

“You sure?”

“Yep!”

Great, I thought, *an intermittent problem.* Those were the worst.
I took the back off the Lyr and ran it through its paces again.

No problem. No DC spike. Nada.

This made no sense! There wasn't anything different about our test rig and the listening setup.

Except—*I'd taken the back chassis off the Lyr.*

Nah. Impossible. It made no sense.

Shaking my head, I put the back chassis back onto the Lyr, and plugged it back into the test rig. I powered it up and waited for the relay to click on.

BANG! A huge pulse of not-just-DC, but multi-megahertz oscillation hit the scope, the moment the relay closed.

"Oh, you've gotta be kidding me," I groaned.

"What?" Tony asked.

"I think it only pops when the back chassis is on."

"Huh?" I nodded. "Exactly." I went back to the bench and quickly confirmed that three out of

three Lyrs all worked fine with the back chassis off, and oscillated with the back chassis on were .

"Why does it do that?" Tony asked. I shook my head. MOSFETs are weird. Something as simple as the chassis being close to the output devices might cause enough parasitic capacitance or inductance to make the amp unstable. But that was extremely unlikely. None of the rest of the Lyr runs had shown this kind of instability. It was a no-feedback amp, it didn't have loop problems.

And, the biggest problem? When you have an oscillating amp, you can usually poke around with a finger to see where the problem lies, or at least the basic area. But with the Lyr bolted all the way together, there was no doing that.

"How are you going to work on it, if you can't get into it?" Tony asked, echoing my thought.

"I don't know." I admitted.

The Brute Force Fix

The "how" of fixing Lyrs that only oscillated in the chassis turned out to be a tedious, trial-and-error process of endless disassembly and reassembly. I did all the normal stuff you do when you have an

amp that isn't behaving—additional bypassing, bigger gate stoppers, etc—and none of it worked.

The only thing that worked was to replace the MOSFETs. And, since many of the "poppers" also took out the servo and part of the Dynamically Adaptive stage through oscillation, we replaced those, too.

That finally killed it—the brute-force replacement of about a dozen critical parts.

Why did it fix it? No idea.

Yes, that's right. To this day, I have no idea what caused the problem. That single run of Lyrs is the only run that ever popped. Nothing before, nothing since. No parts were changed, the boardhouse was the same—it should be absolutely, totally the same. But it wasn't. I have some theories, of course, but they're all pretty iffy.

The most plausible idea is that the output MOSFETs were somehow damaged by static. And that's not saying much. Our PCB assembly house is an ISO-certified facility. I doubt if their assemblers suddenly decided to leave off their static straps and wear angora sweaters to work en-masse. And we usually don't handle the parts outside of their own protection.

Or, it could have been a bad run of MOSFETs. The ones we replaced them with had a different lot code. Again, not very plausible.

Finally, tolerance stacking—the boards, the solder, the parts and the assembly were just different enough this one time to make it unstable. Again, not very plausible, probably the most farfetched of all. I suspect we'll never know what the problem truly was, for the most irritating failure mode in the world.

Final word: Tony. To this day, Tony is our lead tech. He's tested and/or programmed tens of thousands of products. Pretty good after Mike's first intro, right?

Chapter 15

DAC in a Toilet Paper Roll

Shortly after we started shipping Bifrosts, Mike brought me something that would change the company.

"Take a look at this," Mike said, handing me a Bifrost USB card. I didn't think anything of it. We were making Bifrosts, we had tons of Bifrost USB cards, they worked, and that was that. And I knew we weren't going to be suddenly introducing a new USB card, only a few weeks after we announced the Bifrost itself.

"It's a USB card. So?" I asked, not taking it.

Mike waggled the card at me. "Just look at it." I took the card and sighed. And that's when it sunk in: *this USB card had RCA jacks on it.* I looked closer. The USB card also had a few more parts added onto it. I recognized one: an AKM4396 D/A converter.

"Is this a USB DAC?" I asked, incredulous.

Mike nodded his head vigorously and cackled. "Yeah, it is!"

"USB powered?"

"Yep! And it sounds really good!"

"Wait a sec," I said. "I thought you were Mr. Anti-USB?"

"Yeah, but sometimes you just gotta say, 'what the heck!'" (Except without the h and e, replace with f and u.)

Aside: this was really the beginning of Schiit's ongoing "WTF" phase, where we'll try a lot of different things—and, if we like them, make them into products. This is what got us Mjolnir, Vali, the upcoming Yggdrasil analog stage … as well as a shelf full of experiments that may never get turned into products.

Anyway, back to the frankensteined USB card. As I held it, three thoughts immediately came to mind:

1. How inexpensive could this thing be?
2. How do we keep the cost of the chassis from dominating the cost of the product?

3. This is a whole lot smaller than anything we currently make, what would the chassis look like?

"What's the BOM look like, cost-wise?" I asked Mike, using corp-speak for the Bill of Materials, or, in English, *All that stuffs ya gotta put together to makes it.*

Mike cackled again and told me.

My head exploded. If we could get a chassis cheap—I mean, really cheap—we could sell that little DAC for $ 99.

Ninety-nine bones. That's a whole different part of the market. I didn't know how many we could sell, but I knew, even then, it would be a hell of a lot more than Bifrost.

"And we could put it in a toilet paper roll," Mike said.

"What?" I asked, coming back to reality.

"Well, you know, for those guys who want something cheaper than a Bifrost. Here you go. DAC in a toilet paper roll. You don't want to upgrade, you want cheap and disposable, we have cheap and disposable." I had a brief mental flash of a (thankfully) alternate future in which we'd be

selling cardboard tubes with electronics inside of them.

A toilet paper roll wouldn't pass FCC," I said.

"So wrap it in foil," Mike suggested.

"No."

"Or machine it out of aluminum tube."

"No."

"We can cut the little spiral on the outside so it looks like a toilet paper roll."

"No!"

Mike frowned, looking offended. "You don't like my new DAC," he whined.

"Oh no, I like it just fine," I told him. "Assuming it sounds good."

Mike laughed and grinned. "Just plug it in."

So I did. And it sounded good. Really good. I knew right there that this would be our next product. I knew we had to make it. It wasn't a case of 'why,' it was a case of 'why not.'

But not in a toilet paper roll.

The Challenges of a Changed Game

That first modified USB card completely changed Schiit as a company. Arguably, it's the biggest factor in us moving from a "hobby business" to a "real company."

But note when this happened: say, November 2011. Modi (and Magni) didn't show up on the scene until late in December 2012—over a year later.

Huh? Our simplest products took over a year to develop?

What's up with that, you ask. (And some are snickering in the background about how long it took to get a sellable Ragnarok. Hey, bite me. I almost did an April Fools announcement that the Ragnarok and Yggdrasil were cancelled.)

Bottom line: yes, developing even simple products can take a good long time—that is, if you want to get them right. Especially if they're clean-sheet designs that move you into entirely new spaces as a company.

Modi itself was a challenge, on several fronts:

1. **It required an entirely new, and hella cheap, chassis design.**
   When we started the development, I didn't

know if we could meet the very aggressive price point I set—especially without going to China.

2. **It had to be as simple as possible,** which meant surface-mount and very easy assembly. We were already doing surface mount with Bifrost, but it was brand new to us at the time. Chassis-wise, even our insanely simple chassis for Bifrost, Asgard, etc were clunky and slow to put together—we needed something simpler and easier.
3. **It required huge production runs**—in the thousands—much more than we'd ever done before. Huge production runs meant (comparatively) large investment—we had to be ready for it.

So what did we do? Well, let's start with the one thing we *didn't* do: we didn't sit back on our asses and say, "You know, we have a pretty good thing going here. Why take chances? Iterate the current products, keep milking the products for all they're worth, and take the money."

Because, like it or not, that's what most companies do.

They're scared. Risk-averse, they call it, in typical corporate doublespeak. Call it what it is. Scared. You're scared to jeopardize your accomplishments

to date. You're scared of burning your profits on something that might not work out. You're scared to step out of the mold you've made.

Hint: the mold you've made is a coffin. Break out, or die.

"Okay, fine, I get it," you tell me. "Go ahead and write what you did, already."

Fine.

Chassis: Set A Target, Stick To It. On the chassis side, the first thing we did was to actually set a price target. If we could bring it in at or under the target, we had a product. If it came in more than the target, we'd have to think again. This is the first time we'd set a price target on a chassis, instead of just sitting back and saying, "Well, let's see where it lands."

To maximize our chance of hitting the target, we put the chassis design up for bid at MFG.com, as well as with our two local chassis vendors at the time. Note we had two for a time. Note we're now down to one. Corollary: cut your losses quick if things don't work out.

It ended up that several vendors undershot our target price, some significantly. One of the local guys was under the target, and one was over. The

local guy who was under the target price wasn't the least expensive bid, but I knew already not to gamble with quality. We went with the one local guy, even though the chassis could be cheaper. Corollary 2: Don't be frigging cheap. If you set a target price and your most preferred guy comes in under it, don't grind them for the couple of bucks you might save if the long-shot new guy across the country works out. Just place the PO.

Assembly: Simple At All Costs. I'm doing this out of order, because the chassis design comes before the chassis. But you get the picture here. Assembly time is a function of chassis design. The simpler the chassis, the lower the assembly time.

We first experimented with variations on our current chassis— a U-shaped piece of metal wrapped around an inner sled. But, in that case, you were still talking 16 screws or so, between the ones that connected the chassis, the ones that connected the boards, and the ones that, well, connected the connectors. That's a lot of screws.

That's why we decided to have a "sled" design quoted—a new concept that used only a front "L" shape, rather than a U-shape. The main advantage of this is that it eliminated the bottom screws. The end result? 7 screws, not 16. I know,

it doesn't sound like much, but it makes a huge difference in production.

At the same time, we asked our vendor their opinion on aluminum versus steel—we were smart enough to know that painted steel was better, but we needed the vendor to say, "Yes, steel will be less expensive—and it can be repainted if it's damaged."

Refinishing a product was a dramatic new concept for us. As I've mentioned before, our other chassis, at least the aluminum parts, are junk when they're damaged. Being able to refinish the chassis was a big deal.

And that's how Modi ended up in a steel box, rather than in aluminum and steel: simple economy.

Production: Bite the Bullet. Yep, big runs are pricey. There's no way around that. And that took us to another point that could change the company.

When you're talking big production runs, you really have two choices:

1. **Save your own money.** Funny, this is the way that businesses used to do it all the time. Seems it's gone out of fashion today.

2. **Borrow the money.** Go to the bank, get a line of credit, or get a loan against your inventory or receivables. This is what our accountant advised, citing all the normal reasons for getting in bed with a bank:
   a) You're growing fast, this allows you to grow faster
   b) Keep your own money for other stuff, like building spaceship-styled campuses, Porsche GT3s, and vacation homes
   c) At current interest rates, it doesn't cost that much.

Guess what we chose?

Yep, right in one. As in, #1. We chose to save our own money for this, because either (a) we're a little stupid/touched in the head/out of touch with thems modern ways of doin things, or (b) we don't think a bank should pay for our own dice-rolls.

Okay, I'm being flippant. But I have seen what happens when a company gets addicted to bank financing. Once you're on the take, you don't get off. And then the bank comes in and starts dictating what you should sell, and when.

No, thanks.

So, that's my long-winded way of saying that there's no easy way around the capital outlay necessary for large production runs. We bit the bullet and made the investment—an investment much larger than we ever could have made when the company started.

We were growing up.

Ego Talking

But, in my mind, there was something even bigger: Modi needed to make sense … as part of a whole line. Mike's original idea was a tiny chassis (much smaller than today's Modi) that could be used with any of our larger amps. But I had noticed—already—that people were stacking Asgard, Valhalla, and Lyr atop Bifrost … and they looked very good together. Which meant only one thing:

"We need a matching amp," I told Mike.

A matching amp for what?"

"The little DAC. We need a small, cheap amp to match it. $99 as well. A sub-$200 stack of an amp and DAC."

Mike looked thoughtful. “How are you going to do any good headphone amp that cheap?” I waved a hand. “I have a bunch of ideas.”

It turns out I shouldn’t have dismissed the “amp problem,” because that was arguably what set the Modi back a good 4 to 6 months. It took a long, long time, down many dead-end paths, before we had an amp we could pair with the Modi.

But that’s a story for another chapter.

Chapter 16

Growth, Garage Style

You may have seen the video I posted of the early Bifrost era/Lyr debacle era in the garage, and thought "holy moly, these guys are really cramped in there."

In reality, it wasn't so bad. Most nights, there were just two or three of us—me, Eddie, Tony. Lisa shipped during the day, so she usually wasn't in the garage with us in the evenings. And we'd gotten our space efficiency down pretty good, storing some chassis at the Centric office, some in the built-out attic, and some outside on the patio (yes, it doesn't rain much here …)

So we didn't really feel the pinch that much. At the time (late 2011), it seemed like a sustainable business for the mid-term, without taking on separate production space. Still, a little voice in the back of my mind kept whispering, "You'd better start looking."

What did I do? I ignored it.

If I have a failing in business, it's in being too conservative. I'll wait until we're completely overloaded until hiring, until we're so cramped we can't move until expanding, and so on. Part of this is simply being raised by parents who were, well, tightwads/intelligently thrifty/not using the house as an ATM. Part of it is seeing the consequences, first-hand, of what happens when you "hire forward" or "expand forward."

The Penalties of Optimism

In 1999 and 2000, Centric was in the middle of the biggest boom we'd ever seen—the first internet wave. Companies were panicking, declaring, "we gotta get on the web, damn the budgets!" and spending money like, well, like someone had turned on a money shower. We got a ton of work, and expanded like crazy, going from 2800 square feet to 7100, and prepping to hire a lot more people if the boom continued.

Even then, though, I was conservative. We took no business from internet start-up companies that weren't funded (or, in other words, we weren't suckers taking stock options as payment—stock options that soon would be worth less than toilet paper.) We took very little business from internet

companies, period. Because I know what a bubble looks like, and we were most definitely in a bubble. And I wanted to avoid the inevitable bursting of the bubble, when everyone had their website, looked at the cost, passed out, and said, "Never again!"

Hint: in marketing, if corporate says, "We don't care what it costs," there's something very, very wrong—as in, you're panicking for no reason (like the first internet boom, or the social boom, or the mobile app boom, etc) or your CMO is soon going to be shown the door.

Instead, we were doing business mainly with companies that worked on the hardware side of things—from the people who built the process equipment and metrology products to make and measure hard drives, semiconductors and optical fiber, to companies doing optical networking chips and new processes for connecting the "last mile" at high speeds.

> Remember, this is fundamentally the beginning of the DSL and cable era—we were still running a T1 line at the office for internet.

Because we were conservative, and because we didn't take a lot of internet business, and didn't

get big bank loans to expand like crazy, people thought we were, well, kinda stupid. I got comments like this a lot of the time:

> Wait a minute, you're doing web development, and you're only 22 people? What's wrong with you? Why aren't you hiring forward like those guys with 300 or 600 or a thousand people?

I just nodded and explained we were an integrated marketing company that did web development, so we were selective and didn't take all comers, we wanted to grow organically, and we saw the web boom busting soon.

Their response (in 2000):

> What? Are you kidding? Everyone has to get online, I just read a Fast Company article about billion-percent growth rates and so-and-so being bought for a trillion dollars or something like that!

Yeah. When you're in a bubble, most people don't see it. Rah rah, buy everything in site, leverage all, and grow! Because it will never stop.

Uh-huh.

Fun fact: the "you're a web development company, and you're only 22 people, you must be doing something wrong?" comment changed in 2001 to "you're a web development company and you still have 22 people, wow, you're doing well!"

So, yeah. By being picky and conservative, we managed to avoid the web bust almost entirely. We lost only a single client in that morass.

But … we should have been even more conservative. The reprieve was short-lived. Remember I said we were working with a bunch of optical networking and tech players? Well, 2001 wasn't so bad for them, but come 2002, very bad schiit went down. We lost 7 clients in one year, either because marketing budgets were slashed, or they went bankrupt, or simply dissolved. One company had taken $ 200 million in venture capital and never produced a product.

And—these are the "hire forward" companies. These are the guys who had the executive chefs come in and cook for everyone, every day. Who had the fully-stocked juice bars and all-you-can-drink beverages. Who had masseuses. Who had foosball and arcade games and lounges full of multicolored couches and bean-bag chairs.

Who had ultramodern polished-concrete-and-glowing-translucent-walls-of-glass offices. Who had Porsches as the "car allowance" car.

Yeah. *Because it will never end.*

Back to 2011 …

Okay, fine. Enough reminiscing. But you start to see why I didn't want to go out there and sign a lease for office space. Because it could end. Booms die. Competitors come out of nowhere. We could screw up again. It could all just be a fad. And so on.

"What's the big deal?" you ask. "It's just industrial space, it can't cost that much per month."

And no, it doesn't. Not per month. So, it sounds like it's time for a quick primer on leasing space for your business.

Point One: Know what a lease is. Distilling down the 70 pages of legalese, here it is in simplest form: *You will pay us $XXXX per month for XX months on or before this day of the month.*

Point 2: Note the lack of any outs. The lease doesn't give two craps if your business is in the toilet, if your cash flow sucks, if your sales forecast

was wrong, or if you're late on your mortgage as well. Pay us. Every month. Until the end.

Point 3: Subleasing sucks. Someone who's never had a lease before sez, "Well, you can always sublease the space." Yep. Have you ever tried? By the way, you're on the hook until it subleases—and even after. They crap up the place? Your problem, not theirs.

Point 4: You'll have surprises, and they won't be good. Instead of hearing "hey, you get a free month of rent this month," from your landlord, be prepared to hear, "Hey, well, I don't care if the plumbing isn't working, that's inside the building, so that's your problem," or, "We had to refurb all the air conditioning units, here's your pro-rated part of the bill.

Point 5: There's less space than you think. Even before you sign a lease, you'll quickly find that many spaces won't fit your needs. They'll be too big, too small, not air conditioned, in a bad neighborhood, etc. etc. That big long list gets very small very fast—and then forget about leverage on the lease rate.

And—an unwritten fact of life—you'll probably be on the hook for the lease, especially if it's your first business lease. They'll want you to

sign what's called a "personal guarantee." That means that even if the lease is in Arglebargle Inc's company name, you have no corporate shield. Fold up? They come after you. Can't pay? They come after everything you've got.

Leasing a space is very much one of those invisible lines in business. Once you do it, you won't go back. Nor will you back out. So you'd better be damn ready to do it.

So why would you ever lease anything? Because you need the space.

We were just on the edge of needing the space, but we were doing well enough. We could squeak by. And we could rationalize a lot of stuff for the future, like:

- We could move more cars out of the garage
- We could have the PCB assembly house do even more—to the point of completed products
- We could buy a different house with space to build a dedicated assembly area (and we even looked into this, but see "there's even less space than you think" above)

And we had one huge advantage, late in 2011: we ran out of stock all the time. Since we ran out of stock all the time, we didn't have to keep huge stock. That saved a ton of space.

Fun fact: running out of stock and going into backorder became so bad that we actually were out of stock on every single product at the end of 2011. That's right. No Asgards, no Valhallas, no Lyrs, no Bifrosts. The holiday rush crushed us.

Still, I knew the end was coming. Mike and I were talking about the next step-up products, the as-yet-unnamed Mjolnir and Gungnir. I wanted to do a balanced amp, and he wanted to do a balanced DAC. We knew they would be bigger than our current products. Bigger products meant more space.

And—there was Modi, the DAC in the toilet paper roll. And an amp. If I could get one to work, that is. Bigger runs meant more space.

But for the moment, we stayed put, as 2011 ended, and 2012 began.

Chapter 17

Resurrecting the Circlotron and Other Mid-Centuryisms

Fair warning: this chapter's gonna be highly technical. The engineers and technology-minded in the audience are probably going to love it. For everyone else, it may be a little hard-going. However, there's a lot of useful information here that might make some of this "Class A, JFET, circlotron, etc" stuff a little more understandable, so you may want to have a look.

Let's set the stage first. This is the beginning of 2012. We're still in the garage, we're still selling the same 4 basic products: Asgard, Valhalla, Lyr, and Bifrost. Mike has a mini-DAC prototype that I know I need to design a little amp for. But before that amp, I wanted to do a be-all, end-all balanced amp design that was substantially more ambitious than what we'd done before.

"Balanced." "Substantially more ambitious." Yep, that's about as much of a design brief as I had,

when I started designing Mjolnir.

In the real world, a design brief can be tens to hundreds of pages, spelling out everything from measurement and power output goals to detailed feature sets, form factor, cosmetics, producability requirements, and company best practices. But that's a lot like telling an artist, "Sure, do anything you want. However, it has to be done in soapstone in a Bauhaus style, not more than 360 total square inches, with symmetry appropriate to production using not more than a 2-piece silicone mold. And it has to be orange." And, in our opinion, it's why you get a lot of stuff that looks a lot like the last product the company made, with nice fit and finish, but no real surprises. No stunning advancements.

And so, in the spirit of exploration that Mike started with the Modi, we set out to design an end-game-worthy balanced amp with really only five words in mind.

Whys, Wherefores, and Design Goals

"Balanced," someone is probably saying. "Why did Mjolnir have to be balanced? I've heard that single-ended amps can be better than balanced, I was told balanced is a scam perpetrated by

incompetent engineers, most headphones aren't balanced, etc."

Okay. Here's the deal. Everything has advantages and disadvantages. *Everything.* The Ferrari you paid $ 400k for is going to be hard to get into and out of. The clutch packs on the robotic manual transmission will need adjusted every 8000 miles. The brakes will cost $ 20 000. The parking lot attendants will hoon it around. People will make jokes about compensating. Blah blah, woof woof. Guess what? If it's right for you, it's right for you.

> And, you know what, same goes for Lexus, Tesla, Mercedes, or any other car. All have their advantages and disadvantages. There is no perfect solution.

So what does this have to do with balanced versus single-ended? It gets down to tradeoffs. Single-ended has some advantages, and balanced has some advantages, too.

Single-ended advantages:

- Can have lower noise (only one active stage and ground)
- Simpler to design and implement (usually)
- Easier to connect (most headphones are single-ended)

Single-ended disadvantages:

- Ground path management can make or break the design
  (ground is not perfect 0 V—where do the currents run?)
- High rail voltages required for high power (runs into
  device limitations)
- Balanced input is problematic (feedback connects to one side)

Balanced advantages:

- 4× the power for the same rail voltages
- Rejection of common-mode noise on the input (if differential)
- Elimination of most ground path problems

Balanced disadvantages:

- More parts, more complexity (except for maybe circlotrons)
- Gain difference between phases if 2-stage amp is fed balanced or SE input
- Most headphones ain't balanced, duh

But, you know what? The main reason we went with a balanced design is that, in our experience, balanced designs offer better sonic performance than single-ended designs … as long as it is a purely balanced design.

Yes. A subjective reason.

Aside: Note the "a purely balanced design." There are tons of amps out there with balanced input and output connectors that aren't "pure balanced." They take a balanced input signal and sum it back to single-ended, so they don't need a 4-gang volume pot. Or they just hang single-ended outputs on a Neutrik 4-pin connector. In our opinion, these aren't really balanced amplifiers, and performance of a truly balanced system shouldn't be judged by their capabilities.

Escalating the Headphone Power Wars?

Funny. Some have accused us of firing the first affordable shot in the headphone power wars (Lyr) and escalating it with Mjolnir and Ragnarok. But there never really was a headphone power war at all, at least in our minds. I never had a power output goal for Mjolnir, other than, "About as much as the Lyr would be fine."

"Well, the fact is, you do a bunch of high-powered amps, and most headphones don't really need it," someone says.

Yes. True. And that comes down to the First and Second Laws of Audio, as esposed by John Chen

at Grado:

1. You can never have too much power.
2. See the first law.

Now, I'm being flippant here, because the law should read something more like, "More power is always a good thing, as long as there aren't any tradeoffs to get that power."

Yeah. More tradeoffs.

The tradeoffs for more power are usually:

1. **Higher noise** (higher power=higher gain to reach that power output, so higher noise.)
2. **Greater need for protection** (muting, DC detect, automatic shutdown, etc.—although a 100 mW op-amp headphone amp can destroy a headphone if it lets go, something with 100 000 µF of filter capacitance and 50 V rails is gonna be very, very bad news)
3. **Paralleled output devices** (really only in the speaker amp realm here—bottom line, paralleled devices, even matched ones, are never quite the same)

From the start, Mjolnir was going to be a high-power amp, so we were aware of the tradeoffs. #3 didn't apply, but we paid a lot of attention to #1 and #2. Because, after all, we were in the middle of the New Orthodynamic Revolution,

and many of those orthos weren't that efficient. Also, many great headphones were still 300 Ω and 600 Ω models, which need plenty of voltage to run them. Both mean big rails, and a big amp.

Today, orthodynamics are actually becoming more efficient, so the need for extreme power is abating. The headphone amp power war, which never really existed, will probably seem pretty silly in a few years time.

And yet still … I'll take the high-powered amp.

Onto the Circlotron

So why was Mjolnir a circlotron from the start? After all, there are plenty of other ways to do a balanced amp.

Well, a big part of it is simply that I have a soft spot for circlotron, or "cross shunt push-pull" amplifier designs. They're simple, high-performance, and neatly sidestep some of the problems inherent in other amplifier topologies (more on this later.)

Back as Sumo, we made circlotron-style amplifiers, but they were Jim Bongiorno's designs. I'd never designed a circlotron amp. And yet they kept drawing me back in. First, because the topology is so different than anything else out

there. When you first look at it, your natural reaction is "how the hell could that ever work?" Then, when you understand the principle behind it, you think, "wow, that's really elegant. Why aren't there more of these?"

And, another big part of the decision was based on the fact that there were no circlotron-style headphone amps on the market. Period. None. Zero. Nada.

"Well, that's being contrarian," someone says. Yes it is. But I'm a bit contrarian. I mean, hey, look at the name of the company.

But maybe I should explain a bit more about amplifier topologies, so you can better understand why Mjolnir was always a circlotron.

Some Solid-State Amplifier Topologies. Disclaimer: this is not intended to be an exhaustive summary, so yep, if I missed your favorite topology, sorry.

- **JLH.** A great example of an early transistor design by John Lindsay Hood. Underscores two facts about early transistors: (1) They were pricey, so it uses only 4 transistors, (2) The PNP versions sucked, so it used only NPN output in a quasicomplementary arrangement.

This topology is much more like a tube amplifier than a modern solid-state amplifier—capacitor-coupled at the input and output, using a single voltage rail.

- **Lin/Blameless.** About 99 % of all audio amplifiers today are Lin amps. A Lin topology uses a differential amplifier at the input, a second voltage gain stage, and an output stage that is usually complementary and biased with a Vbe multiplier. This describes virtually every speaker amp on the market. Of course, there are endless variations: complementary, quasicomplementary, symmetrical, buffered, linearized, output-inclusive-compensated, etc … but at the heart it's a Lin. This topology offers a lot to like, including easy DC-coupling at the input and output, and a convenient terminal to run the negative feedback to.
- **CFA/Current Feedback.** This relatively new topology dispenses with the Lin's differential amplifier, which is usually the limitation on a Lin design's slew rate, and replaces it with a diamond buffer and current-feedback architecture. This topology can provide excellent performance at low parts count, but its advantages and disadvantages are hotly debated on, say, DIYAudio. Typically has better bandwith and slew rate, and worse distortion than a Lin

design.

- **Supersymmetry.** This is a patented Nelson Pass topology that is inherently balanced, and has a fundamental simplicity that is very appealing.
- **CSPP/Circlotron.** Note that none of the topologies discussed above, except supersymmetry, is inherently balanced. Now we get to the Cross Shunt Push Pull amp, which is a very old topology (from the 1950s, google "Circlotron patent" for more info. First applied in tube amplifiers and named "Circlotron" by Electrovoice, the CSPP topology at first looks like a mistake. Oversimplifying, all of the topologies measured above use the output devices as, well, "valves" that control the flow of current from one or two voltage rails to a single output node. The CSPP uses these devices to "unbalance" the flow of current from two cross-coupled power supplies to two output nodes, one positive and one negative. Thus, it's an inherently balanced design. It will never be anything other than a balanced design. And you will never get anything out of it except for balanced output (more on that later.)
- **Chip/Integrated, or Chip and Buffer.** Today, it's easy to get a headphone amplifier output

chip that pretty much does it all. That's fine, but then you're beholden to what's on the chip—it becomes a "black box" topology, which can only be tweaked via power supply and ancillary components. Boring.

So, we went with a relatively old power amp topology because it was inherently balanced—and also very simple. It uses only N-channel devices (or only P-channel devices, if you swing that way), so there's no worry about the N-channel and P-channel devices being mismatched. It does, however, require a complex power supply—two separate non-ground-referenced power supply rails for each channel. If you look at the CSPP transformer on the Mjolnir (the larger one), you'll see it has about a billion output pins. That's why.

And why did we call our version of the circlotron "Crossfet?" Because it's not really a circlotron, if it's not tubes. We wanted something that expressed "cross-shunt" and "MOSFET" in a short phrase. It wasn't because the MOSFETs were mad.

Amplification Devices Disambiguated. And, with that, why don't we talk about amplifying devices for a bit, because I'm sure that you guys wonder if we engineers just make up silly acronyms like JFET or MOSFET for fun:

- **Triode.** This is a tube with three elements (hence the "tri"): anode, grid, and cathode. A heater heats the cathode so that it emits electrons, which flow across naturally to the anode due to the overall circuit potential (voltage.) You can control the electron flow by applying a voltage to the grid. Triodes are very, very old (nineteen-teens), and, in some forms (such as the 6SN7, 6DJ8, 417, etc) are the most linear amplifying devices out there. Some tubes can do 0.005 % distortion open-loop. So, while many tube amps have higher overall distortion than solid-state, it's not due to the tubes' inherent linearity—it has more to do with their current output capability (low) and the use of output transformers (necessitated by their low output current.) Tubes are also interesting in that there is no physical connection between the elements—the grid, anode, and cathode are all in hard vacuum, so a tube can be considered almost a perfect voltage-controlled device.
- **Pentode.** This is a tube with five elements: anode, control grid, screen grid, suppressor grid, and cathode. These types of tubes have much higher gain than triodes, but are inherently less linear. These were more commonly used as output tubes (driving a transformer)

in audio power amps. Some can be run in triode mode for good results.

- **Bipolar, or BJT, or just "transistor."** The Bipolar Junction Transistor was the earliest successful solid-state amplification device. Formed of three doped regions of silicon (negative-positive-negative, or positive-negative-positive (NPN and PNP, respectively), it uses current input (not voltage) to control an output current. The gain of a BJT is expressed in current gain, or beta. Betas of 20 to 500 are common. BJTs are used almost everywhere in discrete amplifier design, and if you avoid common problems (beta droop, nonlinearity, etc), they make great components. Can be made to withstand very high voltages. Just be sure that you obey their current-drive needs.
- **JFET.** Junction Field Effect Transistor. This takes a hunk of n-doped silicon and uses it to control the current flow through a channel of p-doped silicon. The input terminal is essentially a reverse-biased diode, through which virtually no current flows. However, current flows with 0 V bias from the drain to the source, just like it does from anode to cathode in a tube. And JFETs act a lot like a tube—but, to be precise, it acts more like a pentode than a triode. These can be made

very low-noise, and are frequently used as inputs on an amps differential stage.

- **MOSFET.** Metal Oxide Semiconductor Field Effect Transistor. In this case, the control terminal—the gate—is insulated from the drain and source, so it can also be treated as essentially a voltage-input device. With a catch: most MOSFETs have significant input capacitance, and care must be taken that you have enough current capability to drive them at high frequencies. Like a JFET, these devices operate kinda-sorta like a tube, with curves that look like pentodes. Commonly used as output devices in audio power amps, MOSFETs have a reputation for being noisy that takes them out of the front-end and VAS stages (usually.) MOSFETs come in two basic flavors: enhancement-mode (meaning they do jack squat if they don't have a pretty big bias voltage on them) and depletion mode (which will run current even at no bias, like a JFET or a tube.) MOSFETs can be made very robust and very fast, making them a good choice for output stages (as long as you watch for parasitic oscillation—did I say "fast?")
- **SIT.** Static Induction Transistor. A relatively new device that is has some characteristics of both JFETs and MOSFETs—with an interest-

ing twist: the curves these devices produce look a lot like triode curves! Unfortunately, when available, SITs are eye-wateringly expensive, leading to limited use in commercial applications.

- **Opamp.** The Operational Amplifier is an amalgamation of hundreds or thousands of BJTs, JFETs, MOSFETs, capacitors, resistors, and other devices on a single chip. Available with gains in the tens of millions, these are complete amplifiers on a chip that offer very good performance in terms of traditional measurements. However, if you want to tailor a topology to your own application, or reduce loop gain to ensure constant feedback across the audio band, you're out of luck.

Okay, so what does this mean for Mjolnir?

Well, let's leave the tubes out of the equation. Mjolnir was never going to be a tube amp (though we did have a design for a big all-tube amp in a Mjolnir-sized chassis, which we never did anything with, but that's another story. Mjolnir was going to be solid-state from the start.

So where does that take us? We have BJTs, JFETs, MOSFETs, SITs, and op-amps to play with. We don't have anything against any of those devices. But for voltage amplification, we tend to like

JFETs and BJTs, in that order, and we tend to like MOSFETs and BJTs for output devices, in that order.

Why the hate for BJTs? Well, it's not really hate. Just caution. Current-driven devices are fine, but they need to have a little extra work to make sure you have enough current to drive them, even when they're working hard and beta is drooping. And you have to watch their safe operating area and thermal characteristics a bit more.

What we ended up doing in the early Mjolnir design (and we're talking breadboards here, not PC boards) was trying two different topologies:

1. High-voltage JFET front end and MOSFET output with no overall feedback.
2. JFET front end, BJT VAS stage, and MOSFET output, with local feedback around the VAS and output stage only.

We focused on these two topologies because both were simple, and both sidestepped the "different gain per phase" problem inherent in balanced amps that are driven single-ended.

What do I mean by this? I mean, if you drive a differential amp with overall feedback with a balanced signal, it produces a balanced output. 1 V in, gain of 10 = 10 V on either side.

But, if you drive a differential amp with overall feedback with a single-ended signal, it produces an unbalanced output: 1 V in, gain of 10=10 V on one phase, 11 V on the other.

Huh?

Yep. Look it up in an opamp cookbook. You'll see the different gains per phase and ways to compensate for them.

However, since we wanted to have an amp with both balanced and single-ended input, we wanted to avoid having different gains per phase. That would mean we'd have to switch the feedback resistors (say, with relays) to compensate if a single-ended input was used. No, thanks. I didn't really want to have 10 relays inside a Mjolnir. This was supposed to be a simple, no-frills, performance-is-everything kinda amp.

What was interesting about those two topologies was how closely they measured. We found that by using 95 V rails and a special high-voltage JFET (which I think we own the world stock of), we could get very, very close to the measured performance of the amp with the VAS stage—without any feedback.

This made for a very simple amplifier. The path was set. Mjolnir would be a no-overall-feedback, single-stage amp design.

So Is It Class A?

One of the things we get asked about all the time is "What Class amp is it?" It's a terrible question—not because we hate to answer it, but because manufacturers have mis-applied amplifier classes, especially Class A, to the point where there's a ton of confusion out there. I won't repeat my screed about Class A amps a few chapters back, but I think it'll be useful to go through some common amplifier classes.

Amplifier Classes Explained. While Class is in session, why don't we talk about amplifier classes a bit? This will be fun. Like everything else, this isn't exhaustive—I won't be talking Class C or S or T—look 'em up!

- **Class A.** Class A amps run full-out all the time. The transistors all conduct, all the time. They never turn off. They're hot. They're big. They're heavy. And they are, by definition, no more than 25 % efficient. So if you have a 125 W per channel Class A amplifier, it's going to be sitting there dissipating 1000+

watts at idle. It will get cooler the harder you run it. There are no shortcuts, no excuses, no easy outs. If it's not hot, big, and heavy, it's not Class A. Period.

- **Class B.** This type of amp really isn't used for audio. This is where the output transistors turn off as soon as they cross zero, because they are completely unbiased. The problem with this: huge crossover distortion, as the transistors turn on and off.
- **Class AB.** This is Class B, with bias on the output transistors so they run Class A some of the time. Fun fact: BJTs have an optimal bias for linearity, so "cranking up the bias" doesn't necessarily translate into better sound. MOSFETs don't. Crank them! When a Class AB's outputs eventually shut off, the transition is managed much better than in a Class B amp. This is the most popular audio output stage, because it combines high efficiency (up to 75 % theoretically) with good performance. Mjolnir is technically a Class AB amp. It runs in Class A up to about 1 W, then in Class AB thereafter.
- **Class D.** Switching or PWM amps. These "digital" amps have gone past the "exploding parakeet" stage (based on a comment from Mike Moffat about seeing one of the first commercial Class D amps, which output so

much RF noise that it would probably cause a parakeet to catch on fire), and can provide good measured performance. They can also offer 90 % to 95 % efficiency, so your thousand-watt amp can fit in something the size of a cigarette pack. Still, barf.

- **Class H.** This is Class AB with voltage rail switching. These amps run at lower rail voltages to increase efficiency and reduce heat, then switch to higher rails when output demands. A neat way to get high efficiency without the drawbacks of Class D, but necessarily more complex.

So what does all this pedantic BS have to do with Mjolnir? Think of it as some of the stuff an audio engineer has to hold in his (or her[1]) head as they work on a new design. How crazy you wanna get? How many chances you want to take? Should it be an all-BJT, Lin topology,

1 Too bad there aren't more female engineers. When I was in school, one of my classmates snarked, "You've dated all the girls in engineering." To which I replied, "Yeah. Both of them." Which wasn't far off the mark. Come on, guys can do this. It can't be that hard. And they teach you wayyyyy more math than you need. Don't be scared by nonlinear differential equations. You'll never use them … well, unless you plan on being the Ph.D in residence and presenting papers before the AES. Which is fine. Me, I'd rather blow up … er, I mean build stuff.

Class AB amp, because that's the best-known and most-documented design option out there, or a whackazoid tube-input, level-shifted, DC-coupled hybrid supersymmetry circlotron with a Class H output stage?

Yeah. You get the picture.

Early Adventures with Mjolnir

"Circlotron?" Mike asked doubtfully. "Isn't that something that only 6 people in the world know how to do, and even then they have to chant incantations and swing dead chickens over their heads to make them work?"

"No, they really aren't that bad—" I began.

"Famous last words," Mike cut me off.

"They're actually really simple—"

"Except for the keeping it balanced problem, the voodoo transformer, the eight thousand voltage rails, the weird in-the-air outputs, and making sure some idiot doesn't ground the negative output and blow it up problem, you mean."

"Well, yeah, but—"

"But you're gonna do it anyway."

"I already prototyped it. It works fine." Which was true. Circlotrons are really dead-simple. They just scare people, because at first glance, they look like a very, very bad mistake that will catch on fire and burn your bench to the ground. In reality, a circlotron using enhancement-mode MOSFETs with no bias and no input will just sit there happily and do absolutely nothing. With a decent function generator, you can program in an offset voltage and two out-of-phase sine waves and run it easily. Really not a big deal. But Mike's scared of weird analog things, and I'm scared of complicated digital stuff. So there you go.

"Really?" Mike said, doubtfully.

"Really. It took like an hour to build it up."

Mike sighed. "Ohh-kayy. How are you keeping it balanced?"

A differential servo."

Mike's eyebrows shot up. A differential servo?" I nodded and explained. I'd found a cool way to use a servo to compare the difference between the two outputs, and set the bias on one of the paired MOSFETs to match the other. That way, you just had to adjust bias on one side, and the other would follow.

Aside: "Servo" is short-speak for "DC servo." This is commonly used in amplifier designs to ensure that the DC offset at the output is very low, without having to use coupling capacitors to block DC, or twiddling pots to null DC (and hope it doesn't drift over time.) DC servos are relatively simple and very powerful, but like most things that are simple and powerful, they demand respect. DC servos are not perfect. They inject some audio back into the servo summing junction, so they'd better be high-quality and well-filtered for best performance (or used at a point that's not an input, like I was doing in Mjolnir.)

"And the voodoo transformer and eight thousand voltage rails?" Mike asked.

"Already have one. Got prototypes from the transformer guys. And yeah, it'll have a lot of capacitors in it, so what?"

Mike nodded. "Is this one of those things where the outputs are 40 V in the air?"

"Nope, they're close to 0 V."

"How about the shorting problem?" I frowned. Because Mike had a point. Both of Mjolnir's output phases were active. If someone connected

them together (say, by using a balanced 4-pin to 3-pin TRS adapter), it would be a very bad day. And I knew that no matter what warnings we put in the owner's manual about how Mjolnir was a balanced amp and a balanced amp only, and that they should never, ever, ever, not even on a bet try to use an adapter, someone would do just that. Probably about five seconds after the first one hit a customer's door. Even if we put an electronic flashing sign on the top of the amp saying "NEVER USE ADAPTERS THAT SHORT THE OUTPUTS!" it would still happen. And there would be fireworks and consternation.

"I'll figure that out," I told Mike.

"And how about single-ended output?" Mike asked.

"This is an end-game amp. I figured it would just be balanced."

"And what if these end-game guys have single-ended headphones?"

"Then they'll have to get balanced cables, or re-terminate them for balanced," I said. In retrospect, I should have paid more attention to this. Although we tried to work up a single-ended summer for Mjolnir, we were never happy with

the performance of the design we had. So it never got single-ended output. And yes, I'll admit … single-ended output is very useful. But it'll always be a summed output, if you're doing a circlotron. There's simply no easy way to get a single-ended output from it.

And, you know what? Mjolnir really was a fairly simple amp to get working.[2] Except the protection. Mike called that one right on. A simple output delay wouldn't protect Mjolnir from shorting adapters. A DC detect circuit wouldn't do it, either.

In the end, I came up with a complex analog-computer-style circuit that continuously monitors output current and DC offset, and lifts the output if the current goes over a pre-set point, or if DC offset goes higher than a predetermined limit. I

2 Working, yes. Right, not so much. Mjolnir's first appearance was at the Audeze-sponsored meet in Los Angeles. It was running far too high bias, its protection was only kinda-sorta working, it wasn't thermally stable, and the servo was being, well, very un-servo-like (we found out later that it was oscillating.) But it ran through the show and didn't blow anyone's headphones up. Unfortunately, it sounded very fat and strange. It wasn't until May that we had a real, production-intent prototype that made us happy to listen to it.

think it uses more parts than the original Asgard gain stage.

But it has saved our butts many a time. When we get an email that goes like, “Hey, I plugged in the Mjolnir and it sounded a bit funny, then it went “click,” in the middle of a song,” we immediately ask, “Are you using an adapter to plug your headphones into the jack?” The answer, unsurprisingly, is “yes.”

Whew. Big chapter. Lotsa tech talk. But hopefully illuminating.

Next up: Gungnir. More tech incoming. You’ve been warned.

Chapter 18

The Pinch-Off Problem

Late in the development of the Gungnir analog stage, I was sitting in the living room, talking to Mike on the phone about a problem we'd been having on the prototype boards.

To me, this was "just another day in engineering." Weird schiit happens. You gotta figure it out. So I didn't think anything about our conversation … until Rina walked into the room, laughing so hard she could barely stand.

"What?" I asked her.

She just laughed harder, holding on to the kitchen counter to keep from falling over. Literally.

"What's wrong with you?" I continued.

"You—" she said, gasping and pointing. "You—you—"

"Me me what?"

"You have a company—called 'Schiit'—and you're talking about your—pinch-off problem!" Rina said, through gales of laughter. I stopped dead. Then I started laughing, too. Schiit had a pinch-off problem.

Background: Pinch-Off is an Engineering Term. Really.

Well, more accurately, it's an old-guy engineering term, like "plate" instead of "anode" for tubes. What it describes is the voltage it takes to turn off a JFET (or other depletion-mode device.) As described in the previous chapter, a JFET runs current through it from drain to source as soon as it's connected to a voltage. But, if you lower the voltage at the gate of the JFET, eventually it won't conduct at all. That's known as $V_{gs(off)}$, or, in old-dude speak, the pinch-off voltage.

So Mike and I had been sitting there, talking about our pinch-off problem, and that's what Rina had walked in on.

Yeah, I know, stuff like this surely doesn't happen at Sony …

Now, as far as why we were discussing our $V_{gs(off)}$ problem, it was simple. I'd built some

perfboard versions of the Gungnir analog stage, measured them, and been very happy with the result.

Then we got the PC boards, put the parts on them, and suddenly they were running 50× the distortion of the perfboard versions.

Yeah. Stuff like this happens all the time in engineering. The real world isn't the same as simulations. PC board protos are different than built-in-the-air protos. Manufacturers change processes and parts don't work the way they used to. If you want a simple life, consider a career as a fisherman in Costa Rica or something. Maybe. Who knows. I've never actually done that, so it might be as bad (or worse) than engineering.

So, after spending a night with Mike and Dave trying to chase down the distortion (looking at compensation, oscillation, badly routed PCB traces, bad solder, wrong parts, bad parts etc—all the obvious stuff that shows up on prototypes), we were all baffled. The circuit worked, but it didn't work well.

So I went home and slept on it. Problems that seem huge at night sometimes become really obvious the next day.

But this one wasn't. I went back to it that evening, swapping parts and measuring. And a parts swap did make things better. Just not enough better. So I swapped parts again, just for the heck of it. And it got worse.

That was when the light bulb went off. The in-the-air prototypes weren't built with surface mount parts. They were built with through-hole parts. The JFETs we were using on the PC boards were supposed to be a near-equivalent to the through-hole parts … *but maybe they weren't.*

The datasheets told the story: the through-hole parts we'd been using on the perfboards had a pinch-off voltage range that was very small, and spec'd pretty tightly—about 0.2 V to 0.5 V. The ones we were using on the surface-mount PCB? 0.5 V to 6 V.

Yeah. 12× different. Like, duh.

And, considering where they were used in the circuit (as followers), that big pinch-off could cause all sorts of problems. I tacked some through-hole parts in their place, and suddenly the boards were acting (and measuring) just like the early prototypes.

From there, it was only a matter of finding a surface-mount JFET with similar specs, which only took a quick web search. A few days later, when the parts came in, Gungnir's analog stage was working as it should.

But only after a painfully hilarious conversation …

But Analog Isn't The Real Story

Okay. I front-loaded this chapter with a funny story about Gungnir's analog stage, but in reality, that was probably the least interesting part of the DAC's design. When Mike said he wanted to do a much more no-holds-barred design, I knew exactly what I wanted to do on the analog side. That is, a more sophisticated discrete stage, with a better topology, with higher voltage rails, and this time using a DC servo rather than coupling capacitors on the output. Other than the pinch-off problem, the development was relatively uneventful.

With Gungnir, the real story was on the digital side. Like Bifrost, we started with not much in the way of a product brief, except for Mike Moffat stating that he wanted to "do a proper hardware-balanced DAC." Beyond that, nothing. No sizes.

No feature sets. No colors. No 500-page list of specs.

But Mike is very, very specific when it comes to digital. "It needs to be big enough to keep the Hatfields and McCoys out of each other's corrals," he declared. "That means two transformers, one for digital and one for analog. And clock regeneration, we need to look at that a lot harder. And absolutely, positively hardware balanced, none of this single-DAC-per-channel stuff."

Let's translate:

Big enough to keep the Hatfields and McCoys away from each other. In Mike-speak, this means careful segregation of the analog and digital sections. Grounds. Power supplies. Clock routing. Physical space. Mike looks askance at tiny products that mix analog and digital. So, Gungnir was gonna be big.

Clock regeneration, we need to look at that. Bifrost uses a lot of the tricks Mike learned to get SPDIF jitter to acceptable levels, but Mike's work at Theta always featured VCO clock regeneration, and he hated to give that up. Eventually, that grew into Gungnir's unique Adapticlock system, which actually assesses the quality of the input signal (in terms of center frequency and jitter)

and routes it to either a VCO or VCXO oscillator. It also meant that Gungnir needed a much bigger and more powerful microprocessor to do this analysis and routing, for all supported input resolutions and sample rates.[1]

Absolutely, positively hardware balanced. Hardware balancing, or using one stereo DAC per channel, pays off huge dividends. Lower distortion, lower noise floor, elimination of more of the high-frequency noise that comes out of a modern sigma-delta DAC—these are all wonderful things. It also comes at a cost of using two DACs and twice the analog components, plus discrete summers for single-ended output.

So yeah, digital is the real story. And the real story of Gungnir is probably Adapticlock, it's unique feature. That's a Mike Moffat original that he's justifiably proud of. As far as we know, no other DAC tells you if your source is good or bad, and, even if bad, still provides clock regeneration. It

1 Actually, let's talk about that a bit. In the old days, you only had to worry about 16/44.1 and 16/48. Now there's a LOT more variations. And if you're interested in keeping everything bitperfect, that's a hell of a lot of management. Why do Gungnir and Bifrost click when you change sample rates? Because you have to reset the whole system to run at the new clock multiple.

took a ton of code to make that one work—and some very expensive VCXOs.

"If it's bad, we'll light up a front panel light," Mike said. "We could call it the 'buy better gear' light."

And "Buy Better Gear" is what stuck. It's technically the "VCO Mode" light, but that's a whole lot less interesting, right?[2]

In the spirit of the last chapter, let's talk about the parts of a digital audio system, so hopefully all of this stuff makes a little more sense:

Storage. Digital music has to be stored, whether it's on a plastic disk, magnetic disc, or in the cloud. At this point, it is no different than any other data you have. And, like other data, it can be lost if your hard drive is made by Western Digital (er, I mean, when it breaks.) Sorry, WD has had the majority of breakage in my personal experience—this is not a statistically significant result, just a personal opinion. The important

2 And there's not a lot of really bad gear out there, to be honest. Pretty much any computer won't light it. It really only comes on with really, really awful stuff, like satellite receivers and Apple Airport Express sources. And some old CD players that have gone off-frequency. That's about it. Everything else runs in high-precision VCXO mode.

thing is to make sure it's backed up. Or, if you're still a dinosaur (er, I mean, using plastic disks), don't treat them so poorly as to destroy them.

Formats. Most of you guys already know this, but let's go ahead and be inclusive. Digital music comes in tons of different formats. Let's cover three broad swathes:

- **Lossy (MP3, AAC, etc.)** These formats have had a whole lot of bits thrown away. Lossy formats contain only 10 % to 20 % of the data of the original digital file. This means they have to be reconstructed into a semblance of the original signal by an algorithm running on a computer or player, with the losses masked by perceptual encoding. This is in no way, shape, or form a "bitperfect" solution.
- **Lossless (FLAC, ALAC, AIFF, WAV etc.)** These formats preserve all the original bits, even though some use lossless compression to pack them more efficiently. All of these are Pulse Code Modulation, or PCM, formats. These formats are bitperfect—they preserve the original music samples.
- **DSD.** Instead of using multiple levels to represent the music, DSD encodes it as a high-speed, pulse width modulated (PWM) datastream. Theoretically, a DSD recording made on a pure

DSD analog to digital converter and never converted or processed could be reproduced on a pure DSD DAC (or simply via a very good switch and low-pass filter) and retain all of the original DSD "samples"—or, more precisely, information.

Transmission. From the stored digital file, you need to get it to where its going. This should be relatively simple, but sometimes it isn't.

- **Ideally.** In an ideal (PCM) world, the DAC wants to see three things: bit clock, word clock, and data. In a delta-sigma DAC, you can add a master clock to the list. Which means you'd ideally have 3 to 4 BNC cables connecting a digital source and a DAC. However, in the dim dark days of early digital, this was Kevorked for a single-cable solution—simpler, easier, but far more problematic in terms of jitter, since the clocks were buried in with data. Welcome to SPDIF.
- **Internal.** Let's say you want to move data from a CD drive to a DAC internal in the CD player. After recovery of the data from the disk, it's usually moved around via I^2S, which is, hey, bit, word, data. Ideal? Sure. Why don't they use this outside of the box? There's no single accepted standard. But we can dream.

- **Optical.** One of the flavors of SPDIF. This was chosen in the early digital days because it glowed, looked cool, and in general seemed like a nifty futuristic idea in an era when people also thought the 1980s Corvette digital dashboard was cool. It's bandwidth-challenged, though, and can only sometimes stretch to accommodate 24/192.
- **Coaxial.** Another flavor of SPDIF. Looks boring, but generally performs better than optical. BNC connections make it seriously good for high-frequency use. Unless there's a ground issue, of course.
- **USB.** And now a left turn into computerland. I've gone on and on about how much of a pain USB is, but it's here to stay. Comes in many, many flavors and implementations. Different than SPDIF in that it is a packet-based system. Bandwidth isn't a problem with USB 2.0 and above.

Reception. Beyond the digital connection, there's a receiver to process the incoming SPDIF or USB signal. In the case of SPDIF, it recovers the clocks that are embedded in the data. In the case of USB, asynchronous transfer controls the clocks locally.

Clock Management. Okay, so you have digital

data. Now what? Some manufacturers choose to upsample everything to a specific datarate, no matter what's coming into the box. This eliminates the need for clock management, but … you guessed it … asynchronous sample rate conversion, or ASRC, is not bitperfect—it replaces the original samples. So, for a bitperfect DAC, this means clock management is necessary. In short, this is the process of telling the digital filter and the DAC, "Hey, I'm sending you 16/44.1, get ready, 'kay?" Or 16/48 or 24/96 or 32/192 or whatever. Run through the different bit depths and sample rates, and you'll quickly see that there are many different combinations. Clock management isn't trivial.

Digital Filter. Digital filters are where bit-perfect transfer usually dies. Digital filters upsample the incoming data to higher data rates (typically 8×) to reduce the need for analog brickwall filtering. This is handy, but again—what it outputs is a mathematical approximation. That is, unless it is a closed-form digital filter that retains the original samples. And we know of only one of those—which is what's going in Yggdrasil.

D/A converter. From the digital filter (which may be inside the DAC chip itself), the data gets passed off to the actual D/A converter. These

typically come in two varieties:

- **R2R.** This is nearly dead as a technology today, because it's a pain in the butt to implement, especially with bit depths that exceed 20 bit. But it's the only D/A technology that can be bitperfect.
- **Delta-Sigma.** About 99.9 % of DACs (or more) today use delta-sigma conversion. The reason? Because they allow marketers to use big numbers, like 24 bit (or even 32 bit). And because they are inexpensive. However, like ASRC and digital filters, they only output an approximation of the actual data.

Analog. And you thought we were done? No. Some DACs output current, which requires an I/V converter. This is a place where discrete designs, with proper low-impedance inputs, can offer huge advantages over ICs. Some DACs output voltage, but it still needs filtered, and, in some cases, summed. So there is an analog component at the very end—and it is definitely critical to the performance of the entire system.

Although I can go on and on about the engineering side, I'll spare you the huge dissertation, as in the last chapter. Because, the more I think about it, I believe the central thing comes down to philosophy.

"Philosophy? What the heck does that have to do with DACs?" you might ask.

I'll respond: It has everything to do with DACs. And amps. And business in general. So let's move on to that.

Philosophy: Or, Why You Do Something

Okay. Let's say you start a company. Why did you do it?

Let's say that company makes products. Why did you do one, and not another?

Let's say your products are made a certain way. Why did you choose that method?

Let's say your products have certain features, or you choose to leave them out. Why?

If you keep asking, "Why?" you might find that there is a good reason behind all of the answers. Or you may find none at all. And, even if you find there is a good reason, you may not agree with the "why." This is why companies have to have a philosophy—and stick to it. Because otherwise, they're rudderless. Aimlessly wandering. Waiting for a magical hit product to pull them out of the morass. You see this in a lot of big companies—

ones where products are created and approved by giant committees, endless meetings, thousands of hours of "gaining consensus."

Car companies are a great example. How many midrange cars can you simply swap the badges on, and not even know it's made by a different manufacturer? How many are so completely forgettable. How many seem to lose their way every decade or so—and then suddenly release a flood of models that simply copy the one hit they recently had?

There's no philosophy. Just the endless chase of benchmarking and specs-list-stuffing and crossing your fingers and hoping that somehow, something with enough personality makes it through the bean-counters and second-guessers to make a difference for the line.

If you start a company, make a product, decide on features and specs, ask yourself, "Why?" And be specific with your answer. You'll do a lot better for it.

So what do we mean by this? Fine. Let's use Schiit as an example.

Our philosophy is that ***we want to make fun,***

affordable products that are as true to the musical source as possible.

Note what comes first, second, and third.

First, fun. Come on, guys, this business is far too serious at times. Let's have some fun with this.

Second, affordable. The elephant in the room in high-end audio. On a recent panel, some guys tried to make the point that "personal audio" wasn't really any younger than high-end audio, citing examples from companies making $1000 to $5000 products. Like, duh. I reminded them that our audience was much younger than the norm, statistically, because of one simple thing—they're affordable.

Finally, true to the source. In the digital realm, to us, this means retaining the original sampled data as much as possible. In the lower-priced realm, this also means forsaking the totality of this goal, because modern delta-sigma DACs sacrifice the original samples for a mathematical approximation.

Now, this isn't to say these DACs can't sound good. And, congruent with the third part of our philosophy, we choose to preserve the original

bit depth and sample rate as far as possible down the chain, and to minimize errors caused by jitter. That's the best we can do with delta-sigma.

Now, very soon, we won't have to compromise on this, but that's a substantially more expensive product—though still much less than most megabuck DACs. And that takes a digital filter that runs an algorithm that retains the original samples—like the unique one we have in store. But note that this is congruent with our overall philosophy and goal.

"Wait a minute," you say. "You 'want to stay as true to the source as possible,' but you do tube amps and stuff like that. What's up?"

What's up is that tube amps don't necessarily have to be high distortion, and, even then, relatively high levels of low-order THD aren't correlated highly with audibility. And a lot of people think tube amps are fun. So we're hitting the second part of our philosophy. If we were saying, "You must use tubes, and a great tube amp is $ 50 000," that's antithetical to our philosophy. But fun, accurate tube amps are not at all.

"But I don't agree," someone is saying. "You might have a philosophy, but I don't agree with it."

Yep. And that's the thing. You'll never have 100 % consensus. If you don't agree, simply move on, and find a company that meets your own goals. That's why there isn't one *The Amp Company* out there to rule them all. It's a wonderfully wide, varied world—and that's a very good thing.

What we can say is that we do actively ask ourselves why, and bump up the answers against our philosophy, as we develop products. Let's see how that works:

Hey, a D/A converter manufacturer brought out a new delta-sigma device. Let's build a new product!
No, sorry. Not congruent. Just because it's new doesn't mean it's better. Is it a meaningful upgrade? Inexpensive? Then maybe.

Hey, a D/A converter manufacturer brought out a new R2R device. Let's build a new product!
Hell, yes! R2R DACs are bitperfect. So let's go for it—well, unless it's crap or a billion dollars.

You know, some people really like the old, lush euphonic tube sound, let's throw a really nonlinear tube in our DAC!
Nope. Although it can be pleasant, and although

some prefer it, it's not staying true to the original, is it?

You know, we can make something even better and more fun, but it's even better and less expensive than our top-of-the-line products. You bet, let's do it! Everyone benefits!

Did you see that there's a new Class D module that puts out 300 W per channel from something smaller than a cigarette pack? Let's throw it in a DAC and make a power DAC. Not gonna happen—throwing away the signal for a nonlinear-control-system approximation of it? Not our bag.

Anyway, perhaps you see where I'm going here. Philosophy is key. Whys are important. But by sticking to them, you won't please everyone all the time.

But then again, trying to please everyone all the time won't do that, either.

Chapter 19

Every Road is a Dead End: Early Adventures with Magni

This chapter is a lesson in hubris—and in the value of chucking it all and starting over.

It happens to every company, I'm sure. There will always be a time when things are going well, reviews are great, and new products are flying off the shelves. We literally couldn't keep Bifrost in stock. Asgard, Valhalla, and Lyr were all doing well. I'd just been contacted by the Arizona Audiophile Society, where the Bifrost had beaten all the other DACs in their blind listening test (with retail prices up to $7500.) We had working prototypes of Mjolnir and Gungnir, and were looking forward to their launch. And we'd just started looking at space so we could move out of the garage (more on that later.)

This is the time when you start thinking, "Hey, this is going pretty well. Man, we're really tearing it up. Wow, maybe we actually are pretty good at

this!"

This isn't good. In fact, this is the time you should be the most terrified.

Now, to be clear, we didn't go completely over-the-top on the narcissism. We didn't do anything truly stupid. Nobody bought a Ferrari. None of us went out and bought $1000 bottles of Scotch. None of us created audio product derivatives to sell to Wall Street. And none of us rode to work in a sedan car borne by a dozen acolytes.

But, this run of good luck was enough to have me thinking, "Heh, a little amp? How hard can that be?"

As it turned out, it was damn difficult. Remember, 13 months from Modi proto to launch? At least 6 or 7 of those months were spent running down the wrong paths on Magni.

Philosophy Can Also Be A Prison

I spent the last part of the previous chapter nattering on about philosophy. And I truly think that every successful company should have a well-thought-through, concise philosophy that informs everything they do.

But a philosophy can also be a prison. If it's too specific or too inflexible, you won't be able to change when you need to. You won't be able to adapt to new needs, new markets, new competitors. That's also necessary.

It's also why our philosophy is pretty broad. And if I'd done nothing more than apply that philosophy to Magni, I probably would have been fine.

Aside: our philosophy is to "make fun, affordable products that are as true to the musical source as possible," in case you skipped the last chapter.

Instead, I larded on a bunch of additional "wants" to Magni's initial design brief. Some of these were based on market reality. Some of them were sheer fantasy.

Let's start with the ones based in reality:

1. **This amp should be versatile enough for most any headphone.** We already had some very specific amps, like Valhalla for high-impedance headphones and Lyr for power-hungry orthos, but this should really be a do-all amp, since it would likely be a starter amp for many audiophiles.

2. **If we couldn't do better than the inexpensive amps already out there, why bother?** To me, "doing better" was a mix of more power and sonics, in a simple, attractive package.
3. **This amp needed to hit a very aggressive price point**—a price point unimaginable when we started the company. We had to be careful about design, construction, features, reliability, etc. I had $ 99 as a target, to match the Modi.

And now, the fantasy:

1. **The topology should be as simple as possible**—insanely simple, just a few transistors and a very simple power supply, almost like a solid-state version of a tube amp. That's probably what would sound best, I thought.
2. **To keep costs down, we'd use a switching wall-wart** to generate a single DC rail. Switching wall-warts are so cheap they show up in Cracker Jack boxes these days. I mean, unimaginably cheap. We'd have to use a switcher to keep cost down.
3. **It should be a neat, unique topology.** I'd messed around with two-transistor gain cells. Maybe that would be cool. I'd also played around with the old JLH topology, which I remembered sounded good.

Looking back, all those extra fantasy items are really funny. Simple amps usually have to resort to Class A output to get them linear enough to work well—and Class A was absolutely out. Magni's tiny chassis wouldn't be able to dissipate the heat. Switching power supplies are cheap, but absolutely scary in terms of power supply noise—not to mention the fact that a single rail would mean we'd have to use coupling capacitors at the input and output of the amp. And a neat, unique topology? Yeah, there's a reason those are scarce. The "cool" stuff I'd played with in the past simply had too many limitations—not enough voltage swing, not linear enough, not stable into a wide range of headphone loads, etc.

But I'm An Idiot

So, of course, the first thing I had to do was to try a JLH-style amp with a switching wall-wart that I bought off of eBay. I think it was $3. Which meant, in production quantities, it could easily be a $0.50 part. Think of that—a cord, plastic chassis, PC board, switching supply doing 24 V at 0.5 A for half a buck.

And yeah, it was about as good as you'd expect for that price. It was so noisy that it made the JLH

amp oscillate constantly at full voltage, without any input. I'm talking full-scale noise at a couple of megahertz.

To translate: instant headphone fry. Assuming the output stage lasted that long. I tried a couple of other switching wall-warts, but they really weren't much better. So I tried filtering them. Which doesn't work so well when you have half a volt of noise on the ground (the engineers here are cringing).

Finally, I gave up and simply hooked the JLH topology up to our lab power supply. Now it ran fine. No oscillation. Which is what you'd expect from a clean supply.

There were only two problems:

1. The JLH topology really, really doesn't like to be transformed to a Class AB design. It's very nonlinear, with high distortion.
2. It sounded like ass.

I mean, it sounded awful. As in, 1960s solid-state awful. I'd forgotten how bad solid state could be. Bright, nasty, confused, muddled … it simply didn't stand up to modern designs. Not even an opamp-and-buffer design. Which I also didn't want to do, because that's been done to death.

Yeah. Hubris.

After that failure, I tweaked around with the circuit for a while, and ended up with something that sounded kinda decent. But by this time, the original 16-component-per-channel design had ballooned to over twice that. It was more complicated than some of the 60- and 100-watt speaker amps I'd designed. And that was really stupid.

So what did I do? For a while I just gave up. I had a non-optimal topology and an unworkable power supply. I'd wasted a couple of months getting exactly nowhere.

The Non-Lighting Light Bulb

Sometimes when you walk away from a project, the insight will come when you least expect it. You'll wake up one morning and have the answer. Or you'll be driving into the office and it'll hit you so hard you'll say, "Hell, why didn't I think of that before?"

In the case of Magni, walking away didn't work. As Mjolnir and Gungnir moved towards production, and as we started our first move out of the

garage, I had plenty to occupy me. I could forget about it.

But the answer didn't come.

Not that I didn't try. Sure, I put together a half a dozen neat circuits. JFET-MOSFET gain cell. Simple current-feedback amp. Etc.

But all of them had at least one fatal flaw. And all of them still wouldn't work with a noisy power supply. Even if one had worked, the supply still killed it.

So I wasted more time—drawing up chassis for the Modi and the non-existent amplifier, trying still more wall-warts, tweaking circuits and hoping that something would work out. Nothing did. And I was starting to sweat. Any day now, Mike would ask me how the Magni was going, and I'd have to tell him. And he'd say, sarcastically, "I thought you said it would be easy, Sparky!" I didn't want to have that conversation. I didn't want to say, "You know, an opamp and a buffer wouldn't be so bad."

But, you know, even if I'd done an opamp and buffer design, it probably wouldn't have worked because of the power supply.

In the end, I was sitting in the garage one weekend, staring at the perfboard mess that should be a Magni. And I suddenly remembered that one thought I had: *Hell, this thing has more parts than some of the speaker amps I designed.*

So what if I just did it like a speaker amp? I wondered. That would eliminate the topology problem. Lin topologies could be very low-distortion—and Class AB—and direct coupled—and very, very robust.

But that was crazy! A full Lin topology for our least-expensive amp?

What the hell. I opened the schematic capture program and drew up a simple Lin amp.

It was simpler than the mess I'd designed.

But … a Lin amp really needed a bipolar power supply—that is, both positive and negative rail voltages. They didn't like to hang halfway between a single supply and ground. That meant caps in the feedback loop, input biasing, and other ugly stuff like that.

So what if I just said, "the hell with it," and did an AC wall-wart (basically a transformer in a box) and a half-wave bridge to create both positive and negative voltages?

Half-wave bridge, barf, I heard Mike's voice in the back of my mind.

But I didn't care. Maybe this was the way to go. Maybe a full Lin amp with a bipolar supply—and, what the hell, a DC servo too, might as well go crazy—maybe this would work. Maybe modern surface-mount manufacturing would make this feasible.

Aside: I really had no idea. I'd never done a surface-mount board before the Magni.

I built the Lin circuit that night and ran it on the lab supply. The damn thing worked first shot, as if to say, "Why didn't you just do this from the start." And it measured well. Not just well, but spectacularly. And with a few tweaks, it was running almost rail-to-rail.

Aside: "Rail to rail" is important for efficiency—a very important part of a Class AB design.

Now, I was excited. This was getting somewhere. If we could get a power supply put together to run it, we might have a product!

Except—I had no idea what a linear wall-wart would cost. They're pretty scarce. Most people have gone over to switchers these days.

But again, like I said in the beginning, most answers are not much more than an inquiry or two away. Since I knew we were shooting for minimum cost, I wasn't going to be able to get it from a US manufacturer. So I turned to a new source—one I'd never used before—alibaba.com.

Yes, that Alibaba. Chinese manufacturing. Now, there's nothing wrong with that. But it was different than anything we'd done before. Luckily, Alibaba has a pretty good feedback system, so you have at least an idea of the companies you're working with. We quickly had quotes from a half-dozen manufacturers, all at amazingly inexpensive rates. Not as inexpensive as a switcher, but still well within the envelope of a $ 99 product.

But what would they look like? Would they be any good? Even if they were, how well would Magni perform on a smaller power supply (smaller than the lab supply). I ordered some samples and sat back to wait.

In a week, I had my answer. They looked like standard cheap wall-warts, the kind you see on dozens of different products. But these had one big difference: they were AC wall-warts, delivering 16 V AC to a half-wave rectified supply

running MC-series regulators. I did a version on perfboard and verified the performance—and sat back in shock. The Magni prototype delivered nearly 2 W into 32 Ω at clipping, and distortion was less than 0.004 % at 1 V RMS (a much more typical headphone load). And this was from the wall-wart. 60 Hz hum from the half-wave supply was over 100 dB down from 1 V RMS.

It measured better than anything we made.

Still, what did it sound like? That took more waiting. Because I usually don't listen to breadboards or perfboards—I just build a single channel and then get into the PC board, then listen to that.

Into Surface Mount

Before I did Magni, I'd never laid out a surface-mount board. It was a profoundly alien experience. I wasn't used to the parts. I wasn't familiar with the best way to route them. And, most of all, I still wasn't confident it would work. The lesson from Gungnir's pinch-off problem was too fresh in my mind. What gotchas would we find when we went to surface-mount? Would the equivalent parts even be available?

Parts turned out not to be a problem. In fact, they were a real eye-opener. When you hear someone say, "They don't make great audio devices anymore," and wax poetic about the glory days of Japanese transistors, they don't work with surface mount parts. They don't know all the cool new stuff that's available right now—and the majority of it is in surface-mount packages. I learned a lot throwing that first board together. But, because I wasn't confident it would work, it wasn't a full design. No muting relay. No servo. Hell, it didn't have a power switch. But I wanted something we could try. Something we could listen to, and decide if it was good or bad.

In a few days, I had PC boards to play with. I threw one together and measured it. It ran pretty much the same as the prototype.

After that, it was the moment of truth. I grabbed a set of Grados and took them out to the test bench. The little Magni prototype drove them shockingly well.

But it should also be able to do better than Grados. It had tons of power. I decided I'd bring it to its knees with the LCD-2s.

Magni laughed at the LCD-2s. No problem. No big deal at all. It would easily go to ear-bleeding

levels. I sat there, laughing at the spectacle of this tiny little amp driving the LCD-2s. It looked absolutely ridiculous.

Rina came out to see if I'd lost my mind. "What are you laughing about?" she asked.

"Magni. The little amp."

She saw that I was holding the LCD-2s. "On those?" I laughed again. "No problem."

"Really?" she took the headphones out of my hand and put them on. "Play my song."

Aside: Rina has a specific song she uses to evaluate new headphones and amps. It's not what I'd call hi-fi, but she's heard it so many times that it's a perfectly good reference for her. It's Enigma's Seven Lives (Radio Edit.) Yeah. I know. Talk to her. Hey, those of you with the earliest Asgards had them listen-tested to New Kids on the Block, thanks to her. Think about that the next time you listen.

I played her song. She listened for about a minute, poker-faced. I frowned. What did that mean? Did she like it? Did she hate it? Was it really crap? Was I hearing things?

Eventually, she took the headphones off. She shook her head sadly and looked at me.

"So what are we going to do about Asgard?" she asked.

Chapter 20

The HOA Problem

While we were working on Magni and Modi and Gungnir and Mjolnir—through all the troubles and triumphs and setbacks and workarounds—I couldn't fight a growing unease.

Unease about how big we were getting.

Yeah, I know, laughable now. Hell, I'll be signing a lease to increase our total square footage to over 8000 square feet on this week of our 4th anniversary.

But that's now. As of early 2012, I had very good reasons to be worried. Rina and I live in a neighborhood with a home owners association, or HOA. For those of you outside the USA, that's where a bunch of insanely picky buttheads get together so they can determine the acceptable colors of each others' homes, send nastygrams about cracked concrete or broken bricks, and generally act like a bunch of old-timers with bad "get off my lawn" syndrome.

Of course, I'm being flippant here. My business partner at Centric specifically avoided buying a house with an HOA neighborhood, and later had a neighbor paint their house in day-glo sky blue. Yeah. And, in our case, the HOA does serve a useful purpose, since it maintains the greenbelt that separates us from the hillside brush. Which I was very thankful for in 2003, when brushfires literally came within 50 feet of our house—and were stopped by the greenbelt.

Protip: if you're having trouble with your HOA, the easiest way to get them off your back is to learn Morse Code and get your ham radio operator's license. Hams are protected by the FCC as part of the critical communications infrastructure of the USA, so if you wanted to, say, put a 90-foot-tall radio antenna in your front yard, the HOA can do absolutely nothing about it. Threatening to do that will shut them up good and fast.

Anyway, what does an HOA have to do with getting too big?

Plenty. Most HOAs prohibit operating a business out of your home. Now, this isn't usually strictly enforced. They don't care if you have a home

office or a studio, or if you're shipping some eBay or Etsy stuff out of your house, or selling a few hobby things you make in the garage.

But when you have two employees coming and going every evening, 7 days a week, with garage lights blazing and music blaring … well, that's a different story. It was only a matter of time before the HOA would start complaining. And since the city doesn't really like businesses being operated out of homes, they could absolutely enforce it.

So we had to look at moving … and soon.

7 Figures in a Garage, and the Reality of Having Your Own Business

What's the limit of a garage business? A lot higher than you might think. When I finally started looking for space, we were still only using 1/3 of a 3-car garage for "production and shipping floor." In that space, we'd just cracked seven figures in sales. A note on numbers: Schiit is a private business, and I usually don't discuss revenues, because (a) it's gauche, (b) we are not required to, and (c) it's really not that important.

However, in this case, I'm using a specific number to illustrate what you can do with a self-funded,

home-based manufacturing business. Remember, we started this with $10k. 18 months later, we're into 7 figures annually. In a garage.

This isn't intended to be bragging. This is intended to be inspiration for you. Starting your own business is absolutely do-able—without taking loans, leasing tons of space, hanging your ass out for bankers, gambling on delivering a crowdfunded product on time, or otherwise betting big on getting big.

But that brings us to what I call *The Reality of Having Your Own Business.*

In my opinion, the difference between working for someone and having your own business comes down to a single phrase: *When you have your own business, you can't say, 'that's not my problem.'*

Sounds too simple? No. Sit back and let it sink in.

When you're working for someone, you'll usually have a fairly well-defined role. No matter if you're a clerk or an engineer or a COO, in all cases, you'll know what's expected of you. A clerk isn't expected to design a new product—that's not his problem. An engineer isn't expected to write the ad copy for the new line launch—that's

not his problem. A COO isn't expected to stand up in front of the press when the firm does a billion-dollar acquisition. It's not her problem.

With your own company—especially a small company—*everything* is your problem.

"Well, I'll hire people to take care of marketing, production, operations, etc," you say.

Yes, and hiring them is your problem. As is budgeting for their salaries. And keeping them motivated.

And who do they report to? You.

Everything is your problem.

And I mean, everything. In this book, I've covered only the top level stuff—engineering, putting things together, dealing with production problems, space. But let's look at a bigger list of things that will be your problem, if you start your own manufacturing biz:

- Engineering Documentation
  - Prototyping
  - Schematics
  - PC board layout
  - Testing
  - Compliance (FCC and the like)
  - Production procedures

  - Testing procedures
  - Documentation and archiving
  - Managing changes
  - Pricing
- Functional Items
  - Product manuals
  - Product description
  - Product setup guides
  - Driver installation guides
  - Packaging
  - Packaging testing
- Design
  - Industrial design
  - Packaging design
  - Collateral design
- Marketing
  - Website structure, functionality, copy, design
  - Product “sell”copy
  - Press releases
  - Photography
  - Informal communications (blogs, forums, etc)
  - Advertising
  - Review samples
- Sales
  - E-commerce capability, functionality, testing

  - Sales through other sites
  - Sales through dealers and distributors (ack)
  - Accounts receivable (if you are dumb enough to give terms)
- Purchasing
  - Vendor interface
  - Purchase orders
  - Accounts payable
  - Account setup
  - Wire transfers
- Customers
  - Incorrect orders
  - Technical questions
  - General questions
  - Customer service
- HR
  - Hiring
  - Firing
  - Vacations
  - Benefits
- Shipping
  - Choosing shipping providers
  - Negotiating rates
  - Setup for e-commerce
  - Mis-shipments
  - Lost shipments
- Accounting and Taxes
  - Bookkeeping

- Forecasting
- Federal taxes
- State taxes
- Sales taxes

Yep, it's a long list. And it's not complete. Not by a long shot.

Now, it may seem I'm relatively cool on the prospect of starting your own business. No. Just realistic. It's a ton of work—but I'll re-iterate. It's the biggest, most satisfying thing you'll ever do. At least for me. But I'm weird.

Again: *it's always your problem. No hiding. No passing the buck.*

If you're cool with that, go for it. Create your own business. Own it. Grow it. Enjoy the great times—and there will be plenty. And work through the tough times—those will happen, too. Hire and build and make it easier on yourself.

But don't do it because you think it's going to be easy. Or because you think you'll have more freedom.

Because, if it's your business, it'll always be your problem. Until you decide to sell it and get out. And when you're small, everything is your problem.

Including looking for space.

Which is another of those invisible lines that, once you cross it, you won't go back. It's a big step that takes you from "This may be a hobby," to "OK, we're committed, this is a real business."

Our Advanced Search Technique

When I started looking for space to lease in early 2012, I employed a proprietary algorithm using non-Fourier wavelet mathematics, known as DBLFSN, as well as a well-known methodology called BSLA. DBLFSN and BSLA rapidly narrow the lease candidates to a handful of locations best suited for a business' use.

Yeah, I'm having fun with you:
DBLFSN: Driving By Looking For Space Nearby
BSLA: Borrowing Someone's Loopnet Account

Here's the deal. I was still working daily at Centric, which is located in downtown Newhall. Being lazy, I wondered if there was anywhere nearby that we could use to house Schiit. So that's where I started. And, being observant, I noticed a building only a block away that might be a good candidate. Looking it up on Loopnet

confirmed that it was about 1800 square feet, but didn't give a lease rate.

Now, 1800 square feet is kinda big for getting started with a small company. But it's not really that pricey. At a standard industrial lease rate of, say, $0.60, plus $0.10 CAM, you're looking at under $1300 a month. Plus electricity, water, gas, etc.

And this space had one thing going for it: the building was a schiithole. It was a mixed siding-and-stucco one-story building with holes kicked in its sides, dry-rotted eaves, peeling tan paint that revealed 5 decades worth of colors underneath, a nasty potholed blacktop parking space or two …

In other words, it would probably be almost free. I was thrilled. It seemed like the perfect place. I called the realtor listed on the sign. "Hey, that building on the corner of Railroad and 6th, is it available?"

"Yep," said the realtor. "What are you planning to use it for?"

"Um … I have an audio company, we'd be doing some light manufacturing—"

"Manufacturing, nope," he cut me off. "It's not zoned for that."

My heart sank. “We’re not talking machine tools and stuff. Just assembly.”

“No can do,” the realtor told me. “You start doing manufacturing in there, the city’s going to come in and shut you down. The whole thing is a mess. We’re now technically a flood plain, and the whole Enterprise Zone thing is screwing everything up, we don’t know exactly what we can do with that building, period.”

“There’s no way we can make it work?”

“Nope,” he said. “Sorry.”

Stupid Rules, and Types of Business Space

Okay, if you’ve never leased business space before, you’re probably shaking your head, wondering what the hell is going on here. Let me ’splain.

What’s going on here is zoning. As in, a city divides its space, and decides what you can do in each place. Some common zones will include:

- **Residential:** You can only build houses here. This is where you live.
- **Commercial:** Only shops, thank you. Get your $10 fancy ice cream here.
- **Industrial:** That’s where them crazy people build things. Like Schiit amps.

Now, some cities are crazier than others about zoning. Where Mike Moffat lives in Agua Dulce, you could probably build an 80-foot-tall purple freeform house with a 10 000 square foot manufacturing facility and a Mexican restaurant, and bring in 15 employees to run it, and the city wouldn't blink. (I'm exaggerating, of course.)

Santa Clarita, which encompasses Valencia, Newhall, Canyon Country, Saugus, etc, is about 1 000 000× opposite. In fact, Valencia is entirely master-planned. Which means it has big business areas full of offices and warehouses and manufacturing firms—and not a single restaurant. Everything must be in its place. And they are famous for sending their snoops around to make sure you're abiding by the rules. Yeah, we live in a wonderful place.

And the problem was, that beautiful Schiithole was zoned for commercial use, not industrial. Commercial meant that we could have an ice cream parlor or an office there, but not a manufacturing floor, nor a warehouse.

So did we look for other space? Yeah, sure. We looked in the Valencia Industrial Center, which is where we are now. But I really, really didn't like the idea of driving from Centric to Schiit to get

things done. Remember, we had no operations people at that time, so it would be me overseeing things, Rina shipping and doing listening tests, and Eddie and Tony putting things together.

Which brought us back to that crappy little building in Newhall. The question was, how the heck could we get our hands on it?

It took two more conversations with the realtor to get him to agree to let us take a look inside the place. He still didn't think he would lease it to us, but I guess he was either (a) too bored, or (b) pissed that the space still wasn't moving. I still don't know why he did it, but it was that meeting that turned the tide.

And, man oh man, it was at least as bad inside. The tile floor had been cut open to do plumbing work, then roughly cemented over. Every thing was covered in dust. A contractor was using it to store parts and equipment, and it was packed floor to ceiling with all kinds of stuff. The heating and air conditioning were nonfunctional. The layout was a weird L shape around a dirt-and-concrete yard. It was next to a barber shop and a thrift store. And the hammering of passing trains across the street made bits of the acoustical ceiling fall like snow.

“I guess this wouldn’t be all that great for an office, even,” said the realtor, looking at it with new eyes.

“It’s perfect,” I told him. “If we could just screw a few things together here.”

He shook his head. “This is commercial. For shops and such. Places that sell things.” A sudden idea hit me. “But we do sell things,” I told him.

“Yeah, but you’re a manufacturer,” he said. “I’m talking about shops that sell to the public.”

“But that’s what we do,” I told him. “We sell direct to the public.”

“Hmm,” he said, rubbing his temples. He clearly wanted someone in the space that would pay more than the guy who was using it for storage, but he was still scared of the city. “So people could come in here and buy one of your products?”

“Theoretically, yes.”

“Theoretically?” he looked doubtful. “How do you sell your stuff now?”

“Online, mainly.”

“Hmm. But you could sell something to someone who came in?”

“Yes.”

The realtor nodded. “Huh. Well. Let me see what I can do.”

“Great!” I said.

He frowned. “No promises.”

But somehow I knew: *this would be Schiit.*

Chapter 21

You Catch a Cold, We Die: Bigger Products, Bigger Problems

The first and second quarter of 2012 weren't just the beginning of our look into really, truly moving out of the garage—they were also the ramp-up to Mjolnir and Gungnir, our two most ambitious products to date. I've already covered some of the engineering challenges presented by these products, but that wasn't where the pain ended—not by a long shot.

First, Gungnir and Mjolnir broke our chassis design. Somewhere in January 2012, I submitted drawings to our chassis provider. Like all of our products, they were simple two-piece designs—an outer "U" and an inner sled. No problem, right?

Wrong. A couple of days after I submitted them for quote, I got a phone call from Russell, the guy we work with at our sheet metal fab shop.

"We can't make these parts," he said.

"Which parts?"

"The 01-25 and 01-30," Russell said, which was code-speak for the Gungnir and Mjolnir outer aluminum chassis, respectively.

Aside: Parts numbers, internal and otherwise. Okay, it's time for a lapse into engineer-ingland here. If you're going to start a company that makes any kind of custom parts—chassis, transformers, knobs, bolts, whatever—you're going to have to get used to part numbers. Suppliers don't take you seriously without them. They are simply not comfortable with saying, "Hey, the Gungnir outer chassis has a problem." They'd much rather say, "Hey, the 01-30 isn't producible, can you make these changes and send us a Rev B drawing?"

Ah hell, let's talk about parts numbers *and* revisions for a bit. Both are important. Because if you order an 01-18 Rev C when you intended to order an 01-18 Rev F, you're probably going to be boned.

"So what's this weird crap you're talking about?" You ask.

Let's break it down:

Parts numbers. When you create a custom part, you should assign a unique part number to this. Now, this doesn't have to have an bazillion-digit code like a UPC, or be done in hexadecimal or Klingon. But it should have a part number. Some companies break down their internal numbers with a prefix and suffix, like this:

- 01-XXXX: Wire
- 02-XXXX: Screws
- 03-XXXX: Sheet Metal Parts
- 04-XXXX: Milled Parts
- 05-XXXX: Cast Parts
- 06-XXXX: Transformers
- 07-XXXX: Capacitors
- Etc.

Now, that's pretty helpful to you, if you're looking for a part and you forgot what it was. But of course Schiit wasn't so organized. All of our custom parts are simply 01-XXXX. Knobs, chassis, transformers, whatever. It's relatively simple, and we don't have that many parts (though, to be fair, we just did drawings for the 01-132, so maybe we should think about segmenting it.)

But we probably won't change.

Why?

Because that's a whole new bunch of pain, because we'd have to re-educate our suppliers on the new part numbers, which would probably result in some mis-orders. And mis-orders mean backorders for you. No, thanks.

Revisions. These are what happens when parts change. If you're betting you'll get the drawings right the first time, you're probably wrong. (Though, admittedly, the first articles of Ragnarok—yes, Ragnarok, with about ten billion holes and super-complex PC board layouts—fit the first time. This is not called "mad skillz," this is called "damn f'n lucky." Of course, the finish work was so terrifying it was unsellable, but that's another matter.) If you think you won't have to make changes over time, you're wrong.

This means your drawings should have a revision level, usually specified as a letter, like "Rev A" or something like that. So when you change the location of the indicator dot on the knob, it's now "Rev B." And when you find that it doesn't fit the shaft of the pot, and you have to change the drill size, it's now "Rev C." And so on.

Revisions should be specified:

- **On the drawing.** So you know what it is, duh. Protip: add a "Revision Notes" panel to

remind yourself what you changed with each revision. Trust me, the vendors will thank you.

- **On the file name.** Seems basic, but you'd be amazed how long it took us to figure it out.
- **On the purchase order.** This may be less obvious. But if you end up with a whole boatload of parts that don't fit because you specified the wrong rev, you're going to eat them—and customers are gonna be howling about the backorder.

Learnings, or Why It's Not Always Bright to Think Everyone's Like You

When I was at Sumo, I thought all this part numbering business was a gigantic pain in the ass that made it impossible for people to know what the hell they were doing. I mean, why call a 121 Ω, 1/4 W resistor an 05-1262? Why not call it what it was? Wouldn't that be a lot easier?

Turns out not so much. By thinking "this is a pain, people won't know what part it is," I was actually thinking, "I, as an engineer, think this is a pain, because of course I know it's a 121 Ω, 1/4 W resistor, like duh, hell, you can see the stripes on it."

In reality, the people putting the products together (or, today, the robots) don't care what it's called. An 05-1262 has no more or less meaning than a 121 Ω, 1/4 W resistor. And when you get into chassis or custom parts, something like "the new, non-screwed-up Gungnir tops" is a whole lot less descriptive than an "01-31, Rev F."

So, if you're going to be starting a business with custom parts, I'd recommend the following:

1. **Set up a parts numbering system** that covers, at least, every custom part. It doesn't have to be complicated, but it should probably be segmented. Especially if you plan to produce more than a couple of products. This will get you taken more seriously by your vendors, and (believe it or not) will save you pain in the long run.
2. **Document all of your revisions**, and do everything you can to label revisions correctly. There *will* be changes. Yes, even on 3D CAD pre-fitted, pre-qualified-with-the-sheetmetal-module files. Your vendor will need to know what changed between revisions. And you won't want to be ordering 1000 pieces of a wrong Rev that doesn't fit anymore. Because those go straight in the trash can. And your vendor will be more than happy to point at

the revision level on the purchase order, and say, “We’re very sorry, but it’s your own stupid fault. Want to place an order for the right part?”

I guess what I’m saying is that working with external suppliers, and working with external assembly, is kinda like writing code. You want to be very explicit, and make sure your syntax is right.

Now, some vendors are gonna be really good, smart, and on the ball. But I still wouldn’t want to tell them, “Hey, make this part this new way,” and expect that a verbal change will filter through to the final delivery.

Use part numbers. Document changes. Pay attention to Rev levels. You’ll thank me for it.

Fun fact: Schiit is up to Rev H on some parts. Yes, even our simple stuff. Revs happen. Keep them straight, and your life will be a lot easier.

Back to Russell …

Wow, that was a hell of a diversion. Let’s get back to Russell and the non-producible chassis.

“Why can’t you make them?” I asked. “They’re just like our other parts.”

“They’re too deep,” Russell said. “We can’t bend something that’s 8 inch deep on both sides.”

Crap, I thought. How the hell were we going to do Mjolnir and Gungnir, then? A boring conventional chassis with a front panel? A front U and extensions? My mind quickly started running through the variations. (No kidding, I frequently think about how to put chassis together in interesting ways. Yeah, not exactly Running with the Bulls, but I think it’s fun.)

“Can you get a bigger brake? A different tool?” I asked Russell.

“No can do. The problem isn’t the depth so much, it’s actually getting the tool in there once it’s that deep. If it was a J-bend, sure. But a U-bend won’t work.”

Note: By J-bend, he meant a piece of metal where one side is much shorter than the other. So, when seen on-end, it looks like a J.

But a J-bend wouldn’t work for us. We had to transfer heat from the bottom to the top of the chassis. That was the beauty of the Asgard-style chassis design. It was a heatsink as well. But to work as a heatsink, it had to transfer heat. A

J-bend, with a break in it, wouldn't transfer heat. The Mjolnir would cook.

"Let me see what I can come up with," I told Russell. "I'll see if I can get you a new revision tomorrow."

Then I got to sketching. As in, with a Rotring pencil and big eraser. You kiddies can laugh at the dinosaur now. But I don't think you can iterate ideas any faster in 3D CAD than you can with a sketch. I took Russell's idea—the J-bend—and sketched up ideas that would allow us to transfer heat between the bottom and top chassis. The first sucked—a 1/8 inch thick bracket to attach the two pieces. But that was a whole nother piece. And it would have to be tapped, or have PEM nuts inserted into it. Which would drive up cost. I could do a joggle bend, of course, but that would look terrible. The inner chassis wouldn't be able to hide the joggle …

… unless I did only a partial joggle, and left the outside flush.

There we go.

And that's how Mjolnir and Gungnir got a three-piece chassis with a joggle bend hidden on the

bottom. Because it couldn't be done in a single piece.

"Well, I'm sure you could have found someone out there who could do it in one piece," someone is saying.

Yeah. Maybe. And maybe they would have cost 5× as much. Or delivered crap. Or a thousand other things. Believe me, in sheet metal, the devil you know is usually much, much, much better than the one you don't. If you're contemplating a vendor change, do it when:

1. You don't need them.
2. You're very happy with your current supplier.
3. You have a ton of extra time.

> Or, in other words, it'll probably never happen. But that's another story.

To make a long story short, I got new drawings out to our metal guys, they nodded approvingly, and they got started on building the first article metal.

On First Articles and Cheapness

First articles are the first "proof of concept" metal from your chassis supplier. This is what they make so you can:

1. Check fit
2. See what the finish is like
3. Make crappy-looking prototypes and show them to people who won't understand they're prototypes, no matter how big the signs are.

First articles you pay for. No metal supplier is going to do them for free. At least not for a small company.

Because we were cheap, we got first articles that were unfinished. As in, no paint on the steel, no screens, no anodize on the aluminum. This saves quite a bit of money.

But, unless you're using the first articles ONLY for fit, it's best to get them done all the way. That way, you can see what the finish is really like, if the screen lines up with the holes, and you can take it to shows and tease people with it, if that's what you're into. This is what we do today. In the past, we were cheap—which resulted in overheating Mjolnir prototypes (unanodized aluminum is a really crappy heat radiator) and a crap-looking Ragnarok that wasn't ready for prime time.

On Metal, Transformers, and Announcing Early

By the time we were ready to think about announcing Mjolnir and Gungnir, we were smart enough to know that pre-orders weren't a good idea.

"We'll do an interest list instead," I told Mike.

"An interest list?" he said.

"Yeah. We'll put complete product info up, but instead of taking orders, customers can leave us their email and check which products they're interested in."

"Oh, so like a pre-order, but with no credit card," Mike said doubtfully. "So you'll still have to answer all the questions about 'hey, when's this gonna be out?'"

"No. It's just a list. And we'll set availability at, say, 60 days out. That'll give us plenty of time to deal with any glitches."

"Hmm," Mike said.

"What hmm?"

"Glitches always take twice as long as you expect. That's the Second Law Of Vendors."

“There might not be any glitches.”

“Right, and the Pope might convert to Judaism,” Mike retorted.

“Look, I think we’ve got this figured. Lyr was out ahead of time … ”

“And Bifrost wasn’t. And I know how crazy you get when everyone’s hounding you.” I sighed. “I like to think positive.”

Mike shrugged. “I prefer to be realistic. But if you say we’re doing an interest list, that’s what we’re doing.”

“It’ll be fine, you’ll see.”

“Uh-huh,” Mike grumbled.

Of course, you know how this goes. Mike was absolutely right. In fact, he was actually thinking positively when it came to glitches.

Because, a week after the interest list went up, we got the transformers in (01-21 and 01-22, for the parts-number-centric out there.)

And the 01-21s hummed like refrigerators. I mean, if we’d used those transformers in Mjolnir, it could have been a headphone amp/massager product.

Now, of course, the prototypes didn't hum. But production did. And this time, it wasn't our fault. The Rev was correct. But the vendor had built an earlier rev. And they were junk.

Okay, not the end of the world. We still had 50 days or so left. Transformers made in the USA can be had in 3 to 4 weeks, no problem.

But when you factor in another round of prototypes, and a short run to make sure they were really, really quiet, well … the time stretches out. We ended up getting the replacement transformers only about a week before the initial release date.

Which still would have been fine, except for one small thing.

About 30 days after the interest list went up, we got the metal. And it was crap. The graining was completely random, with skips and hops and slips like tire tread. Completely unacceptable for a relatively expensive product. I got on the phone again."Russell, what the hell happened?"

"Ah. Yeah. Our timesaver bearings are maybe a little woppity. We need to replace them. But it's a custom part. It'll take a while."

My stomach sank. “And you thought these were good enough to send to us?”

“We knew you needed them fast.” I groaned. Of course. We’d been pressing them to deliver on time. It didn’t excuse the quality, but I could understand why they did what they did. They didn’t really know what kind of finish we needed.

After I explained this to Russell, he said, “We’ll new ones to you as fast as possible, but it may take a while for the bearings.”

Luckily, they got the bearings in within a week, and promised new metal in 3 more weeks. Which meant we could still make it.

But I could see Mike smirking in the background. And when the metal came in—on time—I had no reason to gloat. Because the bottom chassis were still crap. All of them had a big gouge on the back of the chassis.

It was only days until the release. I called Russell. “What the heck happened? They’re all messed up the same. There’s this big mark on the back!”

“Oh, yeah, that’s where they have to clip them for anodizing.”

“Then clip them somewhere else! We’re a week away from launch, and on the second set of metal. We can’t ship these! Pick ’em up and fix them.”

“Will do,” Russell said.

“How long until we get them back? Like I said, these are promised in a week. Boards are at the boardhouse. This is the only holdup.”

“I’ll see.”

“Make it fast. Please.”

“We’ll do it as fast as we can,” Russell promised.

Well, ‘as fast as we can,’ ended up being about 5 weeks. Those of you who remember the Mjolnir launch remember the delay.

And so, even with the best planning, even with a nice big buffer between announcement and scheduled ship date, even without taking pre-orders, the launch was still a bust. Mike had been absolutely right. We should have just kept our mouths shut.

That one glitch pretty much wrecked the early summer. If we didn’t run with substantial cash reserves (we are *extremely* conservative), very, very bad things could have happened.

Which is maybe the most important lesson. When your vendors catch a cold, you get sick. When they have a problem, it's your problem. Your customers don't care about excuses or The Reality of Making Things Today. They want their stuff. When it was promised. Period.

After Mjolnir and Gungnir, it was clear what we had to do: *never pre-announce a new product, ever again.*[1]

1 And yeah, yeah, we did talk about Ragnarok and Yggdrasil before they are available, and they're still not available, and we're massively late. When they are both shipping, I'll feel free in a way I can't describe to you. Because then, nobody will know what we're shipping next. And there will be exactly zero pressure to ship a partially-worked-out product to an artificial schedule.

Chapter 22

Introducing the Schiithole

Okay, the last time I talked about moving out of the garage, I left you hanging—with the realtor saying, "Well, maybe we can find a way for you to do your light manufacturing in a zoned-for-commercial space, but no promises."

To cut to the chase, we got it. What finally sold it was probably three things:

1. Persistence.
2. Not appearing too flaky or insane.
3. Willingness to take the property as-is.

Believe it or not, #2 matters quite a bit. There are plenty of flaky, insane people in commercial real estate. Landlords want nothing to do with them. And #3 is also a big deal. Landlords don't want to do a bunch of custom buildout—even if you're signing a long-term lease.

And—one other thing: the willingness to take a little bit of risk. Because we were in a commercial space, after all. Not industrial. The city inspector

could conceivably come by, decide we weren't conforming, and shut us down.

Yes, it was a risk.

But it was a risk worth taking, because it got us an inexpensive space near our other office. And the risk, we told ourselves, wasn't high. After all, we were retailers. We sold direct to customers. And if 99.99 % of it was shipped via FedEx and USPS, did it really matter, as long as we had a place that someone could theoretically walk in and purchase something?

Perhaps. But that would be up to the inspector. If they ever came by. I crossed my fingers, hoped they wouldn't, picked up our first liability insurance*, and signed the lease.

And that's how, around March 2012, we got the worn and pitted keys to our first Schiit building.

* Hidden Expenses and DIY Dreams

I precede this aside with a * to connect it to the liability insurance mention before. I do this because this is a great corollary to "everything is your problem" reality of having your own business. Well, here's the second harsh reality of your own business: *there are a hell of a lot of*

hidden expenses ready to jump up and bite you in the ass. I sometimes get taken to task by DIYers who say, "I can build something like that for a lot less than you're charging."

Sure.

Well, maybe.

No. Wait.

Actually, they can't, because the price on single pieces of stuff like transformers and chassis will make any one-off a budget-busting exercise. Even if you're talking off-the-shelf transformers and project boxes, it really isn't going to be that much cheaper. And that's not factoring in the time the DIYer spent building it, nor the time it took to learn their construction skills, nor the cost of their tools, nor the cost of the new tools they had to get while making it. You get the picture.

And, just to be clear, I love and support DIY. As far as I'm concerned, we should all be juggling soldering irons and dropping them in our laps, grabbing on to 120 V (or 230 V) once in a while, putting transistors in backwards, watching capacitors explode, spending endless hours wondering why the new prototype doesn't work quite right, getting excited when the new PC

| boards come in …

But DIY isn't production. It's not production in a garage, and it's not a business with all sorts of crazy expenses. Expenses like:

- Liability insurance
- Product liability insurance
- Workman's comp insurance
- Facilities lease
- Facilities upkeep
- Facilities changes/expansion
- Equipment cost
- Equipment upkeep and calibration
- Bookkeeping
- Local gross receipts tax
- Sales or VAT
- Business licensing/registration (if applicable)

And this is on top of the normal, fun stuff like local, state, and federal taxes, payroll, parts cost, shipping, assembly cost, etc.

Yep. Tons of fun.

Mike's Perspective

Mike, of course, saw right through me, as soon as he drove by the place. He took one look at Rina running a shop-vac over the cracked and dusty

floor (and vacuuming up big pieces of ancient tile in the process), and told me:

"You got it because it was cheap."

"Right," I told him, not even hesitating.

Rina and Eddie were also arguing over space for shipping versus space for production. Eddie was arguing that we should take all the used, battered Ikea office desks (that I got from Centric's storage unit) back outside and blow the dust out of the space with a compressor.

"There's about a hundred pounds of dirt per square foot up there," Eddie said, pointing up at the sprayed acoustic ceiling. "That crap's gonna fall down if we don't blow it out."

"It may fall down if we blow it out anyway," I told him. The ceiling didn't look too robust. We could be looking at sheets of acoustic cottage cheese if we started blowing on it.

"But it was cheap!" Mike said.

"Yes," I snapped. "And convenient. And it keeps the HOA from shutting us down."

"Until the city inspector comes."

"If he comes."

"I'm going to need more space for shipping," Rina interrupted, indicating where she wanted her finished-goods racks placed.

"That's two-thirds of the building," I told her. "Eddie and Tony need more space."

"We need more space for shipping!" she insisted.

Eddie shook his head. "And we should stop this crazy vacuuming and blow this place out. This dust is gonna get *everywhere!"*

"But it was cheap!" Mike added. I groaned. Thankfully, Tony wasn't there, or would have probably had a comment or three as well. Suffice to say, it wasn't our finest day moving in. The place was really a mess.

Maybe I should step back and describe the Schiithole (the name Rina dubbed it with on that first contentious day—and it stuck.)

The Schiithole was an old, L-shaped stucco-and-siding building on the corner of 6th and Railroad in Newhall. Railroad is named because of, well, the railroad that parallels it. This railroad was instrumental in making Newhall one of the first boom towns of the late 1800s (together with the discovery of oil.) It now carries mainly Metrolink traffic. Many times a day, trains rattle by, shaking

the old building. Cars rush by on the 5-lane street outside at all hours. With no insulation, it was a hollow, loud, booming space.

Outside, the stucco was fading to an off-white from what had once been a taupe color. Several large holes had been punched in it, whether from frustrated passerby, or by some other mechanism, I don't know. The siding was peeling paint, and much of the external wood was collapsing into dry rot.

Inside, the floor was uneven and patched crudely with concrete. And by uneven, I mean "like, one side was a good foot lower than another." Some traces of fiberglass tiles remained, but they were rapidly flaking off as we cleaned. The sheetrock walls were relatively unmarked, but the bathrooms were only partially functional, after the building had been stripped of plumbing and wiring by thieves and only makeshift restored. There was no hot water. No heating. No AC.

Out back, it had a dirt-and-concrete yard full of knee-high weeds and assorted detritus. Of the four doors, one looked out on the back area, two were on the Railroad side, plus a roll-up door the realtor warned us "never to leave open, because that's a sure sign you're making things

here," and a door on Market Street. None of the doors matched. None of them were handicap accessible. I didn't know it at the time, but it had once been the home of the Daily Signal, Santa Clarita's newspaper that still survives (barely) to this day. We were later shown the room where they used to melt down the lead plates every night, to re-cast them for the next day's paper. We were using it as a storeroom.

But, as Mike said, it was cheap.

So, yeah, a fun move. But by the end of that first day, we had test equipment plugged in and running, Eddie had his assembly bench set up, and we were moving things in from the garage and the rest of the house.

We had a home away from home. The Schiithole.

Interlude: Business Space Philosophy

Okay, so why did we take such a crappy place (because it was cheap) and tolerate the noise (because it was cheap) and the dust (because it was cheap) and the lack of any typical niceties, like heat and AC (because it was cheap!)

Well, yes. But also because of an observation I've had about business space:

The moment you build a palace is the moment you die.

Now, it may take many years for that palace to kill you. You may end up with some very good years there. It may serve as a very useful way to awe and astound customers or clients that are easily impressed by such things.

But the moment you start focusing on business *wants,* rather than *needs,* you're dead.

It happened to Sherwood. It happened to Marantz. It happened to dozens of ad agencies I've seen come and go. It's happened to scores of clients who spent their startup money on nice offices and celeb chefs and foosball tables and lounges. We'll see how Apple's "spaceship" campus does for them, but I'm betting right now I know how it ends.

Here's the deal. I business, there are certain things you need. These are things like:

- **A functional space.** That is, one large enough to contain what you want to do. Trying to put a full wood shop in a 20 foot×20 foot space isn't gonna work out very well. Depending on what you're doing, you may need no more than a single office—or a home office—or a whole lot more space for stock, finished goods,

production area, machinery, etc. But first and foremost, you need a functional space.

- **Effective places to work.** This may mean an Ikea desk and a Wacom Cintiq and a WiFi router and nothing more. Or it might mean heavy-duty workbenches with static mitigation and 10 outlets for work on sensitive electronics. It it may mean self-built 2×4-and-fiberboard tables with static mats for accomplishing the same thing. Plus chairs and such.[1]
- **The equipment you need to do your work.** This may be nothing more than a laptop, or it could be an entire suite of test equipment. Or CNC routers and laser engraving machines. Bottom line, stuff like this is critical, if it is critical for doing your work. Don't skimp.
- **The right connectivity for your business.** This may mean nothing more than a simple DSL line for basic spreadsheets, web surfing,

1 I have never seen such a rip-off as high-end office chairs. Please don't get me started on this. Might as well buy audio jewelry lovingly carved from hunks of solid titanium by master craftspeople living in Monterey. Just go to Office Depot, plop your butt in a bunch of chairs that are $150 or less, and pick the ones that are most tolerable and cheapest. Our creative director once tried to talk me into getting Aerons for the whole office. I swapped his chair for a steel folding chair the next day. We bought sensible chairs.

and ordering, or super-high-speed optical cable with dedicated symmetric lines for a phone bank to serve a creative shop with outbound calling.

Please notice that none of the above includes things like:

- **A cool-looking building.** Please. Who cares? Save money for your house.
- **A cool-looking office space with polished concrete floors, $16 000 European couches, and iDevice-controlled programmable LED lighting.** See above. If stuff like that impresses your customers or clients, they're not analytical enough to understand:
- Every cent of that came from their pocket.
- You're making far too much money.
- You're not good with impulse control and probably won't be around long.
- **Corner offices for everyone.** Geometrically impossible, anyway. Plus, let the infighting begin.[2]

---

2 Let's talk about this a little bit more. The biggest fights I've ever witnessed amongst employees was regarding "who gets what office," or "who gets which desk." Honestly, this is completely useless and divisive stuff that you really don't need to deal with. Start-ups probably shouldn't have private offices, period.

- **Nice private offices for everyone.** Yep. Let's give them reasons to close the door and hide from problems. Not a great idea, especially in a startup.
- **Lounge/recliner/videogame/relaxation/informal meeting areas.** Yeah, you love your employees. But it's more important they love you and believe in what you're doing, without having to be tempted by silly perks. The best person for the job is one who wants to work for you above all others—because of what you're doing, not because of the icing on the cake. If nobody wants to work for you, start looking in the mirror—and hard.
- **A big sign in front of the building.** Might as well advertise "there's expensive electronics in here," in our case. In most cases, this is nothing more than ego.

And that's why we ended up with a space that was really nothing more than a large production floor, with no offices, in an ugly, run-down building. Because it had what we needed. And nothing we didn't.

And it was cheap.

The Earliest Days of the Schiithole

Going to the Schiithole in the early mornings, shortly after we moved in, are some of my most vivid memories of Schiit, beyond the early start-up phase. Why? Because we were finally in our own space—and that opened up so many new possibilities.

And because I was crapping-my-pants busy.

The move, as simple as it was, put us back a few days in production. And in those times of "build tonight, ship the next day," that means that we were very behind. Tony and Eddie would come in during the evening and build the boards that Jaxx delivered, but I was doing all the sound-checking.

So, I'd come in before going into Centric (5 AM to 6 AM), take the plastic sheet off the pile of finished units on the burn rack, and run them through sound check. If any failed, I'd note it and put it on my desk for later that evening, when I'd come back in from Centric at 6 PM to 7 PM, and fix whatever didn't make it through burn.

Rina would come in during the afternoon and ship, but her time was starting to come at a premium—her own business, Twilight's Fancy,

was taking off. She had to spend more time there and less at Centric.

And I was quickly burning out from the long days. You can do 14-hour days for a while, but they'll eventually kill you (if your significant other doesn't do so first.)

Tony wasn't going to be able to take over shipping—he had his hands full, especially with Mjolnir and Gungnir imminent, and Bifrost flying off the shelves. Eddie couldn't do shipping, either—he liked to work at night, when it was cooler and quieter in the shop.

Yep, you know where this is going. It was time to grow again.

And this time, we needed to look at it more strategically. Sure, we needed someone who could ship. But what we really needed was someone who could do a lot more than that—someone who could grow to be in charge of operations.

Luckily, Rina had the "perfect" candidate …

Chapter 23

"I Didn't Know People In the Private Sector Were As Lazy and Incompetent As the People In Schools"

I've already mentioned some "invisible lines" in business—hiring your first employee, getting your own production space, etc. Well, here's another: hiring your first employee who won't directly be producing products.

Yes, your first manager. I can hear the groans now. Well, I hate to break it to the fans of completely flat organizations, but yes, some management is necessary. That is, if you ever want a snowball's chance in a blast furnace of ever having a life. If you ever want to take some time off, have a vacation, or do any of those things that non-workaholic people do, you'll need management that ain't you.

Why? Because, believe it or not, you can't do everything.

Repeat that again: *you can't do everything.*

No, let's turn that up: *you **shouldn't** do everything.*

Now, before you all think I've become some tool relaxing on a beach, espousing the benefits of outsourcing every part of your business and life so you can sit on your ass some more, building your empire on a cloud of of a temporary situation in which some places in the world offer very low-cost labor, relax.

Similarly, if you're thinking I'm going to be retiring to the CEO's giant office, reclining in a $6000 Italian-leather office chair, foregoing the day-to-day work, and focusing on creating dream-team org charts while the company is mismanaged into non-existence by the direct reports under the dream team, oh *hell no* to that as well.

No. Let's be clear. We're not talking about a billion-dollar "hey, I won the VC/startup/etc lottery" mass-market business here. We're talking niche. We're talking niches where people get up in arms about the tiniest decisions you make. We're talking an environment where micro-social will dissect everything we do.

Get out of your business and you kill it.

Hire a "dream" CEO to run your brilliant idea, and it's dead meat.

Ignore the nuts and bolts and focus only on a "vision," and you might as well throw in the towel.

But, at the same time, this doesn't mean you have to do *everything*. Which means hiring people with responsibility, intelligence, and the ability to make a decision. Call them "managers," "smart dudes," "team leaders," or what-have-you, but you will need them to do things you can't do, don't have the time to do, or don't want to do.

But this is a massive change. Once you've hired someone in management, you've now put on direct overhead. You have a salary that isn't totally dedicated to pushing products out the door.

It's a gigantic change. To put it simply: *once you've hired a manager, it's a real business, not a hobby.*

Jason's Thoughts on Business Management, or, Saving Face Before I Get My Ass Kicked By Flat Organization Fans

Now, before we get any further let me be perfectly clear about how I feel about management layers in a business: they're a necessary evil. Extra layers should be avoided at all cost. Like nuclear waste, you don't want to get too much management on you.

The problem is, the vast, vast majority of businesses in the USA have far too much management. Instead of treating it like nuclear waste, they treat it like an all-you-can-eat free buffet … the more, the merrier.

How much management is too much?

- How about a 150-person company with 17 vice-presidents?
- What about an 80-person company with 65 in management and 15 in production and shipping?
- How about a 200-person company with 6 layers of management: C-Level, EVP, VP, Director, Senior Manager, Manager—with anyone below the EVP level unable to make decisions on expenditures that cost more than $1000?

Yep, been there, seen all that. And lots more.

Why do companies have too much management? Three reasons:

- **The 20/80 rule, or hiding in a crowd.** When a company reaches a certain size, and the founding management checks out of the day-to-day reality, it's very easy for groups to form where 8 out of 10 people are simply hiding in the crowd, and only 2 people are actually doing the work. Solution: add a manager. All too often from the 80 %, though.
- **The "ain't my job" CYA factor.** When something new and scary is thrown at a company running the 20/80 game, the natural reaction of most 80 %ers is to run away. It might fail. They don't want to get that on them. (Seriously, at my first corporate job at Magnavox, I was pulled aside by the Senior Team Leader and told very simply, "Never volunteer. It makes us look bad. And you're working too fast.") Good thing there was a manager and senior manager above him, before the department head and then the "real" engineering management.
- **Title inflation.** How do you make someone feel good if you don't want to give them a raise? How about a fancy title? Yeah. It really

is that simple.

Mike and I took the titles of “Co-Founder” when we started Schiit, rather than “CEO/President” for the simple reasons of:

- They describe accurately what we are
- They set no limits on what we can or should do
- They are about as anti-inflationary as they can get

So, to this day, and into the foreseeable future, Mike and I are as hands-on as we want to be, or need to be. Mike directly designs pretty much all of the digital products. I do pretty much all of the analog products. When digital problems come up, Mike solves them. When analog problems come up, I solve them.

Will we some day have engineers working with us doing some of that? Maybe. But I don’t see us getting out of those arenas entirely, or anytime soon.

Similarly, I do the marketing and the product cosmetic design. Will I get out of that? Probably not completely. But will I hire a marketing person to deal with the mechanics of ads, shows, etc, and to beat off the ad salespeople with a stick? You bet. But not now. Not yet.

So what don't we do?

Well, Mike and I are both what you'd call "administratively challenged." We're not good at detail. Rely on us to make sure the parts show up on time, that we build the right mix of 115 V and 230 V units, or even that we get to the right show on the right date … hmm. Maybe. Maybe not.

Or to put it in more typical business language: we clearly need help in operations. And that brings us to our first management hire …

Let's Meet Another "Ideal" Candidate

Okay, let's first set the scene. About the time I start thinking, "Hey, I need someone who can do more than test and assemble," we're not even yet in the Schiithole. It's very close, though. Imminent. And I'd really like to have someone in there, ASAP, someone who isn't Rina, Mike, or me.

At the same time, we're starting to think really hard about the Mjolnir and Gungnir launch. Which means the workload was going to get even bigger. The days, longer. Compounding this, Rina's own business was taking off, so she had less time to devote to Schiit.

We needed to hire someone. At least to help with shipping. But I really wanted more—someone who could grow into our lead operations guy. Call him the Proto-COO, or Future Director of Operations, or whatever title sounds appropriate. But I knew we needed someone who was much more than a clerk.

That's when Rina had a brilliant idea.

"Hire Alex," she said. "He could run this whole business."

"Alex?" I asked.

"Jen's husband." Jen was Rina's friend from way back. They've written books together. Currently, they're producing a web series (necrolectric.com). I knew Jen. I'd met Alex before, but I wasn't intimately familiar with his background. I remembered something about computers, but that's about it.

"What does he do?" I asked.

"He'd be perfect! And it would be great. And they've always wanted to move out here."

"Wait." I remembered that Jen and Alex lived in Hesperia. To those not familiar with California,

that's a hell of a commute to Newhall—about 90 minutes one way … with no traffic.

"But he could start part-time," Rina pressed. "He wouldn't have to come out here every day … " I shook my head. "Part-time," and "operations," don't really go together. And anyone who was willing to do part time …

"What's he doing now?"

Rina frowned. "I told you. He was an assistant principal at—"

Rina kept talking, but I didn't hear her. Every negative stereotype of someone who'd worked in the school system came clamoring to the fore, shouting for attention. *They were overpaid. Lazy. They could never make it in private industry. They'd fold in a second and run back to the land of fat, tax-funded pensions.*

But at the same time, I remembered talking with Alex—and he was sharp.

"Um," I stalled.

"He really would be perfect!" Rina said. "I know it!"

"But … "

“And when they’re living here, it’ll be even better.”

“I really don’t know about this,” I told her.

“You should give him a chance … ” I knew where this was going. I didn’t want to argue, so I pretended that I had something to do (or maybe I really did have something to do) and begged off of making a decision.

Out of the Frying Pan …

But Rina kept hounding me. And that’s how, shortly after we took possession of the Schiithole, I found myself talking to Alex over breakfast with Lisa and Jen.

And … you know what? Alex was sharp.

What’s more, he didn’t flinch when I spelled out exactly how it had to be, to start: he’d be a contractor, not an employee. The hours might be a little sketchy, depending on demand. The place he’d be working in was not a palace. He was going to die in the summer heat. And he’d be having to learn the ropes of audio—an insanely picky, somewhat neurotic, completely unique market that he had no deep experience in.

No. Instead, Alex asked reasonable questions, like when we'd end up switching to salary (soon, I was already getting nervous about the "contractor" definition, and I knew it had to go), opportunity to move up (there is no ceiling, I told him the story of Mike and Centric and creating our first interactive department from nothing). And the story of Theta and its bonuses.

And, during that breakfast, I began to get the feeling: Alex may just be one of those diamonds in the rough, like Eddie and Tony … but on a much more functional frequency. He was clearly smart enough to get in and understand our operations.

But was he? Those old prejudices about people who worked for schools kept nagging.

It didn't matter. It wasn't like Mike and I were going to be organized enough to put together a help-wanted ad and screen the applicants. It would be easier to do the "sink or swim" thing again, just like we'd done with Tony and Eddie.

So what did we do? I said something like, "You want to get started now?"

That afternoon, Alex was helping pack and ship Valhallas at the garage. From there, he helped us move all the crap out of the garage into the Hole,

then trained with Rina on Mondays, Wednesdays, and Fridays for the next couple of weeks.

After that, it was pretty much the Alex Show.

Alex determined that we needed more racks for better segregation of finished goods—so he went out and bought them. He decided to reorganize the packing and loose parts. He cleaned up the Schiithole to the point where it was much more livable. He oversaw the installation of a security system, told us when we were running out of space, found a place with mobile storage units, and added one in the back. He let Eddie know what to build and Tony what to test. After some short training with me, he took over a lot of the part ordering. He didn't complain, even as the days got hotter and his hours got shorter because things were slow.

And he did this all while he and Jen were looking for a house out here, closing the deal, and he was spending his evenings renovating the house, as well.

And that, my friends, is what initiative looks like.

Alex was the perfect candidate. Rina was entirely right.

Now, Alex got to interact with a lot of our vendors—the metal house, the PCB house, the electronic parts suppliers, the packaging guys … and I think this was a revelation to him in more ways than one:

- **Revelation One:** he didn't have to fill out 50 pages of paper to buy $500 worth of stuff. I'd just say, "If you want it, get it."
- **Revelation Two:** his own assumptions about private business were about as lacking as my own about people who work in schools.

One day, as we were finally shipping Mjolnir and Gungnir (after we'd gone through a couple of rounds of metal incompetence, and Alex had helped us work through them), Alex said something that still stands to this day as The Greatest Schiit Quote of All Time.

He said:

> *"You know, before I started this job, I had no idea people who work in private industry were as lazy and incompetent as the people who work in the school system."*

I busted up.

But I couldn't help thinking: *wow, that's a turnaround from my own prejudices.*

"When I was working at the school," Alex continued, "That's all they ever told us: that we'd never make it in the private sector. That we were so lucky to have these jobs. If we ever lost them, our lives were over."

"They actually told you that?" I asked.

Alex nodded.

"Wow." I shook my head.

"But I get out here, and what do I find? The same bozos." I laughed. "Yeah, being part of a private company doesn't mean anything. Especially in marketing. You know what we used to say about marketing?"

"No," Alex said.

"You go into marketing because you're not smart enough for science, and not ruthless enough for sales."

"Ouch."

"Yeah. And probably about as true as everyone in the school system being lazy and incompetent. The thing is, everything's a continuum.

There are sharp people and not-so-sharp people everywhere."

"I'm still looking for the sharp ones," Alex told me.

"Good. We'll need them."

And that really is *the* challenge, more and more: not letting your company fall to the 20/80 rule, where a handful of good people do most of the work. Not letting the bozos on the bus.

And, to this day, Alex has helped us do that. He oversaw the move into, and build-out of the space we're in today. He spec'd the racking and forklift we needed. He set up the shop layout so it would be more efficient. He's helped us work through half a dozen new product start-ups, and made decisions to hold back on shipping until he was happy with what we are making. He brought on the extra people we need for customer service and for shipping. And he's overseeing the buildout of extra space right now, as I write this. All transparently, without hours of micromanagement from me and Mike, and without us second-guessing everything he's doing.

In short, without Alex (or someone very much like him) we would not have been able to grow

as effectively. Schiit would have been much more of a burden on the day-to-day side. We probably wouldn't have been able to launch as many new products as we have, nor the flood of new stuff that's coming.

That's why you hire (good) management. Because it makes things better all around.

Note the (good). Because accepting mediocrity isn't acceptable. Nor is sitting back, crossing your arms, and saying, "You know, we're doing pretty well."

Bottom line: you can't tread water. You can't stand still. You have to sacrifice your babies. You need to look straight-on at cannibalizing your own products. You always have to be asking, "What can we do better, less expensively?" Even if it lays waste to your entire lineup.

Because, you know what? *If you don't do it, someone else will.*

Bonus Chapter

Schiit Happens (and Keeps Happening)

Hey all, something happened today … something that will tear me away from the torture/self-immolation that is learning more about 3D CAD. (Yeah, we're making that transition—all the new products will have chassis drawn up in real parametric CAD, which should speed up the prototyping cycle—yeah, you can start laughing now.)

So, about 3:00 this afternoon, I get a call from Alex. "Um, hey, I hit a sprinkler with the forklift. Can you give the building owner a call and see if they have a preferred plumbing contractor?"

> And yes, I was out of the building, and not planning on coming in that day—see my comments on "you can't/shouldn't do everything."

Hmm. *Yeah, that sucks*, I thought.

I'd had a sprinkler shoot out of the landscaping on the side of the building and land inside one of

my cars (a convertible), so I knew how that could suck. And the forklift—an electric model—wasn't so hot outside anyway, having about 1 inch of ground clearance. I told Alex I'd call the building manager, which I did. They didn't seem overly concerned, and said to use whatever plumber we wanted. I relayed this to Alex.

Crisis averted, right?

Well, no. What I missed from the original message was that the forklift had hit a fire sprinkler. As in, *inside the building*. As in, *above the Schiit assembly area.*

Yeah. A little different.

Rina texted photos to me about 4:15 or so, showing the shop floor covered in water and the fire department working in the background. That's when the light went off.

Oh, you mean ***inside*** *the building,* I texted back.

Like, duh, she sent back.

And sent more photos. This time, I saw what we were really looking at. The forklift had tweaked one of the fire sprinklers that was directly above the racking where we kept some chassis parts, and where Eddie and Miles worked. Imagine

turning on a proper 1950s-style shower (before they did the damn low-flow thing that makes you have to drill out every showerhead you buy) above a bunch of products and workbenches.

No, imagine it raining indoors. Raining hard. That's what happened.

Luckily, Alex, Rina, and Tony were on it. While Alex was calling a plumber, Rina started organizing everyone to move Schiit out from underneath the deluge. Tony called the fire department (they were fire sprinklers, after all.)

Fun fact: The fire department laughed and asked why we hadn't called 911. Tony said, "Well, it isn't really an emergency, sorry for bothering you." But they laughed and said "Hey, this is the most exciting thing to happen today. We'll be right out."

> If the hillsides had been burning, I don't think they would have been so bored … but hey, glad they could help.

By the time I got there, most of the affected parts and product had been moved into Unit B, which we're still in the process of moving into. Everyone was busy unwrapping stuff to let it dry. And unboxing products that got soaked.

Yeah. Products. All the Ragnaroks.

> Sorry, just kidding. I really wanted to see if you would have a heart attack, after all the hyperventilating in the Ragnarok thread.

In reality, the only things really affected were:

1. **Aluminum chassis** for a handful of products, which aren't watersensitive, but they were sorted and unwrapped as necessary.
2. **Steel chassis** for a handful of products, which should not be watersensitive, since they are powder-coated, but all were unwrapped and dried to be sure.
3. **Lokis.** Oh darn, DSD suffers another blow. But perhaps we should ask why there were a ton of Lokis up on the top rack, boxed and ready to go. The brutally honest reason? Because they don't move, even if we put Ex-Lax in the boxes. Lokis are a painfully slow seller. On the other hand, the DSD furore has not put a dent in the increasing sales of our other DACs, hence Mike's recent pronouncement that we won't be doing much more with that, unless Sony opens its vaults for real and we suddenly have 4000 albums to choose from, and not at $ 45 each—which Mike and I think is about as probable as Neil Young personally

delivering a palletload of Ponos to me before I finish writing this.

So what do we do with the Lokis? We dry them out and see if they work. If they do, we'll sell them as B-stock with full warranty. If they don't, they don't. Maybe we should have a special Schiit Happened Deluge Sale Loki model at a reduced price. I don't know. Help me out here.

In the end, yeah, it sucks, but schiit happens … and keeps happening.

Special Bonus Section: Q&A

Q: So will this affect the Ragnarok delivery date?
A: We recommend you switch to a decaffeinated brand.

Q: No, seriously!
A: Or take a long constitutional in a relaxing, low-stress part of the world.

Q: Why can't you take me seriously? OMG-WTFBBQ! I'm gonna explode if I don't get a Ragnarok in the next 30 seconds!
A: Or seek psychiatric counseling.

Q: Does this affect my current order for any shipping product?

A: No. Even if you ordered a Loki, we have tons of them on shelves not affected by the Great Schiit Deluge of 2014.

Q: Wow, how can you be so nonchalant about this?

A: Because:

a) I know I can't do everything. Nor can Mike.
b) Everyone performed above and beyond the call of duty. They even ordered in pizza (Alex, please expense this.)
c) Schiit happens, get over it.
d) I am probably (meaning certainly) drinking now.

Bonus Chapter

Perspective

Author's Note

Hey all,

Thank you all for the condolences. I wasn't in a particularly good place last week. This week is still a morass of things to deal with and unexpected surprises, but I am getting back into the swing of things.

So, what does this mean for Schiit Audio, Ragnarok, Yggdrasil, and everything else in general?

First, the Ragnaroks are not in the dumpster, and Mike and Dave continue to move ahead. I'm going to shut up on pronouncements about when you're going to see them … but you will. Mike (Baldr) can provide additional commentary, if he'd like.

Second, Mike and Dave are also moving ahead on Yggdrasil. Same here, not gonna promise dates anymore. Ask Baldr.

Finally, everything else? Everything else is moving ahead, much more rapidly than you might imagine! Everything else is fun. Everything else has no pressure of artificial launch dates on it, and we can take our time to get it right. It's a great time to be in engineering at Schiit … because you ain't seen nothing yet!

Yes. It's a bright future. And I'm going to keep it that way.

And so, instead of the next chapter I'd planned to write, I'd going to do something completely different. Here's a completely unplugged, unfiltered treatise on perspective. Perspective, both in business and personal terms. It's far too easy for everything to devolve into shouting matches about meaningless stuff like formats and technologies, while the big issues get a pass. So let's pull the elephant out of the corner and talk about what might be the hardest thing of all: maintaining perspective.

All the best,
Jason

A Prescient Comment

After a few pages of the DSD/PCM/provenance/upsampling/etc debate, do

you wonder why so many people think us audiophiles are a little bit nuts?

Time to listen to some music, I think. Does it sound good? Yep. Then all is well.

From the Schiit Happened Thread
Posted by Jason Stoddard
July 13, 2014

The Audio Biz and Loss of Perspective

Okay, let's get some stuff out of the way. In my opinion, we work in an industry with some profoundly broken corners. I've mentioned that Mike and I got out of high-end largely because we didn't want to chase the then-new trend of "superprice audiophilia." The price escalation for the sake of price escalation, with no new ground broken in terms of technology—that wasn't for us.

And today, it's a hundred times worse. People argue over $20k+ DACs. Reviews of $40k preamps are common. There are dozens of speaker models with retail prices over $100k. I was told that a "moderate price" system was $250 to $500k at a recent show, by a guy who said it with no trace of irony in his voice.

Let's be clear. This is insanity.

Obsessing over $ 250k system is out-and-out nuts, no matter how much you make. Period. Get out. Buy a Ferrari. Get laid. Listen to real music. Start a band. Travel the world. This is what people do when they haven't lost perspective.

Similarly, producing products that cater to this uber-priced segment is nuts. It just fuels an additional "my price is bigger than your price," escalation—and this escalation usually doesn't result in pushing the limits of actual audio performance, except in a handful of cases where implementation is astoundingly challenging (I'm thinking of discrete R2R DACs, and, to a lesser extent, turntable designs.)

Yes, I'm indicting an entire sector of the industry, but that is my honest opinion.

And … this is why I like headphone audio. By and large, the "price-is-everything" attitude is much, much less. Yes, there are expensive products. But not orders of magnitude more expensive. And, with a more tight-knit community, products that offer poor value are usually deconstructed pretty quickly. There's still a reasonable amount of perspective out there.

I'm really hoping we can keep our perspective, as personal audio grows up.

So how do we do it? Hell, I don't have all the answers. But I think I can at least outline the signposts on the path to lost perspective, and maybe, just maybe, help some companies and individuals avoid the loss.

Business Perspective, and Avoiding Devolution into an Algorithmic Robot

This is every businesses' dilemma: how do you avoid devolution from a human business, into a soulless robot driven only by algorithms and metrics?

Most businesses start human. People usually don't start a business only to make money and screw people on customer service. They usually start a business because it's something they love, and something they believe in. They put their soul into it. And this love and soul is reflected in everything they do. When a business is small, everything happens at a human level. When you talk to them, you're talking to a real person. When something goes wrong, it's a personal failure for

the business founder, and they scramble to make it right.

Then they encounter their first insane customer. Yes, they exist. And I'm not talking about insane in terms of "they didn't like the product, and returned it," I'm talking insane in terms of "they scammed you," or "belligerently tried to take down your business," or something of the sort. And they are out there.

That encounter chips away at that love. That soul.

Then they get another. Then they get yelled at for something beyond their control. Then they get dismissed as a hack or incompetent. And every one of those encounters wears away a little more of that love and soul.

Warning sign: if you ever start saying, "It's good enough," start worrying. You're starting to lose the love. Yes, even if you have to ship something late.

This continues as the business grows. Businesses get scammed by people with bogus credit cards. Or their dealers don't pay, if they're silly enough to have dealers and give them terms. They get people who make it their mission to take them

down, over some imagined (or sometimes real) slight.

And that leads to the first of three phases in becoming a soulless, algorithmic corporation.

Retraction phase. This is when you stop caring. When "Good enough!" becomes your mantra. When you start saying things like, "There's nothing we can do about it." You're pulling back from your business, removing the love and soul. When you start bitching about your customers in front of other people in the business, and spreading the contempt. This is a disease. It has to be killed before it spreads. Because, you know what? It's your humanity that separates you from the rest.

Codification phase. If you don't actively stop the retraction phase, you'll end up spreading it throughout the company, where it can end up being codified. Your production line will take "Good enough," and run with it until it isn't "good enough" to compete. Your engineers will stop caring about what they should be doing, and start copying other company's designs that are "good enough." Your customer service will stop answering inquiries quickly, and put in their own rules of "Well, 24 to 48 hours for response is fine," or "Put it in the ticket system, we'll

get to it." Once the company has reached this point, you're in big trouble. The founders may be celebrating success by buying Ferraris or taking long, expensive vacations, but the end is in sight. Soon, you'll end up like Time Warner or AT&T, where people only "Like" them on Facebook to bitch about their poor products and abysmal service.

Algorithmic takeover. Once it's been codified, it's time for the professional managers to move in. These are the number-crunchers, the bean-counters, the benchmarkers. They'll bring in data about how you're doing relative to other companies in your industry, so you can "improve" your processes. What this usually results in, unfortunately, is usually the same-to-the-penny offerings and beyond-abysmal level of customer service. Because nobody else is doing any better. And it's very easy to look at a CEO dashboard that says, "Hey, we have a 39 % higher customer satisfaction rating than our competition," without revealing that your competition's customer satisfaction is at 9 %. At this point, congratulations. You're not a person anymore. You're a robot, moved only by algorithms. If you're lucky enough, you may be able to move fast enough to survive.

So, how do you avoid this fate?

First, by recognizing why you started your business: your love of music, or cars, or code, or whatever—and never forgetting it. If it works to put a banner up on your wall that you see every day as a reminder, do it. Otherwise, make sure you have enough time to sit back and remember.

Second, by defining a philosophy. Your philosophy should be a natural outgrowth of what you love about your business. It will help keep you on track. Need another banner? Add it to the wall.

Third, by active reflection. Remember the great customers, the wonderful accolades, the moment you first held (or heard) a new product. Take time for yourself. Don't pack your schedule so full that you don't have time to sit back and put your feet up. Because your business has many wonderful times. Don't forget them in the rush to do "What's next!"

Or, as I've said before: *stay small, stay human.*

Personal Perspective, and Avoiding Mutation into an Inflexible Ass

This is a common personal dilemma. Someone buys an expensive product, and is told by someone that it "really isn't that good," which then turns

into a shouting match about minutiae that no sane person really cares about. Or they buy into an ideology that must be The One Right And True Way, and begin inculcating everyone around them into that Way.

Add instant, anonymous communication into the mix, and boom! You have a recipe for transforming an otherwise sane and rational person into a didactic, inflexible ass. Now, this doesn't happen all the time, but sometimes I wonder if the ability to disagree congenially is on life support.

And yes, I understand that this inflexibility can be an expression of personal philosophy, or love of an object or idea. That's cool.

But … ask yourself one question before you hit the keyboard: *does it matter?*

Most of the time, the answer will be "no."

Yes, I know, when you love something, or believe in something, it's easy to take any negative comment as a personal attack. And that might get you going to the point where you want to "educate" the attacker.

But will you convert them? No.

Some people love French wine. Some people like California wine better. Neither is going to convince the other with words. Some people love Corvettes, and some people love Porsches. Same deal. Some people love Schiit, and some people love other products. Same.

DSD vs PCM? Same.

Rock vs classical? Same.

Objectivist vs subjectivist? Same.

You know what, if you love it, then that's fine. It doesn't matter what other people think. And if you think you're going to convert them by pounding a keyboard on an online forum, or writing a book, or yelling at them in person, you're profoundly overconfident. Or a much better writer or speaker than I am.

> And, here's the thing: the other guy might have a point. Try something new. You might surprise yourself.

So, it's simple. Ask yourself: *Does it matter?*

And, if you want to get a little introspective, ask yourself: *What do I really know?* If you're just parroting marketing blather or some "consensus"

opinion you derived from reading two posts, well … you may not know as much as you think.

And, if you're interested in broadening your horizons, remind yourself: *Maybe I should try some new things.*

That's the path of sanity. Everything else, not so much.

And now, time for some music. Probably 16/44. Probably forgot to turn Bitperfect on. Who knows? Who cares? It sounds good. I like it.

And that's what matters.

Chapter 24
Getting Our Schiit Together

Astute readers will notice I swapped a new chapter in here. As to why, it's simple—I realized that the "Song of Ten Thumbs" chapter actually came a lot earlier than I had it pegged in the timeline … and that it really isn't just a story of incompetence, but a story of making big changes that set the tone for our company well into the future.

But don't worry—Ten Thumbs will make an appearance here.

Let's define the game first, though. The summer of 2012 was one of the most transformative times at Schiit, and it was largely due to Alex helping us pull our heads out of our rear ends. It was the first time we started actually trying to predict demand and buy ahead to meet it, rather than simply looking at the shelves and saying, "Hmm, looks like we're out of Asgards, we'd better order some more." It also meant major changes on the employee side, the shipping side, the operations side, and on the facilities side. In many ways,

it was when we moved from being a “hobby business” to a real company.

In the Beginning (of Summer, That Is)

At the start of summer 2012, we were still operating pretty much as we always had, except for the fact that we were operating out of the Schiithole.

Now, we only had part of the Schiithole (about 1000 square feet) at the time, not the entire building. It was just as crappy, dusty, and miserable as I’ve described. And, as an added bonus, the Santa Clarita summer was on us, bringing 100 °F to 110 °F temperatures during the day to a building with no air conditioning and no insulation.

Yeah, Alex pretty much cooked. Eddie, Tony, and I only came in during the evening when we could prop the doors open and run some big industrial fans to circulate the cooler air outside. Alex stayed there pretty much all day, to keep the shipping going.

Alex tried to combat the heat with a portable air conditioning unit, but it really struggled in the uninsulated space. We mostly used it to keep the products on the burn-in racks from going full

Chernobyl. Indoor temperatures of 90+ were common.

It was bad enough that we ordered a larger wall air conditioner, but it was delayed, then cancelled for some unknown reason. The delay was good, though, since that was when the landlord came by.

"I'm thinking of sprucing this building up a little," he said. "Putting in heat and AC, painting, doing the floors, stuff like that."

"Great," I said, a little suspiciously. Because landlords don't do things like that just for the fun of it. Usually such pronouncements are followed by, "And we'll bill you $X thousand dollars for it."

"It'll make the building more attractive to tenants," he continued. I sighed. *Here it comes,* I thought.

"And I was wondering if you wanted to take the rest of the space," he said, looking at me expectantly.

"At the same rate?" I asked.

He laughed. You know, that landlord laugh. The one that says. *Oh, you silly boy, do you think anything in this world is free? Why don't you just*

buy a bunch of buildings and sit back and collect money like I do?

Aside: I have nothing against landlords in general. I'm just not cut out to be one. I've got enough on my plate without having to deal with crazy people on a day-to-day basis. No. Wait. Nevermind.

"So what are you looking to get for the whole thing?"

He named a price. A surprisingly reasonable price. A silly cheap price, in fact.

Now, I know that in Biz 101, they tell you to never look impressed, to press any advantage you have and try to grind people down to the lowest possible price, but, you know what? That's also called "being an ******."

Actually, let's expand on that a bit. Haggling to get a lower price is one of those things that I've done in the past—and I was pretty good at it. But I never liked doing it. It felt bad. Wrong. Slimy.

Because when you haggle, you're saying, "I think what you're doing isn't worth what you're asking."

It also says:

- I don't trust you

- You don't know how to price things
- You're incompetent in general
- I think you're trying to screw me

And it says:

- I'm a cheap-ass
- I don't have a lot of money
- I'm going to be a difficult customer
- I might not pay you

When a customer starts haggling, a truly good and competent supplier does one of two things:

1. They reiterate the value of what they're offering, and stand firm.
2. They walk away.

Yes. That's right. I'm actually advocating for *non-negotiation*.

Well, that's nuts, some of you are saying. *There's a ton of companies out there who are set on screwing you, and will quote silly prices. How the hell do you know if you got a good deal?*

First, by knowing the general price of the product you're looking for. If you're looking for a knob you've bought in the past for $2.25, and you're quoted $1.80 by someone who's genuinely eager to make it, guess what? That's a fine price. If you're looking for an industrial space and Loopnet says the average rate in your location

is $0.65 per square foot, and you're quoted $1.10, you're probably getting screwed.

Second, by doing your research if you don't know the general price.
This could involve putting parts up on MFG.com to see what kind of range you get, or getting multiple quotes from different distributors.

Third, by knowing if something is either "too good to be true" or "worth paying a little more for." If those same $2.25 knobs were quoted by a brand-new business for $0.35, guess what? It may be too good to be true. But, at the same time, if your reject rate on the $2.25 knobs was 35%, it may be well worth your time to take a $3.10 quote.

Fourth, by knowing when to say when. If the quote you get meets what you expect to see on your bill of materials, and it's from a trusted vendor, stop. Don't grind. Just take it. You've accomplished your goal. Put your feet up, relax, and listen to some good music.

Fifth, by being up-front. Okay, so let's say you have a new product that's very cost-constrained, and you've been working with a vendor you're very happy with, but you're unsure if they can meet the cost constraints. Tell them what you're

looking for. Don't make it a 16-quote guessing game. Their time is valuable, too. And if you have to look elsewhere, you'll know right away. Also, give them a chance to comment on the part design—they may have a much better idea for meeting your price.

So what did I say to the landlord when he quoted a very reasonable price for more space that was literally next door—space that was going to be all shiny, polished, and air-conditioned? I said "yes."

Like, duh.

Enter Ten Thumbs

Unfortunately, our landlord was probably one of the biz school 101 guys who never grew out of the "grind until you're sure you screwed the bejeezus out of your vendor."

Why? Because he hired a guy for the renovation work who soon came to be known as "Ten Thumbs." I bet he got a great deal. Or at least I hope so. Ten Thumbs didn't look much different than your usual itinerant handyman, but he was put in charge of the renovation of the entire building—interior and exterior, structural, AC,

and electrical. This is something like giving a kid who's just built his first CMoy set the task of designing the next Audio Research tube preamplifier. Not just scary, but possibly dangerous.

We should have known what was coming when Ten Thumbs started on the outside of the building. The old structure featured quite a bit of dry rot on the wood beams—something you'd normally remove and replace. Ten Thumbs? Nope. Bondo would be fine. He bought it by the five-gallon bucketload. Similarly, you'd normally relay the scaffolding under the holes in the the stucco. Ten Thumbs? Fill it with newspaper and wall putty, then stucco over it. Windows? Don't bother masking them, scrape them off later. Roof? Well, you can't see it, so why bother? Paint? Hey, it looks like Home Depot has a remnant sale. Let's mix them all together to get something approximately the color of sick-baby poop. Spray that stuff everywhere. Only add a little lighter green and white accents when the landlord comes by and complains it's ugly. Street numbers? Don't need those. The few decorative trees outside? Cut them down. Trash? It can pile up in the back. If nobody can see it from the street, who cares?

You get the picture.

That's when Ten Thumbs moved inside. Luckily, he started with the area that we weren't using. In came spackling paste in 5-gallon drums. In came cheap white primer paint. Use a keyhole saw to cut spaces for the AC vents? That's crazy talk. Just use a hammer. There's plenty of spackle. And nobody will notice they're about 5 degrees off true. Doors into our current space? Well, we had some leftovers that don't really fit the hole that was cut, but if you sand about an inch off of them, they kinda fit. Of course, they aren't hung on any beams—they're just bolted to sheetrock. Floors? Well, the white paint can go all over it, because it's going to get epoxy-coated. What, you say epoxy doesn't stick to painted floors? Ah well, not Ten Thumbs' problem. I shook my head. It was crap work, but the space was still cheap.

That's when we had to move into the finished space, so Ten Thumbs could work on the space we were already in. They told us it'd be done in a couple of weeks, so we went ahead and executed the lease and moved into the finished space, in expectation that we'd have the full run of the place on October 1st.

And, you know what? We probably would have had it done in that short time. Except for Ten Thumbs' little mistake: when he did the epoxy

floor in the other half of the building, he didn't mix the hardener in with the epoxy.

Those of you who know epoxy are groaning now. For those of you who've never used epoxy, *it doesn't harden without the hardener.*

So yeah, we had a big, sticky mess of a floor. Just walking over it caused the unhardened epoxy to lift off and reveal nasty, untreated floor.

After a humorous episode where the landlord tried to blame Alex for the floor not hardening (saying that he washed it down, as if water would change anything), he learned that "no negotiation" doesn't mean "pushover." Because if there's one thing I will absolutely stand my ground on, it's defending our employees.

The landlord caught me on the phone and tried not only to place the blame on Alex, but implied that Alex was sabotaging the job on purpose.

"I don't believe that," I told him.

"And you're just going to take his word for it?" the landlord asked.

"Yes."

"Are you calling me a liar?"

“No.”

“Then what are you trying to say?”

“I’m saying that you don’t know any more about this than I do. Neither of us were there. Alex was there. He saw what was done. Plus, the way the floor is behaving is what you’d expect if there was no hardener in the epoxy.”

“How do you know Alex isn’t deliberately trying to mess up the job?” he shot back, getting frustrated.

“Well, besides the fact that he’s packed in there like a sardine, and he needs the space, and we’re already a week late on completion, I also trust him implicitly.”

“So you’re not going to investigate this?”

“No.”

“I don’t believe this!” he thundered. “I demand you look into this for yourself.”

Aside: several clients have tried to pit me against my own employees during my time at Centric. They got the same response. Usually shortly before we fired them. Newsflash: in real, productive companies, you hire competent people, support them, and let them shine. If you have trust issues, you have the wrong person.

“Okay. I’m going to say this just once more. No. I will not spend a minute more on this. I hired Alex to run the show, and I trust him. What did or didn’t happen isn’t the issue. The issue is that you’re two weeks late on a lease that we’re already paying.”

After some more bluster and spluttering, we agreed that Ten Thumbs would finish at all speed, and we’d stay out of his space. Two weeks later—exactly one month late—we had the full space, complete with a handmade soffit for a new air conditioning duct with runout that could be measured in inches, if not feet.

But the air conditioning worked …

… and it only took two more weeks to fix all the screw-ups in the electrical system

… and since it doesn’t rain much in SoCal in summer, it was a couple of months before the roof started leaking.

Ah, bargains.

The Real Work

Okay, so Ten Thumbs makes for some humorous anecdotes, but let’s talk about what we really

accomplished that summer—taking the first steps towards acting like a real company.

And, for all the hellish renovation work, we ended up in a space that (except for the remaining sticky places on the floor) looked reasonably like a real company for the first time. People could come by and see us, and we wouldn't be entirely embarrassed. It was still far from a palace, but hey, at least we weren't cooking anymore.

But what changed most was on the process side. Now, Process is another scary word that conjures up images of big corporate flowcharts, meetings, and managers pounding tables while yelling, "We've always done it this way!" But, in reality, every business needs processes. They just have to be flexible enough to change as the business grows.

In Schiit's case, Alex shored up five major things:

1. **Planning and scheduling.** Like I said, the old method was "wait until we're out of something, then start ordering again." This really isn't all that bad of a method when you're very small. It conserves cash flow, and helps generate a perception that things are special, and in-demand. That's why we used to go into back order so often. Cash flow is why we didn't

really bother addressing it for a good long time. But as we grew, we realized we had to look at what we were selling and place larger, timed orders that would help ensure that we had better availability. The only problem with that turned out to be rising demand. It's hard to project forward when demand keeps growing. And that's why the endemic out-of-stock problem lasted long past summer 2012 (in fact, it was so bad that we literally had zero stock at the end of December 2012, except Magni and Modi, which had just been introduced two weeks before. And those went out of stock just a couple of weeks later.) But planning is absolutely critical if you want to grow out of the one-guy job-shop model, and it's worth trying to figure out what you will be selling, and placing larger orders with timed shipments from your suppliers. You'll get a better deal—and you'll have some wiggle room if you end up slipping your schedule a bit.

2. **Facilities layout and production flow.** It's one thing to have a small truck delivering a few boxes to your garage for one guy to test and another to work on. It's a whole different thing to have multiple vendors delivering many different products at the same time, and having to move them through the shop. Alex

set things up so we had a logical flow through the shop, including all the hardware needed to make it happen (you know, "dumb" things like carts, which aren't so dumb when you're handling pricey products, like the then-new Mjolnir and Gungnir.) He also got us set up to handle the anticipated flood of high-volume stuff, like Magni and Modi, and made sure everyone knew what they needed to do.

3. **Shipping logistics and relationships.** If you are starting a business that will do volume shipping, you owe it to yourself to be very close with your shippers. Every shipper (besides USPS) offers significant discounts for their volume shippers, and they can tailor your discounts to what you're actually shipping. Alex got us in contact with FedEx, our shipper of choice, and began the process of negotiating better rates … rates that we then could pass along to our customers. If you notice our FedEx international shipping is frequently only a tiny bit more than USPS (or, in some cases, even lower), this is why. We should have been in touch with them from the start. But Alex fixed this, and continues to be their interface to this day.
4. **General operations and vendor communication.**

Before Alex, vendor communication was pretty ad-hoc (meaning, I'd talk to them in the spare moments I had between running Centric and writing books.) This isn't great. Especially when you consider the relative laggardness of many metal, transformer, and board vendors. An astounding amount of them still really only respond to phone conversations or in-person meetings. Email is treated as a secondary communication method, and can easily be ignored. Today, they're paying more attention to our emails, as we grow and become an important customer, but having real phone time or face time helps. Alex provided a focal point for the vendors to talk to, and he knew much more about our day-to-day needs than I did in short order.

5. **Employee management, specifically hiring and training.**
   Finally, Alex started laying the groundwork for both me stepping back from hiring, and for the coming Salary Apocalypse. Yeah, that's right. As of summer 2012, everyone was still a contractor. Since nobody had set hours (and everyone came in at fairly bizarre times … hell, Eddie frequently worked from 1 AM to dawn), we could kinda squint and get away with it. But I knew that, very soon, someone

would say something about it, and we'd be in trouble. We also really, really needed to provide healthcare and all the normal job perks, since we were starting to have employees with families and such. Sounds simple, doesn't it? Well, stepping into full-boat employee status, tax reporting, and benefits is a very big step. We didn't make it until January 1, 2013. Alex also started hiring as we needed, bringing in someone to help him with shipping (another underappreciated job).

The sum total of this was a company much more ready to take on the coming Magni and Modi avalanche. How ready? Well, we've only had one or two short backorders on Magni and Modi since their introduction in December 2012. Which is pretty good when you consider that Magni/Modi volume was about 5× to 10× what we'd ever planned for in the past.

But all was not well on the planning side. Because sometimes you undershoot the needs … and sometimes you get a big surprise.

Like the one that came next.

Chapter 25

Dead Media Ain't Dead: NYT Strikes

Look. Nobody is perfect. No matter how many degrees they have, no matter how high they scored on their IQ tests, no matter how many years of experience, no matter how many companies they've launched. Period. Puffery and pontification about the One True Way and instant dismissal of any alternative viewpoints may be the sign of great learning—but it's also the sign of a deeply insecure, lazy person.

But, at the same time, a learned caution and reliance on "what you know" can sneak up and bite you in the ass as well. It's not a light-switch change from being young, enthusiastic, and open to new ideas, to a grumpy, set-in-your-ways know-it-all. It's a continuum.

And I was at least thirty percent down that road to the grumpy know-it-all when we had our butts handed to us on a platter late in 2013.

The further irony? I got bit not because I was

old-school conservative, but because I'd drank the new media kool-aid for a bit too long.

What am I talking about, you ask?

Hold a sec, and let's talk a bit about pushing the limits, common wisdom, experience, and the fiction of modern marketing.

Lessons from the Leading Edge of Marketing

I founded Centric, my marketing agency, in 1994, the same year that Netscape Navigator was released. Now, these two occurrences were not coincident; Netscape came on the world about 10 months after I started Centric.

This is important because, over the next few years, the world of communications saw the growth of its most important medium, ever.

Yes, *ever.* Don't argue with me. I'm at least 25% leading-edge to this day. And if you want to argue with me, name one other medium that disintermediated the distribution of content in the way the internet did, and one medium that resulted in the significant shrinkage of other media. The common wisdom of my college years was that all new media was additive—that is, radio didn't replace newspapers, as TV didn't

replace radio, etc. I argued vehemently that this wouldn't always be the case, even way back in 1988 … but my professor was less impressed about my gut feeling than his statistics. Hey, bite me, prof. Who's right today?

But it wasn't just the appearance of the internet that drove Centric to always be on the leading edge. I'd been doing marketing for some time before Centric, and had always relied on "desktop publishing" programs, as the entrenched typesetters and color separators of the time diminutively referred to them as.

"We do *real* typesetting," I was told, with a sneer, more than once, when I asked about camera-ready or film output from my files.

Yeah, and you'll be out of business in a few years, I thought. The economic benefit of computer typesetting, and, later, computer-based full layout and color separations was too compelling to ignore. And those early desktop-publishing programs opened up the field to a lot more people—it democratized design, in a way never seen before.

If you're too young for this dinosauric crap, just consider that the way you used to lay out

brochures included sending your copy out to have typeset and output on an optical typesetter, then actually pasted to an art-board, which you sent to the printer along with any color separations you had done—from film—to make your brochure. The first brochure I did for Sumo had a hand-painted gradient that was sent out for photography and color separation. Think about that next time you use Illustrator and create one in 2 seconds.

In any case, by the time the internet came around, I was absolutely ready to jump on its leading edge. When I first saw it in early 1995, I literally got chills. *This will change everything,* I thought.

Which is why Centric had one of the first ad agency websites online (27th listed on Yahoo at that time—there were 8 when we started development.) It's why we did some of the first e-commerce work, including the first online leasing system for Compaq. It's why we built some of the first customized product configurators. It's why we built our own content management systems. It's why we embraced search engine optimization and online marketing back when you could only buy keywords from GoTo.com (which Google later copied and turned into Adwords, the powerhouse that drives most of its revenue today.) We did

some of the first database-connected personalized banner ads. We did some of the earliest Flash sites, games, and COPPA-compliant virtual worlds. And we did some of the first social marketing stuff out there for Warner Brothers, Cotton, Inc, and other big names. Hell, we built HP's outpost in Second Life, and created a virtual environment for David Rumsey Maps that was profiled in MIT Tech Review.

What I'm trying to say here is that, in terms of marketing, our leading-edge cred is way, way up there.

Which is perhaps one reason I dismissed the power of the conventional press when I started Schiit. We were *so* past that. It belonged to the era of paste-up and color separations.

And, to be fair, it was starting to look pretty unhealthy, on lots of fronts (hell, we had one industry—data storage—in which the industry print publications all folded up before the turn of the century, so we knew where things were going.)

So, when Schiit got rolling, we pretty much just ignored the conventional press, except for sending them press releases when we launched new products. Yeah, we were happy for the coverage in

Wired and such, but I was far more jazzed about Gizmodo, Engadget, and TechCrunch. Those were publications of the present day, reinventing press as we knew it on a real-time basis. They were what mattered. We'd eventually get money to do some Adwords, and roll it into more online, measurable ads as time went by. And that's how you did things in the Shiny New World of the 21st Century.

Or so I thought.

But Social, Oh, Social, and the Fiction of Online Marketing

You may have noticed that I haven't mentioned one of what's considered the most "leading edge" marketing vehicle out there, though: social media.

That's for very good reason. Even in 2010, we were pretty much done with social. After 4 years of pimping MySpace, Facebook, Twitter, etc, we'd learned some hard truths. Namely:

- If you're an entertainment company, social marketing is the greatest thing since sliced cheese. It should absolutely be front and center in your plans. Every entertainment

social media program we did produced 10× to 100× the results of an equivalent investment in conventional media

- If you're not an entertainment company, social marketing is really, really dumb—easily the biggest time-sink and resource-eater out there, with returns 1/10 to 1/100 of an equivalent investment in conventional media[1]

"Wait, what?" you might be saying. "I still hear that talking to your customers on Facebook, doing videos to put on YouTube, and Twittering your latest office party pictures (and Instagramming, or Pintresting, or whatever) is the way to have an authentic conversation with your prospects and create engaged brand ambassadors."

In short, no.

Sorry, guys. People are on Facebook to talk to friends. Not shills.

They're watching YouTube for funny cat videos, not smooth-talking tours of your factory set to some hip music.

They're on Twitter to get celebrity tweets.

---

1 I can't tell you how much hate mail I got when I "came out" against social media. It's nice to see some recent studies corroborating our experience.

Et cetera. If you want to talk to your prospects effectively:

1. Clearly communicate the unique benefits of your products on a good, easy-to-use website.
2. Have a memorable brand.
3. Provide fast responses to any inquiries.
4. Take care of customer service before it spreads to Facebook.
5. Make sure the press (online and off) know when you have something new and cool, but otherwise stay out of their face.
6. Invest carefully in measurable marketing vehicles such as Adwords, reinvest in successful vehicles and revise or discontinue underperforming ones.
7. Continue improving your product so someone doesn't have a clear, unique benefit over you before you know it.

And that is that. Social media will take care of itself, at that point.

"But wait, does that mean we can pretty much ignore social media?" you ask.

To be blunt: *yes.*

This "ignore social media" advice is even more relevant if you are a business-to-business company—that is, selling products or services to businesses.

Do not spend a single second on social media. Concentrate on the 7 points above. Don't dismiss 1 and 2 because you're B2B. And you're done.

Sure, there are gray area companies (such as audio companies like Schiit) which have rabid fans and may get some small benefit from social media, but your marketing dollars are better spent on more conventional advertising and PR. And then there's micro-social (like this site), where there's a very concentrated niche audience that is absolutely relevant to your products. Then, be there. Yourself. Regularly. Don't leave it to an agency. Don't use paid shills. And don't be a dick.

Still not convinced? Fine. Cool. You want to get into social media? Answer these questions:

- **Who's going to create the content?** Note: this is not just posting funny pictures on Facebook. This is about creating an article every week for a blog, or a video every month for YouTube, or doing cool imagery for Facebook. Are you going to do it? Or are you going to hire an agency to manage it for $200/hour? Yeah. There you go.
- **Who's going to respond to comments?** Social media is social. You'll get comments. You'll get hotheads. Someone has to manage it. Is it you? Or more $200/hour time?

- **Who's going to decide what's OK to say?** In larger companies, this is a HUGE problem, usually involving a team of lawyers. But even for smaller companies, this is no joke. What will reflect well on you? What will reveal competitive info you'd rather keep under wraps?
- **Who's going to measure and manage it all?** Social media is still evolving. Who's going to be the oversight to determine what to spend where, based on the sales you're making? What metrics software will you use, and what metrics will you track? What will you do when Facebook changes its algorithms yet again?

Yeah. Have fun. Unless you want to cheat. See below.

A Bold Prediction

Now, I haven't scratched the surface on social. Because it's a lot more than companies posting contests on Facebook these days. Social is, more and more, a bot-infested, paid-propagation-driven exercise.

Which, in English, means:

- Much of the content you see on social media wasn't posted by a human at all, but by software—a bot
- Much more of the content you see on social media was posted by offshore shill accounts paid a pittance to say, well, pretty much anything
- And even more of the content was posted by people just like you who want to get something for cheap (or free) and are doing so to get referral credit towards their shiny new object of desire

So, social media is less and less about "an authentic conversation with real people," and more "prospect-powered advertising." Or, to be more topical, crowdsourced advertising.

And in an environment where anyone can be a shill, and their financial motives aren't known, credibility disappears.

Which leads me to my bold prediction: that in the next decade, we're going to see paid, conventional advertising in big-name venues become the most credible source of information, and word-of-mouth the least credible.

"Wait, you're saying that we'll trust paid advertising more than our friends?" you say, aghast.

Yes.

Why? Paid advertising in big-name venues:

- Requires a large investment that the company wants to pay off, meaning the company is at least successful enough to make the investment.
- Subjects the company's claims to scrutiny by individuals and by regulators, which the company can be held accountable for.
- Has significant repercussions if those claims don't gel with the customers' experiences—as in, class-action lawsuits, etc.

With that on the line, paid advertising could very well be more accurate than the offshore shell account shilling the latest hot gadget, or even your friend, who may just be parroting fake claims in the hopes of getting something he wants.

Crazy? Perhaps.

But we'll see.

On to the Old Media

"So, after all this blathering, are you actually going to tell me what happened?" you ask. "I'm dying of suspense!"

Yep. Here's what went down.

In summer 2012, I was contacted by a writer for the New York Times, who asked if he could get a Bifrost to review for an upcoming story on DACs. And let's be clear: no matter how new-media biased you are, here's what you do when the New York Times asks for a review sample: you say, "Yes."

Same goes for if the national TV news wants to talk to you, or if the Wall Street Journal wants to have a phone chat, or Wired wants to do a profile on your company, or, hell, if Ladies Home Journal wants to try some of your cables. It's big exposure. You say "yes," and make it happen.

Aside: well, actually, first you do a quick Google search on the writer's name and the publication title, to make sure he's really who he says he is. If it checks out, then get that product out the door.

So I sent out the Bifrost to our guy Roy at NYT and kinda forgot about it. Because sometimes these articles don't happen for a long time. Or ever. And that's fine. Editors have their own agendas. And you'll live through a missed opportunity or two. I forgot about it, at least, until early in September, when Roy gave me a call.

“Wow, this is a great DAC,” he said. “Really, really good.”

“Cool,” I told him, or something lame like that. I’d probably forgotten who he was, and was trying to put the pieces together.

“It was the crowd favorite,” he enthused. “But I’m not sure if the Times editors will let me actually say that. I want to, but there’s limited space, and, you know, stuff like this is a little controversial anyway.”

Aha. Times. Everything clicked. “Stuff like this?”

“High-end audio. Subjective reviews. It’s kind of, well … ”

“Too much voodoo?” I prompted.

Roy laughed. “Voodoo. Yes. But it’s real voodoo. It’s just, well, we have to be careful not to go too over-the-top. But great gear really deserves more coverage.”

“Roy, I spent 20 years in marketing. Believe me, I understand.”

“No, it’s not the advertisers or anything,” he corrected. “It’s just, well, everyone seems to have a price in mind for any piece of audio gear, and if something goes over that price, but the specs

aren't any different, well, it's hard to explain how and why it's so much better. I don't understand why some companies can make great-sounding stuff, but so much gear out there just sounds, well, awful."

"But measures good," I added.

"Exactly." I agreed it was unfortunate, and let it go at that. But if I'd had a few drinks, I probably would have told him something like, *It's because nobody ever got fired for engineering a product that meets specs, but sounds bad. And it's because most people don't really care that much.* The first is understandable in a highly regimented, performance-review-oriented, hypercorporate environment. The second is understandable, too. Some people don't give a crap about cars. Some people don't give a crap about audio. Some people don't give a crap about cuisine. That's just the way it is.

Aside: And, while I'm absolutely for introducing everyone to great sound, we're going to meet plenty of people who don't care. And we have to be careful not to be tiresome proselytizers. Imagine if every Jehovah's Witness suddenly converted to the Church of the Perfect Sound and started going door-to-door with Audeze

LCD-3s and a Mjolnir/Gungnir rig to spread the word. Yeah. Just as irritating.

In any case, Roy and I talked a bit more about audio, about what was different about Bifrost, and that was that. I kinda forgot about it in the craziness of the next week.

Then, suddenly, we got a huge explosion of Bifrost orders, and the email loaded up … not with our normal, relatively technical questions, or questions about "is Asgard or Lyr better for my headphones," but with a ton of super-noobie questions like "Can I hook up Bifrost to my flat panel TV," and "how can I connect Bifrost to an AVR?" and "what's an optical cable?"

An email from Roy confirmed what I had suspected: the article—"A Sound System as Resonant as a Concert Hall"—had gone live.

What I didn't expect was that it would hit the actual print version of the paper.

And I certainly didn't expect the response. No way. No how.

In fact, before the Times article, Alex and I were feeling rather smug. We'd pre-ordered and scheduled a double run of Bifrosts in anticipation

of the holiday rush. We were well-set to sell a ton of Bifrosts. There was no way we'd go into backorder on them.

Until that article. In the space of 3 weeks, our double run of Bifrosts was gone. Completely annihilated. Suddenly we were staring at a big, long backorder, right in the busy time of the year. I scrambled on orders for metal, boards, parts, etc, but there's only so much you can do when your metal lead time is 6 to 8 weeks. You can beg a bit faster, but that's about it.

And that's how we ended up, not just with a single backorder that fall, but also a second backorder at the end of December, as the second double run we'd done also disappeared out the door. A good chunk of that was the Times article. The noobie questions continued well past the end of the year, and we continued to sell Bifrosts into homes that probably didn't even know what a DAC was, before reading that article.

So, if someone says, "Old media is dead," laugh at them. They don't know what they're talking about.

Or, if someone says, "We can't convert new audiophiles," laugh louder. We absolutely can. We just need to get more attention in the mass media.

And continue our inroads on sites where younger people discover stuff, like Reddit.

And that *can* be done.

But it *won't* be done with $40 000 preamps, $3000 USB cables, and $500 magic pucks. It won't be done with aspie-level obsessive in-fighting about formats and provenance. It won't be done with religious fervor to spread the word about the One True Sound or the One Perfect Measurement.

And—this is the important one—it won't be done with magic name-branded processing tricks, or deceptively-named mediocre technology, or the breathless hype of the Truly Established Big Names. Because that's one thing about marketing that's really changing … people today see through the BS much faster. And if they don't, their friends do, and they spread the word. That's the real power of social media, and, believe me, that's something the billion-dollar audio behemoths don't want to learn about.

So how do we break into the mainstream, and introduce more neophytes to great sound? Well, call me biased, call me old-fashioned, but I believe it will be done in only one way: with quality prod-

uct made at a price that's fair for its performance, construction, and looks.

And that's what we're focused on. It's astounding how many Magnis and Modis we sell to first-time audiophiles.

It's a small step, yes.

But just think: that person could have ended their journey with an iPod.

A Counterpoint

There's a counterpoint to the Times story.

This is a story about a review we got in a major audio magazine (which shall remain nameless.) This was a big deal for us, because they usually aren't so hot on direct-sale product, especially from relatively new manufacturers. And, outside the headphone community, we're not very well known.

It was a glowing review. We got our name on the cover.

And, remembering the New York Times review, we stocked up on Bifrost with another double run (actually quadruple, since we doubled the first run amount as a standard.)

But this time, the sales didn't materialize.

"What the heck?" you're probably asking.

Yeah. So were we. Sure, there was a minor bump—maybe 10 % over our standard run rate—but nothing like the Times review. And there was another bump in emails—but this time they were 8000-word essays about how they got into audio, owned tons of very expensive equipment, and contained 18 questions about tiny, tiny details about the product, usually regarding buzzword compliance with the lastest press-propagated meaningless terms of the day.

Why? My best theory is the old marketing adage of "one ad does nothing." If two other reviews of Schiit products had followed the first one, maybe things would be very different. But that's not likely, given the focus of the magazine. They're more on the two-channel side of things, and we haven't yet made a big dent in that market.

But, in the Times, one review did a lot.

And that's a valuable lesson. That there are still huge opportunities out there … sometimes where you least expect them.

Stay open. To new things. And to old ones.

Chapter 26

Finally, the $ 99 Solution

Back in the early 1990s, while I was working at Sumo, if you told me I'd eventually be designing and selling audio products that retailed for $ 99, I would have told you that you were insane.

And, if you'd told me that I'd be designing and selling audio products that retailed for $ 63—the 1991 equivalent to $ 99 today, I'd think you were even nuttier.

I suspect Mike Moffat would have had the same response. Which makes it all the crazier that he was the one who brought the idea for the $ 99 DAC to me first, and started the whole inexpensive-gear thing at Schiit.

Why would this be so unbelievable, you ask?

It's simple. Back then, it was a different world. Building and selling a product for $ 31 to $ 37 in 1991 (which would allow the dealers to have their cut, in the dealer-dominant marketplace of the time) simply wasn't gonna happen.

Selling direct was, to put it mildly, unfeasible in the pre-internet world. And, back then, we'd probably have to make the products with thru-hole parts, rather than using surface-mount parts (which allow for efficient automated assembly), driving the production cost even higher.

And, early on in the Magni/Modi project, I had doubts that we could do it for $99, even direct.

So I was very, very thrilled when the original numbers came in for the chassis and wall-wart, confirming that we could, indeed, sell gear with the near-unbelievable two-digit price tag.

Once the numbers were confirmed—this was in late August 2012—we started placing the biggest orders we'd ever done to date with our suppliers. We knew Magni and Modi would be big, and we knew how crazy things could get if we went out of stock. So we stocked up, in anticipation of being able to release it at CanJam, well-timed for

the holiday rush that hits every year.[1]

Of course, people who attended 2012 CanJam know there was no Magni and Modi there (well, not out on the table—there were under-the-table prototypes we shared with a few people.)

And, even more sharp-memoried readers will remember that Magni and Modi didn't ship "comfortably before the holiday season," and, in actuality, made their appearance on December 13—deep into the Christmas buying time.

Not ideal, right?

---

1 Audio is a strongly seasonal business. October through March are hot. April, May, and June see falling sales, to a low in July/August. Then it picks back up again as people get into a stay-indoors or holiday-buying mood. This seasonality tracks very well with cold weather in much of the northern hemisphere, where people are more likely to stay at home and curl up with a nice glass of scotch … er, I mean wine … er, wait, warm milk or something is probably a little more PC, but whatever. As the days get warmer and vacation season hits, sales fall. This has been going on approximately since the earth cooled. If you're thinking of starting your own business, be aware if it's a seasonal one or not—planning in the fast season can leave you buried in unsold stock, and projecting forward from the slow season can put you in heavy backorder.

Yep. But much less ideal than setting an arbitrary deadline to be missed, or pushing a semi-finished product out the door.

The Luxury (and Penalty) of Closed Doors

Magni and Modi were the first products we didn't announce beforehand. Indeed, they were the first products we didn't even hint about. This "closed door" approach is great in many ways. Namely:

- There's no chance of missing the deadline, because there isn't any
- If there are problems in production, you have the luxury to take your time and sort them out
- Your competitors don't know what you're doing until you launch, which means it will take them longer to counter

Of course, closed doors have some disadvantages, too:

- With no deadline, you may not work as hard as you need to in order to hold a schedule
- If you have cash flow problems—that is, if you're running a business on a receivables-financing basis, or a assets-financing basis, or if you don't have shipping products to tide you over—the luxury of time is a cold comfort

- If there are significant delays, you may miss your first-mover advantage

The real danger to Schiit's launch of the Magni and Modi was the first point—with no deadline, you might not work as hard to get it out on time.

Now, this doesn't mean we sat around. But when metal was late, and when the board house was slow freeing up their surface mount line for us, we were probably less diligent than we should have been in being "in their face."[2]

So yeah, metal was a bit late. That killed showing Magni and Modi at CanJam. And boards were a little later. Which wouldn't have been as big a problem, except for two others:

---

2 Please note that "in their face" does not mean "being a total whiny bitchy ass at every opportunity." Running a business that doesn't do everything in-house requires great relationships with your vendors. They need to know when things aren't critical, and they need to know that you're not going to hold them to some insane algorithmic standard for to-the-minute delivery on every product. That means, when you do need fast delivery, or a quick change, they're going to be much, much more likely to get that done for you. Because you don't ride them on every little thing. So, "in their face," means "calling them more than once every three weeks to see how things are going."

1. The pots (engineer-speak for 'volume control') we planned to use were late—even later than their 6-week lead time would suggest. Which hung up delivery of the boards, since a headphone amp without volume control is, well, pretty much useless. Especially when there's a big hole in the chassis it was supposed to poke through.
2. When we finally got the completed boards in, we had an engineering "oh schiit" moment when we realized the boards couldn't be inserted into the chassis. As in, at all. The big capacitors at the front of the board hung up between the pressed-in board standoff and the top inner chassis flange.

Aside: and, before you start saying, "3D CAD would have showed this clear as day," well, not really—not unless it showed articulating exploded views. And, these days, finally, everything we're doing is 3D—but not back then.

So, what did we do? Well, I knew at least one fix:

"We could send the chassis back to the metal guys and have them notch the front flange to clear the standoff."

Mike was less than convinced. "How long will that take?"

A couple of weeks."

Mike looked pointedly up at the calendar. It was the first week of December.

"Yeah, I know, I know, we miss Christmas … "

"If it's really only two weeks," Mike said. "We could lay the capacitors down, couldn't we?"

"Yeah, but that looks awful." But then, I had an epiphany—the one I should have had in the first place. "But we can find shorter capacitors, I bet."

> **Aside:** yeah, I know, obvious. But tell me you've never made any boneheaded mistakes. With a straight face.

"Can you get capacitors with the same value?" Mike asked.

"I'll find out. But even if we have to shrink the values a bit, it's pretty overkill already." A quick Mouser search confirmed that we could, indeed, get the same value capacitors, at about half the height. Mocking it up with capacitors of that size proved that we could, indeed, just clear the chassis flange. Done.

Now, the above thing about capacitors of different sizes may seem strange to non-engineers, but it

really isn't. Capacitors of the same value—for the sake of argument, let's say 1000 µF at 16 V—come in a huge array of sizes. Some are large for the sake of higher temperature ratings or lower ESR. Some are large because a lot of engineers think that large caps measure better than small caps. That's not always true, because a lot of manufacturers caught on to this little factoid, and sometimes make electrolytic capacitors that are full of a whole lot of air. You need to look at the specs, and make your decision based on that … and also consider that oddball sizes are the ones more frequently out of stock, and pricier.

So, after the capacitor debacle was resolved, we got some right-size caps in next day, and sent the first article boards back to the PCB house. Luckily, there were only 10 first articles that needed reworked. The rest of them came to us with the new caps in-place.

And, in a few days, we had enough Magnis and Modis on the shelf to announce …

The Real Statement

On December 13, 2012, we announced Modi and Magni to the world. And, oh, what a crazy

announcement that was—and a crazy last two weeks of December!

In terms of mechanics, it was a relatively uneventful launch. We had product in stock. We had enough labor to keep it in stock, and a big enough first run that we didn't hit backorder until January.

In terms of issues, it was also relatively uneventful. Some dead wall-warts (which we tightened up by better vendor communication) and some out-of-spec pots with unacceptable tracking, even at relatively high levels (which we reworked, and re-spec'd the pots for the next run.)

But in terms of what it did to the company … suddenly we were all about $ 99 products. They were the vast majority of sales—easily outstripping all other products combined. We struggled to keep up with shipping and with inquiries, which led to fast changes on the shipping side and slower changes on the customer service side (until November 2013, I answered the vast majority of customer service questions.)

And, to this day, the best comment I think we ever got on the launch was amongst the chatter about the Schiit statement products (yeah, people were

talking about them even back then, the as-yet-unnamed Ragnarok and Yggdrasil.)

The comment was: *"Now, this is the real statement."*

Exactly. Thank you. Making another pricey product—no matter how advanced and innovative—is cool and all, and makes for good ego fodder (that is, at least when you can ship the darn things.) But making a good, solid product that almost anyone can afford, that's a whole 'nother thing entirely. It's wayyyy more important for the Magni and Modi to exist than Ragnarok and Yggdrasil.

And—you know what? It's a lot more fun, too. Magni and Modi will always be special to me, because it proved that we could actually make, sell, and support a no-compromises, all-discrete amp, and a top-shelf asynchronous DAC in the USA for a two-figure price tag.

Nobody was expecting that.

Nobody was expecting that *from us.*

Magni and Modi were the first products that actually made me want to go to CES and do a press tour. To go in front of them and say, "You know, there once was a time in this country when

we actually made things. Practical things, with top-notch technology and affordable price tags. Things without excuses about 'well, we can't really do that in a global economy,' or 'we don't have the supply chain here to do that,' or 'labor here costs so much we can't be competitive.'" And then hold up Magni and Modi and continue: "That stops today. Presenting the only fully discrete, no-excuses headphone amp made in the USA, from predominantly-US-sourced parts, with a 2 year warranty. No excuses. $99."

You know, that would make a decent commercial, if Fiat-owned Chrysler hadn't co-opted the idea. Ah well, what can you do?

And, it's funny. After the launch, of course we got compared to the O2 and ODAC. And some people even opined that we brought these products out as a direct counter to those.

Actually, no. I was thinking bigger. I saw the headphone market growing. And I realized that, very soon, some large entity would wake up and say, "Hey, you know, we should have some of the accessories side of that. Like amps and DACs." Someone like Logitech. Someone like Harman. Someone *big*.

And, with a market fragmented amongst a bunch of smaller players, you know what? It could have happened. They could have gotten some engineers together, put together some of TI's textbook op-amp/headphone-driver solutions, built the whole thing in China, and sold it for a very attractive price. Maybe not $99 through distribution, but dang close.

But after Magni and Modi? A powerful, fully-discrete US-built amp for $99? A USB 2 asynchronous DAC for $99? Well, then suddenly the market doesn't look so attractive. Going up against that is a lot harder than going up against Asgard, Valhalla, and Lyr.

And that's what I wanted to do—to set the value bar, and the barrier to entry, much, much higher. So that any of the big guys out there looking in would say, "Hmm, well, there's some very high-value, well-marketed stuff already established—considering the size of the market, let's take a flyer."

Hopefully, we have achieved a small fraction of that goal.

We'll see.

The Lasting Impact

At launch, I really was all about making an impact—but it took quite a while for me to realize what it really is. Chest-thumping about made-in-USA product is one thing. Setting the barrier to entry high is another. But the former doesn't really matter for much of the world, and the latter really doesn't matter much except to us and the competition.

No. The thing that really matters is that Magni and Modi are giving many, many first-timers their initial taste of very good sound. They're recommended here, and on many other online communities, including the biggest first-timer venue out there: Reddit.

It's become almost a knee-jerk recommendation: "Oh, you're looking at getting into headphone amps and DACs? Well, there's the Magni/Modi … " (and others, of course), but Magni and Modi are usually amongst the first mentioned.

Why? Because they're inexpensive, and because they're good products. I wasn't kidding when I wrote, "They may be the only amp and DAC you'll ever need." I was dead serious. Magni has a ton of power, and it's pretty open and neutral for many different headphones. It's easy to recommend.

Same with Modi. It plugs into a whole lot of different sources, doesn't need drivers, and pretty much just works. Again, easy.

Sure, there are amps and DACs out there that scale up higher, or have more features, or serve the format du jour, or whatever, but for most people in today's largely 16/44.1 based world (and, dare we say it, AAC and MP3), they're just fine. And either can be had for the price of a good steak dinner.

So that's the real impact: bringing that first audiophile experience to a larger audience.

And *that* matters.

Chapter 27

Twilight of the Gods—Ragnarok from 2009 Until Today

Okay, so shoot me. I moved the Ragnarok chapter up a notch. But I think it makes sense, in the context of yesterday's announcement of the Ragnarok public beta. Yes, we are almost there. Soon, we'll be shipping Ragnaroks … for real.

Yes. Stop laughing. Rina still doesn't believe me.

Aside: the reason Ragnarok was delayed (again) was twofold. One was my personal drama. The other was that we found another software bug in the final code—a particularly nasty one that caused production main boards to go into meltdown mode when they were first plugged into the production control boards. It took us a while to ensure that it wasn't something inherent in the amplifier itself, and to identify the offending code. For the technically minded out there, it was related to some debugging code that remained in the software—

it swapped two microprocessor output pins so we could monitor the operational parameters in real time. However, since the pins swapped included the output to the control DACs, and the updating was much slower than production, the Ragnarok microprocessor would end up trying to bump the bias up to unsustainable levels. When it finally updated, bang! Bad news. That code is now gone, and Ragnarok is running stably.

Another aside: Jude's Ragnarok, as well as others in the very small private beta, did not have this software, and hence didn't have the problem. All in all, we think the public beta will be short and uneventful, but in an abundance of caution, we're proceeding with it anyway.

The Ragnarok Saga: Pre-Blab

And … it really has been a saga. Ragnarok was one of the first amps we discussed in Schiit's pre-launch days. I have sketches dating back to late 2009.

So why was it one of the earliest concepts? Because it was a direct outgrowth of driving headphones with Sumo speaker amps. I figured, well,

if a Sumo 60 W amp works with headphones, why not make a reeely reeeeellly ridiculous headphone amp. We could put a warning sticker over the headphone output, saying, "If you remove this, you realize that you can make the magic smoke come out of any headphone, and agree to hold Schiit harmless for any damage."

Yeah, silly. But an interesting concept for a publicity stunt for a then-unknown, unlaunched company.

Of course, reality soon set in—driving headphones with a speaker amp sometimes works out just fine, but most speaker amps are:

1. Relatively noisy (forget your sensitive headphones and IEMs)
2. Not headphone-friendly on turn-on (1 V of DC is no big deal to a speaker, but a bad day for some headphones)
3. Big, hot, heavy, and not desk-friendly

So, the idea got shelved, after I'd spent some time playing with the old Sumo Antares and came to the realization that, like duh, a speaker amp usually isn't a headphone amp—at least not without a whole lot more development work.

But the kept ticking over in my head. And, after our early experiences with planar headphones

and the moar power!!! phenomenon, I decided that yes, it was something that we should pursue.

And, of course, I had to mention it. This was 2011.

Yes, I am an idiot.

The Ragnarok Saga: Post-Blab

Okay, so after I'd opened my mouth and told the world that we were working on a statement amp and DAC, I realized that we'd, well, actually have to deliver something eventually. But I didn't set any timeframe, because, well, I didn't figure there'd be many people interested in an unnamed, un-spec'd, unreleased super-power amp.

But they were. From the first mention that Schiit was going to do a statement amp, we started getting inquiries. When will it ship? What will it be like? What's the price?

Based on super-scientific complete-WAG BS, I started telling people it might be about a year, and it would be super-powerful, and probably about $1000.

Yes. Dig yourself deeper. This is a textbook example of what not to do. If you are silly

enough to pre-announce a product, here's what you should do:

1. Say nothing more, and hope the inquiries stop.
2. If pressed, say you got drunk and spoke out of turn, and the product doesn't actually exist.
3. If presented with documents you accidentally emailed someone confirming the existence of a project to develop said unreleased product, claim that you were either:
   a) Blue-skying it, but it has since been abandoned
   b) Developing a product for someone else.

Or, in other words, *shut up until it dies.* Then, when it's ready, bring it out and let people buy it.

Yes, it really is that simple.

But no. I had to confirm its existence. I even had to opine that it would be a really, really cool product. Again: don't do this.

But … since the cat was out of the bag, I figured that I should at least get started on the design. People who know a thing or two about large power amp design know that you really start with two things:

1. **Thermal design.** How do you get rid of the heat? Running speakers isn't like running

headphones—there's going to be significant thermal load.

2. **Transformer.** How big of a transformer do you need, and does it fit in the chassis you have in mind?

Once you've gotten those two things set, you can then proceed on to the rest.

So, the first calculations I did were for the chassis and transformer. Would our trick of using the chassis as a heatsink work? It turned out that the answer was yes—if we were looking at a standard 2U height rack-sized product. Which gelled with other 60 W amps I'd designed in the past. Transformer? Sure, you could use a 3 inch high product with a thick stack. Done that before, too.

But what about features? The first idea we had for Ragnarok was mighty conventional:

- 60 W/8 Ω, 100 W/4 Ω
- 3 levels of gain switching
- Alps RK27 balanced volume pot
- Nice Grayhill input switch
- 5 inputs, 3 SE and 2 balanced
- Circlotron topology
- 16 inch×12 inch×4 inch chassis

So really, nothing too nuts, except for the gain switching and circlotron topology.

We ran with that idea for a while, long enough to get some switch samples and order the balanced pots. But we never built a prototype of it.

No. Ego stepped in, and started us down the path to insanity.

Fun fact: one of the reasons we did Mjolnir was because the Ragnarok concept had already started to grow … and I wanted to use those balanced pots somewhere. So, Mjolnir was our shot at the simplest, most basic balanced amp we could do. Ragnarok had already been elevated, in our minds, to something that was quite a bit more.

How much more?

Well, it all started with the switches. Grayhill switches are very nice, but pricey. It really wouldn't cost any more for us to switch inputs with relays. And with relays, we could do the switching right at the back of the amp, where the inputs came in—which would save having to run a bunch of traces for each input up to the front switch. This would make things like crosstalk a whole lot better.

So hey, why not use relays?

But then you have 7 relays in the chassis (2 each for the balanced inputs, 1 each for the other inputs.) And, if you have that many relays, why not have a few more? Alps RK27 pots are nice. Sure. But a resistor-switched stepped attenuator is better. And another 12 relays would get us a 64-step attenuator at 1 dB to 1.2 dB steps—much better than you'd get with a front-panel switched attenuator … and tons of control with 3 different gain levels.

So now we're up to 19 relays. Still not a huge deal.

But if we have 19 relays, we should be switching gain with relays, too, right. So add another 4. And, of course, you have to have output muting. So that's another 2.

Yep. 25 relays in that first Ragnarok (there are actually 29 in production, since we added a separate headphone/speaker mute, and changed the way the gain switching relays worked—but we'll get there.)

But 25 relays wasn't crazy enough. I also decided it might be fun to be able to run the amp with either a solid-state or tube input stage. So,

that meant that the transformer would have to accommodate either one. Which meant a very large, complex transformer with like 4 Molex connectors hanging off of it, and a separate gain stage board (one for tubes, one for transistors).

And, of course, with 25 relays, we needed a microprocessor to run them—to handle switching, volume control, etc. So then we needed a shielded front control board …

… and, if we were going to all that trouble, then we might as well have a really cool volume control—that is, one that actually acted like a volume control, with stops at either end, rather than an endlessly spinning encoder. Which meant we needed a fairly beefy microprocessor, so we could A/D in a voltage and use that to set the output to the resistor ladder volume control.

Okay. So why do all this, you ask? For better sound. A better volume control and better switching is a big deal, when you're going all-out. Plus, we'd already begun to believe our own BS, with all the inquiries about the "statement products."

This is how you talk yourself into making something fairly insane.

And that's where the first layout began.

The First Ragnarok

If the first Ragnarok had worked, it would have been a pretty damn good headphone amp, but a fairly crappy speaker amp. This is because I'd become used to doing headphone amps, and I wanted to do a super-simple single-gain-stage topology.

There's only one problem with that: single gain stages have low overall gain—which means you can't use much (if any) feedback to get the output impedance down.

> **Aside:** Yes, Ragnarok is not a no-feedback amp. It is a no-overall-feedback amp, however …

Why is this important? It's important for control over the speaker load. Speakers are physically large. They need a reasonable damping factor. As a single gain stage amp, Ragnarok would have had an output impedance of about 0.5 Ω in high gain—fine for a headphone amp, but equivalent to a damping factor of only 8 into 4 Ω loads. Not exactly ideal.

But, as it turned out, the first Ragnarok prototype was so problematic, we never got it fully operational. The tube stage didn't get enough

power from the transformer. The solid-state stage had so many layout errors that it wasn't really worth fixing them all. And the board changes needed to accommodate both tube and solid-state really made it impractical to have it swappable—it would have to be built one way or the other.

> And, of course, this tube vs solid-state realization came after I'd shot my mouth off about how Ragnarok would be configurable for either. Yeah. Again: just shut up.

We hacked around with it a bit, blew up a few output stages, and never really heard music through it. So, finally, we did something sensible.

We went back to the drawing board.

The Second Ragnarok

The next Ragnarok was stripped down and simplified. We got a new transformer, only suitable for solid state use. We designed the main board only for solid state. We went to a two-stage amp with enough gain to give us the damping factor we needed.

But, it still wasn't the Ragnarok we'll soon be shipping. (Stop laughing.) This Ragnarok still

used DC servos and trimpots to set the bias for the amp, as we do in Mjolnir. This worked, albeit with some difficulty. Unlike Mjolnir, Ragnarok wasn't thermally stable. You'd set the bias, and it would creep up, and up, and up … until bad things happened.

In retrospect, this isn't surprising. Mjolnir uses 1 Ω source resistors in its output, and Ragnarok uses 0.1 Ω. Mjolnir's operating point is stabilized by local feedback from the source resistors, whereas Ragarok didn't have that luxury.

So what do you do in that case? Well, you either increase the source resistors (not a hot idea in a speaker amp), or you use something truly yucky, like a varistor, to try to compensate for the bias creep.

We didn't want to use varistors (which are resistors that change value depending on temperature, which is something you usually don't want—most resistors are designed to be as stable as possible over temperature, not the other way around—and considering that resistors set gain and other operational parameters, varistors are usually pretty bad juju.)

So, it was back to some hacks … more board hacks to fix the servos, and a whole lot of parts

tacked on to try to get it thermally stable … and once again, we found ourselves at an impasse.

And that's when I had my crazy, scary idea. The one that set back Ragnarok for over a year.

The Left Turn Into Hell

"Dave, if we have a microprocessor in Ragnarok, can't we use it to set the bias, too?" I asked, one day we were deep in the Ragnarok Second Prototype Fiasco.

Dave is our software guy. Well, he also does digital and analog design, too, but since he's the only one around here that does software, he's our de facto "software guy."

"Well, yes, but—"

"And if we use it to set the bias, we could get around the thermal stability problem."

"Yes, but—" I didn't hear him, though. I was getting excited now. "And, if we used it to set the operating point on both sides of the circlotron, we could throw out the DC servo, too!"

Aside: this *is* very exciting. There are two ways to get rid of DC on the output of your

products definitively (well, there are three, but I really can't count that one, because, as far as I'm concerned, it's not definitive):

1. **Coupling capacitor.** This blocks DC. DC gone. Done. Easy. But now you have a capacitor-coupled amp. And you get to try to choose the most sonically transparent capacitor you can find.
2. **DC servo.** This is a feedback loop that only operates at very low frequencies (say, 0.1 Hz and below.) This is usually more sonically innocuous than a capacitor, but even the best servos feed back some audio-frequency stuff … and cause phase shift in the bass. Not ideal.
3. **Trimpots.** And, if you're crazy, you can simply trim out the DC. And hope the operational point of the amplifier stays constant. And hope it doesn't drift. No thanks, I like certainty.

"Well, yes, but—" Dave tried again.

"Then that's what we'll do!" I cried, envisioning the truly revolutionary amplifier that Ragnarok could be with full microprocessor control. It would always run at the same operating point … and we could get rid of the DC servo … and it

would be insane! Nobody was making anything like this!

"We can do it," Dave said finally. "But the software will be pretty complicated, and I'm sure we'll run into some issues—"

"But we'll work them out!" I said, not hearing him.

In fact, I don't think I heard anything past the initial "yes."

But it didn't matter. I knew this was the way to go. Ragnarok would be the first truly fully-managed amplifier, with a giant microprocessor brain overseeing everything it did. It would be amazing! I got to laying out the new boards—which fundamentally changed almost everything about the Ragnarok. We had to add control lines from the DACs (yes, there are 2 DACs inside Ragnarok, but not for playing music … they take a 12 bit input from the microprocessor and use it to set the operational point of the amplifier.) We needed to add more sense lines so we could monitor what the amp was doing.

But I got it done, and we sent out for boards. This was it. I could feel it. We were going to have an incredible amp.

It was only later—a lot later—that I remembered Dave's, "Well, but … "

Releasing the Names

So what did I do after all these changes?

Well, I didn't do the sensible thing, which was to shut up. Instead, I believed we were far enough along that a 6-month development cycle would wrap everything up. Which meant we could make some noise. We could claim the names of our statement products (which had been the same since 2009, we simply didn't use them). And we could show off a little.

So, I released preliminary information on Ragnarok and Yggdrasil on Head-fi, together with some product renderings and a very abbreviated spec sheet.

That was more than a year ago.

Yes. Shut up. That's all there is to do. Else, suffer the consequences.

The Trouble With Software—and Hardware

In the early summer of 2013, Ragnarok seemed pretty buttoned down. We got the boards back,

stuffed them, verified basic things like DC operating points, and gave them to Dave to do the software. All that was left was to sit back and await the software, which would ensure the amp always operated the way we expected it, each and every time.

Yeah, and software never has problems. Never has bugs. It always works entirely as expected, from day 1.

Uh-huh. Right. When you get into software, expect lots of fun—and the more fun, the more complex it is.

Actually, a primer on software should go here. Because, Ragnarok's software isn't really soft. Nor is it hard. So, it's what we call "firmware."

Yeah, I know. Sounds silly. But here you go:

- **Hardware:** this is stuff that performs an analog or digital function that is set and defined. If you're talking digital, the hardware is set up to do basic or complex logic in a set pattern, for a set purpose. You don't go and change hardware. That's why it's "hard." As in, "hard-wired."
- **Firmware:** this is software that is usually embedded in a system for a specific use, and is not intended to be updated frequently (if ever.)

Ragnarok's microprocessor uses software that is programmed into it during production, via an RJ-11 jack inside the chassis. It's not intended to be updated, ever. That is, if we get it right. In the past, this might have been known as an "embedded system," and burned onto a ROM. I used to do software like this. You really don't want me doing it again. I'm a very bad software engineer. Similarly, all of our USB DACs use firmware that is programmed via the USB port; theoretically we could change this over time, but if you've ever used the C-Media firmware updater, you'd think twice (at least) about unleashing this capability on the world.

- **Software:** This is code that's intended to be changed from time to time. It's also a very imprecise word, because it can refer to everything from machine-level firmware code to C# computer programs, to apps, to website Javascript and a whole host of code done in various high-level languages. So let's leave this at that, since we have yet to do a true software product.

Okay, so where do FPGAs fit into all this? Well, let's start with the brutally honest: "FPGA," in audio, is simply the latest buzzword of the day. They do nothing that cannot be done with a microproces-

sor and code. In some cases, they may do it faster. But this is at the cost of insane development complexity. We'll stick to microprocessors, thank you.

Smart readers are sitting back now, arms crossed, asking, *So how bad was the first firmware you got for Ragnarok?*

On first glance, not bad. The basic stuff—volume control, input switching, protection—all worked. The bias and DC nulling … well, not so much.

In fact, we only got it back after Dave admitted he blew up the outputs on the main board a couple of times. This isn't surprising in development, but Dave's efforts were hampered by his lack of experience with MOSFET-output amps. Which led to the Ragnarok main board running inconsistently, or not at all.

"But I replaced the parts that tested bad," he told me.

"Dave, Dave, Dave," I chided. Experience with thousands of MOSFET-output amps gives you a certain perspective. This perspective is: *complete and utter paranoia.*

MOSFETs are not BJTs. They are not tubes. When something goes wrong, it's entirely possible for

them to test OK … for a while. It's also entirely possible they took out some other parts in the flames of their demise … which you might miss. When MOSFETs go bad, it's best to simply replace everything. And then start up the amp really slow, in case you missed something. Dave was following the procedure you'd use for a BJT or tube amp … which meant more failures, later. And, for Dave, these were mystifying failures.

And we had other problems, too. The output stage oscillated, and needed to be recompensated. Again, not a big deal for someone who is used to power amp development (hell, I get suspicious if it doesn't oscillate at some point before it's been properly compensated in-circuit.) Problem is, an oscillating amplifier doesn't exactly allow for accurate bias current measurement … which means even the best firmware is helpless.

And—to complicate matters—we had to have the microprocessor compensate for each of the gain stages, which required different voltages for each operational point. The problem was that at higher gains, the steps the DAC had to take were too large for good adjustment of the operating point.

This is where Dave stepped in, and came up with

a new feedback arrangement which allowed:

- The same operational point for all gain levels
- The same amount of feedback for all gain levels

This in itself is a huge breakthrough. As far as I know, Ragnarok is the only gain-switchable amp that uses the same overall feedback for all gain levels—which means that the sonic impact of switching gain is effectively nil.

Big win, right?

Right. I was very pleased with our solutions. Ragnarok was finally running reliably … on speakers, on headphones, etc. Everything looked good. And we were still a couple of weeks away from RMAF 2013, which means we could bring an ugly prototype to the show, and ship shortly thereafter.

Except … a couple of days before the show, I turned on Ragnarok … and it smoked. No input, no load. Just smoke.

This, my friends, is not a good sign.

And that's how Dave, Mike, and I pulled a couple of all-nighters getting Ragnarok back together, right before RMAF.

Problem solved. Right?

Wrong.

Ragnarok's First Appearance

Ragnarok's first appearance was at RMAF 2013. This version sported a very ugly, unanodized, unscreened, unfinished chassis that definitely didn't show very well. We put a big sign on it saying, "This is a prototype!"

Which, of course, was ignored. People were so excited to see and hear Ragnarok that I don't think many of them saw the sign.

The good news? Ragnarok didn't blow up during the show.

The bad news? It didn't work very well. In fact, with the early software, the operating points were anything but stable. It was all over the place. It ran hot. It ran cold. It ran in-between. It sounded pretty good once in a while. It sounded pretty bad most of the time.

All in all, not an auspicious debut. But let's go back to that advice:

1. Shut the hell up.
2. Shutting up also means "don't show it before it's ready."

Yes. Seriously.

Do you wonder why we don't talk about products before release anymore?

Algorithmic Adventure

Once we were back from the show, Ragnarok smoked itself again.

All in all, probably a good thing. Because it was time for new boards to clear up some of the analog problems we'd encountered, and to fix some bonehead mistakes on the control board. So, we lost some time as I re-did the boards and got them stuffed again.

Then it was time for firmware.

This time … this time, it did the same thing. Weird operating points. Unreliable operation. Random blowups. In other words, it was not yet a shippable product.

Which really sucked. Here we were, at the end of the year, and Ragnarok wasn't anywhere near ready. We were going to blow yet another deadline on it.

But why? It was a puzzle. At least to me. An enhancement-mode output stage will do exactly

one thing, in the absence of bias: nothing. It will sit there all day long and not blow up. It can't. It's not on. And it won't be on, unless bias is applied.

And … there was no reason it should be unstable as far as operational points go. The microprocessor was setting that. It should be, well, set.

Which pointed at the software.

"Dave, how frequently are you adjusting the bias on Ragnarok?" I asked, trying to get to the bottom of our mysteries.

"Frequently?" Dave blinked, looking mystified.

"Yes. Once a second, once every tenth of a second?"

"Once," Dave said.

"Once? Once once?"

"Yeah, when you turn it on. Or when you switch gain modes."

"Once?" I cried. "And then you let it go?"

"Yes," Dave said. "But the operational point trends downward with temperature, so it's safe."

“Unless it gets so low the bias is turned off,” I said. Suddenly the crap-sounding Ragnarok started to make a lot more sense.

Dave nodded. “Well, there is that.”

Argh. “Dave, it has to adjust continuously, over time.”

“But then you have to parse out the difference between output current into a load for a music signal, and quiescent operating point.” I nodded. “Exactly. But that’s what we intended to do from the start.”

Dave nodded, but fell silent. “That’s a lot harder.”

“Right.”

“I need to do some more work.”

“Right!” A few more weeks went past, since I was tied up with other products, other problems. Eventually, Dave and I got together with new software … and a surprise from Dave.

“I also added debugging code so we can see the output current and the voltage out from the DACs at all times,” Dave said.

And, sure enough, the Ragnarok control board sprouted a new connector: an ancient DB-9 computer port. This, Dave hooked up to a computer. Soon, numbers were scrolling on the screen, updated about once a second.

"This is the output current," Dave said. "This is the other channel. And this number is the DAC output, from 0 to 4096."

"Neat!" I said. Which is true. This was an insanely helpful tool to debug Ragnarok.

"But, uh, Dave," I added, pointing at the output current, which had risen from 250 mA to 300 mA, then 400 mA, then 550 mA … "

"Oh," Dave said, switching it off. "It hasn't done that before. Did you change something? Maybe it's not adjusting fast enough." I had changed something—the driver current. Which meant they were heating up at a different rate. Which the firmware couldn't compensate for.

"Can it adjust faster?" I asked.

"Sure," Dave said. He changed a few lines of code, and we restarted the Ragnarok. This time, the output current rose over 250 mA, but more slowly. But it still wanted to run away.

"I can change it so that it adjusts faster, the farther it is away from the target," Dave said. A few lines of code later, and we had an amp that didn't overshoot the target by more than 10 % before settling down to a nice, constant 250 mA.

Holy moly, we really have something here, I thought.

Of course, that was before the next day, when Ragnarok smoked again.

"Dave," I said. "Why does this keep blowing up? I thought we put in a current limit that would shut it down."

"I was going to do that last," Dave said, a little sheepishly.

Argh, volume 2. *Do you know how much time we would have saved if this damn thing didn't blow itself up as a failure mode?* I wanted to ask. But I restrained myself.

And, after another board rebuild and some new software, we finally had a safe, operational Ragnarok. This was a day before the San Francisco Head-Fi meet in February. Satisfied, we packed off the Ragnarok to the show.

Where … it acted like a real product for the first time. No drama. No craziness. No too-hot, no too-cold.

We were done, I thought. *Now, all we needed was to order boards and parts, and get to shipping.*

Yeah, right.

Interfacing with the Real World

Back at the office, though, tests told a different story.

Ragnarok wasn't good at differentiating between music and bias, as Dave had solemnly predicted. Compressed, low-dynamic-range music into 4 Ω speakers could cause itself to de-bias the output stage. Translation into English: very bad sound. The DC offset wasn't all that hot, because Dave was only looking at bias, not DC offset, over time. At certain temperatures, it would go into a state where it wouldn't switch into high gain mode.

Yeah. More development. We had to work out an algorithm to help differentiate between bias and music. We had to improve the bias setting over time. We had to work out the non-switching-into-high-gain problem, which was due mainly to

the way we were stepping up and down the bias (and never hitting the target, in some cases.)

Which meant, in reality, more months of software time.

And, when we were finally ready to ship the first prototypes (to Jude and a couple other NDA listeners), we still got bitten. Jude's first Ragnarok still had the "I don't want to switch into high gain" problem, which we thought we'd worked out.

In the end, we changed the entire bias algorithm. It's actually quite a neat bit of code … and something I'm not eager to get into. If someone else wants to do this insanity, let them have the pain, too.

Which put us at about TheShow timeframe. Ragnarok has been acting like a product for a while now. Everything seems sorted. Which made us confident enough to bring it to the show and put up a sign that said, "Shipping in June."

Yeah. Right. Shut up.

The Final Gotchas

With all of our software problems, we were (ahem) somewhat distracted. Which meant that things we should have been paying attention to were passed by.

Things like:

- The entire first run of transformers being the wrong voltage. Scrapped. Our fault.
- The pots being the wrong length. Re-ordered.
- The control board tactile switches being the wrong length. Ditto.
- The board house stuffing a lot of the first boards wrong. Back for rework.
- Our production line being totally unprepared for a product with 5 boards, 700 components, wiring, chassis-mount switches, transformer connectors, specific mechanical decoupling requirements, etc. Training, training.
- Our production line literally running out of space. We've gotten more space.

And I've left out a ton of details in the saga above, details like:

- 6 transformer changes, including splitting the circlotron output transformer and the input voltage gain transformer into to entities

- Adventures in circlotron summing circuits, and the multiple meltdowns that occurred while figuring out the right ones
- The insanity that is a circlotron, and the effects of nonsymmetrical loading on it
- The crazy software just for volume control—it actually mutes the input between levels, to eliminate any popping or big glitches
- Various blowups I'd rather forget

Suffice to say, the Ragnarok story isn't a story at all … it's a saga, befitting of the name. It is, by far, our most complex product to date, and it defines a number of firsts in the amplifier area—most notably as the first truly universal amplifier, using the same gain stage and output for speakers and headphones, and as the first completely managed amplifier, dispensing with DC servos and other band-aids that can affect sound.

As far as the whole saga goes—all the mis-steps, all the gotchas, all the surprises, well … this is the kind of pain it takes to make groundbreaking products.

If there's a lesson to be learned here, it's that blazing a new path isn't to be taken lightly. You should take a long, hard look at your capabilities and resources, and plan for how it will impact everything you do.

And, of course … keep your damn mouth shut until it's ready.

Welcome to Ragnarok. The end of the world.

Or the beginning?

Chapter 28

“You’ll Never Do Any Upgrades Anyway.”

When we first introduced Bifrost as “the least expensive fully upgradable DAC on the planet” in 2011, we had some interesting responses.

Some of them went like this:

“Yeah, but it’s not like there’ll be any upgrades.”

“Not that you’re planning to do any upgrades.”

“I’ll believe it when I see it (with respect to the upgrades.)”

Why the doubt? Hell, I don’t know. Audiophiles can be a morose bunch. Maybe that’s the only reason. Or, maybe there was a manufacturer which promised upgrades and recently went out of business, or otherwise didn’t make good on the promise—and that was coloring their response.

But, needless to say, I was shocked at the amount of negative sentiment we received. Sure, there

was plenty of positive press, but the opinion of "the audiophile on the street" was less thrilled. I didn't worry. I knew we'd have upgrades. It's in Mike's D/A to do upgrades. I just didn't know when they'd come.

First, Let's Talk Theta

Mike's upgrade D/A was implanted at the inception of Theta Digital, I believe. (Mike, correct me if I'm wrong, or if you think this story needs more bat testicles.)

Why? You have to consider the environment. When Mike started Theta, digital audio was in its infancy. Manufacturers were still trying to get the price of a good CD player under $500 (think $900+ in today's dollars.) There were no standalone DACs. Zero. None. SPDIF, as a transmission standard, was brand new.

And Mike wasn't just at the leading edge with Theta—he was bleeding edge. Literally. Before Theta, nobody had even considered making a standalone DAC. And nobody else would have started with a flagship $3000 product, using their own digital filter code running on megadollar Motorola DSP chips.

Aside: Think about that a bit. A flagship product that "only" cost about $5500 in today's dollars? Insanity! Add more CNC-machined parts and a fancier chassis and some custom dampers and heebie-jeebie clocks, until it's $100k+! Yeah, that's the only way to go! But that's a comment on how far we've fallen in the past 30 years. From a $5500 first-of-it's kind product using the latest, greatest, very expensive chips of the time, designed and coded by a team that spanned university researchers, brilliant mathematicians, and top-of-their-field engineers, over a timespan measured in years, all packed in a relatively plain-jane chassis, to beautiful, overpriced audio jewelry powered by voodoo and obfuscation. (And, you think this is harsh? Ask Mike what he thinks.)

But, back to the subject at hand: digital audio was new. Nobody knew what was around the corner.

And things were changing. In contrast to today's relatively stable market, new D/A converters and new digital filters appeared regularly—pushing from 16 bit to 18 to 20, from no oversampling to 2× and 4× and 8×. New technologies appeared as well—the first "one bit" D/A converters. And,

these new chips (with the exception of the "one bit," or delta-sigma" converters) usually offered significantly better measurable performance than their predecessors.

And, manufacturers were learning as well. Mike was the first to measure jitter, opine that it might have something to do with the sonic deficiencies of early digital audio, and devise ways to minimize it. He experimented with the best interfaces for SPDIF, using transformer-coupled coaxial, and later adding AT&T ST-optical glass fiber as an option.

In this constantly-changing environment, bringing a very pricey, first-of-its-kind product into this market with no ability to upgrade it would be suicide. Because, even if you won some first sales, how thrilled would your customers be when it fell behind the latest "latest and greatest?"

In short: they wouldn't be.

Which is why Theta Digital was built, from the start, around upgradability. Theta owners could pay a relatively nominal amount to upgrade their DS Pro Gen 1 to 2, 3, and 5 (there was no 4) as the years passed. Same with Theta's lesser gear. This way, they could keep pace with technology, without filling trash cans with their old DACs.

Fun fact: Theta's first non-upgradable product was the Cobalt 307, which I designed—their first "disposable" DAC, at "only" $599 in 1992. Compare to today's $99 Modi. Yeah. There is such a thing as progress.

Today's "Stability" and Upgradability

"But … but today, digital audio isn't stable!" some of you are protesting right now. "USB changes quite a bit, and there's DSD, and standards for transmitting digital over WiFi and Bluetooth, and high-res music, and all that."

Yes. But it's still much more stable than the 1980s and early 1990s, when everything was changing. Today, most DACs have settled down into a comfortable model: inexpensive delta-sigma D/A conversion from AKM, Analog Devices, Crystal, ESS, TI, or Wolfson, coupled with a USB interface solution from C-Media, TI, Tenor, or XMOS, plus (perhaps) a SPDIF interface from AKM or Crystal … plus, of course, a power supply, analog output stage, and associated interface electronics. Sure, there are some outliers, but they're usually at the scary end of the price spectrum.

Beyond that, let's go through the sources of instability today:

- **USB.** Much more mature, but still improving. When Bifrost was introduced, things were a lot less stable. Today, we have robust, relatively good-sounding solutions. But, we suspect they will get better. So, yes, the ability to upgrade the USB input would be a plus.
- **DSD.** Yeah, it's there, but the floodgates have not opened. In our opinion, best to concentrate on the 99.999 %—that is, PCM—rather than a possible dead-end excursion à la HDCD. Still, if it ever becomes more than 1% of the market, sure, it would be good to have upgradability that would allow DSD decoding.
- **WiFi Digital Audio.** Definitely changing, and may be promising in the future. The potential for uncompressed transmission is definitely there, but the challenge is in the interface (joining WiFi networks means a I that is computer-like in its capability and configurability, unlike Bluetooth.) In our opinion, best to sit this out while it matures a bit. Jumping in now would be kinda like backing a digital transmission method before SPDIF was established as a standard—with the downside of extreme software prowess and support requirements.
- **Bluetooth Digital Audio.** Definitely changing, but too compromised to jump into for the sake of convenience, in our opinion. None to

date is uncompressed (not even AptX.) Much of this is highly integrated, single-chip solutions that deliver analog output (not a digital datastream). Some can get I^2C out. New Bluetooth standards with higher speed may allow uncompressed transmission. Probably too early to build a standard "Bluetooth compatibility module" into a DAC—just use an external solution until robust, uncompressed, I^2C out versions are available.

So, what does this all mean?

When we started on the design of Bifrost, the latter three weren't really a concern. We were worried about the USB input changing over time, and, to a lesser extent, the DAC and analog section.

Which is what drove our decisions—and our caveat. When we introduced Bifrost, we told potential purchasers, "We're not going to have a DAC of the Month club or anything like that. When real, meaningful changes come to USB or the DAC/analog stage, we'll release an upgrade."

Of course, we didn't know when those real, meaningful changes would happen.

But, after Gungnir development, Mike started wondering what the Gungnir analog stage would

sound like in the Bifrost. That led to the first possible upgrade.

And, early in 2013, C-Media laid the second one on us: the CM663 A USB input receiver chip.

Which meant, as we went into 2013, we knew … upgrades would soon happen.

But First, Let's Talk About the Way To Do Upgrades, and The Way Not To

"Real, meaningful changes." That's an important phrase for any manufacturer thinking about doing upgrades. That's the phrase that keeps you from being the Burger King of audio.

Remember: it's not about making everyone *like* you. It's about making some people *love* you. And you won't achieve that if you offer everything to everyone, with no position on what is best.

For example, it would be relatively easy to offer Bifrost with a half-dozen different DAC/analog sections. Just pick the latest chip from AKM, Analog Devices, Crystal, ESS, and Wolfson, re-do the analog section to meet their specific requirements, reprogram the motherboard, and you're

off and running. Everyone could have the DAC of their choice! Have it your way!

Yeah, except:

1. Everyone would also argue about which one was best.
2. If there was no consensus, nobody would have any idea what to buy.
3. If there was consensus, we'd be stuck with a ton of DAC boards that were impossible to move.
4. We'd have to stock 12 different versions of Bifrost (6 D/A options, plus with or without USB.) Total disaster.
5. We'd have to spend 6× the engineering time in development, or shortchange one or more implementations.
6. We'd have to maintain records for all variations, so they'd be serviceable in the future. And maintain the records so they were up to date.
7. People would want to order two, or three, or four, or five, or all 6 and swap them, which would be a disaster, since the motherboard would have to be programmed, and Bifrost wasn't designed to be hot-swappable. Say hello to unqualified people disassembling the product and possibly hurting themselves—

then say hello to a huge lawsuit.[1]

No. Upgrades are not "have it your way." They are a path. A path to a rational future where you're helping mitigate the cost of buying a whole new product.

Which means, quite simply: pick your upgrades carefully, and keep them to a minimum.

That is sanity. The other path, less so.

On USB and Mohammed's All-You-Can-Eat Sushi and Deli

When we got our first CM6631A USB receiver chips, we were both thrilled and cautious.

Thrilled, because there were some notable limitations with the older CM6631 that the original

---

1 This is absolutely no joke. There's a reason that "No User Serviceable Parts Inside" is printed on the back of virtually every electronic product, and why we say, "If you aren't an electronics professional, have us upgrade your product." That reason is: we want to stay in business. If we ever hint that it's OK to take apart a product that's powered by an AC line cord, believe me, that email or forum post will be dug up by a forensic attorney and used to hang us out to dry, when someone decides to take apart their Valhalla 2 … while it's plugged in … and while they're taking a bath.

Bifrost USB board was based on—namely, lack of 24/176.4 support, and an extreme pickiness about the USB interface and cable quality. A tiny percentage of systems simply didn't like the CM6631 interface—on the order of 0.3 %—but, in the time-honored tradition of Murphy's Law, of course 1500 % of that 0.3 % happened to be forum members who complained about the problem.

Cautious, because we knew that "new stuff" isn't always "better stuff." C-Media promised 24/176.4 support, as well as support for even higher bitrates (24/384, specifically), but were mum on the subject of there ever being a problem with the CM6631 with some USB interfaces.

So we sent a few out to our PCB assembly house to have built up on the current-generation USB board. The chips were almost entirely pin-compatible, so that was the easiest way to see what they'd do.

Or, more accurately, Mike and Dave did it. Dave may have even soldered some of those 100-pin QFNs himself. No thanks, not for me.

In any case, we soon had a handful of USB boards with the CM6631A on them, happily programmed and running away.

“Here you go,” Mike said, handing me a little foil bag one day. “Try it out.”

“What is it?”

“The new USB receiver.”

“How’s it sound?” I asked. The first-gen USB board was pretty darn good by USB standards, but it was never any great shakes by good SPDIF standards—hence all our trash-talking of USB when Bifrost first launched.

“You tell me,” Mike deadpanned.

> He does that a lot. He wants to know what you think, not get confirmation of his own notions.

So, I went home, installed the board, and plugged it into my most notorious source—an older Apple MacBook that would sometimes not play nice with the older USB board. It fired right up, showed all the sampling rates, and played fine.

Pretty damn fine, actually, I thought, after a while.

And, it wasn’t glitching like the earlier board.

Hmm. Maybe we had something here. I tried it on a couple of different computer sources, and they all worked fine. I played it for Rina. Her face lit up. “This USB actually sounds good.”

"I thought it sounded better," I told her. "But I was more worried about the functional side."

"No. This is good. Really good."

Still probably not as good as SPDIF, I thought.

But after switching the Bifrost back to optical input, I had to scratch that thought. It wasn't the same as SPDIF, but it was definitely not worse … black and white had been turned into shades of gray.

Aside: Mike still prefers SPDIF. I listen mainly via USB, though. To me, SPDIF still sounds a bit more natural, but slightly smeared. USB sounds more precise, but slightly crispy. Yes, I've never been good at audiophile adjectives. I'd be toast if I got a job marketing audio jewelry.

"So when will we be selling this?" I asked Mike, the next day.

"You liked it, huh?"

"Yeah, lots better!" I said. "I think it might be better than SPDIF in some ways."

Mike recoiled in mock horror. Or maybe real horror. "Have you started listening to techno? Savoring Mohammed's All-You-Can-Eat Barbecue

Sushi and Deli? Drinking plastic-bottle Vons-brand tequila? Huffing Testors paint?"

"No, I'm serious."

Mike shook his head. "Ohh … kay. Remind me to pick the restaurant next time we go out."

"Mike!"

"And make sure there's some real music on the computer when we go to shows."

"Mike!"

"Hell, maybe I should just buy you some Beats right now. USB? Better? For audio?"

"It's shades of gray," I told Mike.

Mike shook his head again. I don't think he'll ever really like USB, but he's largely stopped complaining about it. And that says a lot.

On the Voodoo of Analog

At the same time as the USB Gen 2 board, we were investigating another upgrade … this one pulled straight from Gungnir. The idea was simple: if the Gungnir sounded so much better than Bifrost, what would a Gungnir analog stage sound like in Bifrost?

This wasn't as simple as dropping it in, however. Gungnir's stage was designed for the high ±4 V rails in Gungnir. Not for the ±5 V rails in Bifrost. I had to change every value on the board to get it optimized for its new home.

But beyond that, you guessed it—the Uber Analog board is, pretty much exactly a Gungnir board. It has a much more sophisticated discrete topology, and a DC servo to eliminate the coupling caps in the original Bifrost analog stage.

But it's the same DAC—the AK4399.

Why not a different DAC? Well, we hadn't found any we liked better than the 4399. It's that simple. And, to this day, we still like the AK4399 and AK4396 better than the newcomers—they are unique in their implementation of switched-capacitor filtering to lower high-frequency noise, which seems to give them a more natural quality than other DAC options.

> Or it could be all in our heads. The voices, the voices!

If you'd told me that our first DAC/analog upgrade would use exactly the same DAC chip as our standard board, I would have laughed …

… until I heard the new Analog board.

It was a big step up. At least as big as the USB upgrade. Bifrost was no slouch in standard form, but with the new Uber board, it was in a different class.

From there, the hard part began.

On Naming

So, why "USB Gen 2" and "Bifrost Uber?" Simple: to differentiate two very different products.

"USB Gen 2" was a product that could be used in both Bifrost and Gungnir. And it was only the first of a line of USB upgrades, we expected (yes, someday there will be a USB Gen 3 … or maybe even something else to plug in there … but as to when, no idea …)

Bifrost Uber, because it really took Bifrost up in performance, and we wanted a way to differentiate it, without going to a whole new model. When, some day, there's a new DAC/analog board, we'll figure out what to call it then. Uber 2? Who knows? But again, no idea on timing.

And … that's also why you won't be seeing a Bifrost 2 or Gungnir 2 … perhaps ever. They're so modular that they can keep going, with relevant

updates, pretty much indefinitely. Or at least for a very, very long time.

The Logistics of Upgrades

You know how, when a new car model comes out, a whole bunch of people have to have it RIGHT NOW?

Yeah. That's the problem with upgrades. When you announce an upgrade, and you already have a significant number of products in the field, the first months are gonna be crazy. Everyone's going to want it, and want it RIGHT NOW. That could easily bring your service department to its knees—especially if everything all comes in at once. Then you'll also have the fun of trying to explain to customers why it's taking two weeks to turn it around.

Luckily, we already knew this from the Theta days. Unluckily, we also knew that our current systems would never survive the onslaught.

Which is one reason we completely rebuilt the website—throwing out it's creaky old taped-together platform for a new, custom-from-the-ground-up development that was tailored exactly to our needs.

This change allowed us to put in place a queuing system to help manage the updates. With the Schiit Upgrade Queuing System, customers could buy the upgrades, and we could tell them when to send them in so we could guarantee fast turnaround. It also kept everyone in the loop—sending automated emails when the product was received, and when it was shipped out again.

But then, there was the question of self-upgrades. As I mentioned before, the "No User Serviceable Parts Inside" disclaimer was no joke. We seriously considered not allowing self-upgrades.

But that would inconvenience people who really could do it themselves—and it would especially inconvenience international customers, who would have to hope their distributor would be able to do it.

So, a week or so before we announced the upgrades, we reached a compromise: purchasers could choose to self-install, but only if they said they were a "professional electronics technician." The copy is still in place to this day on the upgrade product pages.

"They're gonna break them," Alex said, before we launched.

“We’ll see,” I told him.

“They’re gonna screw up their Bifrosts, and we’ll have to fix them.”

“Let’s see how it goes.” I reassured.

Alex shook his head, but agreed, yes, it was worth a try.

So how did it work?

Well, it was a damn good thing we had the queuing system in place—early orders could have easily swamped us, and even with the system in place, we were sometimes doing 20 upgrades a day.

And, as far as self-install goes, it turned out to be a non-problem. Maybe one or two boards were destroyed by careless individuals, or maybe they just died in transit—yes, even after testing, it could happen. But there wasn’t a flood of bad boards, or bad Bifrosts. So, our compromise ended up being the right thing, in the end.

And those early comments about, “Well, you’ll probably never do upgrades anyway?” Yeah, those ended.

Instead, some people said, “Wait, that’s not fair! You mean you’re gonna milk us every year, and have us re-buy the DAC we just bought?”

Yeah. As if throwing away an entire DAC is a better solution. And as if an upgrade suddenly makes their current Bifrost or Gungnir unusable.

What was I saying about some audiophiles being a morose bunch?

Chapter 29

Worst. Customer. Ever.

Okay. This'll be interesting. I'm about to embark on something that's probably akin to a soldier mooning a sniper, a politician shaking a baby on national TV, or a sports coach making racist remarks on YouTube.

Because, well, you know, *The Customer Is Always Right.*

Yes. In title case and with italics for emphasis. Because this is what we're taught. The customer is god. The customer has all the power. The customer, no matter what, is always right. Always. Without exception. No debate allowed.

So, if I start talking about less-than-ideal customers, I'm grabbing a third rail. I can hear the shrill panic of politically-correct sales managers everywhere, echoing around in the back of my mind:
Nooooo! Don't go there! Never insult the customers! You don't know who they know! It's

Armageddon, end of days, the heat death of the universe!

But this is an important discussion. Because, if you're going to get into business selling product direct to customers, you need to know two things:

1. You're gonna get some buttheads.
2. You're not gonna make everyone happy.

Note: the above two rules apply even if you are *giving away* free Ferraris with sales and use tax prepaid, or a free magic rejuvenation pill that takes 20 years off your age with no side effects, or free 6000 square foot homes in Malibu. It'll be too hard to drive to the store, or your young visage will make you less respected at work, or the home will be too ostentatious and the wrong color.

So what do you do? Well, I think the poster set above Alex's desk sums it up perfectly. It reads:

> We bend over backwards for our customers.
> But we won't be bent forwards.

Love (Most Of) The Customers

Okay, now, let's be clear, though: no business selling direct can be successful without loving and caring for their customers. Period. If you are so cynical that every interaction with a customer is a war, don't bother starting a business selling direct.

And we do really enjoy the vast majority of the customers that contact us. Many of them use the same humor we do—Schiit puns, jokes, etc. Many of them are very complimentary regarding pricing or performance, or the made-in-USA aspect. Many of them just need a simple question answered—something we may have forgot to cover, or were unclear about.

Aside: if they're asking about something that's unclear, make sure you fix it right away. If they're confused, lots of other people are confused, too. Remember Amazon's rule of customer service: *If you have to contact us, we've failed.*

And that brings us to some business-y stuff we should get into, before we go into the Worst Customer Ever. Business-y stuff like, "How do you maintain a high standard of customer service?"

Well, that's a bigger question. Because it gets into things like:

- How do you architect customer service to help ensure satisfaction?
- What do your customers consider a high standard of service?
- How will your customers contact you to get service?

Let's break these down:

Architecting Customer Service. *Wait, what?* Some of you are saying. How can you architect customer service? Isn't it just, people ask questions, and you answer them? They have problems, and you fix them?

In general, yes, but you can do a whole lot of things architect your customer service so that you can offer much faster response—and minimize the personnel and time you need to give great customer service.

Note: Yeah, I know, "minimizing personnel and time," may seem like, well, not the most customer-centric thing you can do. But, in actuality, it is. The smaller the customer service team, the more consistent and high-level answers you can give. And the less time spent on each inquiry means that your response rate

can get really, really small.

So what can you do to architect the customer service experience to make it better for both you and the customer? Several things:

1. **Prohibit the "hard sell."** If your customer service team has a dual duty of solving problems and selling product—especially if they are measured and rewarded on their sales—you're in deep, deep trouble. This encourages the typical "audio fellatio" with promises that you're going to hear god without drugs, or be transported into a magical land of unicorns and DSD—and, of course, the more you spend, the more magical it gets. This gets in the way of honest answers, it makes promises that may not be paid off, and it takes a ton of time to do, especially if you're talking about expensive gear.
2. **Ban discussion of other manufacturers' products.** Same deal. If your customer service team is expected to give value judgements on other manufacturers' products, and upsell, wow … now, not only have you opened the gates to infinite time and unfulfillable promises, you now have started building a reputation for trash-talking other people's gear.

Yes, I know, lots of companies do it, but it doesn't mean it's right.

3. **Don't do promos, points, or sales.** Want to really open the floodgates on customer service? Start doing promos, customer loyalty points, or sales. Then you'll get to hear nonstop about things like, "Hey, I just bought this yesterday, now it's on sale," or "When's your next sale, I'm waiting for that," or "I can't figure out this points thing, how does it work?" This will absolutely eat your customer service alive—it will become all they do.

This is why, almost from the start, we put into place three policies that really, really simplify customer service. These are:

- No sales
- No hard-sell/audio fellatio
- No talking about other manufacturers' products

About 15 % of our emails are of the "can I get a better price" variety.

With no sales, no promos, no discounts, no loyalty program, the answer is easy.

In addition, we don't have to have any staff to manage the sales, promos, discounts, loyalty program, and the resulting refunds, exchanges,

special deals, stacked offers, etc. that go with it. The result is simple: less complexity, less staff, and lower prices for everyone.

Wait a sec? Lower prices? You're asking. *What about the sale prices?*

Yeah, and what about the necessarily higher price you need as a baseline if you're going to do sales? Or the necessarily higher price you need for the staff you need to deal with sales, promos, and loyalty programs? Lowest complexity equals lowest price for everyone—and nobody thinking they got screwed because they missed the sale. That's a win-win—lower cost and higher customer satisfaction.

Aside: with respect to sales, I can't say this more strongly: DON'T. EVER. Or you'll become addicted to them. They will never end.

We get more emails about "will I hear a difference from my receiver/computer/etc" or "how does this compare with the Arglebargle XYZ?"

Again, our policies make the answer really simple. We don't know how you hear, or if you think the difference is meaningful. And we never discuss other manufacturers' products.

Which is really interesting. While we look at this being an honest, upstanding manufacturer, this policy really lights some people on fire. They really really really want to be sold, and they get majorly pissed when we won't engage in the usual circle-jerk about how our stuff is the greatest thing on the planet.

But … ask yourself two things:

1. Do you really think we've heard everything on the planet?
2. Do you really think we have the same sonic preferences you do?

Fact is, we probably can't make a lot of those comparisons you're so eager to hear our opinion on. Ask us about what amp we think best for, say, LCD-2s and K701s, sure—we can answer that. But comparing and contrasting our amps with products we've never heard before … um, no.

Maintaining a high standard of customer service. Okay, this is where we have to get subjective. Because we don't have all the answers. We haven't done extensive customer satisfaction studies, and we haven't tested a system architected from those studies.

But, we think that a "high standard" can be defined relatively simply.

Why? Because we have ample examples of what people hate. When was the last time you talked to your cable company? Or your cell service provider? Yeah. Endless trouble-trees going through all the stuff you already told them and email responses that stretch into days. Nobody likes that. It says, loud and clear: we don't care. You're not important. You're part of the little people.

So, a high standard of customer service starts by inverting their model.

Which is what we try to do. Here it is, in one sentence: *put enough information up about the product so most people can make their own decisions, but when they contact us, make the answers fast and simple.*

Read that again. *Fast and simple.*

Read it again, stopping at the first word. *Fast.* I can't stress this enough. Fast response to customer questions is key. Many businesses promise 24 to 48 hour response to email. This is woefully inadequate. There's no way someone can make a decision, much less troubleshoot a product with a 24 to 48 hour response time. It's like getting customer support from Pluto.

During regular business days, 24 to 48 minutes is more like it. Or even 2.4 to 4.8 minutes. We aim to keep this as snappy as possible. Which is why you're usually looking at minutes for an email reply during the week, and we check email at least 2 to 4 times over the weekend on non-business days as well.

Is this the full answer? Probably not. But it seems to make most people very happy.

Choosing contact options. And, with this, it's a good time to talk about the kinds of customer contact options you have. Because, if you're starting a business, you have a ton of options these days. Which means that most companies start off by checking every box on the options list:

- Email
- Chat
- Phone
- Skype
- Facebook
- Twitter
- Mail

Congrats. You just screwed yourself. Who's going to be immediately available for chats? Who picks up the phone? Shouldn't they be making

something? Skype rings on which computer? Why are you paying attention to Facebook if the rest of your customer service is working? Twitter, are you kidding? Mail? This is the 21st century, what are you going to do, send out brochures?

Here's what we chose to do as an experiment, when we started Schiit:

- Be really fast at email
- Never pick up the phone, but call back in 48 hours or so
- Add additional services if this didn't work

We've never added additional services. Why? Because email response is so fast. I was talking to the owner of another audio company when the subject of customer service came up.

"How many phone calls a day do you get?" he asked.

"Um … average?" I asked. "Maybe 1. Maybe 2."

"What the hell?" he exclaimed. "How do you keep it so low? I got three people working full time on the phones." I grinned. "Easy. We tell people not to call us."

The other audio company owner's eyes bugged out. "You … what?"

"It says right on the website: email us, we're really fast. Call us, and we may get to it eventually."

He shook his head. "And that works?"

"Apparently."

And it does work. It's called, in corp-speak, *setting expectations*. Expect fast email response. Expect really, really slow phone response. Which do you choose? There you go.

But I want to call, some people are saying now.

Yep. Got it. *Want to pay 25% more for your products?*

Thought not.

Fact is, phone support—that is, good phone support, from a real audio guy who can really help you with your question or problem—ain't cheap. It eats up pretty much a person's entire day, because the guy answering the phone doesn't know if he's going to get tied up on a one-hour call covering the life story and audio adventures of someone, say, interested in maybe buying a Magni.

Even on email, it gets interesting. How about 73 emails with 85 questions from one person … on Magni and Modi. Not kidding.

If we invited phone calls … if we were *good* at phone service … we'd probably have at least 2 to 3 more salaried staff working full time on it. Nick can handle all the email we get in a day—questions, support, etc—while still being a tech.

Wait, wait, wait! Some of you are crying. *Are you essentially saying your support options are pretty much email-only? How the hell do you get away with that?*

Simple. With fast answers. If our ran our email like Comcast, well … things would be very different.

So How Bad Was the Worst Customer?

Okay, so was the worst customer someone who didn't like our products?

No. There are plenty of people who don't like our stuff. That's perfectly fine. Send it back, get a refund, no harm, no foul.

Or someone who expected 10-minute response on a holiday weekend, and sent 12 emails complete with onomatopoeic descriptions of what his product was doing?

Again, no. Maybe we should be more clear that we don't work weekends, nor on holidays, so our fast email response may be, er, a little slower during that time period.

Was it the guy who got a new Magni and Modi, threw a temper tantrum when they didn't work (or he couldn't get them to work), so he beat the crap out of the product and sent an email photo of it in his trash can?

Again, no. Though perhaps some anger management is in order there.

Was it the guy who sent 73 emails with 85 questions asking about Magni and Modi?

Again, you guessed it, nope. He never bought them.

> **Aside:** although we haven't done big customer satisfaction research, we have run some statistics that are very interesting. One of them is that anyone who emails us before purchasing is 8× more likely to return the order. 2+ emails takes it up to 30×. But again, are these bad people? Not at all. Merely indecisive.

No. The worst customer ever wasn't just bad. He was criminally bad.

Here's what happened.

After work on Friday evening, I decided to check the customer service email. Until December 2013, I was the primary guy who answered customer service email, so this in itself wasn't an unusual event.

What I found was disturbing, though—an email from a very, very irate customer who had ordered a B-stock Mjolnir. Back in those days, we sold B-stock manually, by individual inquiry. If someone wanted B-stock, they had to email us, we told them the price, and if they were interested, we sent them a PayPal invoice.

> **Aside:** today, all B-stock is sold through Amazon.

Apparently I'd sent them an invoice, and they'd paid for it. However, we didn't ship the Mjolnir the same day, as requested. This also isn't unusual, since we quote a 1 to 3 business day shipping time on in-stock items that aren't ordered with expedited shipping.

But that didn't matter to this guy. He was livid. I mean, full-boat, cartoon-steam-whistle-out-the-ears, screaming red-faced rage. In acidic sentences strung in all-caps, he told us what a terrible

company we were for not shipping it right away, expressed his extreme displeasure with our customer service, questioned our competence in an overall manner, and made various other personal assertions relating to our lack of professionalism and discipline.

And, to top it all off, he told me that Alex was the worst person in the universe, he didn't care about him as a customer, and had never returned the emails he'd sent earlier in the day.

In a perfect, algorithmic, Mr. Spock-driven world, I would have tweaked an eyebrow and said, "Curious," then investigated this incident in a dispassionate manner.

Humans don't work this way, though. I was pissed. I'd been called an incompetent idiot. Alex had been called much, much worse.

So, I bit back my first response and emailed Alex, asking if he'd replied to the guy's emails.

Alex sent me a long string of increasingly irate emails, beginning at 10 that morning—all responded to in less than 10 minutes by Alex.

Okay. That's all I needed to know. Screw Spock and dispassionate logic. This guy was a butthead

of the first caliber. What could we do? I called Alex. "What can we do about this guy?"

"If it were me, I'd give him a refund and invite him never to be a customer again."

"Can we do that?" I asked.

"I can have FedEx re-route his shipment back to us." I only had to think for about a millisecond. "Do it. I'll refund his money, then he's a non-customer."

"Done," Alex said, and went off to do what he does with shipping. He came back a few minutes later via email. "Done and done."

Cool. I went into PayPal and refunded all of his money. We'd be out the shipping and rerouting fees, of course, but that was a small price to pay to be rid of him.

Aside: seriously, I am saving your mind by not posting the emails here. They were seriously, pathologically disturbed. This guy was, no crap, going to lose his mind because his amp shipped a day late.

There we go. Package rerouted, money refunded, done. Right?

Wrong.

The guy came back to me about 10 minutes later on email, even more livid than before. He'd noticed that we refunded his money, and wanted to know what was going on. (But with about 10 000× more expletives and rage.) I sent him a pleasant email in return, saying something like:

> Dear Butthead *(actually his real name)*,
>
> We have refunded your purchase in full and re-routed the shipment of your B-stock Mjolnir. We have done this because you are so disappointed with our service to date. If you are this unhappy now, we have no confidence that we will ever be able to make you happy. We believe this parting is for the best, and wish you luck in finding the perfect component to meet your needs.
>
> Sincerely,
> Jason, etc.

Oh, boy, was he ever pissed. After four or five more irate emails, Alex and I seriously wondered if we'd meet a guy with a lead pipe at the office on Monday morning.

But Saturday was quiet—no emails.

And Sunday was the same.

And nobody was waiting to jump us on Monday.

So, end of story, right?

Oh no.

About a week later, Alex starts wondering where that Mjolnir went. It had never come back to us. And it had only been shipping within California, so it should have come back quickly.

He checked the shipping record, and quickly found the problem: the guy we'd shipped it to had called FedEx himself and rerouted it to a FedEx office, then picked it up.

Yes, the guy we refunded and finalized the transaction, so we couldn't charge against his card again.

"What do you want me to do?" Alex asked me, his eyes dark and murderous.

"Whatever you want," I said.

Luckily, Alex is very good at internet forensics. Through this guy's multiple email addresses, he was able to track down his LinkedIn, his business website (yes, he had his own business), and his

Facebook page. On his Facebook page, in plain view, was a photo of the Mjolnir.

Alex sent an email (and a registered letter) to the guy's main business address, demanding payment for the Mjolnir within 72 hours—or a visit from the sheriff's department.

He tried to play it off:

> Thank you for the gift of the Mjolnir, in compensation for your poor customer service

the smug bastard emailed.

We reiterated that it was clearly not a gift, and repeated the timetable. Pay, or see how having a criminal record for grand theft felt.

Over the next few days, we endured various emails about what terrible people we were, our relationship with our mothers, how we had small body parts, etc. None were responded to, save to remind him of the time ticking away. I really thought we'd have to get the police involved, but on the last day, he blinked. He paid for the Mjolnir.

And that, really, was that.

However—if he ever orders anything else, it won't be shipped. If that Mjolnir comes in for service, he will be getting a check for full value in return, and we're keeping the amp. We don't need customers like that. Ever. For any reason.

Starting a business? Working with customers? Repeat after me: *not every customer is worth having.*

Bonus: How to Get Great Customer Service From Humans

One of the problems with customer service these days is that most of it has become by-the-book and algorithmic. *Choose from these available options so we can route it to the right department. That problem wasn't found. I'm sorry you're having trouble we value your business, your expected wait time is 50 minutes.*

Yep. Endless trouble-trees, backed by least-experienced customer service personnel who ignore the long list of troubleshooting you've already done. Four hundred words of boilerplate about what a special customer you are to them, and how they're truly so sorry they're going to

self-immolate. Ticket systems that promise transparency and continuity, but don't deliver when shared by a team of 150 people.

The result? Everyone knows that when a big company says, "We value you as a customer," it's 195 % BS.

Which means it's open season, guys. Get out the 12-gauge! Give them both barrels! Let 'em have it! Because they aren't really human, and they're not telling the truth!

Is it any wonder that yelling and screaming at large company customer support personnel is almost, well, accepted?

Because if you make enough noise, you might get somewhere. You might trigger the Rage.2.Uplift and get someone who knows more about what they're doing. Or you might trigger the PITA.Refund.1.1 and get your money back.

Big companies are making a very prickly bed with this combo. When nothing but anger works, they're going to get nothing but anger. And then it doesn't work.

Companies like Schiit are a little different. Hell, I bet most audio companies are a little different. Hell, most small companies, period.

Which means if you come in, guns blazing, things may end up very, very differently than you expected.

Why?

At Schiit (and companies like us), you're talking to humans.

And humans have emotions. They are not slaves to a script or to a corporate customer service code of What Can And Cannot Be Done.

Not only that, at Schiit, you're talking to fully empowered humans. Nick, Laura, and Alex all have carte blanche to give you anything they want—or nothing at all.

Now, this doesn't mean that you need to suck up to them. In fact, that can be just as irritating, or more so, than anger. But you should be aware that you are talking to humans that can—and do—bend over backwards. But if you come in hot, that willingness to bend over backwards diminishes.

Instead, if you come to us:

- In a concise manner (we do not need Your Life Story With Audio)

- With clear questions (you'd be amazed how many emails we get where we ask, "Was there a question here?")
- Using complete sentences at least some of the time (no kidding—and phone autocorrect is the worst)

You'd be surprised at the fast, helpful answers you'll get. And you'll be shocked at how much we'll bend over backwards to make you happy.

Aaaanddd … I'll bet any other company with human-powered customer service will work exactly the same way.

Remember, you're talking to people, not machines.

Don't be the next *Worst. Customer. Ever.*

Chapter 30

Death of a Product

In the spirit of full disclosure, I probably should have put this chapter before the previous one. The events outlined here occurred in the early to middle part of 2013, whereas our "worst customer ever" came a bit later. But when I originally outlined this book, this chapter didn't exist at all—mainly just because I was dumb and forgot about it.

However, this doesn't make this chapter any less important. It covers something that any business will have to face at one point or another: the death of a product.

Of course, most product deaths aren't usually deaths, per se—they're more a phoenix-like event, where a new and even shinier product rises from the ashes of the previous one. At least hopefully.

But … that leads us into the first question. When a product gets long in the tooth, do you update it, or give it the full Kevorkian treatment?

Good question. And, with that, let's talk about product life cycles, and product life cycle management. Yeah, good boring corporate stuff. But I'll try not to make this too tiresome.

Product Life Cycles, AKA the Game of Update, Assassinate, or Cannibalize

Okay, let's start with the basics:

1. No product is fresh or competitive forever, especially a technology product; the competition, and the market, can and will change—sometimes in new and unpredictable ways.
2. Because of this, you have to think in terms of product life cycles—or, in regular English, how long a product will be a good, solid competitor in its market.
3. You should determine (at least) a guess as to how long your product life cycle is, so you can be working on updates or replacements before the end of its life.

And, the bonus stuff that most companies ignore:

1. Killing your babies is perfectly OK, if updating won't make them a good product for new market realities—you have to be ready and willing to do this.

2. It's better for you to cannibalize your own product lines, rather than waiting for someone else to do it.

The above is why you typically see an iPhone every year. It's why most other flagship phones are on the same life cycle—the technology, software, and market have changed enough in a year that a new, fresh product is required to stay competitive.

It's also why you see new laptops and such on timeframes dictated largely by the release dates of new chipsets from major manufacturers like Intel—the introduction of the new chipset changes the game enough so that new products need to be introduced.

Cars? They have longer life cycles, typically 2 years between minor refreshes, 4 years between "making it look new on the same platform" and 8 years between moderate to major platform changes.

Same goes for a lot of less techy stuff—appliances, etc. Their product life cycles can be much longer than a year.

Audio? Hmm, now that's a conundrum.

On the mass market side, the major manufacturers of "bulk" gear like receivers and such have

been chasing a 1-year product lifecycle for a very long time—but the new products that come out frequently aren't anything more than re-badged and slightly de-contented versions of their predecessors.

Aside: My theory is that the ongoing decontenting allows them to maintain arbitrary price points (determined by copying their competition). After all, it's much easier to follow somebody elses' rules than make your own. The latter might require intelligent marketing to show how your products are, well, *actually* different from, and better than, the competition.

Another aside: And that's assuming they *are* actually different and better—which may not be a good assumption in the mass market.

In high-end, product lifecycles are all over the place. Some companies make essentially the same products for nearly a decade. Some make changes every couple of years.

What's right?

- **Some use the old metric of "when sales slow down, it's probably time to update."** But this is an astoundingly bad metric. When sales slow down, it's probably too late. When

sales slow down, you're under pressure to come out with something, fast. And when you're under pressure, you may miss a critical feature—or not do your best work.

- **Some use a fixed schedule:** "We have decided our product life cycle is 3 years, so we will begin working on significant updates 24 months after launch." This is better, but what happens when the market undergoes rapid changes and your sales fall off a cliff at 18 months in? Do you simply sit and wait the remaining 18 months to launch a new product?

When we started Schiit in 2010, if you'd asked me what I thought our product life cycle would be, I probably would have shook my head and said, "I don't know. Two years, three years? Let's see how it goes."

And, to this day, I can't really put a number on it. We've decided to set our product life cycle on a more flexible metric than falling sales or dates on a calendar. If I had to put it in words, it would be something like this:

Our products are updated or discontinued when significant positive changes can be made, or need to be made, at a time not disruptive to customers.

Note the specificity: updated *or discontinued* (it's okay not to keep a zombie product alive), significant *positive* changes (not just a small tweak, and not de-contenting), at a time *not disruptive* to customers (updating a product 3 months after release, for example.)

In the case of Asgard, early in 2013, those significant positive changes *needed* to be made.

Why? One word: Magni.

The Death and Rebirth of Asgard

Before Magni, we never really thought about updating Asgard. It was a great amp, a strong seller, and sales continued to pick up. By the "wait until it slows down" metric, Asgard was doing fine.

But as soon as we heard Magni, we all looked at each other, and said, almost in unison: "What about Asgard?"

Magni was just too close in performance to Asgard. In fact, it was more powerful than Asgard. It was quieter. It ran cooler. About the only thing Asgard had going for it, objectively, was a much better volume pot—better tracking, better taper—you can't beat large pots for that, no way, no

how. Aesthetically, it was a much more elegant-looking piece, but elegance only gets you so much. Sonically, we believed it was still ahead of Magni—but not by enough.

So it was time to look at the end of Asgard.

But did we kill it, or did we update it? That was the first question.

Deciding to update it was really easy:

- Asgard was still selling well after Magni—in fact, sales continued to increase until it was discontinues—so there was clearly still a demand for a step-up amplifier.
- Magni had taught us a lot about surface-mount parts. Surface-mount parts, applied to Asgard, would actually reduce its manufacturing cost due to robotic assembly. This would open up putting the money into other areas.
- We regularly got requests for two features Asgard didn't have: gain switching and preamp outputs. Adding these made a lot of sense.
- Assembling Asgard was a real pain, with four separate MOSFETs, plus thermal pads, screws, lockwashers, and electrically insulating bushings for each—I'd been thinking about a new way to do it with one simple bracket, two screws, and a single custom thermal pad. We

were now of a size where we could do the large custom orders that made the switch viable.

- Asgard's gain stage was pretty basic—it didn't even have a current-sourced front end, and it had coupling capacitors at the output. Mjolnir had taught us a lot about how good a high-voltage, cascoded, current-sourced front end could be—and going to surface mount meant we could afford to do it—together with a trick DC servo that connected to a sonically innocuous (as in, non-amplified) terminal to eliminate the coupling caps.
- And while we were at it, we put the rest of the money saved by surface-mount to change the transformer and power supply from unipolar to bipolar on the output, plus added an 80 V rail for the HV front-end like Mjolnir.

That's a big list of updates, guys. Different board, different topology, different power supply, different chassis, different transformer, different feature set—the only thing that our proposed Asgard 2 had in common with the first generation was the same chassis styling, the same pot, and the same basic connections.

So how did development go on this radically updated Asgard?

Almost comically boring. I built a perfboard proto one evening and verified that the basic concept worked—including delivering measured distortion performance that was almost 10× better than the original Asgard. Total time: a couple of hours.

From there, the first boards we got fired right up and worked. They even fit the revised chassis just fine. All the pain of your typical product development—all the tiny little nits and problems—were conspicuously absent. I changed a couple of resistor values from the initial calculations, but that's about it. Total time: maybe 4 to 5 hours.

From there, we took the first in-chassis prototype inside to the listening area where I had Mjolnir and Gungnir set up. I plugged it in, put on some headphones—most likely HD 800s, because they're great at revealing what's wrong with an amp—and sat back.

Crap, I'm still running the Mjolnir, I thought, as soon as I heard it.

But I looked over … and the headphones were plugged into Asgard 2.

Asgard 2—sounding like a Mjolnir? *No way.* Not believing it, I switched the gain to Low, and the

volume decreased, as expected.

Holy moly, I *was* listening to the Asgard 2.

Still not really believing it, I swapped back to Mjolnir—and, yeah, Mjolnir was a step up, but it wasn't leaps and bounds like the original Asgard. Of course, Asgard 2 didn't have the raw power for, like, HE6s, but it was *very* good—much, much better than the original.

Next, IEMs. Dead silent on low gain. I grinned. This was gonna be a winner. I decided to call in Rina and demo it for her.

"I want to hear the Asgard 2, not the Mjolnir," she said.

"You are."

She went through the same rigamarole as I had—switching gain, unplugging headphones, comparing to Mjolnir—before believing it.

"This ain't no Asgard," she finally pronounced.

Mike? Pretty much the same reaction. We had a winner.

And, you know, sometimes things just work out. And sometimes, things work even better than you expected. In product development, this is known

as a “gimme.” Also known as, “oh crap, watch your back.” I should have watched my back.

The Asgard 2 Launch Debacle

The runup to Asgard 2 launching was filled with the same little delays that happen with every new product—waiting for parts, waiting on chassis, etc—so, in that respect, there was no clue that we had any nasty surprises waiting in compensation.

But when we launched the product, word quickly came back—some of the Asgard 2s hummed like a refrigerator. Mechanically. As in, you could hear the transformer humming with headphones on. Closed headphones.

But that made no sense—none at all. The prototype hadn’t hummed. And we hadn’t heard any hum in production.

“But we wouldn’t necessarily hear it,” Alex said, as a big train went by outside, shaking the paper-thin stucco of the Schiithole.

“Crap,” I said, realizing for the first time just how loud it was in our building. It wasn’t just the trains—it was the constant traffic noise from cars passing on the 4-lane road outside.

Late at night, we confirmed it. Many of the transformers did hum. And the prototype didn't.

Why?

This is known as a "production surprise." As in, "Surprise … although the transformer meets specs, we decided to make them a little differently … and that difference transformed your product from a headphone amp into a combo headphone amp/massager."

The transformer manufacturer was apologetic, and promised new samples posthaste. But that didn't fix the humming Asgard 2s in the field—now the entire first run. We'd been so deep in backorder, we'd sold out in just a few days.

So what did we do? The only thing we could do: accept the returns on the Asgard 2s that hummed, or swap them out as we got the new transformers in.

Aside: gimmes are dangerous, guys … be suspicious, be very suspicious, of something that is too easy.

And that's why we started a new policy after Asgard 2: multiple prototypes, multiple listeners …

and multiple locations, some of which we knew were quiet.

Yeah. There you go. But that's also why the launch of Valhalla 2 and Lyr 2 were, well, relatively uneventful. Almost boring.

And, in terms of "production surprises," that's exactly
where you want to be.

Chapter 31

R&D Sometimes Means, “Try It, See If It Works”

Sometimes R&D is extremely focused. You know exactly what you’re shooting for, and you apply the collective smarts you have in a concerted effort to hit—or exceed—the mark.

This was certainly the case with the previous two products I talked about (Ragnarok and Asgard 2.) We knew (pretty much) what we wanted, and set about to do achieve it. In the case of Ragnarok, it was an all-in deal with me, Mike, and Dave all contributing—and a long and winding road to the end game. In the case of Asgard 2, it was just me—and, as I said before, R&D-wise, the product was a gimme.

But I strongly believe that R&D shouldn’t always be so focused. There’s value in making sure your engineering staff has time to play with crazy ideas.

How much value? Consider this:

Without play-time, Mike would have never put together the micro-DAC that became the first Modi. Which also would have meant:

- There'd be no reason to develop Magni, ever
- We may never have investigated low-cost products at all
- We may never have discovered how cost-efficient all-steel chassis could be
- Which would mean that we also wouldn't have Vali, Sys, Loki, or Wyrd
- Our lowest-cost product today might still be Asgard (because why bother updating Asgard in the face of rising sales?)
- Valhalla 2 and Lyr 2 might not exist, either, for the same reason

So what would Schiit look like today, if Mike hadn't had the wild butthair to develop a DAC orders of magnitude less expensive than any he'd ever done?

Well, we'd certainly be a lot smaller. Magni, Modi, and the sub-$150 products are the majority of our sales.

But at the same time, ironically, we probably wouldn't be much farther along, if any, on Ragnarok and Yggdrasil—the low-cost product line

barely impacted Dave's software development time at all.

So, it would easily be possible for us to be, say, half the size. Maybe still at the Schiithole in Newhall. And still not have Ragnarok and Yggdrasil out.

Give your engineers some playtime. It pays off.

Those Tempting Tubes

Sometimes it's funny what sets off those "I wanna play" moments. In the case of what would become Vali, it was eBay.

Yes, eBay.

In this case, it was my fault. I keep an eye open for bulk tube deals. We kinda have to. Valhalla 2 and Lyr 2 production chew through an incredible amount of tubes. And good tubes are getting scarcer.

Why? Simply because they aren't making any more 6N1P or 6BZ7 tubes, just to name a couple of NOS tubes we use. And, while there are some good new-production tubes, they tend to be eye-wateringly expensive. So, we prefer to work with NOS tubes, at least as long as we can.

Note: I'm not worried about Valhalla 2 and Lyr 2 tubes, at least for a couple of years. We have some good suppliers we're working with, and we have yet to be held up in production for lack of tubes.

Anyway, my search for bulk tubes sometimes takes me to tube resellers in Russia, sometimes to surplus warehouses where piles of tubes are forgotten for decades, and sometimes, yes, to eBay.

And what came up—what started the whole Vali thing—was an incredible bulk of NOS Jan Raytheon 6088 tubes, at a very attractive price. I mean, so many tubes that it would take us through half a decade of production, even if the amp sold like Magni. Truly crazy numbers.

And those numbers—and that attractive price—got the wheels turning. What could we do with this? Could we make an amp with them?

An amp, maybe, at a near-Magni price?

There was only a whole buttload of problems with this scenario:

1. 6088s were pentodes. Using pentodes in pentode mode for gain is pretty barfy. And you don't know what the triode-strapped curves

will look like until you run them on a curve tracer. Would they be linear enough to use?
2. We'd never used subminiature tubes, so we had no idea how best to use them, nor any specifics of their care and feeding.
3. Even if the thing worked, how would we run the high-voltage supply (because, as Mike says, running a tube from a low-voltage supply is for "children and amateurs.")
4. And, what about the output stage? What exactly could we bolt this tube to?

And it's doubts like that which kept me from simply clicking "buy it now." Because if I just jumped on it, we might end up with a whole bunch of useless tubes.

But I kept coming back to *all those tubes*, at *that price.*

After a couple of weeks, nobody else had jumped on them, probably because they had the same doubts. I did some research—had anyone used the 6088 for audio? A couple of DIY projects popped up, but they were the most basic and simplistic things imaginable—nothing that would be able to drive a wide range of headphones, nothing that could be sold as a commercial product.

But still, those tubes …

Screw it. I contacted the seller and purchased 100 tubes just to play around with. In the process, they confirmed that they had 5× the amount of tubes they had listed on eBay actually available—a truly eye-popping number.

Which meant if I could do something with them, then we could have a real winner on our hands.

The Road to Vali

The tubes came in a couple of days, and I set up a quick breadboard circuit to see how they performed. I was just interested in the basics:

- What plate voltage did they run best at?
- What plate load did they like for lowest distortion?
- What was the distortion like in triode mode?
- What did the operating current and heater current look like?

Why these basics? Because, based on these measurements, I could make a go/no-go decision on purchasing the tubes. Or so I thought.

It was interesting, running the early tests on those tubes. They were different than any others we had ever worked with. How so?

- Maximum plate voltage was only 60 V—most tubes are in the hundreds of volts
- Plate current was 750 µA to 1.5 mA for typical loads—much lower than most tubes
- The cathode was directly heated, rather than indirectly, like most tubes we've worked with
- It only needed 1.25 V heater voltage and 20 mA heater current—again, far lower than the 6 V and hundreds or thousands of A for most tubes (wonder why tubes run hot? Simple—they have a heater in them that usually dissipates a few watts of power.)

So, how'd they do?

Not so hot, at first. Nearly 1.5 % distortion at the first plate voltage and load I tried. But, by tweaking the plate voltage and load, I was able to chart where the tube was the happiest. THD was still high, by our current standards—0.3 % or so—but it was mainly second-harmonic distortion, and the distortion profile was nice, with 3rd 20 dB down and 4th almost at the noise floor.

But—0.3 % was still pretty high. How would it do, with an output stage bolted to it? What would it actually sound like?

That was beyond what the breadboard would tell me. I needed to lay out a board, and see how the

amp really would do.

Laying it out on a Magni-sized board would be easy … except for the fact that it was a totally different amplifier, with radically different voltage requirements. (Remember I mentioned those high voltages for the tube? 60 V is pretty low in the tube world, but it's still a far cry from the ±15 V we were running in Magni.

At the same time, I wouldn't want to run a solid-state output stage at 60 V. That's pretty, ah, adventurous, especially since the standard TRS headphone jack shorts the output every time it's connected or disconnected.

And, we needed a regulated heater voltage, too, at 1.25 V.

Oh yeah, and 5 V for the relay.

And it would be ideal if we could get all those voltages from Magni's standard 16 V AC wall-wart.

Sounds impossible, right? Actually, far from it. With AC input, you can run a voltage-quadrupler and easily get 60 V after regulation for the tube. Half of that circuit gives you DC voltages that can be regulated to 30 V. And the 1.25 V and 5 V requirements are low enough that you can bring them down from the standard rectified output.

But (you knew this was coming, right?) …

But there's always a but. Voltage quaduplers also aren't very good for high-current output, and have significant ripple. But the half-wave rectification we'd used in Magni (effectively a voltage doubler) worked well enough. For a couple of A going to tubes, run through a voltage regulator, it should be fine. Or at least that's what I told myself. I laid out the board in a couple of evenings. Everything fit really easily, including the tubes and an output stage that kinda started as something out of a Magni, but morphed into a pretty cool design that used a phase flipper to level-shift the output of the tube for DC coupling from the front-end, plus LED biasing for the output devices, which ran in Class AB.

Then it was just … send out the boards, and wait a week.

Two Big Surprises

The first Vali boards that came back weren't perfect, but they weren't bad. They needed a couple of hacks to smooth the input to the 30 V regulator, and additional bypassing for the 1.25 V regulator, but that was about it. After a couple of small tweaks, they were up and running, with the LED

biasing simulating the glow of a traditional tube heater.

> Subminiature tubes with 1.25 V heaters don't glow. They hardly run hotter than room temperature, in fact.

What was even better was the fact that it was running on the standard Magni wall-wart, without any signs of strain. Vali does draw a bit more current than Magni, so this was a welcome sign.

> Thermal design really is the starting point for any amp. Get that wrong, and you're in a world of hurt.

So, what did it sound like? I took the prototype from the garage (where I still did most of the design and tweaking) back to the listening couch, where I kept Mjolnir and Gungnir.

Okay, I'll admit it—I used Gungnir for those early first listens. Overkill, yes …

On first listen, I was a little surprised. Vali didn't sound like I expected it to. I thought it'd be more tube-y, with the more typical euphonic colorations of an inexpensive tube design (rolled off highs, syrupy midrange, tubby bass—that kind of thing.) I had every reason to expect it to sound this way.

The distortion profile suggested it. The single-supply output stage suggested it. The fact we were using (horrors!) coupling capacitors for the output suggested it.

But it didn't.

In fact, it sounded pretty neutral and transparent. Maybe even a little bright. And, it sounded pretty darn good. I sat there for a while and just listened, which is usually a very good sign.

But was I hearing reality? Or was I just full of it? That's always the designer's dilemma—being too close to something, and losing perspective. I gave it to Rina to have a listen. Her eyes widened. "Wow," she said. "When do I get one?" I also gave it to Mike. He listened for, like, 5 seconds in the shop, then picked up the prototype and put it in his bag.

"Hey!" I cried. "That's the only one!"

"So build another," Mike said.

"I will, but … I was listening to it!"

"Don't be lazy," Mike said, and left the building with it. A few days later, he called me. "How many of these tubes are there?"

A metric buttload," I told him.

“Get them all. This is good. Really good.”

That’s what I’d thought, but it is good to have some confirmation.
I contacted the eBay seller and cleaned up on the tubes.

Now, the only problem would be telling Alex he had to find more space—again—for the pallets of tubes that would be coming in. He was a lot more happy about it when I gave him a prototype Vali to play with.

And the early accolades kept coming. At a big head-fi meet, several of our golden-eared friends (including some who have given us, well, brutally honest feedback) proclaimed that it was better than 95 % of the tube amps there. And people really flipped when we showed it at Can-Jam.

So, everyone’s happy, right?

The Catch

Nope. Of course not. There’s always a catch. And in Vali’s case, the catch is directly related to those great-sounding tubes: tube microphonics.

What are tube microphonics? They’re noise that’s generated from tapping or jarring a tube. Some

tubes are pretty non-microphonic (especially the 6N1P we use in Valhalla 2), and some are very microphonic (like the tubes we use in Vali.)

In the case of Vali, microphonics sound like a little "ting" sound that takes a long time to decay. It's like a delicate silver bell. It's actually a neat sound …

… that is, if it isn't interfering with your music.

And these tubes were microphonic enough to ring when you first turned the amp on, and when you plugged or unplugged headphones. We warned everyone about this, of course, but it wasn't enough. It turned out that some amps were microphonic enough to be set off by typing on a keyboard, or simply rang all the time.

So what did we do? We started doing an extended burn-in on the Valis, and checking them when they were still warm, to weed out the self-ringers and over-ringers. We also had input from a very helpful customer, who came up with one idea we hadn't thought of for reducing microphonics (specifically, damping the PC board itself, as well as using sorbothane pads under the tubes.) Those two changes have brought down Vali failures to fractions of a percent.

So What Did We Learn?

Take time to play. Even if that play comes from seeing a pile of tubes, and wondering, “What can we do with these?”

Chapter 31

Name Me One Non-Standard Format That's Succeeded, Ever, Or, A Trickster Cometh

The DSD surge started in 2013, shortly after the announcement of the DOP (DSD over PCM) USB protocol.

It started pretty innocuously. Starting in early 2013, we started receiving a few emails asking if we were planning to add DSD decoding to our products. It was a literal handful to start, but as we got into spring, the inquiries started coming faster, as a number of companies introduced DSD-compatible DACs.

As the company's marketer, I wondered if this surge in inquiries would become a movement, so I asked Mike about the possibility of adding DSD to the existing Bifrost and Gungnir.

Mike groaned. "DSD. Argh. No." I waited for him to explain, but he didn't go any further.

"Why not?" I asked. I actually knew some of the technical reasons, but I wanted to hear it from Mike.

"DSD requires completely different filtering," Mike said. "It's essentially wideband noise. You want that going to your amps?" I shook my head. Running ultrasonic noise into an amp is a good way to test it to destruction.

"So, we need way more aggressive filters to get the noise out," Mike said.

"But, technically, the AKM DACs do DSD, right?"

Mike shook his head. "Technically, yeah. But *doing DSD*, versus *doing it right*, are two different things. If we put in the DSD-appropriate filtering, we'd be compromising our analog stage performance for PCM. And it's not as simple as switching it in and out, because that would require more space on the analog board, and I don't even know if we have the hooks on the analog board input, anyway."

"So no DSD," I said.

“Not without a lot of changes. For Bifrost, we’d need a new USB input board, a new main board, and a new analog board. Technically, yeah, that’s just upgrades—”

“—but it’s essentially a whole new product.” I finished for him.

Mike nodded.

“So what if DSD gets big enough to matter?”

Mike laughed and waved a hand. “Remember HDCD?” I nodded. HDCD was a technology of the early 90s that was supposed to be the One True Savior of digital, allowing more dynamic range to be encoded on special disks that could only be decoded by a specific digital filter.

“HDCD almost took down Theta,” Mike said. “We got in screaming arguments about it. My marketing guy said the same thing you did: ‘What if it gets big? Everyone else is doing it. We’re going to lose sales if we don’t have it.’”

“I didn’t say those last two things,” I told Mike.

“Yeah, but just asking about DSD implied it,” Mike said. “You’re worried that we’ll lose sales, or we’ll miss out.” I shook my head. Though Mike was right, in a sense. If DSD became big, we’d be

vulnerable to other products that offered DSD playback.

"Stop worrying," Mike said. "Where did reel to reel go? Nowhere. Where did quadraphonic go? Nowhere. Where did Elcassette go? Nowhere. Where did DAT go? Nowhere. Where did minidisk go? Nowhere. Where did HDCD go? Nowhere. Where did SACD go? Pretty much nowhere. I expect DSD will pretty much go the same exact place."

"But what if it doesn't?"

Mike groaned. "These special formats all end up the same place, because there's no software for them. When there are more DSD downloads available than SACDs, let me know. Then I'll start worrying."

But the Inquiries Kept Coming

In fact, they intensified. As the press flogged the new shiny thing known as DSD, we began to get several inquiries a day—on slow days.

"Mike, we should do something about DSD," I told him, finally.

“Ignore it,” he said. “It’ll go away. It’s just the press. They’re so monumentally bored, they’ll talk about anything, including a non-starter like DSD.

“But what about, just, you know, as a CYA.”

Mike sighed, and was silent for a long time. He knew we were getting inquiries. He knew some people really wanted DSD. And here he was, between his partner’s paranoia and his experience with dozens of nonstandard formats that have come and gone.

“You want me and Dave to divert time from Yggdrasil to work on this?” I crossed my arms. That was the ultimate threat—taking time away from a product that was literally the antithesis of DSD, and which we believed would help redefine the digital market in toto, to work on something that could be a passing fad.

Mike laughed. “You do.”

“I just think it would be safer—”

“To do what everyone else is doing,” Mike finished. “To jump off the cliff, just because everyone else is doing it.”

“Can we really not afford to take a look at it?”

Mike looked thoughtful. “Okay. Fine. I’ll think about it. That’s all I’ll say right now.”

Mike’s Thoughts

Time went on. DSD inquiries continued. I watched our sales cautiously, but they kept increasing for all the DACs—definitely not an indication that DSD was a gotta-have thing.

But the press kept flogging it, the articles kept coming out, and rumblings of lower-cost DSD DACs started to surface (prior to this, DSD DACs were pretty eye-wateringly expensive.

Eventually, Mike came back to me, grinning like a fool.

“Okay. Here’s what we do. We make the least-expensive DSD DAC on the market.” I blinked. “What?”

“If they want this format to succeed, they need wide adoption. And you ain’t gonna get wide adoption for a grand and a half.”

“So it replaces Modi?”

Mike shook his head. “No. It’s a standalone DSD-only DAC. That’s why, even though it’s gonna be the cheapest DSD DAC out there, it’s still going to

sound insanely good. We'll do the filtering right. Just for DSD, and only for DSD."

"But what if you want to play both PCM and DSD?"

"If you're so into DSD, convert it on the fly," Mike sniped.

"Seriously."

Mike looked thoughtful. "Put a switch on it. Then you can run the output of your current PCM DAC through it." I sat straight up. "So you can use it to add DSD capability to any DAC!" I cried.

Mike nodded, looking very pleased with himself. "Exactly." I nodded. That was a perfect fit with Schiit's ethos. *Keep your existing DAC, add DSD, see if you like it, then go from there if you do.* Instead of throwing your existing DAC away to get a DSD-compatible one.

"How cheap are we talking?" I asked Mike.

Mike grinned. "Not much more than Modi."

Okay. Now this was getting good.

Tech Challenges

There were just a few problems with this plan—starting with the fact that we didn't have any DSD-capable USB receivers. The CM6631A we were using didn't accept DSD streaming or DSD over PCM. C-Media was planning a CM6632 for later in the year, which would be DSD-compatible, but late in the year was too late for our plans.

Enter Dave.

> **Note:** when I say, "It's in Dave's hands now," that means it's somewhere in complicated software/firmware land, from which it will hopefully emerge with working software/firmware at some future time.

Dave's plan was simple, but somewhat insane: use a 32 bit Microchip microcontroller to do our own unpacking of the DSD-Over-PCM standard, and then send that along to the DAC. Yeah, quite a programming feat. But he did it, and soon we had a prototype that could play native DSD, using a Crystal Semiconductor DAC.

There was only one problem: it sounded like crap. Dynamically compressed, soft, boring, and lifeless. Yeah, I know, it measured fine, so it should sound fine, right? Not in this case.

“Why’d you do Crystal?” I asked Mike, one day when Mike, Dave, and I were together at the Schiithole.

“Crystal will do 2× DSD,” Mike said.

“But the Microchip controller would have to be faster to unpack it,” Dave said.

“So we can’t do 2× DSD right now?”

Dave nodded. “But with a faster processor, we could.”

“But not now,” I confirmed.

“Right.” I frowned. “Why don’t we just use AKM, then? We know they sound good.”

Mike shook his head. “We have no idea what they sound like when they’re fed DSD.”

“Isn’t it worth a shot?”

Mike and Dave looked at each other. Dave shrugged. Mike sighed.

And a few weeks later, we had another prototype—this one with an AKM DAC. And it sounded worlds better. It still measured pretty much the same, but it had a lot more life and energy. It had dynamics and pace. Both Mike and Dave smiled when they heard it.

“But AKM doesn’t do 2×?” I asked.

Dave shrugged. “It might do it undocumented, but—”

“—but we don’t know,” Mike finished for him. I sat silent. Should we wait to add 2× capability, for the literally 20 to 30 recordings there were out there done in 2×? It would mean another prototype cycle, and maybe different code, and maybe some unforeseen problems.

“And it would take different filtering,” Mike said. “This is as good as it gets for 1× DSD. Throw 2× in there and we start having to make some different decisions.”

Still, I sighed.

“Let me propose a solution,” Mike said, as he usually does when I’m hesitant about something. “Let’s bring this to market, see how it does, and if DSD really takes off, we can work on a 2× solution, or whatever we need.” I nodded. That made sense.

“Good.” Mike said. “Though I doubt if we’ll ever have to do any more work … ”

Peak DSD

Mike seemed confident that DSD was a non-starter, but his comment seemed to be out of sync with the public at the next TheShow Newport, which we attended a couple of months before introducing Loki. At TheShow Newport 2013, literally every other question from passerby was, "When will you support DSD?"

Mike still didn't look too worried.

And, although I didn't know it at the time, that was the absolute peak of DSD.

Loki Cometh

We introduced Loki under the banner of "Add DSD to any DAC for $149." At the time, the least-expensive DSD-capable DAC was $849, so this was quite a coup.

Or so we thought. It turned out that the idea of a DSD-only DAC and switching system was a little more challenging than we thought. Some people thought we were converting DSD to PCM and running it to the main PCM DAC (why, when you can simply do it in software?). Some people thought we were taking the analog output of their DAC and converting it to DSD (fat chance on that

one.) Some people really, genuinely wanted to throw out their old DAC, rather than run DSD through it.

And lots and lots of people didn't like having to run two USB cables (one to their main DAC, and one to Loki), and switch between the two on their playback software. For a main DAC fed by SPDIF, the switchover was easier (and seamless if they were using a different player, like a CD player, for their PCM content), but it still wasn't something that most people wanted to do.

That, combined with the appearance of new, inexpensive DSD/PCM DACs, quickly cooled Loki's sales. Mike will still argue that doing PCM and DSD in the same DAC is a compromise, and the math (and measurements) are on his side, but convenience usually wins out over sonics when you're playing at the lower end of the market.

Still, these wouldn't be insurmountable problems if we wanted to do, say, a Loki 2 with automatic interface switching. It could then interface seamlessly with a PCM DAC. But it would be significantly more expensive, especially if we added DSD 2× (or 4×, or 1000×, or whatever the latest unicorn format is today.)

And I suspect that's the way we'd end up going if we were to continue pursuing DSD—not adding it to our current DACs, but making a seamless, dedicated DSD DAC with interface switching.

But as of the time of this writing, I don't think it'll happen.

DSD Today

Today, Mike has crossed his arms and declared, "No more DSD development, unless something really big happens."

Why?

Because, from our point of view, it looks like we're past the peak. Despite dire pronouncements from other manufacturers saying, "You can't move a non-DSD DAC with a boxcar full of Ex-Lax," we haven't seen it. Cases in point:

- Sales of our DACs continue to increase—and to accelerate
- Inquiries at TheShow 2014 included literally two (that is, 2) half-hearted questions about whether or not we were going to support DSD, in stark contrast to the literally 200 questions the year before

- Sony's presence at TheShow (featuring DSD prominently) was a ghost town
- Email inquiries have fallen from a dozen a day to maybe one or two per week
- The predicted "opening of the vaults of DSD" hasn't happened—there are still only a few hundred recordings available, many of which have questionable provenance (more on that later—and even SuperHires's announcement about Warner probably won't answer even a tiny fraction of the questions on provenance, and prices remain TBD)

Sure, there are plenty of DSD-capable DACs out there, including some that do 4× and 8× DSD … but where's the software?

Let's face it:

- The bulk of the industry remains compressed streaming—and that ain't DSD
- Pono is another question mark in high-res—and it ain't DSD
- Apple could come out of sleep-mode on high-res if Pono or other products prove there's a demand—and if they do, they define the market—and they won't be using a Sony format, I bet

And, the elephant in the room:

- The most important part of a recording is the master—paying attention on that side will reap benefits beyond any format

So What About the Future

Okay. Let's say the next Sony reorganization (they ain't exactly healthy these days) doesn't kill DSD, but results in them releasing 20 000 DSD recordings of popular artists, all with DSD-guaranteed-from-the-start provenance, for, say, $5.99 an album.

Would this result in a whole lot of DSD out there? You bet.

Would it be a game-changer? Absolutely.

Would it have us dusting off plans for a Loki 2, or working on ways to include DSD decoding in our DACs without compromise? Yeppers.

But I think that scenario is about as likely as the disembodied head of Steve Jobs giving the next Apple Keynote.

What's more likely is this:

- DSD recordings will continue to be a small part of the market

- DSD recordings that actually start as DSD, or were converted direct to DSD from master tapes, will be an even tinier part of the market
- High-rate DSD with the same provenance will be even smaller
- DSD recordings will continue to be very expensive
- Some people will continue to really like DSD, and will flip us off as they pass us at shows
- More people won't care, as long as the music sounds good and doesn't cost a fortune
- Even more people won't care if it's DSD, PCM, or compressed, as long as it's available to download at a good price
- And the vast majority of people will never have any idea what the hell us crazy audiophiles are talking about, as they happily stream compressed music for a small monthly fee (or free)
- About 10 years from now, a new quantum-based encoding format will come out so that everyone can buy their music again …

How about we deal with the elephant in the room, before worrying about formats, hmm?

Chapter 32

No Sample Left Unchanged: Digital Today

Okay. Let's follow digital with more digital.

That's cool, though Mike is also doing his own "story of the Yggdrasil," here now (search for Baldr's posts.) This chapter was originally going to cover Yggy, but I'll let Mike do that now in detail.

Instead, let's talk business. As in, business cases, business philosophies … and, yeah, Digital Today (and yesterday).

However, let's start with some fun facts about Yggdrasil:

- **Yggdrasil is really the first DAC Mike wanted to make,** when we started talking DACs in 2010. Yes, the concept is that old. Considering that it will now ship on the eve of 2015, Yggy has been in the oven a long time.
- **One of the reasons it has taken so long to develop Yggy is simply that we didn't have**

the infrastructure to build a multibit audio DAC today, as we did in the past. In 1993, you could pick from a dozen multibit chips, many of them very good (PCM63, etc), interface it with a standard digital filter, and be off and running. Today, the only real multibit chips are either (a) Almost dead (PCM1704), (b) only 16 bit, or (c) not intended for audio applications, difficult to interface with audio data, and eye-bleedingly expensive. And yes, there are also a few discrete resistor ladder DACs out there, but multiply "eye-bleedingly expensive" by orders of magnitude, if you're talking about something that will be reasonably stable across a range of external temperatures (think oven-controlled temp for the module, etc.).

- **The reasons we came out with less expensive DACs before Yggy are simple:** First, we needed inexpensive DACs to mate with our headphone amp line (and to be used in affordable high-end speaker systems, of course.) Second, we were satisfied with the price/performance ratio of our inexpensive DACs. Considering that Bifrost soundly trounces the Cobalt 307, Theta's first inexpensive DAC, in both audio quality and features, for less than 1/2 the cost in constant dollars, we think we've

achieved our objective for these DACs.

- **Yggy may seem expensive at a projected price of $2299,** but if you compare its price to the Theta DS Pro Gen V in 1993 dollars, it looks like a screaming bargain. The DS Pro Gen V balanced in 1993 was $4500. Yggy, in 1993 dollars, is $1399. Sorry, we don't have any of those 1993 dollars available.
- **To clear up any misconceptions, Yggy will be a real multibit DAC,** running the only closed-form digital filter, with no asynchronous sample rate conversion or other intrusive/destructive technologies to maintain the original bits from the input all the way to the DACs. At its essence, it is similar to the Theta DACs from days past, but with much, much greater DSP horsepower, as well as some really tweaky analog tricks like choke-input, shunt-regulated analog power supplies that even Theta didn't use. Mike will reveal the actual digital to analog conversion specifics when he is good and ready; beyond that, he has sworn me to secrecy.
- **If you are at RMAF, you will be able to hear a very-near-production Yggy in its full chassis,** with full feature sets intact. Other than a few PCB changes and some firmware tweaks, what you hear is what we'll shortly

be shipping.

Now, on to the business of things.

Business Cases, Standards, Licensing, and Assorted Fun

Here's where we go back to being a business book. Because Yggdrasil is an interesting case that illuminates some of the problems that manufacturers have to deal with, in the arena of standards. With standards, you have a choice of "going along and getting along," or "forging your own way."

"Standards, what the hell are you talking about?" you might ask.

Okay, let's take an extreme case: surround sound processors, AV preamps, and AV receivers. Today, these types of gear are on the bleeding edge of standards compliance. A typical AV preamp today must:

- Decode all the Dolby surround standards
- Decode all the DTS surround standards
- Decode any other bizarre surround standards that are en vogue at the time
- Be compliant with the current, and changing, HDMI standard

- Probably also accept SPDIF and USB standard audio inputs
- Maybe accept Bluetooth audio standard
- Maybe accept Apple Airplay standard

"Well, who cares," you might scoff. "Meet the standards, and you're golden."

Yes. Except that the ongoing turf war between DTS and Dolby means that there's a bazillion surround standards out there. I've stopped counting, but I have been told that the standard test disk for surround modes is now up to 1200+ tracks. You run that disk, connect your system to an 8-channel Audio Precision, and cross your fingers. The surround standards guys will tell you if you failed. Then it's back to the drawing board if you did.

Plus, new standards pop up all the time. The new one is Dolby Atmos now. Oh, it's not Atmos compatible? Say bye-bye to sales. You don't control when the new standards appear, so you're always playing catch-up.

Same thing with HDMI. HDMI 2.0 is here. Kinda. Sorta. Well, not really. Because you can be compliant with one part of HDMI 2.0, but not another. And, by the way, 2.1 is coming. Will it

work with your current system? Maybe, maybe not.

Same with Apple Airplay. First step to become an Apple hardware licensee is, literally, "Have your lawyers contact our lawyers." Being of the opinion that lawyers are kinda like raw plutonium—very useful in some specific applications, but not something you want to get near very often—you can see how we feel about this.

Same with Bluetooth. We're on Bluetooth 4.0, and it still can't do uncompressed audio. Want to bet on 5.0, and when it comes out?

Fun fact: the HDMI consortium has meetups called "Plugfests" so that manufacturers can see what they're compatible with, or not. Yes, even they don't fully know. It's up to you as a manufacturer to figure it out. You can't make this crap up.

And, to make all of the above even more fun, guess what? You get to pay some exceedingly non-trivial licensing fees for the privilege of putting those standards' logos on your box.

Bottom line, if you're going to be standards-compliant, you're always going to be at the whim

of the standards-setters. You're not fully in control of your own destiny.

So why do they keep changing these surround-sound standards? Three reasons:

1. To improve performance (higher bit rates and sample depths, more/optimized speaker placements, new algorithms, etc.)
2. To increase revenue from the licensees' additional licensing costs.
3. To drive a continuous upgrade cycle—buying newer gear to unlock new capabilities, upgrading cables for the latest HDMI standards, and re-buying content mastered to these standards.

If you were cynical, you could say that 2 and 3 combine to form a perfect "devil's bargain," where, if you keep spending on licensing, they'll continue changing the standards to keep your market coming back for the latest and greatest. Of course, that's a very cynical viewpoint. And, it only works for so long. When a true home theater enthusiast doesn't know what the latest Dolby Super HD Wowiematic Ultra Extra Fine And The Kitchen Sink standard is (and doesn't care), the whole thing comes crashing down.

So, why do I bring up these "standards?"

One, as a thought experiment. Imagine a surround processor that didn't have any Dolby Digital or DTS logos on it, running its own non-standards-compliant decoding algorithms … and sounded good with most surround-encoded materials. Is this something that would sell? It's an unknown, because nobody has done it yet, probably due to the threat of possible infringement suits from Dolby and DTS (and maybe HDMI, if you don't pay their license fees as well.)

Before you start the heavy breathing, we are not working on this. It is simply a thought experiment. Would potential buyers be OK with a product that didn't have 73 logos on the front of it, and the comfort of the Dolby and DTS stamp of approval? I don't know. But it's fun to think about.

The second reason is more pointed. In audio, we have few standards, and virtually none of them are licensed. USB Audio Class 1 and 2 are standards, SPDIF is a standard, PCM is a standard, and DSD is a standard. Okay, you can also throw in fringe stuff like I^2C over HDMI, as long as you're OK with paying the HDMI licensing fee … but then you're getting into licensing … and technically, Plugfests. Shudder.

So, in audio, why would you change standards?

#2 from the surround example is out (to increase licensee revenue), but seeking to improve performance, and to drive upgrade cycles, are both relevant. We don't get many new standards in audio, so the excitement around something new and shiny is much higher than it is in surround.

So, let's do another thought experiment.

Let's say DSD wins as the next audio standard—it's recognized as a significant upgrade from PCM, and it is embraced by enough users that every manufacturer has to support it, and support it well. What happens?

- Recording studios have to convert to DSD workflows. Manufacturers of pro gear celebrate.
- Every listener must go out and buy a new DAC. Manufacturers of consumer gear rejoice.
- Every listener must go out and buy new DSD recordings. The record industry throws a huge party.

On the other hand, let's say DSD fades away, and PCM continues as the reigning audio standard. What happens?

- Recording studios continue to do what they do. No impact on manufacturers.

- Listeners don't have to buy new DACs. No impact on manufacturers.
- Listeners don't have to buy any new recordings. Big sad face for the recording industry.

Something to think about, hmm?

Digital Yesterday: Steady Progression

When digital audio was new, you could pretty much chart the steady, linear progression of the technology for about a decade. From the first 14 bit multiplexed non-oversampling DACs in CD players in 1982, to the fully realized, 8× oversampling, 20 bit ladder DACs in the top DACs of the early 90s, there was clear and steady progress:

- 14 bit multiplexed D/A converters in CD players, no oversampling, brickwall filtering
- The first 16 bit converters, still with no oversampling and brickwalled
- 16 bit converters with 4× oversampling, to eliminate the brickwall filter
- Standalone DACs with 18 bit converters and 4×/8× upsampling
- Standalone DACs with custom DSP filtering, 20 bit converters and 8× upsampling

And, along the way, you could chart the course in measurements. D/A converters got more linear,

less noisy, and achieved higher performance by every measure. New versions of the old products performed better, because the multibit technology behind them was improving. Publications like Stereophile started measuring jitter, which raised awareness of its importance and led to jitter numbers steadily decreasing.

The result? By the early 1990s, it was possible to get 19+ bits of linearity out of multibit converters—a huge leap forward from the 13 or so bits of early CD players.

Progress wasn't only made on the playback side, either. Mobile Fidelity contracted Mike Moffat (yes, our Mike Moffat) and Nelson Pass to create their GAIN system, an insane recording chain with a real 16 bit oven-controlled multibit DAC that output linear PCM with no missing codes up to 500 kHz rates. This multi-chassis product took up almost a full equipment rack … but it was what was necessary to do good 16 bit ladder analog to digital conversion. Arguably, it still is.

Now, of course, there was only one problem with all of this progress: price.

Check the historic price of a PCM63 D/A converter, and you'll quickly realize that it's something that will never appear in an iPhone (nor would it fit.)

So, what to do? D/A chip manufacturers came to the rescue with products based on 1 bit sigma-delta modulation. These products were less expensive, easier to use, and more highly integrated. And they measured pretty well.

Another leap forward? In one way, yes. Without sigma-delta D/A converters, we wouldn't have the wide range of DACs and ADCs we have today. Your smartphone has a DAC in it with specs we would have killed for in 1990. The analog to digital converter inside it may even output 24 bit samples, at higher sample rates than we would have ever imagined.

And we can't underemphasize the impact of sigma-delta technology. It has allowed us to create more DACs (and ADCs) more inexpensively, with higher performance than we would have guessed, 20 years ago.

But we did lose something in all of this progress.

Digital Today: The Lost Decades

Today, it's largely a sigma-delta world.

- **Recording.** Most recording studios use analog to digital converters that employ A/D chips that use an intermediary multibit sigma-delta

format before their PCM output. Note that this isn't DSD. And note that even sigma-delta can have shades—single bit, multibit, etc.

- **Mixing.** From there, the PCM output is mixed/mastered in PCM (pretty much all mixing and mastering is in PCM … yes, even recordings that end up as DSD.)
- **Playback.** From there, it's typically going to be compressed and downloaded or streamed to a player using a multibit sigma-delta D/A converter.

Or, in the case of some crazy audiophiles like us, it's stored uncompressed, maybe even in high-res, before going to a DAC with a fancy multibit sigma-delta D/A converter.

Or, in a literal handful of cases, it might go to a true multibit R-2R converter, just like the old days. But that's a fraction of a fraction of a percent.

"So, who cares what it is, I just want good sound!" you say.

And we agree! We're far too wrapped up in formats. Take that format-proselytizing energy and aim it at the studios. Lobby them to produce better recordings. That will produce greater benefit than any format "regime change."

But … here's the deal (and here's where we get philosophical.) In today's sigma-delta world, we've lost something that *we* consider important: the original samples.

They're destroyed by upsampling, they are destroyed by asynchronous sample rate conversion, they're destroyed by sigma-delta D/A ICs. What you hear is an interpretation, a *guess,* at what the original content was (they don't call them successive-*approximation* converters for nothing.)

"But this can't possibly matter, it's hard to measure the distortion of your typical ASRC, for example," some will say.

Hard to measure doesn't mean *it isn't there,* we say.

Bottom line, it's a mathematical fact that samples that have passed through a digital filter, an asynchronous sample rate converter, or a sigma-delta modulator *are not retained.* There is no closed-form solution to the math.

"And why should that matter to me?" you ask.

Maybe it doesn't. Maybe the approximation is good enough.

But maybe it isn't.

And this is where we get to the core of what Yggdrasil is about: *what if we haven't been hearing everything PCM is capable of, because we've been hearing it on delta-sigma technology that throws away the original samples?*

Yeah. We know. We're crazy.

And perhaps we are. Perhaps it will make no difference at all. Perhaps it won't be important to anyone other than us. But the fact is: we have a solution to retain the reproduce the original samples, without the drawbacks of a non-oversampling design. It is in Yggdrasil. And we'll see what you think, very, very soon.

And that is the absolute core of our digital philosophy: retaining the original samples, all the way through to the output.

"But, It Doesn't Matter, Because … "

Because this position, this philosophy, is so counter to the currently accepted wisdom, I've prepared a quick discussion of possible objections to it, for your convenience.

"It doesn't matter anyway, because everything comes from a delta-sigma ADC these days. Do you have any original bits at all?"

Actually, this isn't entirely accurate. There are still multibit ADCs out there, though they are probably thin on the ground. There are also plenty of recordings made with multibit ADCs, including Mike's GAIN system. They don't disappear when new technology appears. And, you know what? Instead of being fatalistic and negative, we'd like to consider the best-case scenario: that we actually push PCM's capabilities forward to the point where new multibit ADCs appear.

"But how can those old DACs possibly perform better than the best of today? They're only 20/48. We have 32/768."

Going from 16/44 to 20/44 actually makes more difference than anything else, when it comes to digital. Why? Although the Nyquist theorem says you can perfectly reconstruct a waveform from digital with 2× the sample rate, it assumes an infinite-bit ADC with no quantization error. The more levels, the less the quantization error. 16 bit=65 536 levels, 20 bit=1 048 576 levels. 24 bit is 16 million+ levels, but nobody has ever achieved 24 bit linearity, period. The best DACs are about 19.5 bit to 20 bit, even after 20 years of "progress." (Hence, "the lost decades.") Higher sample rates are nice for analog filtering, but limit the amount of horsepower a digital filter

can bring to bear … and it takes up more storage space. So that's a tradeoff. And "32 bit?" LOLOLROFLCOPTER. There will never be any 32 bit music. Because physics.

"It doesn't matter anyway, I'll buy whatever sounds best."
Yep, absolutely. That's what everyone should do. No argument there. Have a listen, and decide for yourself.

"But I really like the sound of DSD, I want it to win."
That's totally cool as well. Just don't make it out to be anything "magic." It is simply different. As is analog. Which can be very, very good as well. Treat us like the crazy uncle who's a little touched in the head, and continue enjoying your DSD. After all, even if it "loses" as a format, it's not like the files will disappear.

The Summation

Here's what we propose: let's see what we can do with the huge amount of music we have in PCM format.

Can we make it better by retaining the original samples? Can we get out of the performance

we've been in the past 20 years? Can we bring this technology down to lower price levels? Can we change the future by picking up where multibit left off, 20 years ago?

Maybe. Maybe not.

We'll see …

Chapter 33
Black Friday

Business. Let's get back to business.

There's been a lot of technology and philosophy in the chapters above, and that's fine, because in a tech company, there's gonna be a lot of tech, and in any company, there should be a defined philosophy (or mission, or whatever other corporatese you'd like to call it) to define *why* the company does something, because otherwise you'll just be a reactionary company, and that usually doesn't work out to well. Companies without focus, who respond by trying to please everyone, will usually find that they end up pleasing nobody.

So, this chapter is about choices. Let's start with the title subject, specifically two business choices centered around the new American "tradition" of Black Friday.

Setting the Scene

Okay, at the time of these choices, it's October–November 2013. We're still in the Schiithole, but we won't be there for long (that's the next chapter.) Vali is getting ready to launch. We're deep into Ragnarok, and beginning to realize that it will be a lot more difficult than we first thought, after a not-stellar showing of a rough prototype at Rocky Mountain Audio Fest. Yggdrasil is still going through questions about D/A architecture, and we're beginning to realize that there may be some deeper questions we have to answer before we get that product into prototype form. At the same time, several other products are in the hopper.

It's a busy, even somewhat insane time, and it's the first time that Alex decided to load up on ordering for the holidays, so we'd be able to avoid the Schiit 'completely out of stock before the end of the year' syndrome. Alex threw away our predicted numbers and ordered what he thought would truly reflect sales—and, in the end, he was very close to keeping us in stock on everything.

The Schiithole, at this time, is insanely overstuffed. We'd brought in a container so we could store a lot of packaging and other bulky stuff outside

in the beautiful patio area, but one container only went so far. Metal, packaging, boards, and finished products were stacked to the rafters, and more was coming in.

And, in the middle of all of this, we got a shipment of Valhalla, Lyr, and Bifrost chassis. Cool.

Except they're painted black, not brushed clear-anodized aluminum.

Black Schiit?

I suppose we were at least partially to blame. We'd talked to the supplier about alternate finishing options, and we'd had a couple of silver-painted, and black-painted chassis made as prototypes.

But what were we going to do with a whole bunch of black chassis—especially since we'd always told prospective customers, "Nope, only available in silver?"

Note: This is the first choice. We chose silver only for reasons I've gone into before—we liked the look, and we didn't want to stock an infinite amount of variations of our products. Consider Bifrost: it would have a total of 8 variations, if we offer black and silver. It's already painful

enough to have Uber, Standard, USB, and no USB variants.

So what do we do?

“Sell ’em,” Alex said. “They look good.”

And they did. The black powder was very nice—completely unlike what you’d expect to see on a painted product. It wasn’t grained and anodized, but it had a very finely pebbled and consistent texture that looked nice.

“And charge more,” Tony quipped. I groaned. “And then what happens when we have to stock black and silver forever, into infinity?” I asked Alex.

“Make it clear it’s just a one-time special,” he told me.

“If that works,” I said. Visions of Theta, where everyone ordered silver when we had black in stock, or vice-versa, came back to me.

“I don’t want to stock it as an ongoing thing, either,” Alex said. “But we have it. Why not sell it? The guys who have been asking for black will be thrilled.”

Alex was right. We should just sell it and go from there. It would make some people very happy.

And, I thought*, It'll give us a better idea of what demand is like for black products.*

This is the second decision: to run with what we had, and use it as a learning experience.

Which is how, in mid-November 2013, I ended up announcing a limited run of black Schiit products on head-fi.org.

And We Waited

When I made the announcement, I also made it clear that this was a limited run, and when they were gone, they were gone. For all the emails we'd gotten asking for black products, I thought I'd have to sit over the website admin, ready to turn off ordering in very short order.

And I waited.

And waited some more. A few orders trickled in. But not the anticipated flood. I went on head-fi to answer a few questions about whether the black products were going to be an ongoing thing, and made it clear that, with the current response, they certainly weren't going to be. A few more orders trickled in. But again, no gigantic wave. I was really surprised. For all the emails we'd gotten, the insistence that black was "make or break,"

the people who'd actually taken our products and anodized them black themselves, sales were slow.

How slow? It easily took 3× longer than our normal sales cycle to sell out a very limited run of black products. Now, this isn't to say that some customers weren't delighted with their black products. And I'm thrilled we were able to make them happy.

But, in the end, the noise around "gotta have black" eclipsed the reality of the demand. And that's a learning experience (AKA, 'you got boned, remember not to do this again'.)

And that lesson, I think, is one of the most important ones for any company, start-up or not: ***the clamor doesn't always equal the demand.***

Why is this? Well, I think for stuff that's mainly cosmetic, the reality is that if you have a great product at a great price, people are going to buy it anyway. If you like, say, Jura espresso machines, you may want one with the red side panels, but if it's only available in black, that probably won't break the order. Or vice-versa.

Same thing goes for stuff that is very niche. Niche features or functionality can evoke a lot of

passion—and, while that passion may translate into many emails, it may not translate into sales.

If we had simply said, "Yes," to all the requests we get (or tried to accommodate them), it's quite possible that, say, Lyr 2 would be:

- A DAC/amp rather than an amp
- Available with an external power supply that costs more
- A balanced amp as well, and have balanced inputs too
- Available with several tube options, some costing +$300 over the stock ones
- Capable of all kinds of data rates that you can't buy content for
- Orderable with fancy power cords and fuses
- Delivered in at least two functional levels, one more expensive than the "standard"
- Available in at least 4 colors

It would also:

- Cost a lot more
- Make the possibility of shipping screw-ups much higher (oh, you wanted black/uber/power supply/etc)

- Be much more difficult to service (need to maintain docs
  for all variations)
- Be more confusing to buy, possibly leading to paralysis by analysis

And so on.

So, after our brief black experiment, we made another choice: that we weren't going to add black to the line as a standard option. Nor would we speculate on whether or not we'd have black again.

Now, this doesn't mean we won't have it again in the future (black, we realized, is a great way to refinish chassis that cannot be re-grained). But, as before, we won't speculate when it may be.

So, What's This Got To Do With Black Friday?

Well, other than the fact that the black chassis were introduced shortly before Black Friday, it's all about choices. Yes, more choices.

Many companies, both retailers and manufacturers, choose to participate in Black Friday through special deals that start the day after Thanksgiving. Some do a lot more than participate—they actively flog the upcoming deals and whip people

up into a buying frenzy so people can ruin a good chunk of the spare time they might otherwise be spending with their family, or simply dozing in a good turkey-coma.

Of course, this is US-centric, and perhaps other countries around the world don't have this yearly buying orgy. If so, you're fortunate.

Because it wasn't always like this. The stampeding, deal-crazed, sometimes murderous customers-trampling-customers thing is really recent. As in, last decade recent. And each year, it seems to get whipped up more and more.

Why? I suppose the theory is that "if ya ain't got customers right away, y'aint gonna do well this season." Or something like that, translated into corp-speak.

The real reason why? Because companies *choose* to participate.

That's a choice.

Companies (like Schiit) can also choose not to participate. And we don't. But one company did us even better last year. And because of this, they have my ultimate respect.

What company? Cards Against Humanity. Instead of offering discounts, or remaining neutral, they actually *raised prices* on Black Friday.

Nicely done.

Sometimes the hardest thing to do in business is to not just opt out, but stand against a trend. Kudos to them. It's a choice—a powerful choice.

So what am I going to do this Black Friday? I don't know. Turning off ordering is tempting, but I'm thinking of simply doing a banner like this:

> Really? Shopping today?
> How about spending some time with the people
> you care about?
> After all, the prices will be the same tomorrow.

Coda: More Choices

It's funny. I just came back from Rocky Mountain Audio Fest. It was a great show. But I'm simple. Any show is a good show when we remembered to get our hotel rooms, ship all the right products, and our products performed as expected and no

smoke came out of the latest prototypes we were showing.

But it was a great show, because I noticed a trend with some of the visitors who stopped by our booth. Over and over, I heard, "Wow, there are actually prices in the brochure," and "Wow, this is actually affordable!"

And, during the course of the show, we had some pretty big-time visitors, including a prominent audio technology blogger, a very big-name audio design engineer, and a high-level I engineer, amongst others. The audio design engineer and the I engineer seemed very taken with my straight answers on our products, how we achieved the performance levels we did, and the overall look, feel, and construction. That was part of a trend where I noticed that the people closest to the design and manufacturing side were most impressed—from product design and engineering to the DIY community.

"How do you do this … for this price?" they asked.

"Well, we're simple," I explained. "We don't do fancy chassis, because that would dominate the cost." I showed them how we do all the punching, machining and finish work, then fold the metal. They nodded, understanding.

“And we don’t have a dealer network, because we’re not living in the 1980s anymore,” I said. “As soon as you choose distribution, prices will—”

“Double!” one of them finished for me.

“Or more,” his friend added.

“And we don’t have to advertise very much, because the value is clear,” I said. “If we were doing a bunch of products that were very similar to other stuff at the same price point, we’d have to do a lot more advertising.”

“Or you could just make mega-price gear,” the friend said. “Though that’s getting crowded as well.”

“And then we’d have to deal with the audiophile nervosa,” I said. “We’re not really set up to convince people to buy our stuff. I need to do a shirt that says, ‘Schiit sales departmement org chart’ on it, with an empty box below.”

They laughed. “Yeah, upstairs is about $130 000 turntables with no arms, $30 000 players that are computers in billet chassis, and $50 000 amplifiers that would have cost $3k a couple of decades ago,” the first opined.

“That’s a choice,” I said. “They chose beauty, fine finish, and distribution. We’re a lot more functional.”

They laughed. “And that’s why everyone’s down here in this room!”

We talked for a while longer, they had a listen to Ragnarok and Yggdrasil, and listened to Mike Moffat say, well, whatever he says at shows (sometimes I don’t want to know.)

And in the end, they nodded and grinned, clearly understanding.

“Good choices,” they said.

Chapter 34

You Want to Pay How Much? Or, How We Moved Again

This is probably how a "real" company decides to move their operations:

1. Based on future plans and internal feedback, decides they need to have more space/less space/different space/different location (more tax favorable, etc).
2. Gets input from key management on the kind of space they need.
3. Surveys the available space in their target area with the help of an industrial lessor/realtor.
4. Weighs the options available and decides on one.
5. Plans well in advance for the business disruption of a move.
6. Has any build-out done, and a floor schematic ready, at the new office before the move commences.
7. Moves in to the new built-out, planned office space with minimum fuss and muss.

Here's how we moved, late in 2013:

1. The landlord came to us and said, "Hey, we have someone who wants to buy the building you're in. Can you move?"
2. I laughed and said "Sure, have them pay us a year's rent for the business disruption." Expecting them to laugh in return.
3. Unexpectedly, they paid, and we moved. Total time elapsed: about 2 months.

Yeah. There you go. We don't do anything by the book.

Except there's more to the story than that, which I'll get into. But I want to riff on some of Mike's comments about trusting your employees, rewarding them well, and thereby ending up with a highly motivated, self-policing team. Which plays into the move as well.

Condensed Employee Advice

Okay, you can take any number of shill courses on how best to motivate your employees—the bottom line of most being that "they want to be recognized and appreciated, more than just paid well."

This is 100 % total bullschiit.

Exactly two things motivate high-performing people:

1. Money.
2. Freedom.

That's it. If you meet someone who really, truly believes their "employee of the month" cake is a big deal, or who thinks the free soft drinks are a sign that the company really loves you, or who's a head cheerleader at the company's monthly corporate pride rallies, or who says, "whatever they want to pay me is OK with me," do this: **run.** Fast. This is not a person you want in a high-functioning company.

High-performing people know they're good. They expect to be paid well. Period.

High-performing people also value freedom, such as flexible work hours. Sometimes this can be traded off against overall salary.

This tradeoff works very well in the agency world. Many creative directors would be happy with, say 75 % of a top salary and truly flexible hours (work from home 4 days, come in 1 day, for example), rather than a top salary and the 12-hour/7-day grind of a typical agency.

This tradeoff works less well in a business that has to run during typical business hours, but, truth be told, there are no expected times of arrival or number of hours per day for anyone at Schiit. Our assembly team usually works nights, because they want to. Alex and tech usually work during the day, because they're motivated to. And not vampires. Unlike our assemblers, who I sometimes wonder about.

So, how do you pay good salaries when you're just starting up and money is tight? Great question. Tricky answers, too. Because the first temptation usually is to give away a percentage of the business. Which is exactly the worst thing you can do.

"What?" some of you are yelling now. "Why wouldn't I show my trust in who I bring on by making them a partner? That way, both rewards and risks are shared."

It also makes running the business much more complicated. Co-owners may be on a high for a while—until they find out the business isn't going to do as well as they thought, and their portion of the profits is small … or non-existent. Co-owners also may share with each other their percentage … and it's gonna be a bad day if the guy who

got 2.5 % thinks he's more important than the lady who got 5 %. And when the majority owner (or owners) have to outvote the minority owners on something, buckle up for a significant productivity hit … or even discord that could end up with partners exiting.

And when they exit, remember, you have to buy back their shares—either per your buy-sell agreement (you have this, right?) or by some kind of entrail-reading known as "professional business valuation." Oh hey, the partner doesn't agree with your valuator's figures? Buckle down for a lawyer orgy, featuring painful forensic accounting, and "expert witness" tea-leaf prognostication of the company's future value.

Bottom line: you'll both pay lawyers, you'll both lose, the lawyers have a party.

And I haven't even gotten into the tax ramifications of giving away shares. It's complex, painful, not pretty … and can be costly, for both you and the awardee.

And yes, having an agency for 20 years, and making this give-away-the-farm mistake more than once … it is, as they say, a "learning experience."

Effective Motivation—
Without Giving Away the Farm

Okay, so how do you create a high-functioning team without giving away parts of the company?

First, by realizing the extreme worth and intelligence of a motivated, engaged person. Don't minimize their worth, and don't insult their intelligence. If you don't literally want to create everything by hand, yourself, you *need* great people.

Repeat after me:

1. Don't minimize their worth.
2. Don't insult their intelligence.

Now, say it again. I had a boss in a former life who loved to do #1: go around saying, "Nobody is irreplaceable!" Usually after I offered an unpopular opinion or asked for more money. I had a business partner on the agency side who loved to do #2: "We're doing this for you, we all have to make sacrifices."

Riiiiigghhhhttt.

"Anyone can be replaced," should be amended to "Anyone can be replaced, but there are people

worth keeping, and be very, very scared of who might come after."

"We're doing this for you, we all have to make sacrifices," should just be banned from ever being said. Any person with more than a few active neurons know you're doing it for the company, and that people who are laid off are making a much bigger sacrifice than you are. Just say it like it is, do what needs to be done, and treat everyone like an adult, not a child.

So what do you do?

1. **Tell the truth and keep your promises.** Okay, you're just getting started. You can't pay someone what they're worth on the open market. But if they believe in your company and its potential, they may start for less, and bet on the promise of future rewards. But let them know exactly how it is, and set some dates and measurements on when they can begin seeing the rewards. If you don't meet goals, let them know, and don't sugar-coat it. But, always, always, always keep your promises. Because it's gonna get really ugly if you forget about them and buy a new Mercedes first. Remember: good people aren't dumb.

2. **Provide personal motivation.** And by "motivation," I mean *money*. This is the better way. Start with a livable salary, and add bonuses that are based on visible personal or company metrics. Number of products shipped. Number of products built. A bonus on experimenting with, and getting the company into a new channel. Royalties for products developed on the side. Stuff that can be translated into dollars, without having to reconcile profit, and without having to resort to nebulous proclamations like, "If we're doing well." This is really the best way to do things—and it is the best way to get your company running at a 5× to 10× productivity advantage, relative to the "norm."

Schiit has done both of the above. The result? A recent meeting with the city economic development manager revealed our sales, etc (we're going for becoming a duty-free export zone). He said, "Oh, you must have about 40 to 50 employees, then."

"We have 6," I told him.

The guy nearly fell on the floor. "How do you do it?" I told him (pretty much the same as above.) He shook his head. "That would never work in most companies." I disagreed. It has worked

everywhere I've applied it, without exception. The key is getting smart, motivated people to start with. And that, as I've explained in previous chapters, is less of a "checking off things on a resume" thing, and more of a "gut reaction" thing. Every time I've ignored my gut, I've screwed myself.

So, how weird are we, besides breaking the productivity/employee ratio really badly? Let's look at 2 other figures.

1. **We *don't have a sales department* at all** (something I realized when writing last week's chapter.) There's not one single "salesperson" here. Nobody to pimp us out to dealers or distributors, either. Nor anyone on contract.
2. **We spend only about 0.2 % *of our revenue on marketing*.**
   That's 10× less than the smallest figure advised for startups, and 50× less than what's considered "typical" for a sustaining company.

So much for us being "just a marketing and sales company."

Bottom line: your people matter, *and* your customers matter. Treat them both right. Don't insult their intelligence, and know their value. Do that (for real), and you don't have to do the

normal BS babysitting/micromanaging/spoon-feeding/ infighting/sales/shilling/promo/hype thing.

And, I think, your life will be a lot more sane.

One Other Overlooked Rule

I had the surreal experience of going through the first Great Internet Boom, and seeing what happens when companies get way too much of Other People's Money. Celebrity chefs making lunch for everyone everyday. Any drink or snack they wanted. Playrooms full of arcade games, bean bag chairs, and foosball tables. You know, stuff like that. Perks.

Except, when things started to get tight, when the investors started getting nervous about cash burn, what did they do? They started taking these things away.

And what happened after that? Yeah, you guessed it. The best and brightest smelled blood, and made their exit. The downward spiral in many companies was started by the lack of free Red Bull.

So here's the rule: once you give something, never take it away.

If you have free food and drinks, they have to stay, forever. No matter how bad things are going. If you have free daycare and company cars, they have to stay, forever.

So … the corollary is … don't give them. Motivate your team with individual, easily quantifiable bonuses, and let them make tons of money for their own Red Bull, daycare, etc. That always works. And if the bonuses are based on visible metrics, there won't be any complaints if the metrics end up sliding.

So, On This Move Thing?

Why is the above employee motivational blather important? Well, beyond the productivity advantage outlined above, it also made for a smooth, uneventful transition into the new space. Because Mike and I don't have to babysit anyone, we were involved at only two points:

1. Looking at the available space (with Alex.)
2. Showing up when the move was complete.

That's it. Fully empowered employees take care of everything else—including finding and booking the movers, packing things up, unpacking and getting set back up, arranging things once we were there, buying things we didn't have, etc.

And … they were ready to move. The place was packed and overflowing. We wouldn't have been there much longer if there'd been no buyer. It just moved up the timeline about 3 months.

Looking at the space was fun. We were in the market for, as I told our real estate guy, "Something 3000 to 5000 square feet," which would give us 1.5× to 2.5× as much space as the old place. I expected to see a long list of candidates, since Valencia Industrial Center is the largest industrial development in Southern California, and 3k to 5k is above the starter spaces and below the 20k to 100k foot boxes.

And the list was pretty long … until we got to particulars.

The one particular that broke the back of most of the candidates was "air conditioned throughout." Most industrial spaces have AC only in the office section, while the warehouse remains uncooled (and, in some cases, unheated.) This is less than great in 110 °F summer days.

So our list of about 40 candidate spaces was reduced a bit. And by, A bit," I mean, it went down to 2.

Yes, 2.

Both spaces were in the same building, and right next to each other. I went through both of them with Alex and the realtor. One was just over 3000 square feet, one was about 5300. The 3000 SF space looked pretty much like the ideal place for us, except for one thing: a lot of these concrete tilt-up boxes have office built out in two levels up front, so you have upstairs and downstairs offices. Upstairs is useless to us (imagine carting products up and down stairs all day.) So, the 3k space really didn't net us as big an increase in usable floor area as it seemed.

The 5.3k space? It looked stupidly huge. I mean, ridiculously, cavernously huge. Bigger than Theta was at its peak. Bigger than Sumo. (Though not bigger than Centric during the dot com boom, which was in a converted 7200 SF industrial building. How I wish we'd have stayed there.)

But 5.3k? That was silly. Even with the office buildout upstairs, there was no way we'd ever use all that space. I mean, really, we were a small manufacturer, right?

But it definitely had the floor area we needed, and more. It also had a glassed-off area that would be perfect for tech.

So, after another visit with Mike, we made them a lowball offer, and got it.

And yeah, I've said before that you shouldn't haggle, etc, but this was more in line with bringing the price down to reality. It's a 1980s building, not super well-kept, and the landlords were loathe to do anything at all to it—not even the usual new-paint-and-carpet deal you usually get when you sign a lease. So I said, "We'll take it, no tenant improvement, at this rate."

They jumped on it, and the rest was history.

The actual move took place as Mike and I were at RMAF 2013. We left for Can-Jam, and came back to a different building.

And that's when the reality started setting in. This "huge" space turned out not to be so huge at all. With all the stuff we'd packed into the old building on the floor, there was a lot less space than I thought. It wasn't disturbing, as in "having to look at moving immediately" disturbing, but we definitely didn't have as much space to grow as I'd expected.

So what to do? Alex suggested racking the place out, so we could stack chassis and packaging up three levels high. Shortly before the end of the

year, that's exactly what we did—and that's when we started to look like a real company.

And what would we do with all of that space upstairs? Sure, I could take an office up there, but Mike hated the heat, and didn't want one. The bullpen area would be essentially useless as a listening room.

Rina came to our rescue. Her own business, Twilight's Fancy, was growing too big for our house (now, we can claim 2 businesses launched there—one in a garage, and one in a spare bedroom …).

"I'll take the upstairs," she said. "I'll sublease from you."

"How much of it?" I asked, remembering the adage that "a turtle always grows to the size of its tank."

"As much as you'll let me have."

Hmm. That saying really had me sweating now. Rina's business is space-intensive, and she is, ah, well, apt to, um, disarray. And piles.

"Let me think about it."

After some discussion, we came to an agreement where she took one of the upstairs offices, and much of the bullpen. So now we share our

building with a seller of jewelry findings and ribbon chokers. Makes perfect sense, in a way.

What We Didn't Do

Our new home, or "The Schiitbox," as we call it (appropriate for a concrete tilt-up) was big, but it wasn't opulent. The upstairs is carpeted with somewhat-worn, medium-blue industrial carpet that may have been the cat's pajamas in the early 1990s. The downstairs is pure concrete floor, concrete walls, and sheetrock. The windows don't open, there are no balconies, walls are painted utilitarian white.

In the peak agency days, this wouldn't have stood. We would have ripped out walls, done new carpets, gotten rid of the drop ceilings, bought new desks, and generally gone on a remodeling spree to make it something we could be proud of.

At Schiit, we did none of that. The same ugly carpet, concrete walls, and white paint remain. We did put up some pictures, but that's about it. Alex's "office" is a desk downstairs by the side door (we don't use the front door, that's Mike's office now.) We pulled the brand-new, nice carpet out of the tech area so it wouldn't be

a static threat, and left the bare concrete with glue marks.

We did, however, finally buy a full-sized refrigerator.

And racks. And desks. And more racks, as we expanded. And more test equipment. And shipping tables.

Why the focus on utility? Because we're not a listening room, a lounge, or anywhere you would want to hang out (well, unless you're super-geeky.)

And … because, when you have smart, engaged, motivated employees, it doesn't really matter. They think it's funny. They're thrilled to help us grow. And growth doesn't come from Hermann Miller chairs and Steelcase desks and faux-finish paint and $600 LED lamps.

It comes from, as Mike said, *giving a schiit.*

Coda: We recently took the space next to ours, bringing us up to about 8300 square feet. The landlord was as cheap as ever. The carpet is as ugly as it's always been. We needed the space for the Ragnarok/Yggdrasil lines, and some other future plans …

Chapter 36
A Real Company?

This was originally going to be the final chapter in the book, covering where we were as of December 2013: in the Schiitbox, busy racking the place out, and looking upon a company that was no longer a scrappy garage start-up, but well and truly *real*.

We were also supposed to be looking from a point of view of having been shipping Ragnaroks for 8 to 9 months and Yggys for 5 to 6 months, and you all know how that worked out.

> Cue evil, polite, or disgusted laughter, depending on your own POV.

But things change, and in the process of going through 2014, we learned a few more things, and got a few more stories under our belt. Stuff that I couldn't write about at the beginning—the rest of the Gen 2 products, Mani, Wyrd, Fulla, and a few other things that are still coming before the end of the year I hope), were launched without a whole

lot of fanfare. And, in the product launch game, no histrionics is good—that means we didn't have to make sudden, unexpected changes because stuff wasn't working, or performing as expected. The thing you want the most is a boring product launch.

So where do we go from here? Well, there are a few more chapters to get us up to date, and then I expect I'll continue adding to the saga—maybe not weekly, but we have enough stories to definitely set up for something every couple of weeks, or, at worst, monthly.

So everyone who's enjoyed this story can look forward to more … just on a more sporadic basis.

And everyone who's cringed at the thought of reading more Schiit blather will just have to cringe a little more.

But that's looking forward. Let's look back at December 2013, and take a look at just how well we did on the "real company" front.

By Our Own Rules

When I started this whole thing, I tried to condense down the business rules up front for people who just want to cut to the chase. These were:

1. Shooting to be the next billion-dollar mass-market company is insane—you might as well buy lottery tickets.
2. Niche is where it's at—specifically a niche where people can get in fistfights over the color of a knob.
3. Pick a niche you know and love and do something nobody else can do—"me-too" never works.
4. Be memorable—this isn't about getting *everyone* to *like* you, this is about getting *some* people to *love* you.
5. Go direct—distribution is a poisonous remnant of 19th century economics in a disintermediated world.
6. Run from both conventional marketing wisdom and the social media mavens—both of them are geared towards the mass market with eight-digit ad budgets and multiple decades to build a brand.
7. Don't think this'll be easy—this is hard work, but you'll also be having a whole lot of fun if you're doing it right!

So, how did we do?

1. **Not being the next billion-dollar business.** Yep. Check. Schiit will likely hit 8 figures sooner than later, but 10 is really, really be-

yond the pale. If someone gave us $100 million of VC to come out with the broad range of products needed to get there, I think I'd rather leave than continue. Because, to hit that kind of revenue, you're looking at $99 headphones and $79 bluetooth speakers and $199 soundbars and $149 powered speakers with Class D amps … true mass products. And don't dare think about trying to take a creative approach with something—the path to success would be in low-dollar, high-volume products that are absolutely guaranteed to find buyers. Which means you want them in established product categories. Which means no swinging for the fences, no trying to break the mold … and, in turn, that means that we'd more than likely fail. Billion-$ audio companies got there either through decades of product development and marketing (and a bit of luck)—think Bose and Harman—or through 100% sheer luck, by creating a new product category—think Beats, though they never cracked more than $0.5 billion a year in sales. The Apple acquisition is what sent them over the top.

2. **Niche is where it's at, and fistfights.** Yep, check. We are a dominant player in a niche market (what they're now calling "personal au-

dio.") We also play in the larger two-channel realm, mainly with DACs, and now with Mani and some other products. But we're definitely a niche—and a niche so narrow that it's smaller than the audiophile market as a whole. We're not trying to be another Bryston. And fistfights? Yep. All the action is in personal audio. Traditional two-channel audio has the "Buick disease." It's moribund, almost literally. Over 70, you lose 10 % of your customers a year. Over 60, 5 %. That's basic actuarial table stuff. And you can't make up for the loss by increasing your product costs forever. Eventually, the last 200 people who think a $ 120 000 DAC is a good idea will die off, and you're done.

3. **A niche you love, doing unique, non-"me-too" products.** Check again. I got back into this because I've always loved music, and I found a new and exciting way to express it in the headphone space. It was everything that two-channel wasn't. Mike, who'd checked out at the end of the golden era of multibit DACs, felt exactly the same way. This isn't a field you can get into based on spreadsheets and margins and feature analysis. This is something you have to love. And unique? Yes. We have a unique place—not the cheapest made-

in-China gear, and not the eye-wateringly expensive audio jewelry stuff, either. With unique topologies, unique value, unique sound. It's a great place to be, because I don't think everything should be about price, nor should everything be about bling. I think we've struck a good balance. Others disagree. See below.

4. **Be memorable, it's about some people loving you, not everyone liking you.** Check, check, check. There are some people who have literally bought everything we make, and think we can do no (or very little) wrong. There are people who go out of their way to help us. And, at the same time, there are some people who dislike, or actively hate us. Dislike because of the name, dislike because of some perceived slight, dislike because they got a dead product (hey, it happens), irritation because we stood up for our policies, hate because of some imagined agenda pushing crappy products on an unsuspecting public, dislike because we don't toe the line to being 100 % subjectivist, distaste because we don't toe the line to being 100 % objectivist, irritation because we skewer some sacred cows, hate because we say what we believe, even if it's wrong and goes against the "accepted wisdom." But, you know what? It's the dia-

logue between the unabashed fans and the skeptics and the outright hostile that keep the discussions going—and that's what helps more people find out about us … and make their own decisions.

5. **Go direct, avoid distribution.** Almost completely check. We made the mistake of putting on some dealers in the early days, and some distribution that didn't do us any favors. Chalk it up to a learning curve. Now, with our quasi-direct pricing coming into play in Europe through Electromod, and the clear non-necessity of any dealers in North America, we're 100 % committed to staying direct. Asia is still a conundrum … but we'll see what we can do next year.
6. **Run from both conventional advertising and social media both.**
   Check and check. I mentioned our marketing budget as being a tiny fraction of what would be normal for any other company. Part of this is due to the fact of heeding our own advice. Yes, some advertising is necessary, but we've stuck to just a couple of online venues that are completely measurable and trackable as to their actual results. If they don't pay off, the plug is pulled and the funds reallocated. We have yet to take print pages in any magazine

(and the way magazines are going, I suspect we may never have to—except for one stunt for a new product I'm thinking about … hmm). At the same time, we have spent exactly zero time on Facebook or Twitter. Of course, this isn't entirely fair, because we (meaning I) have spent significant time on micro-social like here at Head-fi. But this isn't just marketing … I enjoy writing, and I enjoy the various arguments … er, I mean *discussions* … that we get into. No ad agency would have proposed this. And most corporate lawyers would go pale at the thought of their CEO going online and, say, calling out things like "the worst customer, ever."

7. **Hard work, but fun.** Yes, hell yes, and yes again. In some ways, this is the hardest thing I've ever done, but I'm having a whole lot more fun actually *making things* than I did on the agency side. But it hasn't exactly been smooth and easy, or all dancing-through-the-sprinklers. In fact, probably not a day goes by without a minor problem (out of stock parts, for example) and not a week goes by without something more moderate (new shipment of boards acting weird, for example.) Nothing is insurmountable, but this isn't a hands-off business, especially when we're rolling out

many new products.

The Perspective of December 2013

December 2013 was the first time I could look around Schiit and legitimately say, "This feels like a real company." Of course, we were still just a handful of people, rattling around in a messy, not-yet-organized big box with ugly carpet and scuffed walls. But Theta was just a handful of people, too, and the facility was never a looker. For the first time, we had a real tech area, a real assembly area, enough space for parts, an office for R&D, a listening area … and, more importantly, things were running smoothly. The team had gelled. Nobody there was perfect, but everyone did their jobs—and believed in the overall direction.

For the first time, it seemed like we'd *arrived.*

But at the same time, there was still a ton of work to do. Ragnarok still wasn't working right, after many, many firmware revisions. We were still going back and forth on DACs for the Yggdrasil, instead of moving into production. The Valhalla 2 and Lyr 2 prototypes had their own niggling problems.

And, worse … I'd gone off on some tangential stuff that looked like it might never become a product. We had one product just sitting on a desk because we were too busy working other things out. And at the same time, Mike was playing with phono stages, Dave was designing Wyrd, and I had a bunch of crazy ideas that I figured we could launch in 2014, as soon as Ragnarok and Yggdrasil were put to bed.

But without the freedom to experiment, without the ability to dream about what might be … many of our products would never have come to market. As I told Mike recently, "It says something that our only truly conventional design is Magni."

Mike laughed. "And it's essentially a small speaker amp, which is insane for a $ 99 product."

Remember: *be unique.*

So Where Do We Go From Here?

Anyone with a brain knows that 2014 didn't solve all the problems we were facing at the end of 2013. But we never expected it to. That's not how business challenges work. When you knock the current ones down, you find new ones.

But, in 2014, we were in a new position: that of being part of "the establishment." This in itself is a challenge. We've noticed a recent increase in threads that go like this: "DAC recommendation (besides Bifrost)," or "Headphone amp (not Magni or Asgard 2). Simply by being part of the perceived "establishment" means that we can be dismissed.

Why? Plenty of reasons:

1. **When "everyone else" has one, it may be enough to simply be different.** I know lots of people who don't want a Camry (or a Corvette) simply because they're "too popular," not based on any objective criteria.
2. **Something new and shiny is new and shiny, period.** There's excitement in being one of the first to discover something new and good. We know this. We also know there's inevitable backlash against a perceived "flavor of the month." People were calling us the "flavor of the month" for years. Some flavors last longer than others. Only time will tell.
3. **We don't make what they want.** Whether it's black gear, or DSD, or power switches on the front. And that's perfectly cool. We can't please everyone all the time.

4. **They heard we're crap.** Yep, that still hangs out there. It'll never go away. We have to just keep on keeping on.
5. **They had a bad experience.** This usually means they got something that didn't work out for them (like, say, an original Valhalla and LCD-2s) or we shipped them the wrong thing (it happens—and even if we fix it right away, sometimes bad feelings linger), or they got a DOA/defective product. And even though our DOA/defectives are less than 0.05 % now, that still adds up when you're moving tons of products. Again, all valid reasons (though really, if you aren't sure if one of our headphone amps is ideal for your products, ask us … we'll tell you honestly, and it won't be "just buy the more expensive amp.")

But, bottom line, we aren't the new shiny anymore. We aren't the One True Challenger who will Redefine Everything, Forever.

So what do we do with this?

Well, the first step to overcoming a problem is to recognize it. We can't dismiss it out of hand. Nor do we think it's right to respond in the manner of many companies—which is to throw marketing money at it. Nor can we simply reinvent ourselves as a company making expensive audio jewelry,

because, well, that's not us. Nor is it fair to our customers.

So here's what we're doing:

1. **Staying where we are, and getting even better at it.** This means staying affordable, and putting the bulk of our effort towards creating truly amazing, groundbreaking, and inexpensive products. This may also entail "killing our babies," as I've said in another chapter. Because the only thing worse than bringing out a product that might cannibalize sales of your main product line is to have someone else do it first.
2. **Exploring new vistas.** This doesn't mean $20 000 DACs or $15 000 amplifiers. This doesn't mean a full line of two-channel products, like conventional preamps and power amps. Nor does it mean crowdfunding or "co-creation" or any of those fancy new models that essentially say, "We have no ideas." But it *does* mean we'll be looking at how we can make a difference in the two-channel world—including in some very, very surprising new directions. But all on the affordable side. And all with some significant advantage. No me-too products.

3. **Less product introductions in 2015 … but with more significance.**
   There are some real surprises on the horizon, no kidding. At least one product … no, well, two products on the horizon for 2015 are real eye-openers. It pains me that I can't really talk about them, because the best thing besides introducing a new product is talking about it. And many of these are really far along. In fact, I'm listening to one now. Another I've been listening to for months (but it is going in for some changes to bring in some very cool trickle-down tech.) What I can say … watch TheShow Newport.
4. **Continuing this conversation, and listening to your input.**
   Yes, we do listen to your input, and it does help when we're developing new products. You'll see some of your own ideas reflected in the coming year. But … and here's the real but … it's just that we also have our own point of view. Call us old-fashioned, but we think that if we can't add something of our own to the product development, why are we here? Why not form a coalition and go to a contract manufacturer to realize your perfect crowd-sourced dream product? If we're right, we'll do well, and if we're wrong, we'll take our

lumps. But in any case, this document, this conversation, and this exchange of ideas will consider, as long as you'll tolerate me.

5. **Continuing to challenge the established wisdom.** Whether it's product design, buzzword compliance, unicorn formats, en-vogue branded power supplies, fancy capacitor types, or any other of a dozen different things, we have a unique point of view. It's far too easy for someone to wade through a forum, look at some product websites, and decide that "well, everyone's talking about ABC, and it's gotta have XYZ, and of course it has to be isolated with SuperWowie technology, and of course it has to support rates up to 64/3472 … without knowing that everyone's talking about ABC because a big name is putting big money behind a proprietary technology, and one person said it had to have XYZ, and SuperWowie technology doesn't work at the rates that highres is available at, and 64/3472 doesn't even have any demo tracks available, much less any music you can buy. We call out this marketing-based "common wisdom" and skewer these sacred cows. Some people like it. Some people disagree politely. And we rub some people completely the wrong way.

Because … (you know it's coming) it's not about

getting *everyone* to *like* you. It's about getting *some* people to *love* you.

And that, in a nutshell, is what we're going to continue doing. It's up to you to decide if we're completely insane.

Or not.

Chapter 37

The Value of Diversions

Going into 2014, we had plenty of stuff to keep us busy—the ongoing decrapification of the Ragnarok firmware, DAC choices and programming for Yggy, and the planned introduction of both the Valhalla 2 and Lyr 2 as the main features.

On top of that, too, I had a few new ideas on the burner, including one with four variations—one of which became Fulla. Mike was playing with a phono stage that would become Mani. Neither of these products had firm release dates, but we were thinking, "Well, these are probably going to be 2014 products, too."

So, yeah, tons to do. Definitely no shortage of engineering work. But …

… I always get ideas.

… and Mike always gets ideas.

And there's usually a few things that, well, just happen from them. This chapter is a story of two

of these ideas: SYS and Modi Optical.

Shoulda, Woulda, Coulda

"Oh snore," you might be saying. "Not exactly groundbreaking stuff there."

Well, maybe not. But, you know what? Both products are interesting in their own right, but they're also more interesting as a signpost to how we work. It also illustrates the "Shoulda, Woulda, Coulda" dilemma that many companies have. If you start your own business, you'll come up against this—maybe even before you sell a single product.

The original Asgard and Valhalla are perfect examples of this. If we'd had unlimited start-up funds (or at least much larger), the original Asgard and Valhalla may have been made in custom extruded cases with a very large extrusion profile (coulda, woulda.)

Of course, if we had done that, we would have ended up with a product that was very difficult to finish, perhaps stuck in a non-ideal form factor, possibly with extrusion consistency problems (large extrusions aren't always exactly to print—they warp and curve.) This might have put us

back 6 months, 12 months, or even more from the introduction date we hit.

So, even if we coulda, woulda done it, it might not be something we shoulda done. Being able and inclined to do something one way (and having the wherewithal to do it) might have actually been bad for us.

"Wait a moment!" you might be saying. "I have no idea what you're talking about!"

Well, there's no reason you should. I came up with the Shoulda, Woulda, Coulda thing only after doing this for a long time. It's not a management thing they teach in business school.

But I bet it's something many engineers know on a gut level.

So, let's break it down:

Shoulda. This is something you should do, even if you have to wait to do it right, or at all. The challenge is seeing what you shoulda done, before you go off and launch a product. Some things are very obvious (high reliability), some are not so obvious (a specific feature, like gain switching), and some only become apparent after the fact or due to changes in the ancillary or competitive environment. But if you know you should do

something, do it—even if it impacts timeline or budget.

Woulda. Here's conditional one. "Woulda" is something that you would do, if you had more resources. A perfect example is Magni. We woulda done something like Magni to start—an amp at a price point much less than the original Asgard—if we had the resources to do high-volume surface-mount production at the time. This one's a judgment call. Do you wait until you find the investment (or natural growth) to do something you would do if you could? Or do you launch, grow, and bring it in later?

Coulda. Here's the dangerous one. You coulda done lots of things. You coulda made the Asgard anodized gold, or added a bunch of switches on the front to do fancy things (and made up crazy names for them), or put it in a fancy retail box, or had 16 variations with different performance levels, or offered custom finishes, or put it in a rhomboid enclosure machined out of a solid block of aluminum. But all of these things affect price, complexity, reliability, and serviceability. Just because you coulda done something, doesn't mean you shoulda.

So how do you tell the difference between a

"shoulda" and a "coulda?" Again, I don't have a 100 % perfect algorithm for doing this. A lot of it is strictly gut. A more generalized way of sorting might be that "shouldas" are usually fairly simple and straightforward, and don't lead to huge complexity in design, manufacturing, sales, or service. "Couldas" are more complex, ideologically based, design-centric, and can have big impacts on complexity of all of the above.

Examples of shouldas:

- High reliability
- High performance
- Easy to manufacture
- Easy to service
- Simple to understand and use
- Good versatility
- Good documentation

Examples of wouldas:

- Using specific technologies or techniques with high step-in cost or high set-up commitment to:
  - Enhance performance
  - Reduce manufacturing cost
  - Streamline production
  - Improve reliability
  - Enhance traceability
- Increase volume manufacturing to reduce cost

- Employ high-volume techniques only available at certain operating levels

Examples of couldas:

- Make it look like a spaceship
- Have switches glow and pulse with soft blue light
- Offer 12 different finishes with custom engraving
- Mandate for “unibody” design
- Eliminate all exposed fasteners
- Sealed for life!
- 83 different features for “enhancing” sound
- 16 different filters for ditto
- A 100-button remote
- A 2-button remote with proximity sensing gestural technology
- An Android or iPhone remote
- No external controls for a clean look
- Logos that shine through the PC board
- A nifty glass window to look at the guts
- 8 inputs, including ones that sound bad (Bluetooth)
- 100 % buzzword compliance

Products that are made with lots of “couldas” usually look good on those long checklists that, say, automakers like to use.

> Hey ma, this here Kerfluppelator A8 has three 8 inch touchscreens and reads your brain waves to unlock the doors, plus it has a 40-speaker sound system with Golby GigaSound, and 17 cupholders, plus it lets you change the engine note with a Dynamically Tuned Induction Resonator! Let's definitely get that instead of the Boretonium LE that just has knobs and crap for the stereo.

But products made with lots of "couldas" can also spell disaster on many fronts:

1. It's so complex nobody wants to, or knows how to, use it.
2. None of the fancy features ever works quite right.
3. Service guys shudder and hide behind the building when one comes in to be fixed.
4. Production lines that are always stopped for the latest firmware, hardware, or software tweaks that respond to 1 to 3 above.

The biggest "coulda" failures I've seen (not just in audio) usually stem from one of two things:

1. **Getting fixated on a design that is not producible.** This is the current bugaboo of the super high end. Milled chassis with hours of

CNC work for “striking” features, made out of a single block of aluminum. Yark. Do you know what billet aluminum costs? Do you know what machine time costs? I’ve seen quotes for stuff (not ours) that went as high as $1600 for a faceplate (yes, one faceplate.) This type of design is simply unsustainable, unless you’re literally producing millions of products, like Apple—and even then it’s pricey. So, hint: just because Apple is doing it doesn’t mean you should.

2. **Getting fixated on a long list of features.** Or, the “Well, of course it has to have this and this and this and this, because Portlefloot and Giggleberry both have them,” excuse. If you find someone making a list of the competitors’ features, and that list is over 5 items long, take it out of their hands and tear it up. You’re comparing the wrong things. Especially if it contains the en-vogue buzzwords of the day.

Okay, So What Does This Have to Do With Two Cheap Products?

Oh, you want me to make a point?

Well, here it is: pretty much all Schiit products start as “couldas.” Some of them never make it

past this stage. Some do. SYS and Optimodi both did.

Why?

Okay, let's start with the story. SYS first. SYS (or Switch Your Schiit, our little 2-input passive preamp) started without a name or any purpose beyond letting me switch between two different music sources on my desk in the garage, where I did most of my engineering work, up until an office reorg finally got me a space in the house.

At my garage desk, I usually use powered monitors (variously, Emotiva, Equator, and JBL—nothing super special.) Powered monitors require a line-level signal, and many don't have convenient volume controls. An Asgard 2 would work for one source, but I didn't want to waste a perfectly good A2 on a system where I don't use headphones much. Plus, I had two different sources I wanted to use.

So, an easy solution is a small passive preamp (or attenuator, to be more precise.) We had a small chassis for Magni and Modi that could easily fit two RCA stereo inputs and one stereo output. We had a small pot we used in Magni, and a pushbutton switch we used in Loki.

Ten minutes of PCB layout, and I had a board that would fit into the Magni/Modi chassis and provided the two inputs and one output I needed.

Of course, there was one small problem. Passive preamps really don't want to see a 50 kΩ potentiometer, like we use in Magni. They'd rather see 10 kΩ, or even 5 kΩ. The reason? Since they're passive, they don't have a low-impedance output stage to drive cables, and some cables have significant capacitance. The result? Rolled off highs with a large pot and long, capacitive cables.

But there are other pots out there, so I picked up some 10 kΩ pots, and everything worked fine.

And that might have been the end of it, except I made up a couple more and brought them into the Schiitbox (the prototype PCB service we use usually gives us several boards.) Alex grabbed one and started using it at his desk with powered monitors. Mike wanted one. Tony took one.

That's when the light bulb went off. Our DACs don't have volume controls. Many people have powered monitors and don't use headphones. Even more have power amps with no volume control, and only one or two sources. There was a business case for making this small, unnamed

experiment into a real product. That's when it turned from a "coulda" to a "shoulda."

Note: it may have stayed a "woulda," if we didn't have the resources to dedicate to making a whole lot of SYSes … and taking the chance on them hanging around for a while if I was wrong on people wanting them.

But SYS wasn't a big investment. So, with a quick chassis modification to Magni, a name, a manual, and some photos, we had a new product, and launched it without any major fanfare other than a press release on our site.

Happily, SYS has found a nice home in our product line. Its sales won't light the world on fire, but it's a steady seller, and a ton of people really like them. I still have one on my desk (now in the house) to this day—and it's even more important inside, where I have less space.

The Optical Modi was Mike's gig. As I've mentioned before, Mike wasn't a fan of USB input. The optical Modi was both his way of saying, "I really don't want to use USB," and a response to various inquiries we'd had about offering a Modi with optical input.

Aside: Mike's softening on USB over time, but the fact is, USB input still has its own collection of glitches (usually related to port power management problems), and optical has a whole different collection of them (like 24/192 being very iffy, except when you have a source that can *really* output it, and a good short USB cable, and you've performed the ritual blood sacrifice of a goat when the planets were all aligned). Mike didn't like Toslink back in the Theta days—look at the default input on all the old Theta DACs—but that was then, and this is now.

Mike's optical Modi was a lot more work than SYS, because, other than the DAC, output, and filter stage, everything changed. Optical doesn't carry power like USB, so we had to add a wall-wart and a power switch (both taken from Magni, of course.) Also, we decided to try out the new, cool AKM SPDIF input receiver, the AK4113, which we hadn't used before, so Mike had to learn its oddities and apply his own tweaks to it.

Still, there came a day when Mike gave me a green-board prototype and said, "Have a listen," grinning. And it sounded very nice, just out of the built-in optical of a MacBook Pro.

"Is it worth making it, though?" I wondered.

Mike shrugged. “People have been asking for it.”

“They asked for black, too.”

Mike groaned.

“But … it might be less fuss than dealing with all this new USB port power management crap,” I said.

“And a lot of people still use CD players,” Mike added.

“Yeah, 7 of them,” I shot back.

Still, what pushed it from being a “coulda” to a “shoulda” was simple: I remembered that Apple Airport Expresses and Apple TVs had optical outputs. A Modi Optical would be an easy way to make a much higher quality streaming system from them. Like SYS, if we were a younger company, or more cash-strapped, though, it may have stayed a “woulda.”

Now, I didn’t see this as a separate product; just an alternate Modi. Want a Modi? Cool. Choose optical or USB. Hence no different name, no separate model.

So we ordered some boards, made some more chassis changes, did a manual, photos, etc, and we were off and running.

It was only after we launched that someone brought up a killer app for the optical Modi: positional audio for gaming. It seems that a lot of computer sound cards can do positional audio, and have optical outputs. However, they can't output via USB. So, after we heard that, I added that to the list of applications for the Optimodi.

Like SYS, Optimodi probably won't set any sales records, but it does sell rather briskly. It's found a home in a lot of streaming setups, in gaming, and even with the 7 or so people who still spin disks. It's a neat little product, serves a need, and will probably be around for a long time.

So, About Those Couldas …

I mentioned that I was playing with some other stuff … one of which became Fulla. All of the things I was playing with were on the "portable" end of things. Fulla I'll cover in another chapter. The other three variants (including, yes, a battery-powered model) probably won't ever become "shouldas."

Why?

It's simple: we didn't really like them all that much. The portable stuff we've played with was

either:

1. **Too close to other products on the market.** If other companies are already doing a good job on portables, why get in just to be a "me-too?" Our key sticking point here is in the switching power supply. Sure, you can take a 3.7 V lithium battery and use a switchmode converter to create, say, ±10 V, which would give plenty of power … but there are companies already doing this. How would we be different? By being fully discrete? Maybe. But would it matter enough? More interesting would be using two 7.2 V batteries and not using switchers. Again, maybe. But the charging issues wouldn't be trivial, and the size of the box wouldn't be especially small. So that's shelved.
2. **Too much of a question mark.** There are plenty of other companies doing portables, with long experience using lithium batteries. They know the ins and outs of them. We don't. And we don't want to find out about them the hard way (as in, after the product is released.) Could we have a flawless launch? Perhaps. But also, perhaps not.
3. **Limited interest internally.** I understand the use case for a portable amp/DAC combo. I understand that not everyone has a house or

apartment where their stuff can sit all day, nor a desk in an office. But I don't need such a device, and Mike is even less interested. If we can't give it our all, why bother?

So, will there never be a Schiit portable? Most likely. But who knows? Maybe I can solve the two-battery problem. Maybe I'll make something that I really love and want to bring to the market. But it's definitely not top of mind.

So Why Waste Time On Diversions?

For us, it's simple: because we get frustrated. Projects like Ragnarok and Yggdrasil are immensely draining, especially when you hit the latest bump in the road. When something goes wrong on a massively complex product, you really have three choices:

1. **Go back into it frustrated and angry.** And risk screwing it up worse. Or going down a bad path that you don't have the perspective to shy away from. Getting right back on it, without a breather, is dangerous.
2. **Go watch some Fast and Furious movies.** Or cartoons. Or Facebook. Or whatever mindless activity you need to reset your perspective and come back at it fresh. Of course, doing

this isn't really moving anything forward (even if it is necessary now and again.)

3. **Play with something else.** You know the old expression, A change is as good as a rest?" Yeah. Put away Ragnarok, and work on something simple. Easy. Quick. See if it goes anywhere. If it doesn't, no harm except some spent time and a few prototyping dollars. If it does, you win. Both Mike and I take this approach most of the time. Sometimes it works very well.

It's simple why you'd waste time on diversions in the general sense, too: because you never know what you're going to find. If you're going to start a business, you should feel free to explore—and to let your engineers and design teams explore. Don't put the shackles on them at the front end, and you may be amazed what comes out of it.

But definitely, totally, absolutely, consider the shoulda, woulda, couldas.

Chapter 38
Wyrd Schiit

"Hey Dave, is there any value in reclocking a USB datastream, like you guys used to do with SPDIF on Theta's Time Linque Conditioner?"

Dave looked confused. "Wellllll … it's not really the same … with USB, it's more the hub chip repeating the output … but the hub chip uses, uh, a crystal oscillator, so it's … well, maybe."

That was pretty much the entire design brief on Wyrd.

We'd been talking about USB as an audio interface, and the increasing impact of USB port power management on USB-powered DACs and USB interfaces, and I'd just had a random thought about a little box to both power and re-clock the USB datastream. I mentioned that idea, pretty much in the same sentence as described above, to Dave as we were on our way out of the Schiitbox.

"So we might be able to do something like a TLC for USB?" I asked.

Mike groaned. “We have enough to do.”

“I know, but, like, for when we get bored (or frustrated, as in the above chapter.)”

“It’s USB. It’s like turd-polishing.” Mike said.

“Mike, in case you haven’t noticed, I’ve been using USB input at shows lately. Our USB Gen 2 input is really, really good.”

Mike grumbled something under his breath. “We still have too much to do.” I sighed. I knew. Ragnarok still wasn’t running right, even after several firmware revisions. I had the melted-into-the-fiberglass-insulator MOSFETS to prove it. And a USB reclocker/clean power-er wasn’t something I could do (or should do.) I had plenty of analog projects to work on. *But what’s the harm in talking about silly ideas,* I thought.

And that, I figured, was that. We’d drop the idea, and pick it up again later, if we had time.

Except Dave.

“Just Listen To It”

About a week after our short conversation about a USB widget, Dave showed up at my place for a barbecue. Mike was coming as well, but he hadn’t

arrived yet. Dave was carrying a small box of stuff. I didn't think much about it, because Mike and Dave are always swapping various weird digital things that I don't know anything about.

But when we were inside, Dave said, "About that USB thing you were talking about … "

"Yeah?" *Maybe we would get somewhere on this,* I thought.

"I made one," Dave said, pulling a small green board out of the box. I stared at it, then blinked a few times. *It wasn't possible he'd already designed and built the one-line "what if" product I'd just been talking about,* I thought.

"What do you mean?" I asked.

"I built one. A low-noise USB power supply and hub." Dave pushed the board at me. I took it and looked at it. The board had Dave's signature hacks that he did as he tested and optimized a new design—a ground lug soldered on, a few capacitors tacked here and there. Trailing off of it was a hacked old-skool gray USB cable, with a couple of capacitors and inductors hanging off of it.

"Does this work?" I asked Dave.

Dave nodded, then shook his head. "Yeah. Mostly. I mean, I tried it on a couple of machines, but I don't know if it's really meeting USB power spec, or if it works on everything—" I waved a hand. With Dave, nothing ever "just works." There are always caveats. If it wasn't for Dave, we'd have a lot more undiscovered problems. As Mike says, he's the "trees" guy. We're the "forest." But, as a consequence, few things are ever "done" for Dave.

"So I can plug it in and try it?" I asked. "Right now?"

Dave nodded. "Sure."

"Does it sound better?" I asked.

Dave broke into a big grin. "Just listen to it."

Okay, big aside here, before I am eviscerated by the objectivists:
I know there is no sane, rational way in which a USB repeater/clean power device can scientifically affect the sound of a system. USB, like all digital interfaces, is, in the end, digital. The packets either get there or they don't. If they don't, they are resent, if possible. If they can't be received in time for a streaming application like audio, you get glitches and disconnections.

That should be it. Period, no exceptions. So clean power and shapely waveforms should, theoretically, make no difference.

And another big aside, for the subjectivists: Yes, I also know that USB, like any other digital transmission system, is essentially using imperfect analog waveforms to transmit very high-rate data (in the case of USB 2.0, theoretically 480 Mbit/s, or about 300× what you need to transmit stereo 16/44 content, though overhead usually makes the actual rate much lower). But the fact is, it should not matter in the least how well-shaped or noise-free the waveform is—the packet is either 100 % recovered, or it isn't. If it wasn't, you couldn't simply copy several hundred gigabytes of data to an external drive reliably.

And a final big aside, for everyone: If insanity is defined by the ability to hold two completely contradictory ideas in your head and accept both of them, then me, Mike, and Dave are all completely bonkers. Because we run into this same objective/subjective thing all the time.

"You're not telling me it sounds better," I asked Dave.

“Just listen to it.”

“But, there’s nothing that should … I mean, how would it change anything?”

Dave looked sheepish and shook his head. “Just listen to it.”

Okay. Fine. I hooked it up between my Macbook and the Gungnir/Mjolnir stack I use as a main reference system. But, before I fired it up, I dangled the weird hacked gray cable in front of Dave. “What’s this?”

Dave laughed. “Oh, that. That’s just an idea. I saw that some people are separating power and data in their cables. I wondered about filtering the power instead. Maybe try that first. No USB board.”

“Does it sound better?”

Dave was deadpan. “Just listen to it.” I groaned. “Now I know you gotta be screwing with me.”

“Just listen.”

Fine. Might as well give it a true test. I unplugged the USB board and went back to the full-stock system (no reclocker/power supply, no hacked cable) for a quick listen. It sounded, well, like I expected it to.

Next, I used the hacked cable. I went through several tracks I use to get a subjective gauge on the sound, trying to hear a difference. If there was any, it was really too slight to notice.

"I don't hear it," I told Dave. "Sounds about the same."

Dave nodded. "That's what I thought. But the level of filtering you can get with that approach is really low. If you have a couple hundred millivolts of noise on your USB port power, this might cut it down by half. Maybe. You can't put too much resistance in the line."

"But the board is better?"

That big grin again. "Just listen to it."

Fine. Now, I put the board in-line between my computer and the Gungnir and sparked it up again. If there was any difference in sound, I figured it would be in the textbook, "Sharper, more resolution" direction. But more likely, there wouldn't be any difference at all. At least not in theory. I clicked on the first track.

And sat back, a little shocked. Because I could swear I heard a difference—and not in the expected "sharper, brighter" direction. If anything, it sounded a bit smoother and more natural. At

the same time, the noise floor seemed lower, like I could hear deeper into the track. I looked up at Dave, my mouth open. Dave immediately doubled over, laughing.

Nah. Not possible. It was a trick. I was fooling myself. Or Dave had somehow put real-time processing into the USB stream (ha, ha.)

Still, as I went through the tracks, the difference was consistent—and consistently positive. Dave kept laughing as I shook my head again and again.

I had to be fooling myself. I unplugged the USB board and went back to just a straight old USB cable. And suddenly it sounded a bit flatter, more one-dimensional. Dave kept laughing. I switched back, and it got better again.

"How does it do that?" I asked Dave. Thinking, at the same time, *Nobody's gonna believe this.*

Dave shrugged and shook his head. The same thing he does whenever confronted with something that goes beyond logic. This is a guy who is much more analytical, and much more structured, than Mike and I combined. This is a guy who has designed uber-high-performance systems that deliver floor-of-the-analyzer performance,

together with ultra-high complexity, for various systems far more ambitious than anything Schiit has produced to date. Of course, this is also a guy who has worked in an environment where they used RF to seal thick plastic, with water-cooled tube outputs and stray field that would light fluorescent lights in open air several yards away from the machine. So maybe that explains some of it.

"But … why?" I persisted.

Dave shrugged. "Lots of reasons. The power supply is super-low-noise. It's using LM723 regulators, which are rated 4.17 µV of noise, plus organic polymer capacitors, plus I'm using a tight crystal for the USB repeater. Lots of little things."

"But it still shouldn't matter," I said.

Dave nodded. "Right."

"But it does something."

"Right." I sighed. *Okay. Run with it.* "So this is, like, done?"

Dave shook his head. "It won't output the full 500 mA USB spec, because you have a 16 V AC wall-wart and it's only rated for 500 mA. We need something more like 6 V AC, maybe 1 A or more."

"We could do 1500 mA in that same package at 6 V AC," I told Dave.

"That would work." A sudden thought hit me. What if someone plugged a 16 V AC wall-wart into this product? That might be a bad day. I shrugged. We could make the pin bigger, so you couldn't use the 16 V AC transformer on this product. The 6 V AC version could plug into the 16 V AC product, but that wouldn't hurt—they simply wouldn't work.

And then I realized … *I'm thinking about this like a product. Like we were already going to sell it.*

But, why not? Dave had already used the form-factor of our small products—the board was laid out to fit in a Magni-sized box. All we needed was a different transformer. And it would solve the USB port power management problem.

And it would sound better.

Which made me remember my first thought: *nobody will believe this.*

Argh.

Truth and Marketing

If Schiit was a purely subjective company, the dilemma of having a product that made stuff sound better, without having any rational explanation as to why, wouldn't be a dilemma at all. We'd wrap it up in nice flowery language, throw in some pseudo-meaningful charts that showed the difference in power supply noise levels, and call it a day.

If Schiit was a purely objective company, the dilemma might not be a dilemma at all. Because we might have convinced ourselves that, even though there was a difference, there really was no difference, and so why bother making something that didn't make a difference?

But as a company that uses both objective measurement and subjective listening, it's not so clear. We could do the pure subjective thing with the words about how you're transported in space and time to a wonderful world where unicorns dance and crap like that. Sure. We could.

But that isn't us.

And that isn't honest. Because, you know what? We're really talking about small differences here.

It might not be important to a lot of people. It can be easily dismissed.

But for other listeners, it might be big enough to be significant.

So how to market it?

"It does solve real problems," Mike said, when I was angsting to him. "I have a laptop with power so bad it's noisy—as in, you hear the noise even through a Modi and Magni. This kills it dead."

"And it does make those weak USB ports usable," Dave said.

"So does a powered hub," I countered.

Dave shrugged. "And it sounds better." I groaned. I really didn't want to get into the "hey, trust us, it sounds better," theme. Especially if people thought we were trying to sell a $ 99 add-on to a $ 99 DAC like Modi.

Aside: that is pretty silly, when you think about it. And less silly, when you think about it some more. Think of the Wyrd as the "separate outboard power supply" for the Modi.

Another aside: it's still pretty silly.

“And what are we going to call it?” Mike asked. I groaned again. That was another question. What was the name? I hadn’t figured that out yet. And that was strange. We usually have the name—and the descriptive copy—done long before the product is heading towards production. And, as we were having this conversation, the USB board was on its second Rev and considered production-ready.

Name. And copy. Usually not a problem. But in this case, it was.

Wyrd Schiit, and Wyrder Marketing

Inspiration struck on the name when I was searching for terms relating to the Norse concepts of “time” or “stability.” This led me to Wyrd, which is a term used for the complex interconnecting web that binds all things, living or otherwise. It is so intricate, it is much more powerful than the concept of “fate.”

In fact, the word “weird” is derived from “wyrd.” And when you pronounce it that way, everything falls into place: the weird fact that it does seem to make a difference, even when it should not. The complex interconnectivity of USB. And the connection to Yggdrasil and other Norse concepts.

Cool. So we had a name.

What about the marketing? Well, after some more thought, I decided that the best way to portray Wyrd was as a product that solved real USB problems, and stay away from any subjective claims at all.

"But it makes a difference," Dave said. I laughed. Dave virtually lives with an Audio Precision SYS-2722 and one of our Stanfords. To hear him say that, with no real reason why it should be so, is very funny to me.[1]

"I know. But let's let everyone decide for themselves."

"Ohhh … kay," Dave said, doubtfully.

And that's what we did. We first showed Wyrd at TheShow Newport in final production form. At set-up, we ran into a problem with a new, low-cost Windows 8 laptop—one that had not one single port that could run a Modi. Wyrd fixed that right up. A perfect application for it.

---

1 Actually, there have now been measurements of lower jitter with Wyrd from independent sources, but we're still talking about levels that shouldn't matter—which is why we never published our own.

Wyrd was one of those nice, uneventful launches—the kind you want to have.

Aside: I'll probably have to add another chapter here sometime about all the pain associated with building Ragnaroks, and the reality of having our first real production-line product in house (to date, everything else is a simple build—drop a board in a box, done. Ragnarok's build sheet is several pages long, single-spaced, and has 6 separate productions stations involved. A few weeks after the show, we started shipping Wyrd to little fanfare. It's a good-selling product, though, probably because it's still the only product we know of that both cleans up the USB power, and repeats the signal with a precision crystal oscillator—and does it for half the cost of other products that are simply power supplies.

And yeah, I still use one with the Gungnir that I have at home.

Annnnndddd … *not* because I have noise, glitches, or USB port power problems.

Chapter 39

Unto the Second Generation

When you start updating your products, brace yourself for the questions. This is a statement that is probably the same in every industry. When Apple goes from IOS7 to 8, people start wondering where OS11 is on the Mac. No. Wait. It's 10.10 now. More proof that someone needs a swift kick in the ass there … as they build their spaceship palace.

But, to get back on track, update one thing, and everyone will want to know when the rest of the line gets refreshed.

In our case, the inquiries started as soon as we shipped Asgard 2.

"So when do Valhalla 2 and Lyr 2 show up?" "Is there a Valhalla 2 in the pipeline?" "Hey, you should add the gain switch to Lyr and call it the Lyr 2."

Yes, we know, I wanted to say.

But of course I couldn't. None of us could. Of course, we knew that a Valhalla 2 and Lyr 2 would be great updates for the line. And, of course, I'd already been planning some of the things we could do for the updates. But there were no prototypes of Lyr 2 or Valhalla 2 at the time Asgard 2 was launched.

And even if we had prototypes, *there was still nothing we could say.* Because once you start selling products, you can't hint they'll be replaced. Not soon. Not someday. Not tomorrow. Not in the distant future. Not at all. Never.

Read that again. *You can't say anything. Period.*

Because the moment you do, you kill the sales of the current line. If we'd had a Lyr 2 and Valhalla 2 in the pipeline, and we'd told people about them, then we wouldn't be able to move Lyrs and Valhallas with boxcars of Ex-Lax.

Yes, I know. This sounds selfish. This sounds unfair. *What if I just bought a Lyr and there's now all of a sudden a Lyr 2?* you might be thinking.

And yes, there's the chance of being caught buying a soon-to-be-discontinued product right before the new one hits. It's unfortunate, but let's be brutally honest:

1. At least the company will most likely be around in the future to service your product, rather than going out of business by having to take write-downs of unsold products they couldn't move after blabbing that the new version was coming.
2. It's not like the new product completely invalidates the old one. If you liked the old one enough to buy it at its current price, that indicates that it provided sufficient value to earn your money.
3. Most companies will offer you a discount on the old product if you purchased it right before the new one was announced, or the option to pay shipping to get the new one.

Yes, this is from the business perspective. Remember, this is a business book. And once the business is past the first scrappy start-up years, it becomes a different business. Once your business is an ongoing concern, discussing new products becomes an even bigger problem than when you were starting up.

Let's turn that up: discussing new products before they're ready might *kill your business.*

Read that again, too: once you're running, the best thing you can do for yourself, your employees,

and your customers is to ensure that the *business keeps running*.

Talking about next-gen products that could kill current sales is not a way to keep the business running. Talking about new products that could negatively impact current sales is not a way to keep the business running.

But, on the other hand, *keeping the business running* doesn't mean *growth at all costs*.

This is a mistake many businesses make. They think, "We have to keep running, we have to keep growing, we have to keep expanding." These are not equivalent thoughts. "Keep running" could easily mean steady-state. It could even mean minor to moderate shrinkage to deal with a downturn in the economy.

"Keep growing, expanding," is very different. Growth and expansion without any restraint is the easiest way to get yourself in big, company-shattering trouble. Mike Moffat likens a business to a plate-spinning act, making sure you keep all the plates in the air. If you work really hard, and you're very good, you can keep quite a few plates in the air and make it seem pretty effortless.

But add one more plate …

Or add a dozen more …

Boom. Everything comes down. There's no time to keep them spinning. All you can do is to watch everything fall.

This is one reason that, even as Schiit continues its strong growth, we have natural restraints in place—self-funding strictly through sales, no hiring forward, long product development times with long private betas (I'm sitting here listening to one right now), no artificial offers/sales to boost results in the short term, etc.

Because the most important thing is to keep it running. To me, it's not plate-spinning. It's a marathon.

Murphy Was A Butthead. But He Was Right.

Okay, two more doses of business reality before we get to the whys and wherefores of Valhalla 2 and Lyr 2. Because I'm sure some of you are asking, "If you're doing so well, why couldn't you announce all the second-gen products at once? Hire some engineering talent and get 'er done! You'd be that much closer to the revenue from the new products."

There are two reasons you don't plan to launch multiple updated products at once:

1. **Known-good, continuously-available engineering resources are limited.** Beyond that, you're using hired guns that may or may not understand what you're trying to do, and who may or may not be available when you need them. At Schiit, the KG-CA-ER is me, Mike, and Dave. These are the three people I can count on to:
   a) Understand what a great audio product is
   b) Give a schiit—really—about the design and performance
   c) Work through the initial prototype bugs without turning it into a completely different product
   d) Be there when the inevitable production bugs surface—see #2
2. **Going into production is an exercise in living Murphy's Law.**
   It doesn't matter how many prototypes you've done, how many dozens of months the product has been running flawlessly, or how small of an incremental change it is—there will be pain. Pain like:
   a) A $0.008 resistor is unavailable in the value and size you need—everywhere

b) A tsunami wiped out the capacitor plant, and lead times are now 18 weeks
c) The distributor sent you the right part—in the wrong size
d) The final Rev to the PCB moved the exact wrong trace, and the first articles don't meet spec
e) The metal house built the wrong metal revision
f) Chassis are late
g) Boards are late
h) Chassis are bad
i) Transformers hum
j) People are sick on the line
k) A new test program needs to be written
l) A new test fixture needs to be built
m) For some reason, they just don't work like the prototype
n) And at least 1000 other things

So, let's imagine a multi-launch scenario with contracted engineering. Let's say, three new products. Each has a half a dozen line problems. Your contract engineering crew is busy working on new paying jobs, so they're not available. You haven't been too involved in the design, so you don't know it inside and out. The lines are stopped. Nothing is being produced. The boardhouse is hard stopped because you can't

tell them what's wrong yet. You're about to lose your dedicated line there. You need to make the chassis drawing changes, but you just want to get something—anything—working.

Fun times, huh?

Updating the Classic Tube Amp: Valhalla 2

Tube amps are great. That is, if you have high-impedance headphones. 300 Ω Sennheisers and 600 Ω Beyers loved the original Valhalla.

Low-impedance headphones? Not so much. Some people liked the original Valhalla with the AKG K701s, which were 64 Ω, and some liked it with Grados (yeah, I know, weird.) But, in general, we recommended it for high-impedance headphones, where it could produce very good results.

The fact was that the original Valhalla was not well-suited for low-impedance headphones. It would quickly run out of power and distort heavily into low-impedance loads.

Why? Because it's an OTL (output transformer-less) tube amp. OTL tube amps can't source much current, and typically have high output

impedance. High output impedance into low-impedance headphones isn't ideal. Nor is low current capability. Valhalla, in its original form, had the following specs:

Maximum voltage output: 30 V PP (into high-impedance loads)
Maximum current output: 32 mA
Output impedance: 28 Ω

Not very impressive. Especially in the face of solid-state amps that can source 250 mA to 1000 mA of current without even breaking a sweat, and have 1 Ω or less output impedance.

But we still had customers who really, really wanted to use a "pure tube" amp to run low-impedance headphones. So, part of the design brief for Valhalla 2 was to make it more suitable to a wider range of headphones.

But how?

There are three ways to get better current capability and lower output impedance from a tube design:

1. **Use an output transformer.** The transformer provides a higher impedance load that the tube likes to see, and at the same time provides a lower output impedance to the headphone.

The big problem with Valhalla is that this would also make it a much larger, more expensive amp. No dice.

2. **Use a solid-state output stage.** Make it a hybrid, in other words. This was also a no-go, since we already had Lyr, and customers were specifically asking for a pure tube amp.
3. **Tweak the output stage as much as possible, and add feedback.** This offers smaller gains than the previous two techniques, especially in the case of current output. But this was the approach we took.

Aside: feedback? Eeeeeevillll feedback? Yes. We added feedback to Valhalla 2, both in high-gain and low-gain modes. But we never said feedback was evil. We have said that we prefer not to use it when possible. But, in this case, it worked out rather well. Read on.

First, though, let's cover the tweaks to the output stage.

Valhalla has always used a "white cathode follower" output stage with 6N6P tubes, to maximize its possible current delivery. White cathode followers allow us to double the available current into a specified load, and the 6N6Ps are some very amazing NOS tubes that dissipate 8 W on

the plate and tolerate very high standing currents. There's nothing like them in new production, except the JJ ECC99, which has lower plate dissipation and a slightly different pinout.

What we did for Valhalla 2 was simple: we increased the current through the output stage and re-tuned the White cathode follower for lower impedance loads. The end result was almost twice the current capability into low-impedance loads, without affecting performance into higher impedance loads (in fact, performance increased across the board.

And then, yes, we added the eeeevil feedback. I started with a no-feedback in high-gain mode, then did comparative listening tests to come up with what I thought sounded best. Then we added even more feedback to create the low-gain mode. In both cases, feedback is minimal (about 6 dB in high gain, 16 dB in low gain.)

"Wait, wait, wait!" some of you are yelling. "What about re-entrant distortion? Don't you want to use a lot of feedback, if you're going to use it at all?"

No. Not if the basic gain stage is very linear—which is certainly the case with Valhalla 2. In fact, we made one other change to specifically

enhance linearity—adding a current source to the front end differential amplifier, which required a different power supply. It also allowed us to eliminate the input coupling capacitors. And yes, we added even more eeeevill parts, in this case transistors for the current source. But this change reduced the distortion of Valhalla 2 by a factor of 5. And, with different plate loads for more overall gain, Valhalla was now a very, very linear amplifier, running below 0.02 % THD open-loop into high-impedance loads. After feedback, the THD was even lower.

The results? Valhalla 2 now measured like this:

Maximum voltage output: 60 V PP (into high-impedance loads)
Maximum current output: 60 mA
Output impedance: 14 Ω in high gain, 3.5 Ω in low gain

Not exactly Charles Atlas, no, but suddenly Valhalla 2 was able to run a much wider range of headphones. High-efficiency, low-impedance headphones could easily be run in low gain mode. And, at the same time, it sounded even better into high-impedance headphones. A home run? No. It's still not our first choice for planars. It's less capable than any other amp we make (including

the upcoming Fulla) into low-impedance loads. But for high-impedance headphones, it's a great-sounding amp … and it will now comfortably run your Grados.

Valhalla 2 was also notable in that Murphy pretty much bypassed its run into production. Two prototype cycles, revised metal, new boards, and a new transformer, and it was on the line. The transformers didn't hum, the boards fit as expected, the metal wasn't late, and everything worked as expected.

Guys, remember this: you don't get many of these free passes. Say thank you … and then start looking behind you for the knife coming at your back.

Taming the Beast: Lyr 2

If the original Valhalla was a flyweight, the original Lyr was a musclebound brute. Big, powerful, gutsy … but unable to keep from speaking in less than a shout, and apt to break the china in polite company.

Yes, it was powerful, and yes, it was a great match for power-hungry orthodynamics … but pair it up with efficient headphones, and you could hear the

noise floor. Forget IEMs entirely. Lyr delivered the goods in the power department, and offered many, many opportunities to tune the sound via tube rolling … but it was not the last word in refinement.

So, like Valhalla 2, one of the goals for Lyr 2 was making it more suitable for a broader range of headphones.

Aside: yes, a common theme. It makes sense. You should be able to use your headphone amp for the broadest range of headphones possible, rather than having individual amps tuned for a specific headphone. So it's not surprising we set that as a goal.

But Lyr's problems were different than Valhalla—and, as a hybrid amp, it offered us more opportunities to change.

While the simple Valhalla 2 still uses through-hole parts, it was clear from the start that Lyr 2 would have to go to surface mount.

Aside: through-hole parts are parts that, well, go through holes on the PC board. This is an older way to make PC boards. (And, believe it or not, when companies started mass-producing

color TVs, we were still in the pre-PCB era—they were hand-wired. Yes. Scary.)

Now, there's nothing inherently wrong with through-hole PC boards, except three things:

1. Through-hole parts are usually bigger, so you can't fit as much on the board.
2. It's very hard, if not impossible, to automate through-hole PCB assembly, so you're looking at a bunch of people stuffing the parts individually into the boards. The upshot being is that they are more expensive to make.
3. Through-hole parts are on the wane. Every year, parts selection is more limited, and the interesting parts come out in surface-mount only.

Because of this, surface-mount is the way we're going with most of our new products, and Lyr 2 was not going to be an exception. The original Lyr was already packed completely full of parts. If we wanted to stuff more refinement and capability into it, the parts would have to get smaller.

The original Lyr's noise problem had many sources. Most were directly related to the lack of space on the board. Lack of space meant we had exactly zero regulated power supplies in the original Lyr. The 200 V rail for the tube, the

−25 V rail for the front end, the servo supplies, even the original 6.3 V tube heater was AC.

All of these supplies were smoothed as much as possible, but they weren't as good as regulated supplies. So, noise made it through to the output. For low-efficiency headphones like orthodynamics, it was inaudible. For high-efficiency headphones, it could be problematic.

With Lyr 2, we attacked the regulation problem with full force. The 200 V supply got regulated. The ±15 V supplies for the servos and front ends got precision low-noise regulation as well. Even the tube heaters got separate, dual-mono low-dropout DC regulation, rather than AC. These changes alone took the original Lyr's noise down by a couple orders of magnitude. But these changes weren't simple—as with Valhalla 2, they required an entirely new transformer design to implement.

But we didn't stop there. The input stage got a working on, as well, to decrease the plate load and increase the current, for a slightly more linear operating point on the 6BZ7 tubes we're using.

Then it was on to the output stage. That's always been the key to the Lyr design—the Dynamically Adaptive stage. It's really the first Schiit-specific

circuit I devised, so it's really a special thing. At the same time, though, it's also a source of questions and uncertainty. As a unique output stage, it isn't well-known, tested or busted to death a million times, characterized a thousand different ways in hundreds of different layouts with different analyzers. The vast majority of audio engineers are terrified to venture so far out into the weeds—much safer to do something like a Blameless design, or just throw an integrated chip amp at it.

So, I had to wonder: is there a better way to do the Lyr's Dynamically Adaptive stage? It was easy enough to tweak the overall design and optimize its operating points, but was it really better than, say, a conventional Class AB output stage? Or even a simple diamond buffer? Or something more exotic, like a triangle buffer or error-correction stage? I spent the vast majority of dev time on Lyr 2 trying these different output stage topologies … and, in the end, coming back to the original Dynamically Adaptive stage. Every other candidate had significant downsides—not able to swing the rails, crossover distortion, etc. I went back to a Dynamically Adaptive stage, but with one significant tweak that improved performance at higher frequencies.

Then it was time for the eeevil feedback. Yes, again. Lyr 2, like Valhalla 2, uses some feedback even at high gain, and more at low gain. The reason is also the same: because it sounds better. As with Valhalla 2, Lyr 2 is an inherently very linear circuit, so low amounts of feedback don't result in appreciable re-entrant distortion.

Add a gain switch, and there you have it: Lyr 2. A sum of many different changes, together with a multi-month excursion into alternate output stages.

Going into production wasn't too painful, either—other than a couple of hard-to-get parts, and a last-minute revision on one of the transformers, Lyr 2 went without a hitch.

Was this a sign we were getting better at this whole engineering thing?

No. Never think that. Run from engineers who say they're good. Start looking up at the sky so you can dodge the lightning bolts if they say they're infallible. Lyr 2's painless launch was just another patch of good luck.

And … it didn't matter anyway …

The Reality: Simultaneous Launch

For all my yammering about concentrating on one thing at a time, the grand irony was that Valhalla 2 and Lyr 2 launched at exactly the same time (together with a pre-production Wyrd, Mani, and Ragnarok, at TheShow Newport 2014.)

Now, to be fair, Lyr 2 and Valhalla 2 were developed separately, so we were following our own advice on that count. But the problem was simple: we realized that if we introduced one, the other was obvious. Introduce Valhalla 2, and everyone will know a Lyr 2 is coming. It would have the same effect as going ahead and announcing it.

So, since both products were working, finished, and on the shelf, we decided to launch them together.

Note those three caveats: *working, finished, and on the shelf.* In fact, we had shelves full of Valhalla 2s for some months before we actually announced. Lyr 2s took a little longer, but they were also available in quantity well before we announced.

Which might make you wonder: *what would I see on the shelf if I were to visit right now?*

Are there other simultaneous launches coming?

What's next?

Sorry, I really can't say.

Chapter 40

Schiit Goes Vinyl

Vinyl? Yes, vinyl. As in those round spinny plastic things that you scratch a hunk of rock across to make music, treat with kid gloves in a HEPA-filtered environment to keep some of the inevitable pops at bay, and go mental with tracking force gauges and alignment diagrams and VTA adjustment to get a few hundred more hours on a stylus.

Sounds sexy, doesn't it?

So why, in this futuristic year of almost-2015, when we should be shopping for flying cars and taking vacations on the moon, do people still mess with this crazy, neurotic format?

And why, in almost-2015, are vinyl records the only physical medium to show any growth in sales? And not just tepid growth … strong growth?

It's funny. I probably shouldn't be writing this chapter at all, because I'm really, at the core, a digital guy. A convenience guy. I rejoiced when

I first went CD (from cassettes, yeah … I was a car audio guy before that), I was thrilled when I was able to put all those CDs on a hard drive, and I was overjoyed when Tidal showed up on the scene with uncompressed streaming.

So blame it all on Mike.

Yes, Mike, the father of "digital audio done right."

Mike Comes Out of the Closet?

Well, not exactly. Mike has always been into analog from the very beginning—his first products included a tube phono preamp (the first to use 6DJ8 tubes). If you've been following his side of the story here, you've heard him say that his first goal with digital was to get it sounding as good as analog.

> **Aside:** those of you who have grown up with relatively mature digital audio may think this means "going backwards to soft, distorted, noisy analog sound," but the reality is very far from that. Early digital had big problems. It really did sound like butt.

And, Mike has always maintained a record collection and has almost always had a turntable

around. So, it's not surprising he'd eventually turn back to analog.

What is surprising is how he did it.

One day, out of the blue, Mike showed up carrying a small green board, the same size as what we use in Modi and Magni. I didn't pay much attention to it at first, thinking it was just another DAC variation he was playing with. When I finally noticed it and asked what it was, Mike laughed.

"It's a phono stage." I blinked, thinking I hadn't heard him right. A phono stage? Like for records?"

Mike laughed. "Like exactly for records."

"Vinyl?"

Mike nodded, even more amused at my flabbergasted reaction. "Vinyl."

"For us?" I was still trying to wrap my mind around it.

Mike crossed his arms. "I knew you'd have some trouble wrapping your head around it. But believe it or not, there was analog music before digital, and no matter what they say, all music starts out as analog."

“Except techno,” I said, unable to think of anything else to say.

Mike made an expression like he’d just bit into a lemon. “That’s not music. That’s someone standing up on stage and pressing ‘play.’”

“I think some people would argue you on that one,” I told Mike.

“Whatever. It’s experiential. Not performance-based,” Mike said, dismissing it as easily as you can shake a cane and say, ‘Get off my lawn!’

“But … a phono preamp?” I said, struggling back to the subject. “It must be an inexpensive one, if you’re using the Magni/Modi size.”

“Yep!” Mike said, grinning. “Under a hundred and fifty bucks, I expect.”

“Discrete?”

Mike shook his head. “No. Ultra low-noise op-amps.”

“Op-amps,” I said, curling my lip.

“I know, I know, you can do better,” Mike said. “But can you do better for a hundred and fifty bucks retail?”

“I—” I began, but Mike cut me off.

"And don't answer until you've heard it." I sighed. *Okay. Fine.* I was playing with some decent op-amps for what would eventually become Fulla. I was willing to admit that, in some cost-constrained cases, op-amps could be OK. But I was also working on an all-new discrete, inherently balanced topology that was unlike anything else out there, and it was *insane*.

Aside: Cease the heavy breathing about possible new products. All-new, never-before-done stuff takes a lot of baking time. It'll be done when it's done. That's all I can say for now.

I sighed. "So let's hear it."

But, for some reason, we didn't hook it up that evening, and we both ended up getting distracted. It wasn't until a few weeks later that I went up to Mike's house for a prototype-swapping session that I actually heard what was going to become the Mani.

And before I get into that, let's talk about the whys and wherefores of a phono preamp.

It's the Curve, Stupid

Well, actually not entirely. It's the gain and the curve. Doing both at low noise is what's important for a RIAA phono preamp.

But I'm getting ahead of myself. Let's first talk about why we're spoiled in this modern age, and then let's get into details about what a phono preamp actually does.

Why are we spoiled today?

1. Pretty much every digital component on the planet puts out "line level" outputs, which is typically 2 V RMS for full-scale output from RCA jacks.
2. Pretty much anything that isn't line-level is close. For example, a smartphone may only put out 1 V RMS full-scale. But this is only 6 dB less than the typical output. Not a deal-breaker.
3. Pretty much everything has frequency response that is flat from 20 Hz to 20 kHz … and frequently well beyond.

The upshot? Pretty much everything is predictable, consistent, and easy to integrate with other components. Nothing requires huge amplification or frequency shaping. Take a line-level

signal, run it through an amp with 20 dB to 30 dB gain, and you're at speaker output levels. Done.

Why is vinyl different?

1. Depending on the cartridge, you're looking at signals that range from the microvolt levels to handfuls of millivolts—40 dB to 60 dB below line level, or, in more understandable terms, 100× to 1000× smaller than that nice standard 2 V RMS. Remember, this is a signal formed by running a diamond through a tiny groove and having it generate electricity via wiggling magnets and coils.
2. Different cartridges can be wildly different. Deccas are very high output. They may need only 30 dB of gain. Some inefficient moving-coil designs could demand up to 60 dB. Want more fun? Some perform better at different cartridge loads.
3. And, to make it more complicated, all phono preamps have to accurately apply an "RIAA curve" to the output … essentially a filter that boosts bass and cuts treble, to compensate for the way the records were cut (with less bass to reduce the chance of tracking errors, and more treble so that the RIAA filter would act as natural noise-reduction.)

The upshot of this is that designing a low-noise

amplifier with 100× to 1000× gain is difficult … and creating an accurate RIAA filter is no easy feat, either.

So why bother with the demands of a format that was engineered over half a century ago to sound its best with the archaic technology of the time?

Because, when done right, vinyl can sound very, very good.

Yes, there, I said it.

It's Missing Lots of Pieces

When I finally went up to Mike's house to hear what would become Mani, Rina came along. This is important because there are some times I think she knows it all … and this was going to be one of them.

Now, I have to preface this by saying that this was the first time I'd heard analog since the Sumo and Theta days. I was only too happy to leave vacuum-platter turntables, record washing machines, VTA alignment gauges, ground wires, and all those other neuroses in the 20th century, where I thought they belonged.

But eventually, Mike had it cued up (using one of his Ariston turntables, I believe, and probably a Nagaoka cartridge) and handed me the headphones. I put them on, trying not to roll my eyes. I guess I expected old-style, rolled-off, fluttery, distorted sound—perhaps I'd been partially assimilated by the objectivist Borg.

And yeah, the first thing I heard was the subtle crackle and pop of groove noise. *Yeah, just like I remembered,* I thought …

And then the music started. My eyes opened up. Because this wasn't mushy, rolled, distorted, or otherwise clearly inferior to digital. It sounded very good, clean, and dynamic. And … it had a real sense of weight and space, something I hadn't heard with a DAC and the same headphones. Something, something almost indefinable, was more real, more alive, more *right*.

Mike laughed as I shook my head and handed the headphones over to Rina. "You have to hear this," I told her.

"I don't need to," she said. "I know what vinyl sounds like."

"Just listen," I said, pushing the headphones at her.

Rina sighed, took the headphones, and put them on. She closed her eyes for a while, then opened them and shrugged. "Yep. Analog."

"But doesn't it sound great?" I asked.

Rina laughed. "Of course. It's analog. Did you guys forget what analog sounds like? Or did you forget I used to do live sound? Analog always sounds better. Digital has all those pieces missing."

Mike and I looked at each other. I swear, in that moment, I could pretty much hear his thoughts, because I was thinking the same thing.

Well, it's not quite accurate that there are pieces missing, because Nyquist … Mike and I were thinking.

And our jaws both dropped open at the same point, because we both had the same epiphany at the same time: *… but Nyquist relies on an ideal brickwall filter, and it doesn't take into account quantization error … so saying "It's missing a bunch of pieces" wasn't so inaccurate after all.* I tried for a while longer to get Rina to admit that we'd just heard something special, but she refused to me moved, like a picky dog that really doesn't want to eat what you're trying to feed it.

“Why is this so amazing to you guys?” was all she’d say. “Of course it’s better. It’s analog. It’s the whole thing, not bits.”

Sure. Fine. Whatever. I wanted to remind her how archaic, how truly stone-age the technology was I mean, really, dragging a diamond over plastic, come on, we have iPhones now,) and how much signal processing it had to go through, and how that it was completely dependent on the engineer who’d cut the record, and how it’d wear out eventually, and how the stylus would wear out eventually, and how we’d all go insane trying to tweak the last 1% of performance out of it.

But you know what? I gotta admit … it sounded *great.*

And that’s how we decided to do a phono preamp.

The Most Boring Production Story, Ever

We decided to do a phono preamp, did two prototypes, and then put it into production, and began shipping in September 2014.

Yes, it was about that simple, and about that boring. The only real delay came in locating some of the hard-to-find precision capacitors for the RIAA filter. Mike was insanely specific

about those, because an accurate RIAA curve really is 95 % of the difference between phono preamps … and Mani has one of the most accurate curves out there.

Aside: Of course, he was also insanely specific about many other things, like topology (a fully passive RIAA network), the power supply, the low-noise op-amps, etc … made even more challenging because this was intended to be a fully surface-mount product for low cost. I'm really short-changing the design here, because I wasn't involved very much with this one. Hopefully Mike will chime in with some details.

And, of course, there were a couple of other delays, too … late metal, and difficulty finding a decent ground post that didn't cost stupid amounts of money.

But other than that, Mani was boring as can be. And, like I said before, that's the kind of boring you like. Because "interesting" product launches are usually only interesting in very, very bad ways.

And, after several easy, flawless launches in 2014, Mani was beginning to seem almost, well, normal. Like we knew what we were doing. Like it wasn't just luck.

But then there was Fulla, but I'll get to that next week.

So Vinyl Earns Another Convert?

No. Sorry. I'm still really, really lazy. I'll enjoy a good record at Mike's place, but it's mainly files and Tidal at home. If this makes me a terrible audiophile, so be it. Feel safe and happy that Mike is heading up the vinyl side of things.

So, yeah, I'll take my bits, please. Even if it is bits, and not the whole thing.

But I'm pleased to say Mani is selling briskly. In fact, we're currently out of stock for a few days as the next run ramps up. Early indication is that it's a winner—but that's up for our customers to decide.

Business lessons and Alternate Histories

So, what business lessons can we learn from this?

That being the part of a revival of a dead technology can be interesting and fun? Sure. Mike did that before. He's one of the first to bring tubes back from the dead, at the height of the transistor era.

But will this current interest in vinyl last? I have no idea. It might fade away once again. Or it might get even bigger. It really comes down to the type of sound you like, and the amount of time and inconvenience that you're willing to put up with in order to get it. I will say that both analog and digital can sound amazing, but I suspect I'll be happier with an Yggdrasil in the end. But I am very lazy, as I've said before. I think the main business lesson comes down to: be willing to be surprised by what an "outmoded" or "outdated" technology is capable of.

What do I mean by that?

Well, it's kinda like this. Today, we look back on the first Apollo landing on the moon in 1969 and think, "Yeah, that's pretty cool." But I think that future historians will look back and say something more like, "HOLY FRIGGIN CRAP THESE GUYS WENT TO THE DAMN MOON WITH SLIDE RULES AND VACUUM TUBES! No nanotechnology, no mature information technology, no advanced biotech, none of what we have now, just THREE GUYS IN A TIN CAN. They were completely FRIGGIN NUTS! Are you KIDDING ME????"

To think that a technology so profoundly ancient and compromised as vinyl records can

transport someone into an immersive environment that gets damn close—or even edges out—the best of our wangle-dangle, fully-buzzword-compliant, step-right-up-and-put-yer-money-on-the-table 160× DSD or 45 bit SUPER REFERENCE DIGITAL recordings is kinda like that. Not just pretty cool, but really over-the-top insane. I mean, yeah, I make fun of it as scratching a rock over plastic, but it's even worse than that. Witness the RIAA curve. The imperfections of the cutting of the masters. The imperfections of the press. The fact that the stylus has built-in tracking error on all but one point on the record. I mean, this is really lowest-common-denominator, built-to-a-price-point, ultra ho-dad stuff. But the performance you can pull out of it speaks for itself.

But there's another alternate history to think about … one that might really bake your noodle.

Fact: laserdiscs (as in, those big shiny 12 inch things that a few stupid people like me used before DVDs) were not digital. They were analog. Yes. Analog video. On a shiny disk.

Question: what if the recording industry, way back in the dim days of 1982, had chosen analog

as the format for their shiny disk format known as “CD?”

Yes. **What if CD had been analog?**

Well, copying would certainly be a ton tougher. If it was possible, it might easily involve generation loss, too. It wouldn’t be just bits that could be stored on a computer … and shared … and downloaded …

Makes you think about where we’d be now, hmm?

Chapter 41

Completely Fulla Schiit

Okay, so now we get to Fulla.

But before I start on that, let me start by thanking everyone who's read this far. I get a lot of email from people thanking me for writing this, but the reality is … this isn't the easiest book to read. There's a lot of dense tech and boring business, and it's wrapped around a lot of stuff that may make you think we're totally incompetent at this whole audio business thing.

But, like I've said before … nobody is perfect. What matters is making it right. Which will definitely weigh on the Fulla story, as you probably already know.

What you don't know is all the other stuff wrapped around Fulla, which includes a host of dead-end products, Mike Moffat's disapproval, and the whole grabbing-around-in-the-dark thing that comes with the first digital product I've designed since the Cobalt 307.

Yes, you read that right. This product is one I took to fruition. Although it's heavily based on Mike's Modi design, I made the decisions about what power supplies to use, what op-amps for gain and output, what form factor, what feature set, and so on.

Why? Mainly a difference in point of view. When it comes to digital, Mike feels that we should always look to lead. To stand apart. To set the standards. This is a very natural point of view for someone with Mike's background and resume. And, to be very clear, he's absolutely right. If we were in the business of simply taking the same D/A chip that everyone else was using, and marrying it to the same USB interface everyone else was using, the potential for real advancement would be much smaller.

So, I understand his point of view. But I also like to experiment. And, even in the most perfect world, even Mike will admit that there's a price point below which multibit technology really isn't viable.

And that's the roundabout way I began working on Valkyrie.

Wait, Valkyrie? I Thought This Was About Fulla!

It is. But to get to Fulla, because to get to Fulla, we have to go through Valkyrie. And before I get to Valkyrie, let's talk about the business case here. Specifically, why did I want to play around with a tiny DAC/amp combo, when we already had Magni and Modi?

It's simple:
A $ 198 combo purchase is still a decent amount of money for a lot of people. Something that could sell for less than half the cost means we can reach a lot more people.

How many more? Based on our experience, halving the cost of a product squares to cubes the volume, depending on the product. So, half the cost means 4× to 8× higher sales. Guys, that's *huge*.

"Oh, well, now you're becoming a big evil company," you say now. "Chasing volume, because that's all that counts to you."

Well, uh, no. I mentioned in the last chapter that every ongoing business must think about continuity. And, in the case of an audio business, we have to think about it more than many other

businesses, because everything in audio is strongly seasonal. This is a fancy way of saying that your sales in December and January are probably 3× to 4× your sales in July and August. Products which sell consistently throughout the year are a big deal when you're planning for a slow summer.

And guess which products are most consistent: inexpensive ones.

So adding even less expensive products is a very, very good idea—*providing they fit with your company philosophy, and deliver a level of performance you're happy with.*

Stop. Go back and read that again. Simply designing to a price point *never* works. But, if you can deliver a great product at a low price, it can be a real winner. It can make a lot of people happy.

That was the big experiment with Valkyrie (and eventually, Fulla)—*to see if we could produce a combined product that was fun, sounded good, and provided huge value.*

But, as far as I'm concerned, there are also a couple of other reasons to pursue inexpensive products on the audiophile side of things:

1. **Affordable audiophile products are very thin on the ground.**
   Seems like everyone wants to do $5000 amps and $25 000 speakers, but when you ask for something that might get your kids interested, they shuffle their feet and look uncomfortable. We need more entry-level products that deliver great performance.
2. **It can teach you things that will make all of your products better.**
   Seeing how much performance you can squeeze out of cost- or space-constrained designs is a real challenge—and, in doing so, you may discover even better ways of doing things, period. This is one reason we say, "we prefer to do things one way," but challenge our assumptions on this from time to time.

"Okay, okay, I'll agree with you if you just shut up," you're saying now. "Get to the story. We get that you like to experiment. What made those experiments end up at Fulla?"

Okay, fine. Because it's a pretty good story. Because, when I floated the idea for the products that led to Fulla, Mike hated them.

The Beginnings of Valkyrie

Late in 2013, new dongle DACs were showing up seemingly by the week, chasing the success of the original Audioquest Dragonfly. This alone was enough to ensure that Mike would want nothing to do with a Schiit version.

But I still wondered … what exactly was the appeal of a dongle-DAC? They were inherently limited, in being powered by the USB port. If they didn't use a switching regulator to create a negative rail, they really couldn't deliver much more power than a laptop headphone jack. And many of them never really moved around very much, being stuck into work laptops that just sat on desks.

So, maybe it was time to look at a combined Schiit product. Something that could be powered by a USB port, for a single easy connection and no power cords. Something that could be one small box, for portability.

But I wasn't thinking Fulla sized—at first, in fact, I figured that just putting everything in a Magni-sized box would do it. Just one USB cable in for data and power, volume control, and headphone out. Might as well throw line outs on the back as well. And, well, maybe an analog in would be

nice as well, so you could use a source besides your computer with just the amp.

But that idea never got beyond the pencil sketch stage.

Fun fact: I still keep paper engineering notebooks … pencil on paper is how most stuff starts at Schiit, even though it ends up in tolerance 3D CAD these days. It's how I figure out pretty much everything, from the slick way to mount MOSFETs in Asgard 2 and Lyr 2, to the Fulla chassis.

Why didn't it go further? Because I thought I had an even better idea. Why not make the device smaller, and sell it with a little base so it could sit vertically on a desk. That way, it could be, well, about the same size as a portable amplifier.

This is what I called "Valkyrie." It was a 2.5 inch×4.5 inch×0.8 inch, round-fronted design that sat vertically like a blade on a desk. It had USB in for power and data, plus analog in, plus a gain switch, preamp outs, and full-size headphone jack. It used a switching rail generator to create ±5 V from the USB 5 V in, and had a discrete output stage.

Or, in other words, I stuffed about as much Schiit as I could into it—high power, discrete, versatile, etc.

But this wasn't enough. I also made two more versions—one using the tubes out of the Vali (and a wall-wart), and one that ran on a 3.7 V lithium battery pack and had all the charge management/etc you need to keep batteries like that healthy and alive.

Yes, you're hearing that correctly. Fulla began as three products, none of which was Fulla.

So where did Fulla come in? As a lark. When I was getting ready to send the Valkyrie boards out to prototype, I wondered, *Just how small could we go with this? Could it be stripped down even further?*

So, over the next day, I pared back the Valkyrie design to its bare minimum—the USB input receiver, a −5 V rail generator, the DAC, and op-amps for gain and output. After staring at integrated volume controller datasheets and thinking, *We ain't got no room for a microprocessor in here*, I found an 8 mm Alps potentiometer that was silly small and stuffed it into the board. The end result was about 1 inch×2.5 inch, far smaller than Valkyrie.

Feeling silly, I screened on the board, “Dingle Dongle, it’s a DAC.” No product name. No idea what I’d call it.

It would be funny, having a dongle-DAC with an actual volume pot, I remember thinking. But I didn’t think much about it, because I was really focused on Valkyrie. What would become Fulla was just, well, playing. It didn’t even have any mounting holes on the board.

We probably wouldn’t do anything with it …

All the Best Plans … Change

When we got the prototype boards for the three versions of Valkyrie and what would become Fulla, I laughed at how tiny they all were—three products that were all dwarfed by Magni and Modi boards. I showed them to Mike. He grimaced and shook his head, still not thrilled with the whole grand “cheap experiment.” He picked up the dongle board and said, “So what are you gonna call this one? Dingleberry?” I laughed. “We probably won’t call it anything,” I said. “I just threw it in to see what we could do with it.”

Mike eyed the board. A volume pot?”

“Yeah.”

Finally a grin. “Now, that’s cool. What does it use for output?”

“ADA4610 for gain, and a DSL line driver for power.”

Mike laughed. A DSL line driver?”

“It has good distortion specs, though,” I said.

“No, no, I love it,” Mike said. “I hope it sounds good. I’d love to say we’re using a DSL line driver for audio.”

Aside: we didn’t end up using the DSL line driver chip. It couldn’t swing the rails, so we were throwing away too much potential output. Hence the AD8397.

Mike was less thrilled with the Valkyries, except for the one with tubes. “Now, that’s cool, too,” he said. “But USB power?”

“Nope, not enough power for that. It needs a wall-wart.”

Mike squinted at the tiny markings on the bare boards. “Who’s going to build the protos?” he asked.

“Me.” A laugh. “Good luck.”

And Mike was right. After an abortive attempt to assemble one of the boards, I gave up and ended sending them to our PCB assembly house. The 0402 resistors and 100 pin QFN got me. Yes, I’m lazy.

And, interestingly enough, it was a good learning experience. Having your board house do prototypes means you have to have the bill of materials worked out (good), which takes a lot more time than you expect (bad), and they will inevitably find some problems with the boards you missed (good), but it may take some back and forth to get all their questions answered, meaning even more time (bad.)

Let’s leave it at this: proto assembly by boardhouse can definitely end up extending your development time.

When the boards came back, the only one that just started up and ran was the Dingleberry. I think that should have clued me in what was going to happen right there.

And yep. The name had stuck. Dingleberry.

What Happened to Valkyrie

So why didn't we pursue Valkyrie? Well, we did, for a while. But in the end, all of them ended up as non-products.

Why? Depends on the product:

Valkyrie Tube: Holy crap heat, needed a wall-wart, too noisy for the intended application. No chance of it ever being entirely USB powered. So cable hell on the desktop, too.

Valkyrie: noisy-as-hell supply (which could have been worked through, mainly), discrete output not swinging rails, so that meant a redesign, pricey chassis to do it the way we wanted to, might end up costing more than Magni/Modi combo.

Valkyrie Battery: Same as above, plus plenty of good portables out there anyway, why do something similar? Better to shelve it. In the round file. I'll repeat this here, because I'm sure we'll be asked, but I'm unsure if we'll ever do a portable. It would have to be very cool/different to get us excited about it.

Valkyrie was a great lesson in modesty. Because 0 for 3 is a pretty crappy batting average.

But then there was Dingleberry.

Snap Goes the Dingleberry

Like I said, the first Dingleberry just plugged into a USB port, programmed, and ran. Of course, it wasn't perfect—a couple of clocks needed hacked, and I'd screwed up the output ground. But with less than a half an hour of work, we were hearing music.

And … it wasn't bad. In fact, it had pretty good promise.

There was just one problem: I took it home, plugged it into my laptop's USB port, sat down … and the captive USB connector promptly snapped right off the board.

"Well, that's that," I said. I'd seen plenty of dongles with the captive USB connector, but ours wouldn't be one of them. The next design used a mini-USB receptacle.

But that wasn't the extent of the changes. In the brief lifespan of the first prototype, I'd seen promise, so it was time to get serious. Serious enough to put mounting holes on the board, and figure out how this thing might actually be assembled.

But to do that, I had to figure out the chassis … and that was a challenge in itself.

Why? Because I knew from the preliminary bill of material that the make-or-break point of Fulla was the chassis. If we could get a chassis made inexpensively enough, it could be a $ 79 product. If the chassis was costly, it could easily balloon into a 3-figure product … and, at that point, why bother?

And, when you're talking small chassis, you have a ton of potential choices … but very few that would fit with our budget, volume, and aesthetic:

Milled aluminum? Yeah, I know, this is the thing that everyone likes to do these days. Unfortunately, unless you're making literally tens of millions of them like Apple, the cost will make you faint. And, looking at it dispassionately, this is a horrendously wasteful process, milling a solid block of aluminum into a pile of shavings. That's why it's costly.

- **Plastic injection molding?** Sure, that's cheap in reasonable quantities, but … it's plastic. And you still have to meet FCC somehow. So the metal shields inside might send the budget over the target. So, nope.

- **Metal injection molding?** Seriously considered, but it's really better at higher volumes, since the tooling costs have to be amortized into it.
- **Extrusions?** Sure, but they are thick and clunky, and I wanted this to be thin. Still, we did seriously look at going this way.
- **Folded sheetmetal?** At first glance, this looked like the least likely option. Fulla is tiny, and requires precise bends. But we lobbed a sketch at our sheetmetal guys, and they turned a quote that was well-within our price range. Suddenly, a $79 product could be an actual reality.

So that's what we went with—more folded metal. Fulla uses a tiny, custom aluminum and steel chassis, together with a custom aluminum knob. To further simplify assembly, it also has no captive fasteners at all—the PEM insets are on the PC board, and screws sandwich the two chassis pieces together with the PC board. It's actually kinda amazing we got it to work in sheet metal so well.

Aside: actually, it's kinda amazing how much unique or custom stuff went into Fulla—from the chassis pieces, to our most elaborate knob design (and also our first 3D-modeled part),

to custom M1.4×10 mm screws for the volume knob, to the SMD-mount 2-56 PEMs, to a 4-layer PC board—but I'm getting really geeky here, aren't I?

With the chassis design in place, I made the necessary adjustments on the PCB and sent it out for another proto run—this time alone, without a Valkyrie in tow. Because this time we were going to do a real run—a full panel of 40 boards. This one was named "SAGA," because I didn't really want to commit to "Dingleberry."

When the boards came back, we programmed them with Modi firmware and gave them out to all Schiit employees and friends who were interested, since I wanted to have them banged around for a while before we committed to production. This is when I started going around the house with a prototype Fulla hanging off the end of a pair of HD 800s with balanced-to-single-ended converter and a 1/4 inch to 1/8 inch converter that were much bigger than the board.

It sounded good enough that I really didn't miss the big iron too much … but it did have problems. Most notable was a very noisy negative power supply rail—which, while the noise was at inaudible frequencies, was not ideal. Also, the DSL line

driver was limiting the power output, because it couldn't swing the output close enough to the rails.

After the addition of a couple of inductors, a lot of ground plane work, additional bypassing, and the swap of the DSL driver for an AD8397, we had something that sounded quite a bit better.

"Better than the laptop output," Mike said. Faint praise, but he was at least grinning now. "And the volume pot is cool."

"So we're gonna do this?" I asked.

Mike shrugged. "For $ 79? If it sounds a lot better than a laptop, why not?"

"You still gonna call it the Dingleberry?" I frowned. "Umm, well … "

"Don't tell me you're chickening out?" Mike said. "The intrepid marketing guy is afraid that it's in too poor taste? For a company called Schiit?" I didn't say anything for a long time. Because that's exactly how I felt. Yes, we were the iconoclasts in the business, yes, we had a crazy name, yeah, we made fun of ourselves … but I also thought there was a line. And I thought that "Dingleberry" was stepping over the line.

“So what do we call it?” Mike asked. “The ZX-01 Interoscillator? Thor’s Nuts?”

“I don’t know, but I’ll come up with something.”

“I still like Dingleberry.” I gritted my teeth and said nothing. Mike liked the name. Rina liked the name. A lot of our employees liked the name.

But I couldn’t do it.

Eventually, I happened on Fulla, a Norse goddess. And it all fell into place: not Schiit Fulla, but Fulla Schiit. Mike grudgingly admitted he was okay with it. Rina accepted it.

Or, in other companies, it got the 56 signatures and legal clearances necessary to proceed. I got to working on the final board revs, thinking, *Soon we’ll be Fulla Schiit.*

Soon is Relative: The Problems You Didn’t See

When I got to work on the final board, it was early summer 2014. I figured we’d be shipping Fullas by late August, maybe early September. No problem, the design was done, the chassis was ordered, there really wasn’t a lot to work out.

Except for a few things:

1. The first articles for the chassis were much too thin. We'd originally spec'd all aluminum for the chassis, which flexed too much. Back to the drawing board.
2. The second articles for the chassis were still too flexy. Time to go to a steel inner chassis. Also time to wait some more.
3. We needed a second run of prototypes to qualify the "real" board. This took longer than expected, with a boardhouse underwater in new products, including the Ragnarok boards.
4. The knobs were late.
5. The chassis were late.
6. The original PEM nuts to hold the whole thing together broke the boards when inserted. Time to find new SMD versions.
7. About 80 % of the original USB mini cables were junk. Time to find a new supplier.
8. Our original chassis insulation plan didn't work. Time for a new custom part to ensure the sparky parts of the board didn't meet the aluminum chassis in bad new ways.
9. The original 2-56 screws were too short for easy assembly. Yeah, not a huge deal, but still a delay we didn't need.

So, by the time we started shipping Fullas, we

figured, *Heck, after all of the easy products we'd launched, we were due for a bad one.*

We just didn't know how bad.

The Oscillator Problem

Okay, all of you out there who want to run your own electronics business, take notes: even after a lot of pain, you can still get bitten. And two days after launching Fulla, we got bit, bad: multiple owners were reporting that the 48k and 96k sample rates weren't working on their Fullas—specifically, they were getting distorted output. 44.1 and 88.2 were fine.

Uh, oh, I thought, when I heard about the second case. Because that meant only one thing: the oscillator for the 48k sample rate multiples wasn't working for some reason.

> **Aside:** I will use this to note that Fulla does have separate crystal oscillators for both 44.1 and 48k clock multiples, just like the "big boys." It does not derive both from the 12 MHz USB crystal—an insanely high-end approach for a $ 79 product.

I took a look at the Fulla boards we had in stock.

The 49.152 MHz oscillators were all in place, and the solder looked good.

But they also didn't look like the 49.152 MHz oscillators we used for everything else.

Some more investigation revealed what had happened: a different part had been swapped into the run at some point, most likely from a "taped reel." (We found out later it was swapped in very early—there were very few working Fullas.)

Now, before you get worked up about how the PCB assembly house messed up, the reality was that they didn't do anything wrong with the assembly. Taped reels are common. You're usually just splicing in the same part. Or one that was substantially similar. This was a case of the latter—they used an alternate part we said was OK (it was still a 49.152 MHz oscillator) but really wasn't. It was the first time we'd used them, so we got bitten.

So how did this make it through to production and shipping, you ask? Here's how: first article approval and boardhouse programming and test.

You see, before the PCB assembly house does a full run, it usually does a "first article" for approval. These first articles (usually 5 to 10 boards) are

run with the same parts and processes as the whole run. So, if they work, the whole run works.

As usual, we received the first articles and ran them through the full battery of tests, including all sample rates on the Stanford analyzers. And they were fine. They all worked, on both 44.1 and 48k clock multiples.

So, with that assurance in place, we gave the assembly house the OK to build everything … and to program and test them as well. This was the first time we'd done programming and testing out of house, but at the volume we're running Fulla at, it made sense.

Except … except they didn't test them at all the sample rates.

With the approval of the first articles, this shouldn't have been a problem. But clearly the first articles used the last of the good oscillators, whereas the majority of the production run used the bad oscillators.

Boom. Big problem.

Making it Good

So what do you do when confronted with a problem like this? When you've already shipped a few hundred products, and you have orders for hundreds more?

- Well, first you shut down ordering, so the problem doesn't get any bigger.
- Then, you set a policy to swap out every problematic Fulla, at no cost to the owner—or issue a refund, if they'd like.
- And then you go to the PCB assembly house and see how deep the mess is.

In our case, we shut down ordering on Friday, and were only able to address the problem on Monday at the PCB assembly house. Then, it took a couple of days to determine how many good Fullas we had (pretty much none) and how best to swap the oscillator on the bad ones. It was Wednesday before we had that all down, Friday before we had good stock, and now, on Monday, we're shipping the replacements and getting ready to open up new orders.

Helluva week. Especially when we knew we had to launch the next-gen Magni and Modi, too …

Chapter 42

One Year, Twelve Products

Okay. Let me get this off my chest now: twelve products in one year is too many. Yes, even when half of them were updates, and two of them variations on the updates.

Twelve? Yes, twelve. A nice, familiar, comfortable dozen (well, unless you like that easy metric-system stuff, where 10 is a better number.)

Let's count:

1. Wyrd
2. Optical Modi
3. Sys
4. Valhalla 2
5. Lyr 2
6. Mani
7. Ragnarok
8. Fulla
9. Magni 2
10. Magni 2 Uber
11. Modi 2
12. Modi 2 Uber

When we introduce this many products in a year, it naturally lights people up with a lot of questions. Why did we go bonkers and introduce so many new products in 2014? Will we do this again? When are the third-gen products coming out?

I'll answer all of these—and get into the whole whys and wherefores of the Magni/Modi 2/Ubers a bit later. First, let's start with what we learned.

Hard Business Lessons

If you're thinking of starting your own business, or running one right now, this is probably the most important part—something that can be boiled down to heeding our own advice about "keeping it simple, stupid."

Because we end 2014 as our most successful year ever—and, at the same time, asking ourselves pointed questions about the future.

Here's what we learned.

12 intros in a year is too many with our engineering resources.
And probably even with additional resources. We have nobody to blame for this except ourselves. Early in the year, I put up a whiteboard in my office and outlined all of our possible/planned

products and numbered them. I came up with 16. Three died in the Valkyrie experimentation, and one missed (Yggdrasil.) Even at that point in time, I knew that an average intro of over one product per month was logistically painful—or perhaps even impossible. We are lucky that only one product had significant issues, and that only one product missed its intro date. So how many products makes sense? I'll get to that.

12 intros is also probably too much for good marketing impact.
If you're hitting the press with a dozen new products—especially when they are concurrent—it makes it hard to digest. The intro of the second-gen Magni and Modi coming right after Fulla was not ideal, but our hands were forced by production delays. The initial plan was Fulla in August/September and Magni/Modi 2 in mid November. Yeah. There you go. Fewer products, more widely spaced, allows the press to absorb, digest, and comment on the new stuff, before you hit them again. What's the ideal timing? I'll get to that, too.

Each product has many more complexities than you expect.
Yeah, it's one thing to have a working PC board, and it's another thing to be shipping. Even getting

to a fully working board can be challenging. But, to get to shipping, you need to:

- Have that fully working PC board—in our case, we will go through 2 to 5 revisions of the initial proto to make it right
- Have the firmware fully functional, tested, and busted—and, in some cases, this means multiple firmwares, as with Modi 2 Uber, which has both a microprocessor and a USB input that need to be programmed. And by "busted," we mean, "hand it to some people and have them do everything crazy they can think of, including pushing all the buttons simultaneously, pushing the buttons on power-up or down, etc.
- Have a working, fitting, cosmetically perfect chassis in large quantities, ready to go—again, we usually go through a couple of revisions on this, though some of our chassis are at "Revision J" by the time they're released to production … but then you're waiting for production to arrive, and crossing your fingers that it's perfect.
- Have any knobs/buttons in large quantities—see above, though it's usually just one Rev and we're done
- Have all the fasteners for that chassis in large quantities—surprisingly nontrivial, especially

in the case of complex products like Ragnarok and custom stuff like we use with Fulla … Ragnarok uses literally two dozen different screws/nuts/standoffs/washers/spacers/etc.

- Have any custom insulators/etc for that chassis in large quantities—see the two products above, which use either custom Sorbothane parts, or custom electrical insulation parts
- Have all the parts for the PC boards (and we can be talking hundreds to thousands of parts, each and every one of which might make finishing the product impossible—talk to Alex about this …)
- Have complete documentation necessary for the PC board assembly house, including any special test procedures—or what you get back from them might not be, well, what you expect
- Have your team ready for any new assembly challenges—especially with complex products like Ragnarok
- Have a shipping box that will protect the product tested, busted, and ready to go—sounds simple, but a cheap, sturdy box is not as trivial as you think … put a prototype in it and ship it across the country a few times to check
- Have a manual for the product fully written and ready to go in the box—no manual, no

time? Then you'd better at least have online instructions. But really, a manual. Seriously.

- Have the technicians and final sound-check guys ready to test the product—if they don't know what to expect, again, brace yourself for the unexpected
- Have at least one staff member ready to answer questions about the product, based on actually having used it—again, sounds simple, but this really isn't trivial

And this is in addition to all the marketing stuff:

- Product description—what is it, what does it do, why is it special, how is it different? Yes, I know lots of people don't read, but you do need this.
- Product photography—and I'm not talking photos with your iPhone. Even if it is an iPhone 6. Get some professional work done.
- Specifications. Testing, retesting, confirmation, and more testing, decisions on claims (even if your products measure 0.0004 % THD, are you comfortable claiming this, in all systems, on all analyzers?)
- Press release—an email release in PR-speak that tells the press you are serious, and covers all the basics about the product, including price and availability

- Community announcements—like on head-fi.org
- Facebook, if you're crazy enough to do that
- Product FAQ—it really does help. You should do one. And sit back the first day and watch the online feedback so you can add to it.
- Excerpts and links to any early reviews—if you have them
- Anything else—stunts, videos, laser-etching your logo on the moon, skydivers, rockets, soda dispensers in bathrooms, whatever, yeah, marketing is stupid, but cover the first stuff before "anything else"

This giant-ass list above applies to updated products as well.
Yes, it does. Don't tell yourself it doesn't. Updates, real updates, give you a shortcut in only one place: chassis. You may be able to use something very similar, and cut down on the Rev cycle. But that's about it. A revised product is a new product, unless you are so stunningly cynical you're gonna reprogram it a bit and slap a "2" on it.

The line may now be too large. When you look at it, we have a hell of a lot of options. Which is really cool, because not everyone likes the same thing. But we are getting into "confusion by profusion," so this is something we'll have to look

at. Don't panic—we're not thinking of killing off anything, but we will be trying to make it more clear how the different options stack up … and, of course, we're going to weigh any additions to the product line very carefully in the future. Again, don't panic—I think we have a pretty good grasp of where the shortcomings are … and this does not preclude moving into different markets (more "traditional" preamps and power amps … though our evolving ideas are very much non-traditional …)

Harsh? No, realistic. Again, want to start a business? You're gonna have to be realistic. You need to step back, look at what you've done, and evaluate it as Spock-like and dispassionately as possible.

Am I down on Schiit? No, not at all. It's the most exciting thing I've done, ever … and the stuff we do have coming is, well, pretty darn wowie. Looking at it dispassionately, even.

Would I have done things differently, introduced less products this year? Maybe. Not sure. Get back to me on that one later.

The Ideal Scenario

"Okay, so what would be an ideal number of products for the year?" you ask.

Cool. Let me preface this by the fact that I'm not a trained, professional product manager, and I never spent any time at a big company that introduces dozens, or hundreds of products a year. So this applies to niche businesses … maybe only niche audio businesses … or maybe only Schiit. I think it's about 4 products per year.

Why?

Engineering-wise, it's a comfortable rate. At least for us. Our engineering cycles usually start with "what ifs" more than a year before the product is a reality. But once we've decided we have a product to do, 3 months between releases allows engineering to concentrate fully on the product, rather than pinballing between several different ones.

Production-wise, it fits very well with typical cycles. Metal has a cycle of 6 to 8 weeks, and finding all the other parts can take another few weeks, and dealing with out-of-stock parts, first articles, and production ramp-up is another few weeks. Releasing every three months allows you

to work through the issues on one product at a time, which is much more sanity-inducing than several.

Marketing-wise, it's ideal. Think about it. You introduce a new product, and everyone's talking about it for a month. Another month, and you get the inevitable "it ain't that great" people coming out of the woodwork. Another month, and it's settled down into the line well, and the press who were excited about it have moved on to different products … so when you introduce a new product, they're ready to hear about it again.

"So you're not going to introduce more than 4 products per year from now on?" you ask. "If you go over that, the world will implode?"

No. Not at all. I'm just saying that from a technical and marketing standpoint, it seems like 4 per year is ideal. The number might even be lower. Or it could be higher.

In short: 2015, we'll see. We have 4 "definitelys" and a couple of "maybes" on the board now … which is a heck of a lot better than 16.

And … if you think we could precisely space 4 products at 3 month intervals, you have a lot

more faith in us than we do. Let's hope all four don't ship December 20th. 2014.

> Don't worry, they won't. And Yggy will be first.

So, About These Second-Gen Magni and Modis

Why did we feel we had to update the Magni and Modi, after only two years in production?

Well, 2 years is actually a pretty long run on the low-cost side of things. And the market changes over time. And, we'd had plenty of feedback about how a gain switch would be a good addition to the Magni. And, we'd also had plenty of feedback about how a Modi with multiple inputs would be a good thing.

So, all in all, it was time to take a look at them, beginning in early 2014.

Yep, the dev time was that long. Not because it was problematic, but because we intentionally held the products back in order to introduce them later in the year. A little later than we would have liked, yeah, but hey, sometimes things don't go precisely to plan.

It started with me playing with Magni. Or, more specifically, Magni. The gain switch was a must, of course. But if I was doing that, I figured I might as well take a look at the whole thing and see if it could be improved upon.

And, like many engineering challenges, the question was, "Where do you stop?"

Staying within the original Magni budget meant fairly limited changes, but it still did allow for some pretty significant work. The gain stage had originally been designed as a minimal-parts design. With the addition of a handful of parts, though, I could turn it into a constant-feedback design that promised to sound significantly better than the original.

Aside: so what's all this constant-feedback blather? It's constant feedback across the audio band. Now, Cordell disproved the old Otala TIM hypothesis, but in my work, I've always noted sonic benefits to having an open-loop stage with bandwidth larger than the audio range. So, I extended the open-loop bandwidth to greater than 20 kHz with some parts additions and other tweaks, changed the operating point of the front end (it now runs quite a bit more current, and has smaller resistor values for even

lower noise), and, of course, put in the gain switch.

> **Aside to the aside:** yes, this is heavy geek-speak. If you know Cordell and Otala and concepts like thermal noise, this makes sense. Though you might not believe what I'm saying about the subjective side of things. That's cool.
>
> > **Aside to the aside to the aside:** bottom line, Magni 2 sounds better. And it's not subtle.

How much better? The Magni 2 was originally called the Magni Uber.

But, while I was doing this cost-constrained version, I wondered, "What would we get if we threw some more money at this?" So, I built another super-over-the-top Magni Uber Squared, which had a whole host of improvements, starting with preamp outputs, better, adjustable voltage regulators and higher rails, a complementary-input VAS to cancel even more noise and distortion from the front end (and yeah, I know Self sez you don't need no complementary VAS, that's cool, maybe we hear differently), much bigger power supply capacitors and a giant new wall-wart, and a couple of other little tweaks. This super-over-the-top

version I figured we might never build, but I had to know how it performed.

How'd it do? So well that we decided to build both of them. One of our listeners said, "Just stop building everything else and do these," when hearing the Magni Uber Squared prototype for the first time.

Aside: don't worry, we're not going to do that. He's insanely cheap, and that explains a lot of his comment. It is really good, though.

So, with that decision, we knew that we'd have two new amps. There was some production budget left over, so I drew up a small solid-aluminum knob and added an aluminum top to the Uber to make it look a bit fancier.

Beyond that, prototyping to production was pretty uneventful. Magni 2 and Magni 2 Uber are two of those "boring" products that you really like to have.

Modi? Modi took some more turns.

Why? Partly because the original Modi was a very nice DAC, for a $ 99 DAC. Mike pretty much thought of every trick he could apply to it. So, that made it difficult to make better.

One obvious thing we could do was to open up higher sampling rates—extend it up to 24/192. Of course, this would require drivers for Windows.

> **Aside:** really, Microsoft? No USB Class 2 support in Windows 10? Reeeaalllly? Android phones do this now, you know. Do you realize how stupid you look to the audio community? Yeah, I know, you don't care. But you still suck.

To get around the "hey I just got my first DAC and I can't install the drivers" problem, we added a switch. In Standard mode, Modi 2 is locked down to USB Audio Class 1, and 24/96 max. But you can pretty much plug it into a toaster and it will work.

> I am kidding. Unless the toaster is also a computer. Which is possible around here.

Put Modi 2 in Expert mode, and you can run 24/192, but then you have to install drivers. But at least there's one mode that works, and we're setting expectations that, well, Expert mode is for Experts.

Plus, we tweaked the parts and layout a bit, and there you have it: Modi 2.

"So why U not change the A chip and do 32/768 and DSD 8× and stuffs?" some people are wondering. Well, it's simple:

- The AKM4396 is still one of the highest performing DACs, ever—check its specs against some of the flavors du jour
- It is also one of the best sounding DACs, ever—and yeah, we know all about the new ones
- It's not even 32 bit internally, so we can't claim 32 bit
- 24/352 and 24/384 content is comically scarce, and let me know when you find some DSD 8× content

Fun fact: I did a custom version of the Modi firmware that enabled up to 24/384 for laughs. I really wanted to introduce this as a 24/384 DAC, but then be the only manufacturer to say, "And we did it only for the numbers, unlike everyone else, we'll tell you that this *really doesn't matter.*"

But Mike talked me out of it. So, you could have had a 24/384 Modi 2. Which would mean nothing. But it would have a big number on it. Maybe screen it on the front and put racing stripes on it, like they used to put the "5 Speed" badges on the back of Hondas in the 1980s.

Aside: don't bother asking me for the firmware. Seriously.

In the case of Modi 2 Uber, though, things were very different. Mike and I had talked about doing a Modi with USB and optical input for a while. I'd even drawn up some chassis for it. (This will come in later.) But when we got down to actually talking about a step-up from Modi, Mike opined:

"We should just do the whole thing," he said. "Optical, coax, USB."

"Just like a Bifrost," I said.

"Yep!" Mike grinned.

"Like a mini Bifrost," I repeated.

"Right!" Mike said, still not getting it.

"Don't you think that might cut into Bifrost sales?"

Mike laughed. "Who cares? Plus, if they don't get the whole upgradable thing, or don't want it, then you have a heck of a stack for $300."

Hey, I'm not one to argue. So that's why Modi 2 Uber is pretty much a mini Bifrost. Mike also lobbied for the Bifrost pushbutton on the front, too, so you can thank him for that.

But Modi 2 Uber is a very different product than a Modi 2. Because it has optical and coaxial inputs,

it needs a power supply. It can't rely on the USB input alone for power.

Which gives it a big leg up on Modi 2 right there—a linear, low-noise, multiply regulated power supply. With a wall-wart, yeah, but a linear supply will beat USB power anyday.

It's also different in another way: we have to switch, and manage, multiple inputs. Like all of our other DACs, Modi 2 Uber doesn't use asynchronous sample rate conversion to make all the inputs a convenient single bit depth and sampling rate—which means we have to detect the bit depth and sampling rate, and change clock multiples on the fly for multiple inputs.

This means that Modi 2 Uber has a microprocessor. Which has to have its own firmware. So, when we're producing Modi 2 Uber, we install 2 firmwares, one for the USB input, and one for the microprocessor clock management.

So, yeah—just like a mini Bifrost.

Metal Hilarity

Well, it's funny in retrospect. While it was happening, not so much. Modi 2 Uber was a bit of a

pain on the metalwork side—more so than ever before.

And no, not because of fit and finish. Because of revisions.

Remember, I said that I'd done drawings of a 2-input Modi 2 Uber? Yeah. Those got sent out to quote. Then, when I changed it over to 3 inputs with an aluminum top and button, I sent those out to quote, too.

Then, thinking we might be able to show the products at Can-Jam, I asked for a finished first article of Modi 2 Uber, complete with the Rev level of the 3-input product: Rev D.

Then, when the metal guys asked for the silkscreen, the problems started.

"Hey, this screen doesn't match the metal," they told me.

Weird, I thought, and checked the dimensions. But it was fine, it should fit. So that's what I told them.

"It doesn't fit, like, at all," they told me.

"What?"

"It's like it's for a different product."

Ah, crap. What had I done? But I checked the PO and the drawing, and everything matched. A few emails later, and the mystery was solved: they'd built the older "Rev B" drawing by mistake. At least they didn't build a few thousand of them. But it still set us back.

Then, when we got the PC boards in, again we had problems—they didn't fit. And again, the metal guys were the culprit. They'd built the Rev D drawings, rather than the Rev E I gave them after we got the (corrected) first articles.

With a mistake of this magnitude, you have two choices:

1. Stand on ceremony, return the metal, and wait 6 to 8 weeks.
2. Change the board to match the older revision and throw away the boards we had. This was fairly simple, and probably something we should have done in the first place.

6 to 8 weeks, as of mid-November, was a no-go … and we were rapidly running out of Magnis. New boards, on the other hand, could be had in 5 days.

So that's what we did—went with the older metal rev, and changed the board. Now you know why the Modi 2 Ubers are a little late.

What Will We See Next Year?

Let's get this off the table first: Gen 3 products? Nope, you won't see any in 2015.

As I mentioned before, upgrading Asgard 2 to Asgard 3 would be a monumentally different thing than going from 1 to 2. We've addressed pretty much all of the first generation's limitations within the product budget. So, don't plan for an A3 next year.

So what will we expect to see? Less products, but more significant introductions. In fact, I think I can safely say that all of our planned intros will be "one and only" products that will really surprise you.

It'll start with Yggdrasil, of course. Q1 is the target. With any luck, it'll be mid-Q1. But we'll see. I don't think we need a beta test in this case, either, so when they're available, they're available.

And that leaves three mysteries.

Mystery 1. Expect to wrap your mind around a couple of new concepts with this one. It's really hard to say anything more about this … but consider that one component of it will be compatible with some of our other products.

Mystery 2. I've said I've been working on a new balanced topology, and that will be incorporated into one of the planned mystery products. It does not replace anything we're currently selling, but it does take us in some surprising new directions. And that's all I'll say about that.

Mystery 3. And expect something that nobody is expecting from us. I can't say more.

And, of course, there are a couple of wild cards, too … we'll see what happens with those as we get further along in development.

And with that, we are fully up to date, and this book is complete …
but I hope you put up with me posting new stories from time to time …
about every other week.

Mike

Motivations, Musings and Opinions

I write this not to engage anyone—I am not interested in a debate. I only write this as testament to experience.

Digital Filters

With respect to digital filters, it is important to note that there are two separate aspects which may cause confusion to some.

One aspect is the calculation of the coefficients used in the filter program itself—derivation. The other is the implementation of said filter in whatever environment (Purchased digital filter chip, or as a part of a reduced bit DAC or ADC chip). In our case a Schiit written software program runs in a special purpose DSP chip—execution.

Now when the derived filter is running—whether it is optimized for frequency domain, low-pass, high-pass, bandpass, or in our case time *and* frequency domain that *execution* is indeed closed form.

The derivation is quite another matter. While researching digital filter derivation before I founded Theta Digital, I found that all frequency domain filters were based upon non closed form math and successive approximations. I looked at many of them none were completely solvable in mathematics. The time domain optimization that Schiit uses was originally posited by Bell Telephone Labs with a publication date of 1917. This is well before there were computers to implement them. This led to a long journey which ended up with the tools to derive the coefficients that generate a filter which not only has closed form solvable math, but which also preserves all of the original samples. I have spent much time researching digital filters, and I am confident that this is the only filter like this. Period. I may now be from a small town, but I have been to county fairs and have actually seen three story buildings.

Now, why do this? It is a lot of work. Five years of my life. Simple. It sounds better. Almost all of the high priced competition uses chips bought from a variety of manufacturers which implement the digital filters. They come with instruction booklets which show you how to hook them up. It's a lot like buying a frozen food entree and reading the side of the box. I tire of the arguments

that since the record side may not be ideal so why bother with good playback. Huh? Let me go back a paragraph. It sounds better! If you agree, then get one. If you don't, then get something else. I have been in this business long enough to know that God could personally appear to me and tell me how to build the perfect audio device and still there will be people who do not like it. There are those obsessed with particular aspects of technology, be they polypropylene capacitors, non-inductive wirewound resistors, FET-Input op-amps, Current Feedback op-amps, IIR filters, sinc filters, VCXO's, tubes, etc. etc. I am sure I could lose some arguments with some proponents of some technologies.

That said, I have probably produced, designed, and sold more DACs to more people than any other living person. All that is important to me is producing technologically sound products that makes people have fun with music. If you are the guy with a music system that cost more than your car and you only listen to four recordings, all of which you hate because they are not perfect, look elsewhere, please. If you want to debate technology, then do yourself, as well as the poor people who surround you a favor—get out of talking and into action and build something.

Building Products

Let me begin by opining that in my last 40 years or so of being involved in audio, that audiophiles are an odd subset of the human race. Well above average intelligence, yup. Sane, well … So that speaks volumes about me, and manufacturers in particular. This one of those audio facts that everyone seems to know, except audiophiles. Don't believe me? Take a good look around you the next time you go to an audio meet or show.

At Schiit, poor Jason has the responsibility to publicly articulate our policy and position on audio issues. That makes him our marketeer, press secretary, public relations spokesperson, etc. etc all rolled up into one. That would be no fun for me, and totally at odds with my audio disease. See the paragraph above.

At Schiit, we actually (for the most part) enjoy what we do. We also work pretty hard on our products. When we build something we really like or have worked extra hard on, we really want to tug on someone's coat or apron and show them what we did. For me that is a whole lot easier than for Jason, who is the visible to the world Schiit figure. What you need to understand is that it is very difficult for myself, and Jason in

particular, to remain muzzled about something we have really busted our butts on and we are really proud to show off. Temper that with the fact that for someone who is a spokesman is really better off saying less than more. That's why we have a fifth amendment in the US. I have probably been shot at dawn in previous lives for talking too much. Too much communication is a dangerous thing. Too little pisses people off. It is a tough tightrope for Jason, so cut him some slack.

Won't do any good anyway. You can always tell marketeers, you just can't tell them much. It really is a burden for them. Jason also is plagued by the disease, or he would be doing other things. See paragraph number one.

So I love building this stuff. I have been doing this a long time, which gives me some basis to properly decide how to go about just that. It gets more difficult the higher end it gets. More options, choices. I was raised to believe that if I out-performed, out-maneuvered, and out-hustled everyone in my game, I could have the house on the top of the hill and everyone would admire me. Now, I am happy with my current house. It is also of vanishing low importance to me how many people look up to me. I need fun much

more than admiration at my age. Admiration is always fleeting anyway.

There was a time in my life when I tried to build the very best preamp, amp DAC, whatever. Now I realize there is no such thing. There is only a temporary very best product for some people, which leads to the new very best product for the same or different people when the first one is revealed to be not the very best anymore, etc, etc. As I have said previously, God could appear to me tomorrow and tell me exactly how to build the best audio device ever, and there are those of you out there who will disagree. I used to know a guy named Andy who would put a little scratch on his newly bought cars so he would not agonize over keeping his new car perfect. I now get that. Meanwhile, I will build you the very best product I can and let you put your own scratches on them. I like everyone (almost—one or two exceptions—they really earned it) in this business. This means customers who will never buy stuff from me, customers who will, competitors who make stuff I like, competitors who don't, industry scribes, hangers on, etc. Why? Because they share my audio malady and all show up and are a part of this audio industry which provides me my biggest arena to play in. I didn't create the

arena—all of you did—just by being here. It has been my home since 1976. It was here before me, and will be here after me as well. I build products and bring them to this arena. If you don't like them, that's the way it goes. The very fact that you are in this audio arena makes it more interesting for me. Now if you do like them, then that is indeed a hell of a bonus! This may sound arrogant, but I could care less what any of you think of me. That is none of my business. I did not show up here to be a personality—I am here to bring the very best products I know how to build. I also only care about what you think of my products to the extent that it helps me build better ones next time. This is how I have fun—it explains why I do this—and if you have fun as well then it is another bonus as well.

So these products never show up on time. Why? Some may think that Jason and myself stay up late at night to figure out how to delay products. Consider this: an Yggdrasil has several hundred different part numbers, with multiple quantities of parts bringing the total number of parts to well over one thousand. If only one of the parts is missing, then how are we supposed to ship them? What do we do when parts show up late from suppliers? How do we build protos and try new

stuff out with late or defective parts? When we are a company with growing pains, how do we balance building of Asgard 2, Valhalla 2, Lyr 2, Wyrd, Modi, Valli, Magni, Loki, etc. to fit in Raggy and Yggy? These are just two areas where unintended consequences and just in general things going really wrong delay products. There are many, many more. Cut us some slack! We are doing the best we can and learning in the process.

I^2S and DSD

What about I^2S, 2× and 4× DSD, new as yet unimagined connectivity schemes and software formats? Well first I^2S. What I^2S components do we hook it up to, and what do those components cost? If someone is buying a Schiit or other I^2S equipped component because they think it is the best in the world, and they have to hook it up to the best I^2S equipped source in the world, then 14 out of 10 of them will soon sell both anyway, as they haven't yet realized that the best possible components forever are impossible. See paragraph number five. I would much rather devote the time and energy to other products that are usable in the real world.

Some people are excited about DSD. A year ago, it was even more of interest. Users were demanding it at shows, the press was wetting its pants, on and on. A year ago, I went against my better judgment and produced the Loki DSD decoder. For a new audio format to succeed, it must have a variety of competing decoding hardware, and a wide variety of software (music) available from a wide variety of vendors. We are not talking releases like The Orkney Island Shepherds chant traditional "Poems of Rapturous Ecstasy" complete with happy bleats in the background. I need Miles Davis, Thelonius Monk, The Stones, Vienna State Opera, and so on to take me home. Now the problem with DSD is that there is a paucity of native DSD recordings that are not of limited interest. The music that many more of us will like that has been released are generally PCM brick-walled remasters. May as well just use a regular DAC on the far cheaper original release. Now the native DSD recordings which are available, do indeed sound good. The problem is that there are not a lot of them, and even fewer at 2× and 4×. So please do not hold your breath awaiting a 2× or 4× Schiit DSD device. Just like our industry did with HDCD, DAT, Quadraphonic records, and several more, we are doing the same thing, i.e. creating a new format with little

music available, and expecting a different result (success!). Now don't get me wrong—I have a few native DSD recordings that I like and listen to. That's why I did the Loki—cheap, sounds incredible for the price, but a little hard to use. If you do not like hard to use, then why are you using DSD anyway? I think you can actually fit 20 minutes or so of 4× DSD music on a 4 TB drive. I can't believe people believe they will be able to download a wide variety of files this bloated from Amazon or iTunes, but I am from a small town.

In conclusion—This is a hobby, so I am in it for the fun. I could be doing a lot of other things or not, but this is what I love. This is exactly why, although I could, I am not retiring—I love this too much. Stop taking this hobby so seriously. Have fun. That's the purpose of this whole deal. Don't assume what you like others will have to. Do your best to please yourself. No one else cares what you like. Don't be like the guys in the higher floors of many of the shows spending money you may not have on equipment you won't like in some vain quest of perfection. Don't get pissed off at us for saying too much, saying too little, late product, early product, surprise product. If you like our products, tell others. If you hate

our products, tell us. Finally, remember you are listening to music—NOT to audio gear. All we do is help you do just that.

Opinions on DSD

I have never published my DSD opinions. Here they are. I say opinions because the design of audio gear should adhere to hard science. The user's response however, is totally in that user's psyche. When I worked in Peru, there were tribes in the Amazon region who spoke in vocabularies limited to grunts and delighted in eating insects they found under logs. Then there are people like myself who prefer meat, coffee, dairy, sometimes things green or fruity, starches, and lots of salt.

In the early days of digital audio, multibit reigned. It was suitable, but expensive, derived as it was from weapons guidance and medical science. Note the use of the word science. Analog numbers were converted to digital, and the reverse yielded the same number. Nothing was averaged, no noise was added, no economic engineering geniuses were allowed to make anything cheaper with smoke and mirrors.

The earliest DACs were pretty marginal, but natural selection led to the Burr- Brown PCM-

63, an amazing multibit DAC, still pretty good today. About that time, Burr Brown was sold to Texas Instruments. There began to appear delta-sigma DACs, which is a fancy name for reduced bit width DACs which used the above alluded to tricks of averaging and noise shaping to make up for the data they threw away. Soon we had TI, Wolfson, Crystal Semiconductor, Phillips, and many more manufacturers of these (now marketed as audio—read dogschiit) DACs. Why stoop to make them? Simple—they're cheaper! Never mind they can't be used in medical imaging or defense applications because of their inherent data loss/hallucination. Too late, the audio customer had far cheaper gear. The chip makers sold lots of parts.

Enter DSD, the ultimate extension of this idea. More noise and less bitwidth. You get for free with the bargain, the elimination of the nasty anti-alias filter effects used in the recordings. Cool, huh. This idea works well just as soon as every recording studio on the planet switches over. When that happens (right), what about the old recordings like all of those from SACD days of yore!! Oops, they are already recorded with the filter in place … Unfortunately, they are the bulk of the current DSD catalog available. Can

you get DSD from iTunes? Download DSD from Amazon? Oh …

What about 1, 2, 4, or 87.6× native DSD recordings. Yeah there's a few—I really loved the Folsom Prison Castrati Singers doing Handel soprano motets. My all time fave is the Orkney Island shepherd's Poems and Cries of Ecstacy with the sheep. The plaintive cries and bleats of all involved were immaculately suspended in perfect panoramic image. Even the subtle sounds of the shepherds gently placing the sheep's rear legs in their boots were clearly audible.

Nobody ever explained to me how to design a multi-rate 1×, 2×, etc DSD DAC without a real expensive adaptive filter. Do you optimize it for 1×? 2×? 5.76×? Trouble is, then all of the other rates are compromised. Maybe the over $10k DACs do that. I haven't figured out how to make an over $10k DAC yet, maybe someone will teach me.

In conclusion—this is opinion, mine with respect to DSD: How can I express just how underwhelmed I am. Adjectives such as stillborn, faith-based, and ludicrous come to mind.

But wait—I actually built the Loki DSD DAC! How can I be such a hypocrite! The answer is that

I will try almost anything once. If I don't like it, I won't do it again. But I could be wrong—if servers ever get big/cheap enough that iTunes and Amazon offer DSD downloads AND major label music providers begin to provide native DSD recordings in substantial numbers—then I will cook and eat a crow at RMAF. Meanwhile, all you DSDers—enjoy the grubs!! Buy a Loki!

Inventions and Applications

All I am and have ever been is a product of everyone I've known and studied. And I'm pretty old, so I have known and studied a lot of materials from a lot of people. I also look in odd places. I work in the engineering field, and have a background of running audio companies, although from that I am retired. (Engineering is more fun). Now engineers are, at the end of the day, little more than very sophisticated tradesmen—especially me. Over the course of my life I have been guilty of two patents, but never figured out how that made it any easier to do anything except sue people. The geniuses of the world invent and derive stuff wearing white lab coats, I imagine. They write papers and go to hard science trade shows, twaddle on in intellectual fiddlesticks to impress each other, and apply, innovate, and

build nothing. The far East and the US academic community are full of these sorts of folks.

So back to what I've done and not done. Did I invent the vacuum tube, the transistor, the jitter, digital to analog convertor, etc., etc. … Hell, no. I have manged, over the years built some abortive messes. Stubbornly, I kept on—here is a list of some of the things I have tried that worked: Designed and built the first (as far as I know) applications of: a 6DJ8/6922 audio amp design—the first stem to stern no feedback audio amplifier—the first digital to analog converter for home audio—the first to apply low jitter in an audio design—I could go on.

The point is that I have made a career of building audio equipment which applies technology and available parts in a combination that makes them perform, and therefore, sound better. I believe that the human ear is sophisticated and is capable of discerning much that we do not yet understand. That said, I do not believe in freezing anything, suspending cables, listening upside down, etc, for better sound, except when to do so can improve some tangible measurement. Nor do I believe that the solution to all audio applications is a 5534 opamp, as do some prominent designers and authors. To do so would suck all of the

purpose out of our hobby and leave no room for any our competitors, not to mention ourselves. If you believe that everything sounds the same, just listen to the cheapest stuff you can. Take up stamp collecting and tell your fellow hobbyists all about how much their collections suck.

So what I have to say to the chattering audio geniuses in the lab coats goes something like this: Make something happen; create something which is capable of sonically pleasuring human beings, that is if you are capable. Do something and stop listening with your mouth.

That's what we used to excel at in Europe and in particular, here—invent and build great applications of technology: The Model-T, steam engines, elevators, refrigerators, and so on. That's why Schiit does it right here, in the good old USA. It feels right!

Yggdrasil

The back back-saga

It is spring 1976. I was back from Peru working for Texas Instruments wondering what I was going to do when I grew up. I sank hopelessly back into my audio habit, a mental neighborhood I hadn't visited since before Vietnam. My best audio buddies are Mike, John, and Paul. Mike runs a janitorial business, John and Paul are mechanical geeks—they specialize in making engines for model airplanes that weigh twenty pounds or so that were still legal to fly back then. When they crashed they killed people and occasionally started forest fires. We just finished Bruno Walter's Beethoven 9th on my Quad 1956 ESLs, homebrew/modified Dyna all tube pre/poweramps, and Sony TTS-3000 Rabco arm with Decca 4RC cartridge tracking at 4 grams. I am in a goose-stepping, exited about audio mood.

John looks at me and says "Let's start a company building audio electronics—you know how to make it sound incredible. It just looks like s**t

and I can help you there." Just then Paul raised a leg and anally sang a 30 second ten note atonal melody. Despite my audio reverie, I instantly sensed why he had no girlfriend. Mike, who actually had a modicum of sensibility offered "At least you'll be able to hang out with civilized people" before we adjourned to the garage to escape the stench. Thus was conceived Theta Electronics. (Not Theta Digital—that was almost 10 years later.) I was terrified. I had a wife and step-daughter to support. I thought about what Mike said about hanging with the civilized. I was still naive; I didn't realize that audio geeks were little different from mechanical geeks and I was therefore condemned. What the hell, I'll take the plunge but what will I build? At that pre-digital time sources were predominantly turntables—the best reputation was won by the builder of the best electronics designed for turntables. So the first product was a preamp which, unlike power amps, was a product suitable for all systems, regardless of speaker.

At the time, solid state gear had almost universally replaced older vacuum tube based hi-fi systems. (hi-fi=high fidelity=legacy speak for good sounding audio gear). This was just about the moment that most audiophiles were beginning to realize

that they had been conned. Early germanium based transistor gear made deafening noise like stuck toilets. Low bandwidth solid state power devices sounded grainy and strained when they weren't setting fire to your speaker or stinking up the house. To quote the Japanese, the soul of the music had been destroyed with solid state. The biggest market for the better used tube equipment was, indeed, Japan. That bid up the price of used tube gear, and there was little to no new tube gear available.

Almost all tube preamps used a variation of the same-o, same-o, 2 section 12AX7 active feedback eq circuit. 12AX7 tubes were designed for table radios and those sh**ty phonographs like schools used to use in the 60's–80's because they were cheap. (The only cool thing about older schools were the mimeograph machines which made copies that all of the future audiophiles and drug addicts used to sniff.) Not only were the 12AX7s cheap, but they were also designed with high gain so you don't have to use so many of them. They have terrible curves (read distortion), require lots of feedback to partially correct, and are noisy to boot. I hadn't and wouldn't use these in any of my beloved gear!

So there was this tube called a 6DJ8—It turned

out that the intellectual badasses at Tektronix and Hewlett Packard used a lot of them in their oscilloscope amplifiers. They also had good curves, low distortion, and are 12 dB quieter than 12AX7s. I also lost the active eq, switching to passive, which allowed my Theta Preamp to be the first 6DJ8, passive equalization, NO feedback design.

Before I was done with this company, I designed the first hi-fi hardware app for a 6BZ7 (in a power amp), and the first WE41 A tube design in a headamp (Moving coil pre-pre amp) which was, indeed, the first tube head amp. Why did I do these designs? Because they sounded better!

Now did I get rich doing this? Nope. Did I go bankrupt? Nope. Did I learn? You bet. I also picked up two characteristics, neither of which I realized at the time. The first was an addiction to build audio stuff. The second was a reputation amongst audiophiles for building against-the-grain, maverick audio equipment, even if it appeared to be traditional in nature. There were many other valuable experiences—I met Dave (Yes! *The* Dave that I still work with today that Jason mentions in his writing). I learned how to do tradeshows with hangovers and too little sleep 5 out of the 4 nights. I learned that a huge percentage of the high end audio

exhibitors were playing to their competitors and their dealers rather than their customers. Most tellingly, I learned that a loaded dealer was a loyal dealer, which often did not bode well for their customers.

Be all that as it may, I told myself I was done with audio—another grand adventure awaited me in Japan. So off I went. The trouble was that wherever I went, my audio junkie self was there as well. So I very excitedly bought the first available compact disk player—a full year before they were available in the US!! It was a cool looking thing that played the disk vertically behind a plexiglass door. The first CDs were Japanese pop songs—not exactly my style, but WTH I cued it up and listened. Just then a painful edge distracted me from the music. I wondered if one of those old-fashioned supersonic burglar alarms was in the vicinity. Slowly, it dawned on me that this system really did sound like bats with clothespins on their testicles. The digital audio event horizon deflated with a flatulent roar. Wow, it was as if I just found out that there was no tooth fairy, no more twinkies, and no more LPs or vinyl available because Warren Buffet or whoever cornered the market all at the same time.

Stubbornly I took this player home with me and

took it apart in an attempt to discover why it was so excruciatingly bad. Little did I know that this would begin the Yggdrasil back-saga, which will be continued soon.

Yggdrasil: The back back-saga, Part 1

Here I was, all dedicated to finding out why the CD player I bought in Japan sounded so terrifying. There was at that time, only one problem—I didn't know s..t about digital audio as applied to hi-fi. OK, OK, I had worked with the first, exploring for oil, digital audio recorders in Peru; they were 16 channels of 8 bit log encoding, 5 Hz sample rate machines. They were for measuring ground motion as a function of increasing distance from a subterranean explosion. As far as I knew, digital audio worked great for finding pockets of oil for which to drill. So it would seem to be reasonable that it oughtta sound right for Hi-Fi, no???

So I studied and researched what I could in a pre-internet era. There were a few things that I accepted as creed, chiefly that the more technology changes, the more the requirements stay the same; this explains the syndrome where new doodads promising ever better performance often disappoint; the technology takes precedence over

the application. I learned that the vacuum tube, the transistor, the speaker, and the cartridge, was either invented or perfected by Bell Telephone Labs. The way that it used to be back in the day was that Bell Telephone Labs had a monopoly on all home phone service in the US. In today's dollars, a month's worth of phone service where you had a calling radius of 20 miles or so cost about $100 or so. What that did was subsidize one hell of a think tank that was dedicated to developing electronics goodies; computers, TVs, radios, etc., etc. They made it possible, among many other nifty hardware apps, to hear announcers at baseball games, not only in the ballpark but at home!

Turns out they also developed and propounded information theory and digital audio well before there was hardware to implement same: papers date from 1912 to 1916. This even before the US got sucked into World War I! But I get ahead of myself. After all, you have to crawl before you can walk. I had to look at current CD players. There had to some things one could do to make digital audio palatable.

It turns out that the D/A converter chips are the most expensive chips in the whole CD player. The first CD players in Japan cost 380 000 yen,

or north of a kilobuck! A great way to save money was to use only one DAC chip function for both channels—one stereo channel converted and then the other; no kiddin'! No way for the two channels to ever be in proper phase. That had to be it! So I bought another DAC, duplicated everything else, and fed the DACs (now 2) at the same time to fix the coherence problem. It got coherent! There was an incremental change. I was excited! Mike came by, listened a bit and curled his lip. I put the turntable back on—he was right.

Next I got rid of sample and hold amps and replaced them with brick walled amps. Another improvement! Mike came by again, curled his lip again. Paul showed up, listened for about 10 seconds. He then opinined to no one in particular: "The corn kernels are gone, but it's still a turd." Mike, his sensibilities temporarily abandoned, nodded approvingly.

The solution had to be simple, I told myself. By this time, 2nd generation CD players were out. I settled on a Toshiba Model with 1 DAC chip per channel, did my usual mods and hooked up my scope to the digital end of the DAC. I thought that those were just awful looking fuzzy edges on the digital waveforms. No one was calling it

jitter yet, but it occurred to me that that would convert to some really random wow and flutter (a forgotten analog turntable and tape recorder spec that really is not obsolete in the digital audio era).

So I reclocked the whole deal from the DAC chip back with crystals and plls. The digital waveforms were far less fuzzy. I hooked it up yet again. I listened for two or three evenings before I convened the cast of usual judges. Mike listened quietly, while Paul picked his nose. Both were quiet, requesting more music, inscrutable faces. After an eternity, Mike exclaimed: "There's actually an image". Paul grunted in agreement. "Violins still kinda suck", Paul said, "but its kinda ballsy for rock".

That was fine with me. Progress, I thought. I began to recommend and modify these players for an ever widening group of audiophiles. One of them, Dick Olsher, wrote for Stereophile and gave me some kind words. Everyone basically said the same thing: "Not bad for digital".

Another one of the early adopters was a math professor at the University of Iowa named John Lediaev. Classical music geek and true TWAD (Tree Worshiping Analog Druid). He got inter-

ested in the whole digital audio sojourn and gave great input.

Meanwhile, Phillips came out with a multirate 4× chipset and CD player! (Multirate=different input and output digital frequencies, in this case 4× up or oversampling.) This meant that the signal rate to the DAC was 176.4 kHz instead of 44.1 kHz, allowing a much gentler analog filter than the brickwall filters required before which by comparison rang like old telephones. So I rushed to buy one of these new boxes, hooked it up and listened in amazement at the total lack of consistency in any property at all: image, tonal balance, etc. The only saving grace was a less harsh/edgy top end—logical for the much softer filter. The oversampling approach used a Phillips provided DSP chip. In those days, it was called a digital filter chip.

Reverse engineering of the Phillips digital filter chip, (as well as virtually all other subsequent ones by other manufacturers—up to this day) all utilize a frequency domain (read flat) algorithm which has no mathematical solution, only a succession of approximations. (read close—horseshoes). Worse yet, it does not keep the original samples—it throws them away and approximates new ones. Ugh. I wondered why all of these

filters used this method of derivation. In a flash of brilliance reserved for the obvious, it came to me that this algorithm is efficient (cheap) and gives good measurement (very flat). There were also a variety of toolkits and developmental software available to buy devoted to this algorithm only. This simplified things—making it relatively easy to for anyone to develop such a filter. There was no need to hire and babysit socially handicapped genius savants who needed cages as well as food and water pushed under the door from time to time in order to develop a proper filter. I decided to consult my old Bell Telephone Labs 70 year old at the time references. Wow, not only did they discuss the common frequency domain one, but a time domain one as well! Now here we go! Well, the problem was the time domain one was that the frequency response curve looked like an anthill. (It turned out later that a competitor actually used this very algorithm.)

Out of this knowledge two new conclusions arose. The first was that oversampling digital filters were important—to keep the decoded audio from ultrasonically cleaning windows and loosening tooth fillings. The second was that whatever audio information remained in the original audio samples could not be lost. I called John Lediaev

back and told him we had to combine time and frequency domain optimization for digital filters—he told me it couldn't be done because it required a divide by zero and by the way why did I want to do it anyway. So I told him I wanted to beat the sound of analog with digital audio. Instantly he said "There's no way—it can't be done". A stubborn kernel of determination blew up in me like an airbag. "**** him", I thought, "I'll show him". One of the manifestations of my audio insanity is that if anybody ever tells me I can't do something—I will move heaven and earth to get it done. But there was a problem: intuitively I knew I couldn't get it done without the combination of the two algorithms with the original samples preserved. If there were anyone on the planet sufficiently motivated and capable, it was him. A bonus was I really like and admired the guy and wanted to work with him.

Soooooooooooooo I told him, I bet *you* can't f***in' combine the two optimizations and keep the original samples. Silence … More silence. I had him!! He finally said "Give me some time".

But I was still overburdened; I knew a DSP engine of some kind had to be built and programmed. This was a big deal in the mid 1980s. It was a lot of research and hardware. So a week or so later

I was reading all about I DSP processors when I heard a noise that sounded like a combination of a twenty year old power lawnmower and a coffee can full of nuts and bolts slowly rolling down a steel staircase. The turquoise 1964 Toyota Corona out back with white smoke coming out of its exhaust and black smoke coming out from under the hood confirmed it. There was only one person who I knew that drove cars that were so "****ed-up" that even cops, reluctant to deal with crazies, balked at stopping them. I used to wonder how he got so broke, but later came to realize he was just cheap. I mean, this car was such a loser that you couldn't even give it to a charity, even if you hauled it to them. In any event, out stepped a grinning Dave and I knew the team was complete.

Yggdrasil: The back back-saga, Part 2

So here I was, at home in Sherman Oaks, occasionally modifying a Toshiba CD player, struggling with the necessary space to not only continue modding CD players, but to actually begin prototype manufacturing with eccentric personalities who were either total night-owls used to eating lunch at midnight (Dave and myself), or wannabe farmers who got up at 4 AM (John).

The bottom line was that I needed an office/workspace where any and all had access. At the time, I had some bucks saved, but with little money coming in I was interested in cheap. I found a quaint, windowless, 400 square foot office in Van Nuys for $160 per month. It was 2 blocks from the regional courthouse and cop headquarters, free parking, and was surrounded by a wide variety of interesting business, both in and out of the building! We had bail bondsmen, vd clinics, criminal lawyers, and travel agents with "send money back home" windows in the back. (All

of my time in South America had paid off—I could actually read all of the business signage!) We had a variety of some of the best Mexican food anywhere, offered by dozens of competing restaurants. The neighbors in the building were ambulance chasing lawyers upstairs, sports equipment and computer vendors on the ground floor, and a low budget movie studio in the basement. The neighborhood, although a blatant dump, was quite safe, because of the proximity of the court and supporting police. The local economy was cash based on the transactions on those without proper US documentation—robberies were nearly unknown. Oh, and did I say?—It was CHEAP!

One of our local characters was Russ, the local trash raider/recycler/hoarder guy who lived in the driver's seat of his 1964 Pontiac Bonneville. The car was filled up to the window line with trash except a hollowed out area which was the driver's seat and floor. This car was incredibly degraded, but no worse than the average car driven by other Van Nuys denizens. The cops tolerated him because he helped clean up the neighborhood. Besides, he was prone to loud incoherent babbling and probably too much hassle to mess with. Now, he would grab anything lying the alley. He was a skinny guy, so much so that I was impressed

when I saw him pick up a huge drunken lady who was passed out next to a dumpster. When he was trying to gently place her in his trunk to recycle her, she awoke and took great umbrage at his rescue attempt, running for her orphaned bottle of Night Train Express Wine. Please pardon the nostalgia. The neighborhood was also a learning experience for me. For example, I never knew what a "fluffer" was until I toured the movie studio in the basement. What is even more of a riot is that they are considered part of the costume department! But I digress.

Dave got a day job at this company who sealed together leak-proof plastic liners together that were 1/8 foot thick. They lined landfills, copper mine work areas, and small nations. This was to make the products or services offered at the sites environmentally and legally compliant. So, to seal the plastic together, they used RF power proper to megawatt international radio stations. The voltages on the power tubes were so high that you could literally pull arcs off the sealing machine that were 6 feet long. It was better than a 1950s science fiction movie set! The place smelled like the immediate aftermath of hundreds of simultaneous thunderstorms. Trouble was, all of the cold drinks in the building would heat up.

Dave was fearless.

Meanwhile, back in Van Nuys we started working on this contraption we would eventually call a Frankenstein. The power supply was a collection of surplus power supplies on a rack panel that weighed just slightly less than a refrigerator with the added bonus of sharp edges that could rip huge gashes in your hands. The next thing we needed was a DSP engine to implement the John Lediaev derived filter. The circuit board was laid out by Dave—right after the earth cooled the boards were laid out on transparent plastic sheets with black tape of various widths for the traces. They also used "donuts" of various sizes and different holes specced with "C" holes, "B" holes, and of course "A" holes for the pads. There was a time when we cut the layout all of the way through about 3/4 of the way to the top, cut out another huge section, and taped it back together with transparent tape. Try that with modern CAD systems!! In those days, they still sold blank boards that you could expose to light, soak in radioactive toxic waste looking 40 W chemicals which left permanent stains on your trays and fingers, and drill all of the holes by hand. We did just that with the prototype Frankenstein DSP 2 layer boards!! We drilled our own holes

and put jumpers through the boards to simulate plated through holes—the hardware apparently worked!! We had only three remaining problems—the first was that we had no one foolhardy enough to program the Texas Instruments 16 bit DSP processors in direct machine language. Neither could we afford the thousands of dollars to buy the I proprietary software compilers and assemblers to do it. We were stumped.

Well, my dad used to say that when the student was ready—the teacher will appear. It turns out I had this Toshiba modified CD player local customer named Tom who worked at a think tank in Santa Monica. He had more graduate degrees than I had fingers and was the type who would try any mental exercise at least once. He was blind as a bat, walked around squinting, and smoked like a chimney. I gave him the data sheets and programmer's guide. A few weeks later—voilà! We had a working implementation of the filter that we could look at on the scope.

The second remaining problem was I needed to convert the DSP output to analog—we needed to make another circuit board with the toxic waste chemicals. We did, it got stuffed with parts (multibit DACs—the only kind available then), and we were ready to listen to audio—except for

the third problem. There were no digital outputs on any CD players back then. (This was still the early eighties). What I had to do was hack up the Toshiba Player with the bit, word, and data digital signals and graft them into the DSP board to make it all work. The contraption was complete!

The second remaining problem was I needed to convert the DSP output to analog—we needed to make another circuit board with the toxic waste chemicals. We did, it got stuffed with parts (multibit DACs—the only kind available then), and we were ready to listen to audio—except for the third problem. There were no digital outputs on any CD players back then. (This was still the early eighties). What I had to do was hack up the Toshiba Player with the bit, word, and data digital signals and graft them into the DSP board to make it all work. The contraption was complete!

Contraption it was—you needed a wheelbarrow to carry it around. It was a CD player hard wired to a DSP/converter box, both of which were in turn hard wired to the huge power supply. In a rare moment of genius, I installed connectors in the wire harnesses between all of the boxes—all of the connectors were actually different so that

it was impossible to screw up plugging it together. The device became portable and shippable! The true Frankenstein was born!

So I hooked it up to the rest of the system, and invited all of the usual suspects by to listen. Mike was silent, very silent … a faint nod of approval. Paul also said very little, no disgusting utterances. Dave, who was usually silent unless asked a question, responded with his usual "Weighullllllllllllllll". Tom was mumbling something about fixed versus floating point math implementation in DSP processors. I did NOT hook up the turntable. It was time to build another unit and send it to John Lediaev. A couple of weeks later John had the opportunity to listen to his brand new Frankenstein. He opined what eventually became the four dreaded words: "pretty good for digital". At that moment I realized I needed to listen to my analog system again. He was right. Notwithstanding, I built probably 15 to 20 or so Frankensteins and sold them by word of mouth to my Toshiba CD player customers. Dick Olsher at Stereophile even mentioned them. It was very meager income, in spite of the fact that the team had a lot of time invested in this thing. We'd endured a small office with a freak show in the basement in a dumpy neighborhood for

years. Yet, I was frustrated—there was no way to mass produce them. They had to be hacked in to whatever CD player the user had. I was at an impasse—until—Sony and Phillips actually got together and announced the S/PDIF! The Sony/Philips Digital Interface. Within a few months, CD players from both companies appeared on the market with coaxial digital outputs! Excitedly, I added an S/PDIF interface to a Frankenstein. I now had the basis to build not only a prototype, but a sale-able product. The notion of a D/A converter that could be added to any digital audio device with S/PDIF connectivity was born in my mind! A new product category! Not that I was thinking about it a lot—only while I was awake.

However, I couldn't sell tacos to starving millionaires—I knew I needed a marketeer. So I called an old acquaintance who I knew could sell ice to eskimos, and even better than that, loved audio. I told him what I was doing and could we work together. He said that had just agreed to get involved with a company who was going to make CD players with tube analog. I explained that that was like putting makeup on a really ugly person, they remain ugly, just less so, but he was unswayed.

In those days, the audiophile scene was really

emotionally polarized between analog and digital. Ana-philes and digi-philes were like Democrats and Republicans. The analog guys knew that their stuff really sounded better, but that digital was on the upswing. The digital guys knew time was on their side, and the digital sound could only improve. Kinda like Miley Cyrus and Merle Haggard at the same show at the same time.

So I called another audio/lover marketeer who was starting a tube audio electronics company is his garage. He took the audio as a religion approach with me, preaching that how could a former tube audio product maker like me descend into digital, after all that was why I was located in a s****y neighborhood. I took the high road and avoided telling him that I may be in a s****y neighborhood, but you're in a f****n' garage. I went down the list, but could not find a taker. I felt like the high school senior who couldn't get a prom date no matter what. I knew I had the best working digital product concept, but couldn't move it even with bran. What to do?

Yggdrasil: The back back-saga, Part 3—The Theta Digital Years

So a year or so later, the marketeer I had originally approached, Neil, had once again become available from the Tube CD player Company and gave me a call—I was thrilled! He was the kind of a guy that could get away (in the days before internet and email) a phone campaign to put our product in every dealer in the country that mattered.

You see, back in those days, it was a different world. There was no acceptable infrastructure to launch a direct sales audio company. The systems were more fiddly (think turntable), larger (think big listening room—not surrounding a computer), and far more expensive in 1980s dollars. Also, a dealer served up necessary functions in an era when there were no bloggers. The only data available on audio were in the magazines. There was no convenient way for the ordinary user to contribute opinion except by hanging out at dealers, a few clubs/users groups, and

inviting each other over for the equivalent of what has now become mini-meets. Also, since the users were then as odd as they are today (look around you at any meet), the dealer provided not just the heavy lifting in terms of turntable setup and system install, but also functioned as the therapist/counselor, albeit with biases according to which lines he carried. Even though the very presence of dealers nearly doubled the prices, they were a very necessary component of sales at the time.

My duties then were are they are today: produce and design a line of digital products with proper familial integration. If I were to design the aesthetics, the products would have been as ugly as the original Theta Electronics products so then Neil (and today Jason) made them suitably aesthetic.

But before we built anything, we had to name the company, so Neil suggested that we call it Theta Digital to link my earlier company to the new one. That was simple. Then was the agony of the first product. Neil lobbied we go from the top down in terms of products. The tricky part was to figure out how to configure the first product, difficult since there were no other D/A converter product models. Do we

integrate it into a preamp with digital inputs as well as analog ones? Or do we configure it as a digital in, analog out box required to be hooked up to another preamp. Not being sure, I built both, the Theta DSPre and the Theta DSPro. The DSPro outsold the DSPre significantly, but we were committed—we had to keep both in the line. Further, the norm at the time was everybody building electronics offered both black and silver (clear anodize) finishes. So we started out with four SKUs which was needles to say, increased our cost and decreased our inventory velocity significantly. Of course, this made our products more expensive, but still far less than that of our eventual fellow DAC makers, Wadia within a year, Krell a year or so later, eventually the high end ROW (rest of the world). And so we continued, Neil relaxed, seldom leaving his house, running his phone campaign to put us together a first class dealer network at just the right pace to keep me permanently backordered—it was my job not just to design the line, but to produce the products; to run the factory/operations of the whole deal.

Along the way, I learned quite a few things. The best way to learn them is to ****** up. The first run of Theta DSPros had a 30 % failure rate. I learned all about eutectic solder, rigid static

procedures, and selecting good assembly houses, which Schiit still uses today. I learned about component crib deaths, the benefits burning-in my products. I learned about the tragedy of even one failure of any Theta Product and how much that really costs us. I improved our reliability to the point where we offered free FedEx two day both way shipping for repairs. (FedEx was a very big deal back then—exclusive and expensive) I learned that service needs to be instant; that it also needs to be friendly; and this was the best way to build good will. By this time we were shipping hundreds of units per month. I learned that the biggest mistake I could possibly make was to hire a manger to deal with my assemblers, tech, front office people, and inventory/shipping people. It just insulates me from the people I need the most and causes resentment all around. I put in a bonus system tied to production, reliability, and on-time shipping equal to the manager's salary and learned some more. I learned if you treat a crew like imbeciles, you get the same. If you treat them as self-supporting adults, guess what you get! They even correct problems they see and in one case, even push out those who will not pull their own weight. I learned what employees really care about—its not platitudes or company rah-rah cheerleading or f*****g turkeys or cards

or Sh**ty bonuses at Christmas. It's money … What a shock. When we gave employees a multi-thousand-dollar Christmas bonus (a big deal 25 years ago), they will follow you anywhere, be self-policing, and require little of your remedial time. Moreover, they cared about what we were doing and were proud to be a part of it. I also learned that there was a very small percentage of customers who were abusive; they were the 2 % of our users who were 98 % of our trouble. They would call and have our secretaries in tears. I learned that it was far better to refund them and tell them it was on the condition they never buy anything from us again. We have to be at work a significant percentage of our lives and it should be as pleasant as possible, even with a goal of being fun. The better we do, the better the morale; all I had to do most of the time once the self-policing work force was set up was stay out of the way. Speaking of staying out of the way, Neil, who we seldom saw was marketing us into continued prosperity and never ending back orders. He was amazing!

Oh, and lest I forget, all during this time of Theta's growth there was another audio company in a neighboring building run in a much more corporate manner. In their employ was a young

engineer (late twenties). He was intelligent, street smart, and nearly as irreverent as I am. He hadn't yet learned to take a look around at his audio peers and stop taking himself so seriously, but he's improving at that today. He was still trying to figure out what he wanted to do when he grew up, but then again, so am I as of now. He was a great engineer, but kept babbling about marketing. We began to hang out a lot and he was just like a blotter, sucking up every Theta experience I had. He had a sardonic sense of humor—but, before I get there, there is a back story. At Theta we worked with a recovered cocaine addict named Fred Caccione (name changed). So one day, I walked into the Theta production men's room and looked up on the wall to see a new poster: "The Fred Caccione Cocaine Weight Loss Program" then something about white powder, no fussy diets, appetite reduction, and send $50 000 for your first ninety days supply. I laughed my ass off, and Fred thought it was just as funny as I did. By the way, if you haven't figured it out, that engineer was none other than Jason Stoddard. This was the first of a series of posters from my current Schiit co-founder. He was starting to loosen up!

There were some good times there. We were like

a big family, with enough dysfunction to make it terribly interesting. I had a secretary named Ann (again not the real name). She was one I had rescued from an abusive customer by firing and refunding him. She was in her early twenties, amazingly cute, and had an aura of innocence. I always followed the never schiit where you eat rule and as a result was always well trusted by my female staff.

One day she walked into my office, closed the door, and said to me very seriously, "I have to show this to somebody". She then began to unfasten her pants. My eyes were enlarging as she turned around and revealed a bare rear end so perfect that DaVinci himself could not have done it justice. There was a still healing tattoo of a rose on her upper right ass-cheek. Turning around, she asked me what I thought. I told her that the tattoo was beautiful. Relieved, she pulled her pants back up and thanked me. I could tell she really meant it.

So one day Jason comes over to my house for a batch of Mike's famous margaritas—he mentions that Theta, (which now was a multi-million company with three lines of DACs, and a couple lines of transports, as well as accessories) needed a cheaper series of DACs. He already has a name

(Cobalt), an industrial design, and a price. Even though we were probably drunk (we only got drunk once—it just lasted about ten or twenty years), we worked the electronic design out quick even with jitter reduction to fit the projected $600 price. ($360 direct equivalent) The outcome was that Theta sold not thousands but tens of thousands of those. Lots of money. He was really right.

Even in my mid 40's, I still had the idealistic high-end attitude, kick all of the other high end pretentious ******* companies selling Jewel-encrusted over-machined arbitrary-industro neoart angular looking waaaaaay overpriced crap. I was all about making high performing products with money sunk into performance, not sculpture. Further, in the early 1990s the DAC chip technology had peaked with the PCM63, IMHO the best fu***n audio branded DAC ever made. The new marketing direction was delta sigma DACs, which I have flamed many times before, as being bit non-perfect, full of bad math and noise glare, etc. These were proliferating, with cookbook spec sheets and reference designs published so that even non-engineers could design them. This meant that anyone could build a DAC, and with the fall in prices on the new "audio branded,

sh**ty sounding parts, that everybody and his brother were coming out with me-too DACs with no value other than cheapest. Even big audio corporations that never had built digital products were hiring dish washers and landfill workers to design DACs. What was worse was all of the good DACs were disappearing. Even the new multibit PCM1704 DAC sounded like ass compared to the PCM63. It was time to get out of the DAC business. So I sold my portion of Theta to Neil. I decided once again what to do when I grew up—only to set myself down the most excruciating path I had traveled since Viet Nam …

Yggdrasil: The back-saga, Part 4— The Home Theater Detour

So, in the early nineties I was trying to figure out what I wanted to be when I grew up again. DAC design had become unrewarding as the good building blocks were disappearing from the chip makers, new DACs were beginning to appear as elementary school science projects, due to the cookbook approach of the new "audio DACs", and I was looking for a virgin territory for the high-end.

Home Theater—Perfect! The competition was produced from companies whose priorities were to produce equipment which adhered to Dolby technical standards—not to sound priorities which I had been honing over the last fifteen years or so at that time. At that time videophiles had laser disc based systems (DVDs had not yet appeared). Dolby pro-logic was the encode/decode standard of the day. The competing products were principally analog designs, and all had terrible UIs which required considerable attention

to learn. So, I produced the Angstrom 200 HED (Home Entertainment Director) all digital surround sound decoder/stereo D/A converter. It had such an easy to use I that you could set up the unit intuitively, all digital pro-logic decoding, and was by far and away the best sounding unit on the market. It won design awards, was amazingly well reviewed, and was ready for sale. The problem was that even though I had ponied my licensing fees , I had to await the appropriate licensor test and blessing; I was in a queue and it took a while to get through it. Lesson #1 was that licensed products are never released on my time, regardless of how much I have invested in the parts and how soon I need to turn my inventory. Not to mention how much my customers want it.

Lesson #2 was that high end audio dealers were almost never equipped to handle home theater customer support issues, other than take the equipment to the house, hook it up whether set up properly or not, collect the money, and leave. That left us to do the customer support. Bummer.

Lesson #3 was that there was a whole new crop of installers who had their own trade show who were really the main players in home theater. The customers would give a budget to the installer, and the installer would select and install

the equipment. Totally changed the to whom marketing vector. The show was CEDIA, by the way, and I was the first of the old high end guard exhibitiing there.

So, AC-3 laser discs began to appear—the laser disc solution was to rf modulate a carrier—the first step of the Dolby Digital process was to demodulate the signal. There was a special purpose demodulator chip available to do just that. The problem was that there was no channel through which to buy them in the US. The only way was to send develop our own grey market channel for the parts (Japan or Hong Kong) that was quite an endeavor in the pre-internet era of telexes, and pre-phone deregulation era of very expensive overseas phone calls. Lesson #4 was that this whole home theater deal is really a big corporation's arena. I got the parts (some 6 months after all of the giants)—and came out with an AC-3 (Now Dolby Digital) adapter for the Angstrom 200—then waited for Dolby to approve the product (Repeat Lesson #1) I was a year late, and my customers were pissed. Lesson #4 again. Somewhere along the way, DTS capability was added to the Angstrom 200—I still cared about industry firsts.

About that time, DVDs and DVD players appeared.

No more need for demodulators. Cheaper software. Time to come out with a cheaper model. The Angstrom 100 HED appeared, based on a Crystal Semiconductor Codec and a Zoran DSP engine. I now had a big HED and a little HED. ("So I got a little HED and it totally changed my perspective on home theater"—He'll never admit it but that was totally Jason's idea.)

Months later, Dolby approved the Angstrom 100 and it was released to a user base that largely already purchased other gear.

Shortly after that, the sold Angstrom 100s began to return—with blown up Crystal Semiconductor codecs. I called their engineering support. They insisted that my power supplies were improperly designed. I countered that they were the same power supplies I had used in several other Crystal Semiconductor designs. The returns continued to pile up. I continued to replace the parts until they called me admitting that the codecs they had already shipped me did not work well with my improperly designed power supplies and they were going to replace them with a different date code of parts. I kept on going with the new parts finally fixing the failing in the field units. Lesson #5: If you can't afford a lawsuit, you are screwed. Lesson #6: If you can afford a lawsuit,

you are still screwed. It was too late—I was doomed. I shut it down. I lost a lot of money—but along the way I had still made more.

Let me count the ways why manufacturing to someone's licensed technologies suck.

a) You are doing all of the heavy lifting and taking all of the risk for your licensor.
b) You are entering into a crony based association where the licensor will always do what is best for him in his time.
c) You surrender certain freedoms—inability to ship pending licensor's approval, subject to licensor's priorities and other relationships.

Not to mention lessons #1 through 6 above.

So is Schiit going to do home theater??

Over my dead body.

When hell freezes over.

When shrimps learn to whistle. I know, but just in case I was vague, NO.

Rina

Anecdotes

French fry containers

Funny thing about those french fry containers. When I was gearing up to solder the very very first run of Asgards, I wanted a bunch of parts bins for my workbench (ie. plastic folding table wedged between a car and other assorted garage junk) … and I needed them right now. So, I went to the nearest store that sold takeout food containers and the like. It'd have to be big enough to hold a good amount of parts, shallow enough that I could stick my whole hand in it to scoop out parts and needed either rounded or slanted sides to facilitate that scooping because if I'm gonna be stuffing and soldering, I gotta do it right and not be fighting with keeping all those parts in line. Those darn french fry boats are the best. I could write on them. If they got too dented up, I could just throw 'em away …

4 years later, I think we finally went through that bag of 250 I bought. Best $ 7 I ever spent!

Yes, as my screen name (Soldered1st1K) implies, I soldered the first thousand boards at Schiit. Yes, if you've read previous chapters, you've heard that Jason's wife was part of the story. Well, I've been reading this whole thread since the beginning and enjoying every magical minute of it. Why now have I finally posted? I dunno. Someone saw a french fry boat and I got nostalgic. I still remember it like yesterday. That ton of french fry boats full of resistors and caps and pots and LEDs, etc. allllllll over the workbench and there I was neck deep in schiit, stuffing and soldering and stuffing and soldering and stuffing and soldering. It was fun for me. I love to make things.

The '68 Mustang

There's been some discussion about old cars … I've owned my '68 Mustang fastback since 1993. Until 2001, if she wasn't running, I was walking. I've built and rebuilt her four times now and wouldn't trade her for the world. I enjoy the adventure of driving an old car. Most of all, I enjoy all the stories I get every time I stop at a gas station. Seems everyone's uncle, dad, cousin, etc. had one and wishes they never sold it because it was so much fun to drive. So much different than modern cars. Not better or worse, just different.

A few years back, this very car was what I used nearly every day to drop off amps at FedEx and the Post Office. I'd just load her up and fly. Best, fastest, most fun delivery truck for the job.

One particular Friday in March, I was a little late getting the orders out. There were only 2 or 3 because we were in backorder. The Post Office and FedEx are about 3 miles apart with the old Schiit Shipping Department (AKA the garage) in between. I grabbed the boxes plus the orders from my own little company and hopped in my old Mustang. Didn't bother with a jacket since I was only gonna be gone about 10 minutes. Didn't grab my phone, but so what, I was only gonna be gone 10 minutes. No big deal, right? I made the drop at the Post Office, hopped back in my '68 fastback to fly over to FedEx. Making the U-turn due to the divided road, there was a very loud, expensive sounding k-thunk and my transmission was now a box full of neutral. I coasted the car back around and parked. (Good thing I'm used to manual steering and manual brakes.) I didn't even have the reverse gear. There was nothing but nothing that'd move my car in any direction. I was stuck there, on the edge of nowhere because this post office was closer than then one in town. 10 minutes … yeah, right.

Hmmm … No phone. Still gotta get the 2 amps to FedEx … in less than an hour. How'm I gonna do this?

Luckily, the USPS people knew me by then. I was able to use the manager's phone to call Jason …

Except he wasn't paying attention to what I was actually saying when I said I was at the post office and my car was busted. (Kinda like the Great Schiit Storm of 2014 when it took five mentions plus photos before he understood it was a sprinkler INSIDE the building that was spewing.) I figured I might as well start walking. The icy wind was blowing pretty good through my sweater. No coat, remember … only gonna be gone 10 minutes … But I did have a pair of driving gloves in the console. Hey, better than nothing. They kept out some of the wind as I clutched those 2 Asgards and walked the 3 miles to FedEx because Jason sent a friend to get me at the wrong post office.

That day, I literally walked 3 miles uphill to deliver those two Asgards to FedEx and then trudge back home. Eddie brought a tow strap and we used his El Camino to go rescue my beloved bucket of bolts. Jason paid for the transmission rebuild because he knows the way to my heart is through

my car. After that, I started religiously bringing my phone and, of course, never had another incident.

The real reason I didn't just give up and go back home? The hill I had to walk up was far more steep toward the garage than it was toward FedEx! Admittedly, it wasn't like the Arctic breeze I felt while visiting the Scottish coast, but after I set the two Asgards on the counter at FedEx and walked back, my fingers had yet to uncurl. I had to put them under water to get them back to rights.

Plus, in Jason's words ('cause he's standing right over my shoulder at the moment), "Dude, she's a chick."

Uh-huh. :)

Alex

The Deluge

After such an great introduction on Wednesday (see Chapter 23) I end my week like this. I work with amazing people. Tony was today's hero. He called the fire department and the plumber. I was just the moron who stacked boxes a little high and rushed a little too much. You learn a lot about people when the Schiit hits the fan. Today I learned that the friends and family I have worked with for the past two years and five months are incredible, resourceful, forgiving, and blessed with amazing senses of humor. I already knew this, but it is nice to be reminded from time to time. Eddie, whose work station bore the brunt of the storm, was far more forgiving than I deserve.

Kudos to the Santa Clarita fire department. They came quickly and turned off the deluge before I could finish building my ark out of Ragnarok chassis parts. They even spent a good half hour pushing the water out of the warehouse. If you watched the news recently, these are the same guys who rescued the folks trapped on the roller

coaster at Six Flags Magic Mountain. You know you have a great boss when his response to all this is, “Schiit happens,” and “I’m glad I wasn’t the one driving the forklift.” Then he posts it all up on Head-fi.

The ironic part is the bin of Lokis. They have been sitting on the top shelf untouched since before Christmas. We’ve been joking about strategies for getting rid of the things for months. Buy a Gungnir get a Loki, buy a PYST cable get a Loki, door prizes at every headphone meet in North America, etc. I was even joking with Mike about this earlier this week. The trickster God had his revenge. If you are in the market for a Loki, don’t worry, we still have plenty in the shipping department, another case of assembled units in our backstock area, and a couple cases of boards that haven’t even been assembled yet.

The mystery chassis is from a Valhalla. It just looks bigger because of the perspective. I can confirm this because I rescued it from a watery grave. They hold about a pint of water. Please don’t replicate this experiment at home. It will most assuredly void your warranty.

To answer or expand on some earlier topics. I worked in public high schools for 14 years. The

last 8 years as an assistant principal. My number one job was to make sure that teachers and students had what they needed to learn. I wrote the master schedule, took care of budgets, handled federal and state compliance, supervised the guidance department, ran the state testing program, and handed out more detentions, Saturday schools, suspensions, and expulsions than I care to remember. I miss working with the kids each and every day. I don't miss annual layoff notices that are rescinded weeks before the end of school, ever shrinking budgets, unrealistic expectations from the district orifice, and the unyielding quagmire of laws and regulations created by the California State Legislature and Congress. I had two main rules:

1. Get teachers what they need to do their job and then let them do their jobs.
2. Take care of the kids—whatever they need—corrected schedules, someone to listen to them, someone to kick them in the butt, someone to protect them when no one else will—that is why you are here.

For sale: Crown Reach Truck. Used. Low hours. Small dent in top carriage. Slightly wet.

Answers and Anecdotes

On staff size

We have purposefully kept our staff to a minimum and have only hired when we absolutely had to in order to keep up with growth. Every person who works for Schiit either Jason, Mike, or I have known between 10 and 40 years. We work with our close friends and family. This week I have done something that we have really dreaded. I started sorting applications so we could hire our first outside of the friends and family employee. We have a large pool of applicants (application period is already closed, so please don't ask) and there are some really interesting people that will make good additions to the company. Experience and education are important, but how that person fits in to the company is just as important.

We want really smart people who we can trust to constantly evaluate what we are doing and come up with good ideas on how to improve the company. At the same time, we also have to guard against acquiring bloated corporate or

bureaucratic thinking. I've had enough of that already and Jason and Mike have little tolerance for that. What was probably hardest when going through the applications was passing over people who listed their first qualification as "managed a warehouse staff of 100 people", or "supervised a multidisciplinary team of professionals". I'm not looking for that guy. I'm looking for the person who did all of the hard work and made that manager look good.

They also have to be willing to work for a company called Schiit. As a product name, Schiit has some real advantages. It is unforgettable, irreverent, and does a really great job at weeding out people we probably don't want to talk with anyway. It is a little more difficult when your name is Schiit and you need to get your garbage picked up or establish phone service. I've been hung up on more than a few times. The best part is when you are talking on the phone to a very formal and professional corporation like FedEx or AT&T and they need to verify the name of your company. You can almost hear the sweat dripping down their face as they try to say the company name knowing that "this call may be monitored for quality assurance purposes".

When you work for Schiit you need to have a

certain attitude. We are dead serious about our products and not so serious about everything else. We have a good time. We joke. We tell stories. We even play the occasional prank on each other.

For me, hiring an outsider is a bigger transition than moving from the garage to the Schiit Hole or from the Schiit Hole to the Schiit Box. How do we keep the Schiit culture alive as we bring in more people and by necessity start giving our original employees supervisory responsibilities?

On pricing

Things that need to be considered when setting a price:

1. Cost to purchase the device components
2. Labor to stuff the board, test and program the board, assemble the board in to the chassis, burn in the unit, final testing, cleaning and bagging
3. Cost to package and ship the item
4. Warranty service costs
5. Overhead that is not directly related to the cost of the individual unit: Lease on the warehouse, liability insurance, customer service, advertising

6. Price theory: AKA why does every product cost end with a 9
7. Distribution model: B&M, Direct, some hybrid of the two- each of these models has costs that impact price.
8. R&D (Thanks judmarc)
9. Risk & Profit: Every product that is manufactured at the most essential level is a bet that someone will want to purchase what you produce
10. You risk losing your investment if you get it wrong. There needs to be a reward when you get it right
11. Without this, we'd still be a subsistence level species hunting with our hand-made spear chuckers
12. When James Watt invented the steam engine he wasn't doing it because his wife was out of town and he didn't have anything better to do with his time.

On Ten Thumbs

It is a good but challenging time at the shop. We are busy setting up the production line for Ragnarok, doing show prep for Rocky Mountain, and loading up on inventory for the Christmas rush. I've been interviewing job candidates this

week. Overall, we've been lucky and have some good people to choose from. I am looking forward to getting some more help on board. I am bit late to the party, but I wanted to share a couple of anecdotes about Ten Thumbs.

The first time I met Ten Thumbs was on a hot July morning. The Schiithole had a backyard that was never much use to us. It was accessed through a big chain link swinging gate, big enough to drive a car through. On this July morning, I drive up as I normally do and right away I notice that something it missing. The gate is gone. As I get closer, I see that it is not gone, it is laying on the ground in the weeds. The gate has been pulled off the hinges. Oh great, we've been robbed. I slowly creep around the corner with my trusty box cutter in hand. There was an old, beat-up red hatchback parked behind the shop. I should say mostly red as rust stains and primer spots don't really count. This guy steps out from behind the car. Big. Muscular. Handlebar Mustache. Curly black hair to his collar. Scarred rough hands. Oh Schiit. Is my will up to date? Did I pay the life insurance premium? Then I realize. He isn't putting our stuff in to the car. He is taking tools out of the car. He comes over and introduces himself. He was actually a really nice guy. That's

how I met TT. In 60 seconds I went from certain death to forming a very unique relationship that would last far longer than any of us ever wanted. This was the beginning of the one-month remodel that would take four months. Why was the gate on the ground. He couldn't find the key, so he just pulled it off the hinges.

The air conditioning and the door. It is August. It is beyond hot. I have been stuck in the shop babysitting Ten Thumbs. I would have loved to work in the evenings when it was cooler, but that wasn't going to happen. TT has been at this job for about 6 or 7 weeks now. He had finished with the outside and was now working on the inside of the shop. My shop. Where we build sensitive electronics. Ten Thumbs was anything but careful. For example, when he was painting the outside of the shop, he never taped off anything. The windows, the trim.. nothing. Get paint on it. No problem. Scrape the paint off the window with a razor blade. Paint on the trim. No problem. Just paint the trim again. He covered everything in paint, including my trashcans, my spare pallets, the crate we were going to use for shipping stuff to Rocky Mountain. Now he was in my shop. His first job was making the holes for the air conditioning duct work. A

careful contractor would have taped off the area, used a tool that produced a small amount of dust, and tried hard to keep things clean. We didn't have that guy. We had Ten Thumbs. I had already covered every surface in the shop with plastic. No dust was going to get on our products or work areas. I was wrong.

When making a 2-foot by 2-foot hole, what tool do you use. A drywall saw? A sawsall? He used a hammer. Two days of pounding. Why days. The holes were the wrong size. He had to go back and do them again. "No problem, I'll just patch it up later." The shop was covered in a fine, gritty layer of plaster dust. So I cleaned twice a day. Finally, he was done. The A/C guys came and installed the system and life improved significantly. The temperature in the shop dropped from 105°F at 2 PM to a nice 78°F. Then came the door.

We ended up taking the other half of the building the Schiithole was in. The space between the two parts of the building had been walled off years earlier and now TT was going to put in a door. This time he used a sawsall to cut the hole. The problem is he didn't know how big the door was going to be. No problem. "I can patch it later." So he cut the hole. Too big. Then came the door frame. The frame was not directly attached to

any studs or wood framing. It was shimmed in place with scraps of wood and a crap load of nails. It was crooked. No right angles at all. Then TT hung the doors. They wouldn't stay closed. No problem. Just need to add more shimming. Then they didn't close. Eddie and I would stand there each evening and examine the work and try to figure out how long the door frame would stay in the wall. Then I made my big mistake. I complained about the work. I sent a detailed email to Jason. He sent it to the realtor. The realtor sent it to the building owner. The next day the owner comes in to the shop and just lets me have it. I'm obstructing. I'm not TT's boss. I don't know what I'm talking about. He kept going for a good five minutes. He was beyond pissed. Fists clenched and in my face. I just let him yell. What I wanted to do was pick him up and throw him out the door. I didn't want to put us in a position where we got evicted from the shop. I am pretty certain that battery on the owner is grounds for eviction. One thing I learned working in schools is that when parents are pissed, let them vent. Don't explain right away. Just listen and let them feel like they have the power. Then, when they run out of steam, you tell them that their kid was suspended/expelled/not going to graduate. So I let the owner finish. I told him that I understood

and then he left. I called Jason. Whatever Jason did worked. That never happened again.

Things started moving a lot faster after that. Quality not so much. But things got done. Until TT forgot the hardener for the epoxy. That was my fault too. Just ask TT and the owner. I was the villain in their story.

On suppliers

We get the easy parts (common resistors, capacitors, diodes, etc.) from Mouser and Digikey. We get other easy and some medium hard parts from Mouser, Digikey and other online distributors. The hard parts (volume pots, memory chips, USB receivers, DAC chips, WIMA capacitors, wall warts, transformers, tubes, PCBs, Chassis, buttons, knobs, connectors, etc.) we get through specialized distributors, direct from the manufacturer, custom built, or smuggled out of former Soviet Republics on the back of Yaks. Building a Bifrost for example uses parts from about 19 different suppliers.

Other Ragnarok Facts

The Bill of Materials is over five pages long. The BOM is the list of parts that go in to a unit. For comparison, Magni is around 3/4 of a page. Bifrost is two pages. In an earlier post I mentioned that Bifrost has 19 suppliers. Ragnarok has about 35.

At 38 pounds shipping weight, Ragnarok is by far our heaviest product. The next closest is Mjolnir with a 14-pound shipping weight.

The main transformer is bolted to the bottom chassis and weighs 10 pounds. Jason handed me the first prototype Ragnarok transformer in March of 2012. At the time, our heaviest product was Lyr (7-pound shipping weight). I had no idea how I was going to ship the item this 10-pound weight was going to be attached to. A second smaller transformer sits on a PCB and the PCB is attached to the chassis with standoffs.

The Ragnarok ships in two double wall boxes (outer and inner shipping box) and sits inside of two soft, crush resistant foam trays. The inner box is big enough to fit four completely packed Valhallas. I am a big fan of packaging overkill. I know that your package is going to be thrown, dropped, possibly even kicked. It can fall out of a plane, truck, or off of a conveyer belt. We

design our packaging with that in mind. We get maybe 6 to 10 complaints about items damaged in shipping each year. I don't put fragile stickers on boxes. I also don't put "kick me" signs on my son's back when I send him to school. The two are roughly equivalent.

On bad (and worse) customers

Treat people like you want to be treated. It is a simple rule that has somehow gotten lost in this world of mega corporations and make-a-buck-at-any-cost mentalities. Having said that, working in customer service changes you. You start out fat, dumb, and happy. You are eager to please and strive to do your best. Then that gets chipped away by the angry emails, swearing, threats, outright lies, and fraud. I've been told to ****** off, go ****** myself, ****** you, and I had one very upset gentleman in Thailand tell me to sit on his big … I'll let you finish that one on your own.

90 % of the time I give the customer what they want. Extra time to test their gear because they had to unexpectedly go out of town. No problem. Waive restocking because the customer decided to return the Valhalla and go with the the Lyr.

No problem. 10% of the time, I have to say no. No, you can't pay us Tuesday for an Asgard today. No, I won't send you a free upgrade card for your Bifrost because FedEx delivered your package a day late, but let me see if I can get a refund on shipping. No, I can't make the customs officials in Dubai process your package any faster. Sir, we've already sent you two Valhallas and neither of them "worked" in your setup but both worked perfectly here. No, we won't send you a third but we'll be happy to take that Valhalla and give you your money back.

My favorite thing is when a customer tries to play Jason, Nick, Laura, or I against each other. There are four of us. We have this thing in our company that not a lot of people get to experience. Trust. We trust that each person is doing their best. We trust each person's judgement. We trust that they are going to do the right thing. We trust that they are doing their best. When someone tries to play us against each other, we compare notes, decide how to handle it, and then move forward with one person on point.

On Amazon

We have two distribution methods through Amazon. Our biggest channel, Fulfillment by Amazon is where our products go to Amazon's warehouses and they handle the shipping, returns, etc. We have to pack up our products in large shipping boxes and ship them to Amazon. Shipping costs are $2.00 to $3.75 each for our Magni-sized units, $7 to $8 for the Valhalla-sized units. The products are covered by Amazon rules. 30-day return, no restocking fee, Amazon fee, additional FBA fee, etc. As a result, costs for going through Amazon FBA are much higher than through our other channels. So we charge $10 more to help cover those costs. If you don't want to pay the $10, no problem—you can purchase the items from our website or through our merchant fulfilled Amazon stock. The margins on our sub-$249 products are very slim. If we absorb the FBA costs and don't charge the $10 then after Amazon takes their cut their isn't very much at all left for us.

All of our Amazon fulfilled and most of our merchant fulfilled stock is new a-stock product. The only exception being the listings for B-stock product. We do offer some B-stock on Amazon as it becomes available. The B-stock is clearly labeled as B-stock in the title of the Amazon

listing. I just looked at our Amazon listings. It looks like I need to check the B-stock shelf and update my inventory, there isn't much on there at the moment. When I list B-stock is tends to disappear very quickly. A few weeks ago I put 11 Asgard 2 B-stocks up for sale on a Friday afternoon and they were sold out by Monday.

Last Christmas we started working with Amazon fulfillment and it is by far our fastest growing sales platform. Amazon fulfillment is a program where we ship our products to Amazon's warehouses. Amazon handles ordering, payments, shipping, and returns in exchange for a percentage of the sale. We still handle support and warranty services. Overall it has worked well for us. We usually end up sending them a few pallets of product each week. Number one rule for working with Amazon as a merchant is to remember that you are swimming in their pool and they make the rules.

On Schools and Schiit

As you may remember from several months ago, I worked in schools for most of my professional life. There are two things that working in schools did not provide: freedom and a financial incentive to

excel. In fact excellence in a bureaucratic world usually just leads to bigger, more challenging jobs, and increasing levels of frustration.

When I went in to administration I was usually given the job of managing special programs. What that means is I managed state and federal programs for underprivileged kids, kids with disabilities, English learners, etc. All of these programs had major compliance regulations, required tons of reports, and usually had fairly good-sized budgets. When you wanted to spend money, you had to make sure the item you wanted to purchase was in your school plan and that your school plan had been approved by the school site council and the local board of education. Then you had to do your research and locate an approved source and if the purchase was big enough you had to get several bids and select the lowest one that could fulfill the job. After all that was complete, you sat down and wrote out your purchase orders, attached all of your documentation, took it to get it signed by the principal and then sent it to the district office to get signed by three or more additional layers of bureaucracy. And then you waited. After you got done waiting, you waited some more.

Fast forward to September 2013. We are finally

moving out of the Schiithole. I hated that place. Claustrophobic. Loud. No matter how many times I swept and mopped it was never clean. Then after Ten Thumbs did his work there were almost no right angles. The doors were crooked, the ductwork was crooked, the sink in the bathroom was crooked, even the toilet leaned to the left. The floor was uneven and sticky (no hardener in the epoxy). The Schiit-Box was amazing. 26-foot ceilings, giant windows in the warehouse, skylights, air conditioning, a giant roll-up door—I was in heaven. Except it was my job to move us in to the new place, get the utilities turned on, get it organized and outfitted. There was no school plan, no multiple bidders, and although Jason was great when I needed to bounce ideas, his advertising business was getting busy again and he and Mike were busy getting ready for Rocky Mountain. So I started working 14-hour days and I got us moved. Jesse, Tony, Bill, and Laura all worked their butts off to get us moved. I have probably spent about $100 000 in the last year on furniture, racks, equipment, security, contractors, and electrical and Jason hasn't said no or criticized anything I've done getting the shop going. It is a hell of a lot of freedom and a lot of responsibility. In school if I got it wrong I had five other signatures and a school council to

share the blame. Here it is only me.

In schools if you do a really great job, they usually reward you with more work. The new work is usually more difficult and more frustrating. If you are efficient it means you have more time to work on projects or to cover discipline for another administrator because they have a deadline coming up. If you are a good teacher you get rewarded with advanced classes that require more preparation and grading time or you get the really difficult students which means more frustration, preparation, and usually meetings to discuss academic deficiencies. What you don't get is more money. So when we changed our pay system and Jason told me I was going to get paid salary and a then a bonus for every item shipped I didn't know what to make of it. Then I saw the light. If I did a good job I got more money. If I did an amazing job I got even more money. That sounded like a great idea. It was about this same time that I started ordering parts and working with our board house. So I made the decision that we were never going to run out of product again, ever. I've kept our backorders to a minimum, that has gotten harder with the bigger product line. I've even taken my self away from the fun work (sound checks, assembly, order packing) because

other people can do those things. I can move the company farther by focusing on order processing, purchasing, preparing kits, interacting with vendors, and fixing production glitches. I expect that at some point I will have to leave some of those things behind so I can focus on our next challenges. That doesn't mean I still don't grab a screwdriver or a tape gun when needed. While Jason and the team were at RMAF, I spent three afternoons programming Modis so the night shift would have something to build. Those were nice and easy afternoons. Sometimes I miss the early days when all I had to worry about was processing orders and putting things in boxes.

Between the freedom and financial motivation it has been a very different and exciting place to work. It helps that we have a no ******** allowed policy in the shop. The atmosphere is positive and upbeat. Jason and Mike are the best bosses I've ever had. They trust me to do my job. I trust the people who work for me to do their jobs too.

Schiit Happens: The Ongoing Story of The World's Most Improbable Start-Up

2015

2015, Chapter 1

The Trials and Tribulations of Amazon

Ah, Amazon. One of the biggest opportunities for a start-up manufacturer, especially one that's avoiding traditional distribution. If you can keep it from eating you, that is …

Disclaimer: This chapter will most likely irritate many an Amazon devotee, since it looks at things from the seller side, rather than simply the buyer side.

However, to be clear, this isn't an Amazon-slamming screed. This is a realistic, eyes-open guide to the pluses and minuses of selling on Amazon. This is especially important for companies primarily doing direct sale, and are looking to expand their sales channel. For companies selling through distribution, especially those selling through large chains, none of this will be a surprise. In fact, it will probably seem like amateur hour to them, since they've probably

undergone much larger trials and tribulations than any that Amazon seller can imagine.

"Wait, what are you talking about? Trials and tribulations with big retailers? What are you talking about?" you may be asking.

So, before we dive into Amazon, let's talk about the alternatives first—looking at both the good, and not-so-good aspects.

Dealers and Direct

In the old days, your success was made (or broken) by your distribution channel—or, in English, "gettin into dem dealerz and movin boxes."

I've already blathered on about the disadvantages of the distributor/dealer margin structure, so I won't go into that at length here. Instead, let's look at a couple examples of the dark side of distribution.

Deep breath.

Let's say your grandest dream has just come true. You've built a successful product. Made some good sales. And now, your product has been accepted one of those "big box" national store chains they rail about on TV. They're placing

a huge order with you—bigger than anything you've ever seen before.

You've got it made, right?

Not so fast. Let's look at what might happen.

1. **That order may come in … but with a request for additional margin** due to the large volume buy. Note that this isn't really a request, but more an order. Can you afford to grant this?
2. **Or, that order may come in with 90 or 180 day terms attached.** As in, they pay you in 3 to 6 months after delivery (you hope—more on this later.) Can you afford to do this?
3. **In the fine print of the contract,** there's a note that if the products don't move at the rate the store expects (more on this later, again), they can return all to you at no liability.
4. **As a condition of getting in, the contract specifies the level of marketing** support they expect from you, which can include specific website functionality, colorful multi-lingual packaging, shelf-talkers (marketing wonk speak for stuff to get people's attention in the store, from small tear-sheets to video demos on a kiosk)—and, by the way, you pay for all of this.

5. **Once you're in, you get a call** from the chain's director of sales. Another competitor has shown up and is offering something better/cheaper/more fully featured. They want your margins adjusted down.
6. **Once you've been there a while, you get a report from the sales director** showing how your sales per square foot are less than the average for their store—so they want you to do a promotion to get the sales up "to snuff." More on this, too.

During my tenure at Centric, we saw plenty of this kind of thing. And we've seen it kill companies. So, let's tell some stories.

Takeover by nonpayment. One of our clients got bitten by the 90 to 180 day payment terms—strained to the breaking point by these unnaturally long terms, they patiently waited for the big payday … only to find out that the "write-downs"—AKA "money not paid due to items being returned, etc" made their position non-sustainable. The company was bought out by the chain at a very discounted price, since the only other option was bankruptcy.

Death by a thousand cuts. Another client got stuck building elaborate pop-up displays for his

products, complete with video demonstrations and even more elaborate product packaging, courtesy the largest brick and mortar retailer in the USA. Add in some bad advice from the store regarding product planning, and a "required spend" on advertising, and that one went belly-up, too.

The dreaded zero-price promotion. This one was great. A one-of-a-kind product with over 100 % per year sales growth sounds like a dream come true, right? The only problem was that it didn't meet the national chain's expected sales per square foot. So what did they do? They came to the client and said, "We want you to do a zero-price promotion." As in, the customer gets it for free. Who foots the bill for that? The client, of course. That is, unless they want thrown off the shelves. They did it, of course.

"Well, that ain't realistic when you're talking about audio gear, is it? Audio dealers aren't exactly national chains, right?"

No. Not usually. Audio dealers are usually small, and some are very good, and up-and-up about what they sell. However, some will (unintentionally or intentionally) limit an audio manufacturers' prospects by preferring to sell

gear with the largest margins. Hey, they have to eat, too.

And, believe me, if you're "lucky" enough to get into Circuit City or Fry's, I bet those stories above wouldn't sound off-base.

"Fine, fine, you convinced me," you say. "I'll just go sell direct, like you do."

Yep. Good choice. Lots of benefits. You are in complete control of your prices and your margin … but you're also in complete control of your marketing and customer service as well. You won't have a dealer to be your salesperson, nor will you have them as a buffer for customer support.

So, as long as you can get the word out about your product, and provide good support, you're golden. Except …

> You knew that was coming, right?

… except you're going to be battling two negatives—one relatively big, one relatively small, but only one of which might matter to you.

The big negative of direct. You're not going to be seen by everyone who might buy your stuff. Like it or not, not everyone goes to head-fi.org, or sees your ads on innerfidelity.com, and it

takes some really big-boy pants to do a real ad campaign on sites like Gizmodo. So you're going to be a niche company. I know that people are still discovering Schiit, and I also know that awareness has been a bugaboo with other direct-sale audio companies as well. But, you know what? It might not matter. Companies can get very, very big while being in a niche. This may not be a negative to you at all, depending on where you want to take your company.

The small negative of direct. You're also going to have some losses in direct sales—I'm talking fraud here. It's unavoidable. Your credit card processor is only a partial shield. You have to apply some reasonable intelligence to this to keep fraud to a minimum—like, say, not shipping that order to Malaysia when the credit card is from Tulsa, Oklahoma. This matters to everyone, but it's not too hard to keep the fraud at the noise level. Just don't expect that it won't happen to you.

Okay, so we talked about dealers and about direct. Now, let's talk about the modern hybrid … Amazon.

Not a Dealer. Not Direct. Amazon.

Let's start by getting the positives out on the table. At no time in history has it been easier or cheaper to get started selling on a site that funnels people ready to buy right to your products. Period. This is why you want to be on Amazon.

Go back and read that again. Yes, I am saying that your company should sell its products on Amazon. It can increase your sales—without cannibalizing your direct revenue.

Just be careful not to be eaten, yourself.

I'll get into that later. For now, let's step back a second and take a long-view look at Amazon.

The consumer side of Amazon. Amazon makes it very easy for people to buy—and to get sucked even further into their ecosystem, with free 2-day shipping and other perks on Amazon Prime, Amazon points, gift cards, easy returns, etc. It's a buyer's paradise. If you order from Amazon, you're assured you're going to be getting what you ordered, on time, without any liability. Nothing much to complain about here.

The economic side of Amazon. There have recently been articles on Amazon's new debt offerings, and speculation about how they are

burning more money than they might seem. Regardless, Amazon is pursuing market share at all costs. Consider that Amazon buys Google Adwords for Schiit products—and for Rina's Twilight's Fancy ribbon clamps! This does raise the question of "how long can this go on," though.

The seller side of Amazon. Now, this is where it gets interesting. Because there's a ton of good things about Amazon from the seller side … and some not so good things. To give each a fair shot, let's start a walk through the Amazon process …

So, Ya Wanna Sell On Amazon?

Go ahead. Go to Amazon, scroll way to the bottom, and click on the "Sell on Amazon" link. Pretty simple, right? Just two choices—Sell as a Professional, or Sell as an Individual.

Har, har. No, dude, that's just the entrance to the rabbit hole.

But let's follow it through. Forget the "selling as an individual" thing, unless you have a used Asgard 2 or something that's already listed on Amazon. Yes, that's right. Anyone can sell anything listed on Amazon, new or used. Cool if you're just looking to sell something, but it's of no

use to a business, because you can't create your own listing. You can only add on to an existing one.

But—companies take note about the *"anyone can sell anything on Amazon."* Yes, you read that right. Someone can sell a used Asgard 2 right under your main listing. Dealers can also sell their stock. That's why you sometimes end up with an Amazon listing with dozens of sellers.

And companies take another note—it is entirely possible to lose the "buy box" on Amazon—that is, the first listing that appears for the product—if someone else has a better offer. Yes, you read that correctly, too.

Wait. Stop. I'm getting ahead of myself. It helps to repeat this a bunch of times:

> Amazon has one goal—to sell you something.
> It doesn't matter from whom.
> Nor does it matter if it's what you were looking for.

Sorry, back on track. Let's say you sign up for Selling As A Professional. Now you can list your own products. You're set, right?

Well, um, no.

Signing up to Sell As A Professional is just a gateway to a complex morass of options, some of which are not well-documented, and some of which are evolving in real time. Let's break it down:

Basic Sign-Up and Set-Up Stuff. Once you've gone through the process to Sell As A Professional, which includes typical businessy stuff like providing a Tax ID number, bank account info, and credit card, you can list pretty much anything you like. If you're a dealer, you can simply add your inventory to a product already in the system. If you're a manufacturer, you'll have to create your own listings. Sit down for that … you're talking 6+ pages of info, with over a hundred options to fill out, including make-or-break stuff like keywords.

Self-Fulfilled Orders. Once your listings are live and you get orders, you can fulfill them (that is, ship the dang things) just like you normally do with direct orders. And here's where the gotchas begin.

- **Referral fees.** Yep, Amazon needs its cut too. In this case, it's 8 % for consumer electronics. Better bake this in.

- **Shipping fees.** Amazon is stunningly bad at calculating shipping, and requires you to use a flat fee or zoned approach. Which means you'll lose money on some shipments, and overcharge on others. Choose well, or you'll lose too much money—or irritate too many customers.
- **Amazon policy.** If you sell something on Amazon, it can be returned within 30 days. Period. Even if your policy is "no returns." Better plan for this.
- **No Prime.** You won't be eligible for Amazon Prime if you are self-fulfilling orders. Sorry, tons of people only buy Prime.

Fulfillment By Amazon. You also have the option of shipping your product to Amazon and having them ship it from their own warehouses. This also gets you the coveted "Amazon Prime" logo. Sounds great, right? Send one big shipment, then let Amazon take care of the rest. Well, again, there are caveats.

- **Shipping fees.** Remember, you're on the hook to ship to Amazon. To wherever they want you to send it. Don't forget this cost.
- **A complex morass of referral fees, pick and pack fees, return fees, ad infinitum.** Amazon has seemingly engineered its fees to

be as incomprehensible as possible. Forget that 8 %—it applies below a certain amount, but not above a certain amount, but there's also packing and warehousing and returns and probably fees for the wrong color shoes, as far as we can tell. Plus, the even easier returns through prime mean that your returns are gonna be 10× to 30× higher than you expect. Want a rough number? For consumer electronics, expect your *real* costs to be about 25 % to 30 %.

- **Amazon policy.** Remember those 30-day returns? Now you're also on the hook for return shipping. And when Amazon decides to extend their return period for the holidays, so do you. Tough if you don't like it.
- **Amazon reselling.** So what happens to those products that people ship back? You'd think they come back to us, right? Wrong. Amazon decides if they are resellable—sometimes reselling them as new, sometimes selling them through Warehouse Deals. Makes a great impression when Amazon decides to ship an Asgard 2 in a 12×12×12 box with no cord, no feet, and two paper-thin inflatable bags for shipping protection. Yes, this has happened.
- **Amazon delays.** We're promised that stuff shipped to an Amazon Fulfillment Center will

be available to ship in 72 hours, typically. Note the "typically." Sometimes it's much more. This sucks, bad, when you're out of stock … because if you're out of stock for too long, you may not be the first option that comes up anymore …

Sales by Amazon. Once you reach a certain size, Amazon wants to become your dealer. The plus side of this is that you'll always have the buy box, because your stuff is being sold by Amazon. The minus side is that the margin they want is pretty much the same as a dealer. Beyond that, we don't know anything about this, because we didn't take the bait. But if you read about some of the tricks that big-box stores do to their brands, I think that's a pretty good primer on what to expect here.

Sounds confusing, right? Wait. It gets even more fun.**Amazon Brand Registry.** Okay, let's say you're selling on Amazon … and one day, you find that you've lost the "buy box." That is, you're not the first result listed—when someone clicks on "buy," they're not buying from you. Worse, your replacement is claiming to sell the exact same product you make … but they are not you. They're not your brand. Their product might do something similar, but it's most definitely

not you … and you don't have any authorized distributors.

Impossible? Not at all. It happened to Rina, in her business. Multiple times. She sells ribbon clamps that are custom-manufactured for her, as well as ribbon choker necklaces that she makes, under the brand "Twilight's Fancy." Yes, there are other ribbon clamps and ribbon chokers out there, but they are not hers—and she has no distributors.

So, when this happened, she went to Amazon and said, "What the heck? (but spelled with F and U). Amazon eventually replied that she had to get into the brand registry to protect her brand—something they offered to Schiit at the beginning.

One catch: she had to have a registered trademark.

Yes, that's right: with Amazon, the rules change depending on who you are … and who you are rhymes exactly with "how much you sell."

Well, okay. She went out and got a registered trademark. Her business was already legally sound, as an LLC. She applied for Brand Registry, and got it.

Problem over, right?

Wrong. To this day, she has to chase off competitors who glom on to her listings.

Big deal, right? She's the brand owner and registered trademark holder.

Again, wrong. Amazon doesn't understand—or seem to care—that a manufacturer with a registered trademark and no distributors is the ONLY entity that can list its particular products. It does NOT allow the trademark holder to approve and disapprove resellers of their product.

Go back, read that again. And tell me if your jaw isn't open.

So what does she do? She follows Amazon procedure: she buys the competing product, describes the differences, submits to Amazon, and gets the competitor thrown off.

Yes. On her time. At her cost.

This is wrong, guys. Simple as that.

Experienced legal types may be salivating at the prospect of a class action against Amazon, since this is such a blatant violation of trademark and copyright that it isn't even funny.

So what do you do, if you're a manufacturer and want to sell on Amazon? The only thing you can do today is get into the Brand Registry, cross your fingers, and take the time to defend yourself when necessary …

… and, of course, factor that into the total cost of sales.

And it gets even more fun.

Amazon International Sales. For a while now, Amazon has been promoting itself as a convenient way to reach the international market. To date, this is pretty hilarious. All I can say is, if you think Amazon is your international savior, it's time to think again. Best to simply avoid it for now.

Why? Let's break it down:

1. **Limited international reach.** There are some places where Amazon is strong (some parts of Europe, Japan, Canada), and some places where it's essentially nonexistent (Australia, China.)
2. **In-frigging-comprehensible SKU conflicts.** Amazon international supposedly uses the same database as Amazon US. Not in our experience. Trying to simply add or list products internationally results in a disabling amount of

product conflicts. Want to re-enter everything manually? If you can (see below.)

3. **In-frigging-comprehensible Brand Registry conflicts.** Rina also had the fun experience of being told by Amazon UK that she couldn't list her products because they were already locked up in the Brand Registry. They were. By her. In the USA. Amazon is still apparently working this crap out.
4. **You're still on the hook for shipping.** Remember, it's not like the magic Amazon fairies come and pick up your stuff and fly it over the Atlantic with their levitation booties. You still have to ship it … wherever Amazon says.
5. **If you think FBA fees are incomprehensible, imagine what they'd be like with customs and VAT mixed in.** This is pure speculation, because after #3, we decided to hang it up for a year to see if Amazon got it fixed up.

So, if it's this much fun, why bother with Amazon?

Great question. I was just getting to that.

Why Amazon At All?

It's really simple:

1. **Incremental sales.** Amazon allows you to reach people who you may never have reached

selling direct. There is the real potential for significant incremental sales from Amazon.

2. **Amazon-only buyers.** There are a shocking number of people who won't buy anything if it isn't on Amazon—you'll be reaching them, too.
3. **Potential for positive exposure.** Good reviews reflect positively on you and your products, and get you out of the audiophile niche.
4. **Great way to get feedback on operations.** Amazon does keep you on your toes—if you have customer service issues, or product issues, they'll be fully exposed. This gives you the chance to identify them and improve.
5. **Helps buffer out-of-stock situations.** If you're in a situation where you're frequently out of stock, your FBA stock can help keep things moving.

To be clear: Amazon can be very, very good on a sales front. Don't dismiss it. Don't run away from it.

But, be aware of the not so good stuff:

1. **Confusing and time-consuming.** It's not easy to understand the nuances of the various Amazon sales options. It takes significant time to get your listings right, especially when you

factor in proper keywords and a complete, formatted description.

2. **Complex costs that are higher than you expect.** Amazon's actual referral fees, packing fees, warehouse fees, shipping fees, return fees, etc will be higher than they seem based on the Amazon overview—and good luck figuring out what all the fees will be from the start.
3. **Amazon's own operational problems.** As mentioned, there's no mechanism for a manufacturer to properly manage their own distribution network, nor to easily enforce trademarks. Expect delays on getting product into inventory, and sometimes incorrect products placed into inventory, or lost products.
4. **Potential brand reputation hit from Amazon's policy of reselling.** Whether they try to resell it as new or as a Warehouse Deal, you're taking the chance that a customer is not gonna get what they expect, right down to a naked product carelessly thrown in a box. Several of our negative Amazon reviews are directly related to this practice. And no, they won't automatically send everything back to us at our cost, even though we've asked.
5. **Defects in the Amazon review system.** Did you know you can leave a review for a product

you haven't even bought? Talk about a dick magnet. And good luck getting the trolls removed.

Amazon Abuse

Amazon makes it insanely easy for customers to buy … and insanely easy for them to return. You go into an Amazon purchase with no skin in the game. You didn't spend two hours listening to the dealer blather at you. And, there's no restocking fee if the product goes back. It just disappears. Gone.

This is great for the consumer, but any manufacturer selling on Amazon needs to be aware of it. Once customers discover how easy it is to return products, some of them will abuse it.

We've been watching the stats on Amazon, and have noted an increasing number of serial returners—some of which have bought the same product multiple times, or all products multiple times.

Hey guys—this ain't what Amazon is for.

And that's what manufacturers need to know, if they're getting into Amazon. Take a look at Amazon's acceptable return rate. It is shockingly

high. This is a good indication of where you'll be, once your Amazon presence matures. Better factor that in at the front. We didn't—and we got bitten.

Which is why you're going to be seeing some changes on how we sell on Amazon. Expect higher prices, beginning now. The lowest price you'll pay will be via our site—every time. If you want the convenience of purchasing via Amazon, there will be an incremental cost involved.

An Amazon Strategy

"Okay, fine, I want to sell on Amazon, what do I do to avoid your problems?" you ask.

Well, I can't say this is a universal panacea, but here's how I'd go about it, if I was to do it again. This is based on two years of experience and (redacted) sales—well, just consider that we're big enough that Amazon wants to be a dealer.

1. **If you can, wait for Amazon to contact you.** If you do, they'll help you get set up with Brand Registry. They'll also even list your products (sometimes with quite hilarious mistakes—check them!) Amazon came to us. They didn't come to Rina. We get very different treatment.

2. **If you can't wait, go direct to Brand Registry.** Fill out the Brand Registry form and go through the necessary hoops, or you will have ZERO protection for your brand.
3. **Pay close attention to your listings.** Most of the product listings we've seen are astoundingly bad—limited detail, no keywords, sometimes even placed in the wrong category. Unfortunately, there's no big entity or network you can go to (like a Google Adwords expert) to help with this, but pay attention to the top-selling products in your category, take notes … and be prepared to tweak if things don't work.
4. **Go direct to FBA in the USA.** Seller fulfilled? You're just paying a referral fee to ship what you'd already be shipping, and gambling on shipping costs. FBA gets you into Prime, and even though the costs are more complex, they may not be much more than seller fulfilled. You also don't have to deal with UPC codes.
5. **Raise prices on Amazon to offset Amazon costs.** Yes. Seriously. Don't be shy about this. Rina's prices on Amazon are sometimes 2× to 3× higher than she sells for on other sites—and Amazon is her biggest channel by far. If people have Amazon points, or if they want to be absolutely sure they'll get the product—

and be able to return it—they'll pay more. Don't use Amazon's referral fees alone as a metric. Calculate your costs for shipping and returns as well. These costs can eat you alive if you aren't prepared.

6. **Don't waste time on International—yet.** Amazon is still working on this. Maybe they'll get it worked out. Right now, it's not very workable.
7. **Know that things can change at a moment's notice.** You are not in control of this channel. Amazon's 8 % fee for consumer electronics could be 20 % (like Jewelry) tomorrow. Or, they might roll in the ability for a trademark holder to easily manage their distribution network. Changes can be positive or negative … but you don't get to choose when they will happen.
8. **Keep watching … and adjusting.** Your results, too, will change over time. Don't be afraid to make changes to pricing and availability in order to ensure you aren't caught out.

Now, I know this isn't a complete guide. Hell, I haven't even gotten into Amazon's own advertising platform, nor into how to choose categories and keywords, nor on the need for patience—it

takes Amazon some time to ramp up to maximum effectiveness after you add a new product.

But if this guide helps someone get onto Amazon, make additional sales, and save their own skin in the process, it's done its job.

Jump in … but don't get eaten!

2015, Chapter 2

When To Listen, and When To Act

To cut to the chase:

1. When to listen to your customers? Always.
2. When to act? Only after careful consideration.

And that's really it. Sorry, guys, this is another business/engineering chapter. The fun stuff—on scammers and how to avoid them—will have to be in another 2 weeks. I wasn't able to get with Alex for long enough to hear all his war stories. He's a little busy, since January is a heavy shipping month historically, second only to December (yeah, weird, but maybe we'll get into that later.

If you want to press on and listen to the whys and wherefores of when to listen, and when to act, I'll try to keep it entertaining … starting with an anecdote that serves as a perfect example of how listening can go wrong.

The New Car

One of Centric's employees, and a sometimes-contractor to Schiit, is Denise. Denise and I go way back … all the way back to Sumo. Denise is the consummate planner, coordinator, logistics master, and generally make-it-happen person. Or, in other words, the complete opposite of Mike and myself.

You might have met Denise at the last RMAF, where you also may have noticed that we were significantly better prepared than any time in Schiit history. This is entirely because we brought her on to mastermind the whole deal. Bringing her in meant that we had things like rooms in the show hotel (Mike and I would forget), that we had all the equipment shipped in time to avoid rush charges (again, we'd forget), that we had a much more appropriate space with table-cloths, printed brochures, a backwall, the same computers throughout (little tablet machines—thank Tony for that). Bottom line, she helped us immensely.

But this anecdote isn't about that. It's about Denise and her new car. A couple of years ago, she bought a new car, a relatively fancy one, and it had one of those keyfob transmitter and

pushbutton-start deals that you see more and more these days.

“Hmm,” I said, when she showed me that cool new feature.

“Hmm what? What’s wrong?” Denise knows when I’m not thrilled about something, even if I don’t say it.

“I don’t know about those keyless things,” I told her. We also have a long history of talking frankly with each other, so it’s not like I was hurting her feelings.

“Why?”

“What happens when the battery dies?” I asked.

“Then I get a new battery,” Denise said. I sighed. “What happens when the battery dies in a crappy neighborhood, and it’s midnight, and someone’s been following you, and you can’t get in your car?”

Denise frowned. “You have a bizarre imagination.”

“No, seriously,” I said, thinking about it some more. “Can you even get in the car without the fob?”

“I don’t know,” Denise said, now looking a little nervous.

“Does it have a low battery light on it to let you know when you need to get a new one?”

“I don’t know about that, either.” Denise was looking significantly less thrilled.

“Are the interior door releases electric, too?”

“I don’t know,” Denise said.

We looked. They weren’t. But I knew there were cars that had electric interior door buttons as well.

“So at least you can get out of the car,” I said.

“That’s good, right?” Denise said.

“Yep,” I said. “And it is a great car.”

“But maybe I should buy some extra batteries?” Denise said.

“Or at least find out how to get into the car if it dies,” I told her.

And there you go. A cool new feature. One with lots of hidden ramifications.

Full disclosure: I’m quite familiar with these

kinds of systems, at least from a usage standpoint. The earliest one I had still had a key attached, on a 1993 Corvette that I ordered. It was literally brand new, first-ever tech at that point. I had a fun conversation with a car alarm engineer who refused to believe that such a thing even existed, until I showed him I could walk up to my car and have it unlock itself automatically—no buttons, no consumer input.

And I became familiar with using the key anyway when the batteries died.

And I became familiar with little glitches, like leaving the keyfob in the car and having it lock itself. Shaking the car moved the fob enough to fool it into thinking it was being carried around, and the car unlocked. Still, it wasn't a great neighborhood to have it happen.

And I do have another car with electric door poppers and pushbutton start—but it also has a key that you can use on a hidden lock outside to get in—and if you pull the door handles hard enough, they work mechanically, rather than electrically. Otherwise, I'd make sure I knew where the hidden pull-cables were.

Now, I know why you see these systems more and

more. It's because people think they're pretty cool, and they ask for them.

And there is something neat about walking up to your car, having it unlock automatically, and then just sitting down and pushing a button to have it start up.

But …

> Here it comes, you're thinking …

… but this is a perfect example of listening and acting without considering all of the ramifications. Let's compare a mechanical key and a transmitter/pushbutton start system.

Key:

- **Easy interface.** Everyone knows how to use a key. At least for now.
- **Reliable.** Well, unless the lock is frozen or the key is literally worn out or damaged. Fun fact: I once blew a new notch in a key while using it as a screwdriver on an alternator wire and shorting hot and ground … which meant I couldn't start the car after the repair because the key wouldn't turn in the lock.
- **Simple.** It's a machined piece of metal. Durrr.
- **Cheap.** Lose it? Get a new one at Home Despot for $ 7 or so.

Transmitter/Pushbutton:

- **Easy interface, but with gotchas.** Push and go, right? Sure, if the battery in the transmitter is good, and the car battery is charged, and there hasn't been a mechanical or electrical failure in the transmitter, and if there hasn't been a mechanical or electrical failure in the car. And what happens if it is dead? How do you get in the car? Or out of it?
- **Reliable … maybe.** I'm sure plenty of really good engineers sweat the details on these systems to make sure they are reliable as all getout, because the downside is so significant. However, there's simply more to go wrong.
- **Complex.** Transmitters, electronics, code, batteries, switches, pushbuttons, receivers, pairing, security, etc. You can't deny it's complex. Read some articles on how messy the code is in cars, and you may start getting a very uncomfortable feeling in the pit of your stomach.
- **Expensive.** Junior tosses the fob in the washing machine … go to the dealer and pay $300 for a new one and programming. That is, if you haven't already lost the second fob.

So there you go. A perfect example of a simple system that is being replaced by one that's significantly more complex, costly, and inconvenient

for the customer. You may even be able to make a case that it's more dangerous, too (how do you get out of a car with electric doors that aren't working—yeah, I know there are failsafes, but do you even know where they are … or remember them in an emergency?)

And all of this is thanks to the best of intentions … to listening to and *acting on* customer desires.

The Reality of Unforeseen Consequences

"So is this your passive-aggressive response to the whole 'switchgate' thing?" some wags will ask. "Do you really think keys and transmitters have anything to do with where you put your switches on the chassis?"

Well, no and yes. No, because I'm gonna talk about switches (and about other feature-type stuff) here, and yes, because I do think they are related—both are about replacing a simple system with one that is more complex, and has significant downsides.

First, let's talk switches. Why are our switches on the back? Lots of reasons:

- In many of our products, it's simply impossible to put an AC switch on the front, at least in the

expected left-side location. The transformers take up too much space, and are frequently only 1/4 inch to 1/2 inch away from the front panel. A good AC switch like the one we use on the back needs over an inch of depth.

- It's closest to the AC entry point into the chassis, and minimizes the length of AC traces we have to run around the board. We don't like AC running all the way down the board, as we'd prefer to keep its field away from more sensitive regulated supplies.
- Plus, as Mike said, we like to leave things on.
- Plus, there are these things known as power strips that have switches. Some even have remote controls.

Re-engineering to move the switch to the front could be possible in one of two ways:

- **Very awkward switch position and long AC line runs along the board.** How would you like to have a switch that's 3 inch to 4 inch away from the left edge of the chassis? Nice. Plus, you still now have AC running all over where we wanted to have regulated supplies.
- **Significant price increase—if it fit.** A nifty soft-touch pushbutton controlling an AC relay needs significantly more electronics, plus a keep-alive transformer—translating into

higher costs, assuming the extra transformer would even fit.

So, I think it's a comfortable assumption that the switches will stay on the back.

But, while we're at it, let's talk a bit about some other convenience features. Now, believe it or not, we hear you … we do listen all the time, and use your input to shape future products and product plans, within what is possible in our budget. Sometimes you gotta remember that we're all about bang for the buck.

Here are some other things we get asked about, and some reasons why we're still considering how best to accommodate them … or why we don't think we'll be able to accommodate them in the future.

- **Universal AC voltage.** Oh yeah, this would be great for some customers. The ability to plug into 100 V to 240 V and have it work everywhere in the world isn't a problem for Apple chargers and such, so why can't we offer it? In short, because it is either complex and expensive, or it's a compromise. Apple can do it because they use switching supplies. We won't use switching supplies, because we

believe that they are usually very compromised (in terms of noise) when compared to linear supplies. Linear supplies can use complex switching systems to detect and swap transformer primaries, but again, you need a keep-alive transformer and quite a bit of electronics. It's pretty amazing how big a good auto-switching linear power supply is. It won't fit in a lot of our gear, and even if it did, it would significantly affect the price.

Our solution? Perhaps in the future you'll see a voltage selector switch on the back panel … but suitably protected. If someone switches it to the wrong voltage, it'll be a bad day. And even if recessed and protected, many someones will switch it wrong. We're still debating this one.

- **Fancy screens with nifty graphics.** Yep, they're sexy. They're also quite noisy, as in electrically noisy. They also require quite a bit of programming resources. This would require that we hire additional people, probably full-time. Translation: much higher cost. And what happens when you have to reprogram for bugs?

  Our solution: there probably won't be screens in our future. That's what computers, tablets, and phones are for.

- **Remote control.** Yep, as we get into the speaker side of things, we understand that remotes can be quite nice. They're not as important as, say, getting your grandmother's oxygen tank delivered on time, but they are nice. Of course, remotes are also the thing that is most often broken or lost. Plus, you'll either have to register remote codes (and pay for it), or else take the chance that somebody's surround processor might use the same codes (and you'll pay even more for that), or you'll have to go Bluetooth and do apps for Android and iPhone, and then you're a software company … with all the foibles that entails.
  Our solution: Still thinking on it, but we hear you. I expect you'll see a remote from us in the future. But I don't know what it is yet.

So what can I say here to sum up? It's that sometimes things that sound easy have many unintended consequences. These unintended consequences can be both on the owner's side, or on the business side.

On the owner's side, something that sounds simple (and very attractive on first glance) can increase cost dramatically, reduce usability, or cause inconvenience or even harm.

On the business' side, something that sounds

simple (and that a lot of people seem to be asking for) can wrap them around an axle by diverting efforts to complex, long-term development projects, and add significant cost to the products.

Bottom line … *listen, but dissect.* Ask yourself: *Am I really seeing all the ramifications of this request? Is it really as easy as I think? How much time will it take? Is the company set up to develop this … and support it in the future? And, if we do this, will it make a significant positive difference?*

And listen for what they *don't* say, too … but that's another chapter.

2015, Chapter 3

Our Favorite Scammers

Yes, scammers.

Yes, as in people who are actually trying to get your gear for free.

Yes, as in criminals. People who are committing theft.

And yes, I know, in today's blasé age, it's fashionable to think, "Well, it's a big company anyway, they're just gonna claim insurance, nobody gets hurt, screw the man and the corporations and blah blah blah … "

Well, ****** you. Trying to get something from any company for free, whether it's Time Warner or the smallest startup putting the founders' every personal asset on the line to bring a new, innovative product to market, is stealing. Nothing more, nothing less. Let's not sugar-coat it.

And let's not dismiss it so easily. The reality is that most businesses don't have insurance against

fraudulent charges—or at least not insurance that provides complete coverage. When we get hit by scammers, it costs … and it raises the price of everything we do.

"Wait a sec," you might be saying. "You use a credit card processor. They approve the transaction. Are you saying that some of those transactions they approve are scammers … and you're on the hook for them?"

Yes, that's *exactly* what we're saying.

No processor, no matter how good, and no matter how iron-clad their policies may seem, will take 100 % liability for scammers. Yes, even PayPal, with their "100 % guarantee against fraud." This is the reality. There's a lot of fine print for them to hide behind.

So, let me put this clearly, for everyone who's starting a business and planning on selling direct:

- *Yes*, there will be fraudulent charges that are approved by your payment processor
- And *yes*, you'll be on the hook for them

The Reality: Scammers Are a Law of the Universe

Let me repeat that for clarity. If you have a business selling direct, you'll be subject to scammers. You will be on the hook for some of them. And you'll lose money on some of them. Period. End of story.

So what do you do about it? Hang up the whole idea of starting a business?

No, not at all. You do one of two things:

1. Chalk it up as a cost of doing business, and factor it into your pricing.
2. Get very good at identifying scammers, and work to circumvent them.

We chose the second option. And, in the process of running our business, we've gotten very good (but not perfect) at identifying credit card scammers, enough that we've limited our losses to far less than 10 % of what they'd be if we relied on the card processor alone.

Now, we're not going to outline all the methods we use to identify potential scammers here, because, let's face it, some of them may be reading this. But we can definitely give you a general outline of how to spot the most obvious ones.

But first, let's start with the stories.

The Maybe Scammer, Or Why Pick-Up May Not Be An Option Anymore

Until now, if you'd emailed Alex, Laura, or Amy and asked if you can pick up your products directly from us, you'd usually get a reply like this: *"Sure. Let us know when you're coming by, between 10 and 4, Monday thru Friday, to pick up your order."*

In the near future, that response may read something more like this:
"I'm sorry, we're just a factory, we don't allow customer pick-up."

Here's why:

One Friday, we received an order for a very large amount of gear. Think, several thousand dollars worth.

In a business where most of our orders are for the budget lines, this raises red flag number one. The reality is, if you're gonna scam, you're gonna scam big (usually, unless you're very sophisticated), so might as well go for it.

Second red flag: the billing and shipping address didn't match. Both were to nearby cities in California, but they weren't the same.

Third red flag: the shipping address appeared to be an apartment, rather than a detached house. Nothing against people who live in apartments, but why would someone who lives in a fairly tony area of one town ship it to an apartment in another city?

This was enough for us to do a little more investigation. Google Maps showed the shipping address wasn't an apartment, but actually a PO box in a strip shopping center.

Woo-hoo. No way that was going to ship out.

We sent the buyer a polite note that his order looked a little, well, suspect, and that if he wanted the products, he could pay by a PayPal verified transaction (which would then put PayPal on the hook, rather than us.) As usual with most scam cases, we expected to hear nothing back from him.

Boy, were we wrong. This guy sent us a missive that was the email equivalent of the Tsar Bomba, calling our ancestry, motives, and ethics into question, belittling our judgment, and iterating

what An Important Person He Was and how he wasn't going to lower himself to pay via our instructions, no way, no how.

At that point, Alex responded, citing some of the things that made his transaction seem less than kosher, asked him to look at it from our point of view, and again invited him to pay via a transaction method which placed no liability on any legitimate party.

And that, we thought, was that. We went home for the weekend and didn't worry about it too much.

At least until Monday, when the guy showed up.

Yes, as in drove his ass an hour up to see us and give us a piece of his mind. He spoke first with Amy. Alex, luckily, was in the shop and rapidly intervened, asking the man to leave.

Again, luckily, he did leave.

But he did sit there fuming in his car (with his girlfriend/wife) for a good ten minutes. He was so visibly disturbed that the staff locked the doors, in case he decided to come back.

Again, luckily, he didn't. I keep saying luckily, because it could have been very different:

- Alex might not have been there, and then it would be Amy against a guy who apparently looked like an ex-high-school football player
- He may not have left when Alex asked him, which meant that we'd have to call the police—and be on the hook for assault charges if anyone even dared to touch him
- He may have come out of the car with a gun—disturbed people are, well, disturbed

But it all ended well, except for a large dose of stress we could have all avoided. If we went to locking the doors to the shop at all times, we wouldn't have to deal with it at all. So if there's a policy change, you know why.

"So who the hell was he?" you may be asking. "Some very ballsy scammer or a wealthy prick with an overinflated sense of entitlement?"

Honestly, we'll probably never know. Maybe he was really trying to scam us. In which case, he gets the prize for the most titanium-balled scammer we've ever dealt with. Nobody else has actually shown up on our door.

Or maybe he was just a guy with ample amounts of both money and anger, who wanted to prove a point to the small people in the world. In that

case, he also gets a prize—in this case, the #1 prospect we'd never want as a customer.

Aside (and warning, this may offend some people): Yes, there are people you simply don't want as customers. People who take time out of their day to drive up and verbally assault you are one category. People who have uncontrollable anger issues are another.

And there are more—an unfortunate fact of selling volumes of inexpensive product. We simply cannot provide the hand-holding/hand-jobs/psychological support that some buyers of 5-figure gear expect.

To illustrate, I recently had dinner with the founder of another audio company, one much larger than Schiit. He also makes largely high-value products, and we've shared some of the same business challenges, so we commiserate and share ideas often.

During this dinner, he asked me, "Do you get these crazy emails that are like 8000 words long from guys describing their history in audio and asking 35 questions about things like how the front end biasing is implemented and the ripple noise on the main supply and the listening tests

you've done on different PC board substrates and what selection criteria you used for internal cabling and how long will the product really last outside of warranty and stuff like that?"

"Yes, we do," I told him.

"What do you do about them?"

"We hit *Delete*," I told him.

He just goggled at me a little. "Without answering them?"

"Right."

"What happens when they come back insisting you answer their questions?"

"We hit *Delete* again," I said.

"Huh," he said, rubbing his chin thoughtfully. "I … but … hmm, well, we really can't serve them in the manner they're accustomed to anyway, right?"

"Right."

"And engaging with them takes our engineering resources and won't be to their satisfaction anyway, right?" I nodded. "Correct."

“So it might be impolite not to respond, but better not to engage at all.”

“That’s the reality,” I said.

He shook his head. “Gawd, we spend so much time with guys like that … I think I need to talk to the staff when I get back.”

Fair? Perhaps not to the super-detailed-tech-question-here’s-my-own-credentials-and-history-in-audio guys, not entirely.

But to everyone else, oh yeah.

In the spirit of full disclosure, we absolutely do answer tech questions, including very detailed ones from truly curious engineers. But when it gets to be a dick-length-measuring, futile exercise in “guess the answer I want to hear, ’cause I bet I know more than you,” to 30 different highly detailed questions with ideological landmines for answers, um, well, no.

Tales of Other Scammers

Okay, so we had one angry scammer. So what?

Actually, no, we’ve had tons of scammers. Some of them successful. Some not. But let’s go through some fun ones.

The Amazon shipping scam, or an idiot with balls of brass. Amazon really loves its customers. They love them so much that sometimes they don't care about what's right. Case in point: we had a guy order a Lyr on Amazon, then claim that he hadn't received it. Amazon refunded him in full, case closed. Or maybe not. Brass-balled idiot contacts our tech support a few days later, claiming the tubes he received were microphonic. Nick doesn't know about the Amazon refund, so sends a request to Laura to send new tubes to him under warranty. However, Laura and Alex know about the supposed "lost product," and contact Amazon to let them know that they have a scammer on their hands. Does Amazon care? Not at all. They do nothing. Lesson: if Amazon doesn't care about its own money, do you think they care about yours?

Alex's first scammer. Alex's first introduction to scamming came shortly after he started working at Schiit. This guy used stolen US credit cards and had the products delivered into Russia. Pretty obvious in retrospect, but you have to be burnt a few times before realizing that the "fraud protection" at your card processor is not exactly a 100 % guaranteed iron-plated guarantee. Net score: one Valhalla before we figured him out.

Our biggest scammer. Busy times of the year are worst for scams. We're moving fast, trying to keep up with orders, and we may not take the time to scrutinize them as much as we should. Which was exactly where we were during the holidays in 2013. But even if we'd looked at them in detail, this one would have been hard to catch. The billing addresses and credit card info actually matched (yes, this is possible). But the number of large-ticket orders going into and around one European city eventually tipped us off. Most of the orders we were able to catch and turn around, but we still got taken for multiple thousands of dollars in Mjolnir and Gungnirs. Fun fact: he came back this year for more. Stopped him cold. He got nothin.

The most sophisticated, enduring scammer. This one's based in the USA. Very hard to catch, because it's always different locations, different shipping methods, and sometimes even some "test buys" that go through on smaller products before moving up to larger products.
According to Alex:

> He was very very prolific and tenacious last year. I actually teased him at one point and told him that we had a contest

in our order department that anyone who spots one of his fraudulent orders gets a free lunch. He responded by telling us to enjoy our lunches because we had already bought him lots of great dinners.

Scam so much that your country looks bad. Although we hate to say it, we have to be very cautious about shipping to southeast Asia. From some countries, the number of fraudulent orders is actually higher than the number of legitimate orders. As Alex says about one particularly ballsy scammer:

> At the time we were authenticating identification as a secondary check on orders—a practice since abandoned as one more sophisticated scammer figured out a way to provide fake IDs. This guy flat out refused to provide verification and claimed that he had worked for Interpol and wanted to know how I could be so rude to require such documentation. When I told him that the order would not ship he insulted me for several minutes and then told me to "sit on his fat dick." That ended that conversation. I then discovered that this guy has been

blowing up Nick's email with dozens of inane questions and many dozens of crazy statements. It was probably best that he joined the not-a-customer ranks.

And when paranoia bites you. The problem with scammers is that, over time, you get a little paranoid. Sometimes that can bite you. Here's an example. We'd heard of a scam technique where packages were rerouted to a FedEx office for the scammer to conveniently pick up. We'd also heard about HeadAmp's experience with a high-dollar amp being jacked with this exact method. So when we get a very high-ticket order to ship to a UPS store from a fairly high profile person, paranoia kicks into high gear. Alex lets the person know that we suspect fraud, and that we're reluctant to ship the order. The customer then has one of his references—Justin at HeadAmp—email us saying he was the real deal. Oops. Needless to say we shipped the order.

Selling direct? Repeat after me:

I will be scammed.

But you'll get past it. And you'll learn …

Telltale Signs of Scamming

So you're starting a business selling online. How do you recognize scammers, and how do you protect yourself against them?

Unfortunately, I can't give you a complete guide, because, you know, some of our *scammers* may be reading this, and I don't want to give them a complete paint-by-the-numbers outline of everything we're doing.

In addition, what we're doing has become relatively sophisticated, and is not entirely a human process anymore. So there are things I may not be able to reveal, even if I wanted to.

So, do I think we're safe from future scammers? Not at all. But I do believe we can keep it below the noise level.

> **Aside:** and, if anyone out there is worried about the other side of all this—namely, having their own card number stolen from our online store—don't be. We don't store card numbers at all. Everything is done on the processor side. And if they compromise our processor, well, *tons* of people are gonna be getting new cards pronto, is all we'll say.

So, what are some signs of a scammer?

- **Much larger than average order.** If your average order is, say $300, and you get an order for $6000 from a brand-new customer, look a little harder at the order. Be especially suspicious if it is next-day shipping. They're not paying for shipping, they don't care if it's $400.
- **Billing and shipping address don't match.** Yeah, basic stuff here. And plenty of legit orders have billing and shipping that don't match. But ask yourself how often someone who's living in New York will be shipping to Saudi Arabia. Sure, maybe … but maybe not.
- **Desperate pleas for fast shipping, especially when applied with the above.** It's not beneath a scammer to contact you for fast shipping, sometimes the same day as the order. Why? Because if you wait a day or two, the card may no longer work. They have to move fast.
- **Shipping to a FedEx or USPS facility or drop box.** Again, this may be completely legit, but do you want to take a chance? Take a harder look, especially on larger orders, and especially if the addresses don't match.
- **Attempting to reroute the shipment once it's left your building.**

If someone tries to contact FedEx and reroute, this is a dangerous sign.

There are plenty more, sure, but that'll get you started.

So how do you combat fraud?

Well, we've all seen companies that won't ship to new accounts where the billing and shipping addresses don't match. But that's pretty draconian, and it is definitely inconvenient for people shipping to different addresses for a legit reason I do it all the time, for example … no guarantee that I'll be at home during the day, but I'll most likely be at Schiit, for example.) Plus, believe it or not, the billing and shipping addresses can match … and still be a fraudulent charge.

But you can do a lot simply by not being in such a hurry to ship. Same-day shipping is great and all … but it increases the possibility of fraud. Especially don't be in a hurry for much larger than average orders. And yes, I know, this is inconvenient for customers, especially those who legitimately need it next day—which means you'll have to put those orders first in line for scrutiny.

You can also lock down your FedEx or UPS accounts, to make sure nobody but you can reroute your shipments. This may not work in all cases,

but once it's on record, the shipper is on the hook if they reroute anything you didn't approve.

And then there's a lot of little common-sense things, much of them location-based. A US-based credit card going to overseas address deserves scrutiny. An $1800 order that is going to be delivered to a thrift store in South Central Los Angeles also should be sat on. A Valhalla/Bifrost stack shipping to an empty shack in Texas, same thing.

Want more? How about perhaps the biggest scam prevention method of all? Because, like it or not, most scams aren't very sophisticated.

So what's the method?

Simple. Every business has would-be scammers who claim non-delivery of a product. It's amazing how many of them find the item when they find out we need to contact the police or the postal inspector to obtain a police report for our insurance company. Protip: make this part of your written shipping policy from the start.

To everyone selling online: good luck on catching your own scammers!

2015, Chapter 4

Bridging the Gap

Okay, enough about business. Let's take a look at a much bigger question, one that's been prompted by the furor surrounding the introduction of Neil Young's Pono player.

In case you haven't been following it, the Pono Player has had a slate of extremely negative reviews from the popular online press—some saying it provides no audible difference from compressed files from iTunes, and some even going so far as to call high-res audio a complete and total scam.

Make no mistake: these are not just skeptical or lukewarm reviews. Many of them are downright inflammatory and dismissive.

Why such hate, I wonder? After all, this is just a simple product for playing back music. It's not even horrifically expensive. Why call the whole foundation of high-res into question? Why imply that anyone who hears a difference is deluded

or foolish? Why insult all audiophiles who are looking to improve their sound? Why not just say, "Well, I understand that these guys are trying to offer higher resolution music, and some people may hear a difference, but I don't, let's move on to the latest Android phone."

So that got me thinking, and led to a question: **Can audiophile products ever bridge the gap to a more general audience?** I think they can.

In fact, I think that many of the necessary "ingredients" are already in place. In fact, bridging the gap may be inevitable. I think the market for audiophile products can and will grow.

But I also think that we can speed (or hinder) the process. We can make decisions about communications, features, and price that broaden audiophile products' reach … or doom them to a ghetto.

How? Well, I'm not gonna claim to have all the answers. But I think I can at least post some directional signs to a brighter future.

What Went Wrong (With Pono)

First, a disclaimer: I was not a Pono backer. I do not own one. I do not know what it sounds like,

nor will I speculate on the sound. This commentary is purely on a marketing and communications basis.

With that out of the way, let me take a shot at answering, "Why all the hate for Pono?"

Well, I think it's partly the messenger. And it's partly because humans are wayyyyy smarter, on an instinctual level, than marketers give us credit for.

First, the messenger. I have nothing against Neil Young, but he may be the poster child for the absolute wrong messenger for this product. To the opinion leaders on today's mainstream online media, he can be seen as the product of a corrupt and dying recording industry. He may be perceived as being "rich and lucky," no matter the actual facts. So, having Neil come out and say he's reinventing music may start pegging the bullschiitometer. Especially if he's asking for money up front.

> And yeah, I know, musicians, even very successful ones, won't have the wherewithal to bankroll a rejiggering of the entire music industry, but we're talking *perception* here, not reality.

So let's look at this through a lens of cynical

perception. “What, one of the favorite sons of the hated recording industry is now wanting to make my music better? I bet he wants us to re-buy it all, too.”

Yes, I know. Unfair. But that can easily be the perception of the popular online press.

> It would have been interesting to see what the reception would be if Pono was championed by a young indie band … of course, the economics really don’t work in that case …

Next, the message. Let’s look at what Pono said on Kickstarter, with that cynical voice chattering in the back of your head. Parenthetical content is that cynical voice, other text from Kickstarter.

> Pono’s mission *(Argh, barf, we have a mission statement on the office wall, nobody looks at it.)* is to provide the best possible listening experience of your favorite music. *(Okay, sounds friendly enough.)* We want to be very clear that PonoMusic is not a new audio file format or standard. *(Ah, hell, you mean it will cost more than when I re-bought everything on iTunes?)* PonoMusic is

> an end-to-end ecosystem *(Oh, no, they used one of those big corporate words, let's make sure my wallet is still in my pocket.)* for music lovers to get access to and enjoy their favorite music exactly as the artist created it, at the recording resolution they chose in the studio *(Hmm, sounds interesting, at least.)* We offer PonoMusic customers the highest resolution digital music available. *(Wait, what's high resolution?)* PonoMusic is more than just a high-resolution music store *(AAAHHH! Store! I knew it, it's gonna cost a fortune!)* and player; it is a grassroots movement *(LOL whut, Neil has how many dollars?)* to keep the heart of music beating. *(BS, BS, I'm tuning out.)* PonoMusic aims to preserve the feeling, spirit, and emotion that the artists put in their original studio recordings. *(Yeah, but does it sound better?)*

And then that continues on to the site, from the product description:

> Pono's revolutionary high-resolution portable digital music player. *(What, my*

phone does that, doesn't it?) Designed to be the next best thing to live music. *(Wow, no hype here, huh?)*

The PonoPlayer transports you to a sublime musical experience, from the most delicate passages of a string quartet to the thunderous power of a heavy metal band. *(Huh, you lost me, does it sound better? Stop using them fancy words)* This portable audio player uses circuitry taken straight from Ayre's *(LOL, who?)* own top-of-the-line products, costing tens of thousands of dollars *(OMGWTF-BBQ, what kind of whacko spends money like that for music gear, I must be in the wrong place)*, for unparalleled sound quality and unrivaled listening pleasure. *(Again, wow, you guys really think a lot of yourselves, don't you?)*

Cynical? Yes. That's why I called it a cynical voice. That voice may not be 100 % accurate, but it's impossible to ignore. That cynical voice is that instinctual gut feeling, that "Wait, where's the catch?" doubt sitting in the back of your mind.

Because, even if the message is positive (better music, studio quality, as the artist intended) it's

gonna be tempered by actual experience. And that actual experience with the recording industry has included:

- Having to re-buy everything in stereo (if you're old enough)
- Detours into re-buying in 8-track and cassette (ditto)
- Seeing the first RIAA conniption fit about home taping, and seeing them spend big lawyer bux to try to stop it (ditto)
- Re-buying everything as CDs
- Realizing that CDs are insanely overpriced when it is widely reported that production cost is under $1
- Hearing stories about artists getting boned by the industry
- Seeing articles about grandmothers and teens getting sued for sharing music by the RIAA
- Seeing bands who take money showers in $100 bills complaining about people stealing music
- Re-buying CDs as iTunes downloads (if you're too lazy to rip them)
- Realizing what a crappy master the recording is, again and again

All of that sits in the back of your mind. And when someone from the recording industry says

they have the ultimate solution (yet again), those memories come back … and that little voice starts up.

Fair? Maybe not. But that's sometimes how the mind works.

What Went Wrong (Not with Pono)

Now, to be fair to Neil, it's not like the audiophile industry in general goes out of its way to be accessible, friendly, and in general warm and cuddly. In fact, there are many corners of this hobby that are actively hostile towards newbies, either actively or passively.

Actively, there are some forums out there that go out of their way to discourage what they perceive as neophyte posters—to the point of being openly hostile. It's easy to decide to throw in the towel right then and there. This helps nobody, least of all the general audiophile industry.

Passively, it's a huge shock to pick up many audiophile magazines and see 5-figure price tags and a cult of personality referencing many, many different people and companies that a new listener may have never heard of. Stereophile tries to cover some budget gear, but if a newbie's sole

meaningful content is ghettoized in one column per magazine otherwise filled with megabuck gear, it's not likely they'll continue.

And, even in the larger forums, there's typically no single place a true neophyte can go for friendly advice, simple explanations, and an easy guide to "stepping in" to better sound. It's almost always convolved with discussions of expensive products, heated opinions and infighting, references to industry standards and formats they've never heard of, nightmare stories about incompatibilities and other problems, and general information overload. There's a lot more we can do to make getting started a more pleasant experience.

Plus, the depth and passion of audiophilia is daunting. Endless infighting about the relative merits of digital vs phono, DSD and PCM, sample rates and bit depths, this DAC vs that DAC chip, USB vs SPDIF, ortho vs dynamic, planar vs cones, Class A vs Class D, objectivist vs subjectivist, ad infinitum. Is it a wonder that most newbies see this and run screaming, never to be seen again? It's a hell of a lot easier to buy music on iTunes and listen to it on your phone, comfortably unaware of the fuss and bother.

And, to add to this, the prices of audiophile com-

ponents are downright scary. I've said it before, but I'll say it again: looking at this business from the outside usually results in incredulity. Thinking, "That DAC does exactly what the one in my iPhone does … but costs 5× as much as my iMac," is a little crazy. Especially after you hear someone say, "Well, you have to spend that much to get good sound." (And then be immediately countered by someone else saying that DAC sucks.) It doesn't help that it's easy to bring up crazy, crazy examples like multi-thousand-dollar USB and Ethernet cables, magic fuses, cable suspenders, and other may-do-something-but-there's-no-measurements-to-quantify-it stuff that is eyebleedingly expensive. We can seem more than a little nutty.

And (yes, there's more) the infighting amongst the manufacturers in our little industry sometimes gets out of hand. It's easy to talk crap about other people's products (we get it, you really really like yours), and it's sometimes easier to engage in silly practices like shilling your own stuff on forums. Newsflash: this doesn't have to be played like a zero-sum game. Grow the industry as a whole, and we all benefit.

So what do we do about this? How do we go about moving audiophile stuff towards a more

general market?

Well, I think we start by looking to see if there are some parallels in other industries.

The Whole Foods Explosion

Back in the dim dark days of the 1970s, doctors decried "food faddists" saying that we might want to eat healthier … as in, you know, fiber, whole wheat, green vegetables, yogurt, less of the bacon and hot dogs, etc, and we should be exercising and taking some vitamins. The backlash against the "food faddists" was vociferous and extreme, always taking the party line: a balanced diet is all you need, celebrate surviving a heart attack with a steak (but stay off your feet, you don't want to exert yourself.)

This is especially ironic coming from an industry that *prescribed* smoking for weight loss during pregnancy, not so long ago. Seriously. Look it up.

Now, many of the "food faddist" recommendations they railed against in the 1970s are part of the mainstream. Hell, the brand-new dietary guidelines just swerved away from the high-carb diet recommendations that have been in place

for 30+ years AND began recommending some supplements. Holy crap, dude.

Consider that in 1982, Whole Foods had exactly one store. Now they are a juggernaut, pushing the frontiers of alternative food (and health) in new directions … including into the mainstream.

Now, before anyone gets their panties in a twist, let me say this: I am certain that some of the fringier parts of the alternative food/health industry are 100% bunk, just as I think that some of the rather eyebrow-raising stuff will be part of the mainstream in years to come.

So, things can change.

And yeah, I know, there's still controversy about lots of aspects of the Whole Foods thing, but I think we can agree that if soccer moms are shopping there, they are reasonably mainstream.

But there are parallels. Alternative food doesn't exactly have a ton of studies showing the efficacy of their products. In fact, many studies of organic vs conventional or GMO vs non-GMO have shown little difference. And yet the individual testimonials keep coming. Kinda like audio, where "big" tests like the Meyer and Moran study don't show

a difference, but individual experience stating the opposite keeps cropping up … even amongst the objectivist crowd.

So how did the alternative food market go from the belittled fringe to damn near respectable? With lots of time, plus discontent with the mainstream options, plus a fairly welcoming attitude.

- Time? We got that.
- Discontent? Believe it or not, we have that. Check the study currently on Audiostream—54 % want better sound.
- Fairly welcoming attitude? No way, Jose.

It's funny. Maybe it's because alternative food grew out of the hippie movement, it was more welcoming. It certainly was less combative and presented more of a unified front against the mainstream than high-end audio does today. And even now, go to a Natural Foods Show, and you'll see much more openness and acceptance than you'll find at any of the audio shows I've been to.

And, let's turn it up a notch. Natural and alternative food ain't cheap. It's priced significantly higher than mainstream. Has this been a problem? No. But then again, it's not like organics are 2500× more costly than mainstream, like

some cables. The cost is higher, but it is contained.

And let's turn it up another notch. Are there infighting and factions in natural food? Sure, but it's not even a hundredth of what it is in audio. In general, most people in the industry are more interested in helping, than in propagating a specific mindset.

Maybe this is unfair. Maybe it's easier to get into food because it is, well, food. But the results are clear. This industry has carved out a place in the mainstream. And we, as of yet, have not.

And I think this welcoming attitude has a lot to do with it.

Where We Are In High-End Audio Today

Negative, fragmented, and factionalized are the words that frequently come to mind when browsing online forums or looking at the comments on review sites. But that's not the whole picture. It may not even be entirely accurate.

But … it's easy to get that impression, especially when wading into an esoteric debate about interfaces or formats. It seems like we've gotten lost in the technological trees, and are unable to see

the quality-music forest around us. And that is to all of our loss.

We're also talking, largely, to ourselves. Audiophile discussions can quickly turn into a confusing palette of acronyms, abbreviations, short-speak, insider jokes, put-downs of mainstream companies like Bose and Beats, with references to obscure products and product designers thrown in for extra color. This makes us feel comfortable, perhaps, and it allows us to display our knowledge, perhaps … but it also excludes a lot of truly curious people out there.

And … the niche press coverage tends to go for the most exotic, esoteric, complex and beautiful gear out there. Which means there's an inordinate amount of articles about stuff that people outside the industry will see as ridiculously expensive—or even 100 % BS, like reviews of cables, magic damping pads, or devices that don't even connect to your system. This kind of coverage only reinforces our isolation, rather than inviting new listeners in.

But there's hope. A lot of it.

In audio, we already have a lot of active small communities, busily discussing systems and products, experimenting with tweaks and tricks, figuring

out some of the crazier tech incompatibilities out there (USB audio, I'm looking at you), and spreading the word. Head-Fi.org is the big one in the "personal audio" space, but it's by far not the only one. Look at AVSforum.com in the home theater space for another example. And there are plenty more, from the audiophile community in Reddit to small and scrappy start-ups.

We also have a growing online review community with active commentary. This goes far beyond head-fi, to the more established niche press, tiny start-up sites changing the rules, and even touches some of the larger tech community (though they tend to skew objectivist.)

And there are active audiophile communities in many of the tech giants, if our sales into Mountain View, Cupertino, and Redmond are indicative of anything.

And these communities, in general, appear to be growing. This is not the sign of an industry in decay. This is a very strong indication that high-end audio, or alternative audio, or whatever you want to call it, has potential for the future.

In fact, barring the fact that there isn't any community as strong as head-fi for the traditional 2-

channel home audio market, things are looking pretty good, in a niche sense.

So how do we reach across to the mainstream and bridge the gap?

The Path to Success (?)

Okay. Big disclaimer time. The reason the subhead to this section has a question mark is simple: I'll freely admit I don't have all the answers. My suggestions may be 100 % bunk, and they certainly aren't complete.

That said, what can we do to bridge the gap between high-end audio and a more mainstream audience that wants better sound?

First, I think we need to let go.
All of us. Audiophiles, manufacturers, reviewers. Let go of the notion of getting 100 % into the mainstream. Let go of the idea that Sennheiser HD 600s will replace Beats. Let go of the need to be #1, on top, and absolutely right. Because nothing is ever absolutely right for everyone. Accept the fact that some people will never be open to the idea of better sound. Accept the idea that some people who are open to the idea will like stuff you don't like. Growth into the mainstream

doesn't mean supplanting the mainstream. Don't think Ralphs and Kroger, think Whole Foods … it's a big niche, but it's still a niche.

Second, letting go ain't enough. Relax. Have some fun.
This is also aimed at audiophiles, manufacturers, and reviewers alike. Sometimes this industry seems like it's wound so tight that it's gonna strip the gears. Breathlessly awaiting word about the One True Recording on the Grand New Contender. Reacting angrily about any perceived slight to your Favorite Company or Grand Idea (see Pono … reviewers going insane about the negative reviews and ranting on mainstream sites did not help us, sorry.) Bickering about what is the One True DAC or One True Amp. Belittling other people's opinions when they don't match yours. None of this does us any good. Sit back, push the keyboard away, and take a deep breath. Go listen to some good music. Pet the dog. Play ball with your kids. Go on a vacation. Polish the car. Work on the house. And let the audio world turn. You may find it better when you come back.

Third, be friendly and positive—as much as you possibly can.
If I was somehow elected Grand Sultan of High-End, I'd also add, " … and start with a unified

message." Now, that ain't gonna happen. But we can definitely change the spin on things. If you're a manufacturer, you don't put other gear down. If you're working for a company in customer service, you don't laugh at the curious newcomer with Bose headphones. If you're a community leader, you don't do everything you can to shut people out of the site. If you're a reviewer, well … please continue being honest. But it's not like a negative review has to be written in spite (as the Pono reviews seem to have been.) But simply being friendly, rather than snickering into your coffee when someone mentions a mainstream brand, will go a long way to changing the perception of the industry.

Fourth, work way harder at actively welcoming the newcomers.

Okay, reviewers and community leaders, this one's for you. Why are there no "Getting Started" sections? Why are there no subforums dedicated specifically to entry-level gear. Why are there no dedicated "Hey, I'm new and I want to learn" sections with people who like talking to newcomers? Yes, I know, once you've grown into the great throbbing Donovan's Brain of audio, talking to the proles may be tiring … but you know what? There are plenty of people who wouldn't mind

helping. We can do a much better job at welcoming people in to the party, giving them enough info to get started, and letting them decide if audiophilia is for them. If it isn't, hey, no harm no foul. But if we're passively discouraging them by being so inward-focused, our growth will be stalled.

Fifth, get them together with more small meets and shows.

This one's for the community leaders and reviewers, too. This is something that's already happening. Go to a traditional show (RMAF, TheShow, etc), and you'll notice a couple of things. First, you could probably bowl down the hallways for lack of attendance. Second, you see astoundingly expensive stuff in seemingly every room. Third, you'll see the same morose people shuffling back and forth to the same rooms the whole show, muttering about some real or imagined fault in the system. Damning? Maybe. But we're already seeing the counterpoint to this. Shows are fragmenting and becoming smaller and more focused. There are more small meets. For audio gear that can easily transported, we need even more of these. Get people together as much as possible, let them swap stuff around, listen, and discuss what they hear. Let them

experience more of what's possible, at every price level. Because if we get people together in a small, friendly group, they engage, they learn … and they spread the word.

Sixth, address the elephant in the room. This one's for audio companies. Hello. The reason you're not seeing sales to younger people is simple: your stuff is too bloody expensive. Period, full stop. The world doesn't need another $1500 DAC/amp. It needs stuff that college students can afford. Oh, and by the way, college students, by and large, aren't into the puffed-shirt hyperbole that your 55-year-old marketing guy wants you to use … and many (and I mean MANY) are used to seeing teardowns of mainstream phones, tablets, and other gear that includes an estimated manufacturing cost. They ARE comparing that to your products. And if they see something that looks like it could be made for $80 selling for $2500, they have every right to be skeptical. We *need* more affordable products. And, er, reviewers … we *need* more reviews of them, too—preferably not done with power cords and interconnects that cost 6× more than the products being reviewed. That is a 100% surefire way to be dismissed by the mainstream.

Seventh, don't be afraid to defend your prod-

ucts and beliefs—politely.

This one's for the audio companies, too. Going more mainstream does not mean being eaten by the mainstream. There's a persistent belief that we need to have Bluetooth compatibility and wireless high-res transmission and phone app remote controls and touchscreens to be mainstream. ********. Bluetooth is a compromised product, capable of only lossy transmission. If you're not going to support it, state why in polite and non-confrontational terms, and acknowledge that if someone absolutely has to have a Bluetooth product, it may not be yours, and there are good ones out there. Wireless high-res is still in its infancy, has transmission issues and conflicting standards and really requires an onboard computer for login to networks. You don't need to potentially compromise the simplicity and reliability of your products to chase a standard that may change by the time you deploy it. Again, educate, be non-confrontational, and acknowledge that if you really really want that capability, it may not be from you. And so on. There's plenty of room for everyone. Your approach doesn't necessarily have to change … but you should always be polite.

Eighth, experiment and break the rules—but

not the bank.

If we continue pursuing the same old ideas of magazine pages and banner ads with a side of social media, we're not going to get much of anywhere. Maybe it's time to get together and fund some research to see if there really are golden ears out there. Maybe we need to get together and pursue the research that shows high-res music affects our brain in different ways (look it up, seriously). Maybe we need to have a more integrated approach that engages the popular press in a nonthreatening way, and invites them to consider the fact that there may be people for whom high-end and high-res really matter. Note the non-threatening, though. Press ain't gonna respond to a Pepsi challenge for high end vs low end very well … it makes them look like idiots if they have to reverse their position about not hearing differences, and it makes them look gullible and foolish if they don't. Maybe it's time to encourage indie bands to experiment with high-res with some subsidized bandwidth. Maybe it's time to let an organic message about "using the format the recording engineer is using," develop—without the breathless hype. Maybe it's time to do some more research into how discriminating the human ear is … hell, until a couple of years ago, we thought the human

nose could only discriminate 10 000 odors—now we know it's a trillion. Oops. But at the same time, let's not go crazy and create some grand and uncontrollable industry association that will milk people for dues … and then work only to propagate itself. We got exactly nowhere with the Academy for the Advancement of High-End Audio. We don't need a repeat of that.

And there you go. Speculation on what went wrong with the most visible high-end product in decades. Some parallels to another industry that went sorta mainstream. Some thoughts on where we are now, and what we can do.

Is it definitive? Not by a long shot.

Will my recommendations work? I don't know.

But I do know one thing: we ended up with a real drubbing with Pono. It's in our best interests to ask why … and find ways to change that in the future. These are my thoughts, as penned on one random Tuesday in February.

How about yours?

2015, Chapter 5

A Life in the Day Of ...

So, this is how it goes on a typical day, circa March 2015, at Schiit Audio. It's a long way from the garage operation we were only a few years ago, but it's still a small, lean company, with lots of day-to-day, well, schtuff …

"Hey, everyone, I'm here," I say, as I come in the door, usually around nine o'clock.

Alex's sitting at his desk, literally right next to the door. "Hey, did you order the transformers for the Valhalla 2s? The boardhouse says they don't have any." I frown. "The boardhouse is smoking their lawn."

"I'll have them look again," Alex says, through a groan. "Oh, and they got back to us with some Yggdrasil shortages on the first run, but I don't know if we can use alts to Mike's BOM."

"Send him an email."

"It'd be better to get an answer before he looks at his email at 2," Alex says. I sigh. "I'll look at the BOM, but Mike's the final word on anything critical."

Alex nods. He sends the BOM to me as I head upstairs to my office. Rina's already in the bullpen area, in a jumbled mess of a workspace separated from Schiit proper by alterating teal- and blue-colored trade show drapery. Above her desk is a big banner that reads, "Twilight's Fancy." She subleases space from Schiit—most of the completely useless upstairs section of the company (imagine carting, say, Ragnaroks up and down the stairs all day.)

All around her are racks, piles, desks, and drawers full of billions of shiny ribbon clamp ends, bits of ribbon, half-done experiments, finished products waiting to be boxed, and she's already cursing at Amazon.

"Look at this! They lost my shipment again! And on the other one they're saying that I had eight 10 mm velvet and two satin, but there was 20 of each, gawd, it's like shipping it into the cornfield!"

"Have you talked to them?" I ask.

Rina rolls her eyes. “It’s like talking to the wall.”

“But have you?”

“Of course! It doesn’t help.”

“But they pay you for lost stuff, right?” I ask, trying to calm her down.

“Eventually,” Alex calls, from downstairs. He and Rina commiserate quite a bit about Amazon craziness.

“I gotta check and see if I ordered Valhalla 2 transformers,” I say, heading for my office.

“You just told Alex you did,” Rina says, accusingly.

“Best to be sure.” So I duck in my office and check to see if I ordered the transformers. I did. I breathe a sigh of relief. Alex does a lot of the ordering, but both Mike and I do a lot of the more critical parts, like transformers and chassis and boards. And yeah, we miss things. Most of the time running a small company is about a pile of details. Eventually, the theory is that you can move most of the detail-y, day-to-day stuff on to more focused and competent people, but it always seems like that day is perpetually “another hire or two away.”

And, you know what? This isn't a bad thing. I've seen more companies destroyed by growing too fast than almost anything else. You gotcher venture fundin here, so it's time to go out and get a fancy office and three hunnert employees and everything will work itself out, right? Everyone just needs to follow your brilliant leadership, and presto, it's a money machine! Except it usually doesn't happen that way.

But enough yammering about that. Let's go back a few hours, because my Schiit day actually starts a bit earlier than 9.

Erase, Rewind

I get up around 6 every morning and walk—a fairly vigorous walk, about a half mile down a steep hill to the Chuys by our house, then back up to the house again. I could use the typical rah-rah CEO excuse about how walking clears my mind, allows me to think about what's ahead, plan and strategerize and all that jazz. And sometimes I do. But it isn't really about that. It's about an overall lifestyle change, one that started about a year ago, when I discovered I had high blood pressure.

Yeah, I know, probably not surprising. But it did get me off my ass. And I'm getting near to being able to discontinue my blood pressure meds, due to ongoing lifestyle change.

Yes, you do have choices in your life. This is one of mine.

Anyway, after my walk, the first thing I do is really what you'd call a "survey of the universe." I check in on customer service emails, I drop in at Head-Fi and some other forums, and generally see what's going on in the very small and focused world of high-end desktop audio. Most of the time, I don't have to really act on anything, though sometimes I respond to threads.

> **Fun fact:** If you see a response from me at this time, you can be guaranteed I'm sober, which may not be entirely true in the evening.

Why is this the first thing I do in the morning?

Because it's important.

Because I have to know what's going on in our market.

Me.

Not “the marketing manager,” or “the operations guy.” Me. That’s the reality of a small business. You don’t have many people to push things down on.

This morning, the threads about the switches on the back of Schiit gear have reached a fever pitch. Now some people are accusing us of being incompetent engineers again and saying it’s a safety hazard.

So I step in and post why, although it seems simple to armchair engineers, switches really ain’t gonna be happening on our products anytime soon. Of course, I’m polite, because, you know what, front switches are more convenient, and in some far distant future, they may happen, but when my first thought is “how the hell you gonna get that damn switch in there with the transformer’s fat ass sitting right where it needs to be,” well, there’s reality for you.

Then it’s some emails—order confirmations from new vendors, requests for wire transfers for some parts, a couple of personal notes, stuff like that.

On this morning, there’s an email from a big-name sound engineer—I mean, really astoundingly big, you’d know the studio he works for—asking if he could come by and give Yggdrasil a listen.

Inwardly, I groan. Yggys are at a premium. We have two 0.96 versions, two 0.99 versions, and neither Mike nor Dave nor I want to give them up. And two are going out to early listeners posthaste.

But … the guy was very nice, and not assuming anything. So I send him an email that goes something like:

> Wow, I wish we could, but we really only have two final engineering samples, and they're going out to reviewers shortly. But if you can come by sometime in the next couple of days, we may be able to do something.
> I'm copying Mike to see if we can make it happen.

After that, I do a little engineering at the home workstation—just a final tweak on a new prototype board. The big-time board work really needs hours of uninterrupted time. Mornings ain't that. But you are fresh, and that makes it easier to see mistakes.

From there, it's the proverbial schiit-shower-shave-shinola, then into work …

Back to My Office

On one wall of my office, I have a large whiteboard. On that whiteboard is a list of every product we have planned for the year, and condensed notes about what we need to do for each product to make it happen.

This morning, I start by erasing the Yggdrasil Owner's Manual part of the list. Then I groan as I see the Yggdrasil Photos item. Yeah. We still need photos. And if Big-Name-Audio-Engineer-Guy stops by, that makes the logistics of getting photos even worse. I add that to my current paper list, with a note: urgent.

Yes, I still do paper lists. In some ways, I am quite dinosauric.

On the desk is the latest prototype for the Schiit (redacted). It's one of two prototypes, and a real "production qualifier candidate." It had been working yesterday before I took it home. When I turned it on, though, it wouldn't unmute, and one transformer got very hot (never a good sign).

So I had to fix that prototype … as well as the other, earlier prototype, which had developed a nasty DC offset problem after running it for a few weeks with no issues.

Sigh. Engineering really is about a lot of little things, too.

First, I start working on the later prototype. It had been working, so it was probably something wrong with the solder. 4-layer, thick copper boards really, really suck when it comes to hand-soldering surface-mount parts, and sometimes you don't get everything stuck down to ground … and if that happens to be a bias reference, well, things can really really suck.

And yep, that's it … a couple of diodes had never really been soldered down to ground. They'd made good enough contact so that the thing worked on the bench, then had come unstuck during the trip back to the house. Boom. That had taken out some other parts as well.

After some disassembly, reassembly, and more persuasive soldering, the latest prototype was done and working … just in time for Mike to come in.

"Hey, is that the (redacted?)" Mike asks, plopping himself down on one of the damn-uncomfortable guest chairs in my office. Hey, they were free.

"Yep, it is."

"Does it work?"

“Yep!”

Mike brightens. “Cool! I want something new to listen to.”

“Hey, I haven’t even heard it yet!”

“You’ve been listening to two prototypes of this for the last three months,” Mike grumped.

“Two prototypes that were either supremely messed up, or only kinda messed up, not something we’re actually gonna think about making,” I tell him.

And they had been very screwed up. One had whole sections that didn’t work, and about a half pound of parts stuck randomly on the board. The second one only had a few tacks, but had a transformer that was too weak to get things into regulation.

“Come on, hook a fella up!” Mike says. I sigh. “Okay. Take it. But I’ll need it back. The new transformer is a bit weak.”

“I thought you fixed that?”

“I did. The latest version—the one that doesn’t buzz—un-fixed it.”

Aside: take nothing for granted. Nothing.

"Argh," Mike groans. "When does it come out of regulation?"

"About one-ten."

Mike nods. "That's fine, we're actually a little high on the AC side." He grabs the (redacted) to take it out to the car.

"I want it back!" I yell after him.

And groan, realizing I still have the earlier proto to fix.

Downstairs Rounds

I'm isolated upstairs, with only Rina and Tyler nearby, so I make it a point to go downstairs several times a day.

Aside: I also have a desk downstairs in the tech area that I use for tweaking, repairs, and collaboration with Mike and Dave—though most of the Mike and Dave stuff is moving to Mike's new downstairs office. I am envious of his desks, because they're big heavy industrial stuff you could set engine blocks on.

Before I can go downstairs, Tyler stops me. "I need your hand."

"That could be a really scary statement," I tell him.

Tyler rattles a stack of checks and hands me a pen.

"What, you can't forge my signature yet?" I ask.

Tyler smirks, probably thinking, *Anyone can forge that lame scribble.* I sign the checks and head downstairs before Tyler can come up with something interesting to talk about. That's the problem with (very scary smart) philosophy guys—you can get started talking about the whys of stuff … and get very deep, very quickly.

Tony's at his desk in the tech area, a barely-controlled chaos of boards in static wrappers and stacks and stacks of boxes with more boards piled to the ceiling. On his monitor is a video of some tech chick talking about the latest Android phones. On the desk is a Microchip programming puck.

"What's it today, Tony?" I ask.

"Ubers," Tony says, gesturing at the puck. Modi 2 Ubers have an onboard microprocessor that needs programming. "About a billion of them."

"That's a good thing," I say. We'd just narrowly averted backorder on the Ubers a few days before. There are certain products we simply can't let go out of stock.

"I've been thinking," Tony said. "We could plug in an external drive to our show router and use that at CanJam."

"With the tablets as clients?"

"Right." I frown. "If the hotel doesn't block it, like the last show."

Tony nodded. "Uh-huh. We still have the SD cards. But if we can add a drive, we can have a lot more music."

"Let's set it up here," I say. "See if it works with all 8 clients running uncompressed. If that works, we take it to the show as 'plan A.' If it doesn't work at the show, the SD cards are 'plan B.'"

"Will do," Tony says, and goes back to programming Modis.

And, you know what? Tony will do it. He—and Denise on the Centric side—make sure our shows

go flawlessly. And that's in addition to doing first test on every board that comes through our doors—and many of the repairs.

Yep, everyone wears a lot of hats.

The thing is, Tony really, really enjoys the tech stuff … and he really enjoys shows. Maybe the whole tech thing is really just to fill the time in-between. I head over to the other side, where Jesse is working on Ragnaroks, and cursing. Behind him are the burn-in racks, full of Bifrosts at the moment. To his side are Eddie and Miles' desks—our two main assemblers for the majority of our products. They aren't there, of course. They usually only come in at night. Maybe they're vampires. I don't care. They make good stuff.

Eddie's desk is clean and orderly, with only boxes of boards to assemble on its surface. All around the desk is customized with tiki heads, little weird figurines, old coffee makers, vintage prints, and a hundred other little nicknacks that make it like home. The walls are painted black, because Eddie wanted it that way … and did it himself. Miles' desk is similarly bare, the only ornamentation an ancient Silvertone guitar amp and a vintage Fender. Miles and Eddie both play, sometimes, at night. Miles once asked me if it was OK, and I just

laughed and told him, “We’re an audio company, someone should be a musician around here.”

“How goes it?” I ask Jesse.

“Ragnarok,” Jesse says, shaking his head.

“Yeah, they are a pain,” I agree. And they are. They are our hardest product to build. Which is why Jesse, our Quality Manager, oversees it—and does a ton of the hands-on work. They have to be right.

“When’s the next run coming in?” Jesse asks. “We’re getting thin.”

“Soon,” I say. “I know they’re running some of the boards right now, and they have all the kits.”

“Tell ’em to hurry,” Jesse says. I go forward into the sound check area, where Chris and Olivia are assembling Fullas. Eddie and Miles hate the tiny screws and fiddly assembly, so they do it. Behind them are carts of Magni 2s for sound check and another cart of various returns that need to be re-qualified for B-stock sale.

“How go the wars?” I ask.

“Great,” Olivia says brightly. She always seems happy to be here. Alex found her through Chris—

she's his girlfriend. Despite this, or maybe because of this, they work really well together.

"Can't complain," Chris says. "But when do the new Fullas come in."

"Soon," I say. "Maybe. Ask Alex. It's not my fault."

The pair laugh, and I go to the quietest office in the building—Bill's Zone of Silence. Bill is our original sound check guy, and also perhaps the pickiest human being on the planet. If there's something wrong, he'll hear it. If there's a ding or a scratch, he'll mark it as B-stock unmercifully. He listens to damn near everything we make at Asgard 2 level and above—and much of the Magnis, Modis, Manis, and whatnot. His office is a maze of sources—phono, CD, computer, portable, etc. On either side of his desk are Emotiva Stealth 8s. Above him is a rack of headphones. On the desk are stickers showing headphones killed by defective products. Beyond him is a wall of shame—a rack of B-stock that needs to be sold. I groan inwardly, knowing I need to make some listings for it on the Schiit site. We don't have everything set up for B-stock yet. Blame me.

And, now that I've said it, I'll have it up this week. Promise.

"Hey Bill, how's it sounding?"

"Good," Bill says, taking off his headphones. "Do you need me to do something?" I shake my head. "Nope, just saying hi."

"Hi." Bill slips the headphones back on.

And that's the way it should be. Bill's our resident Card-Carrying Audiophile, really really serious about gear. He's the best guy to be listening. By far.

From there, I circle back to the finished goods area. Alex and Amy are packing the orders for the day. Since it's a Monday, it's an insane time—all the orders over the weekend fill up several large rollable racks, a couple of which are dedicated to FedEx, and a couple more for USPS. During December and January, it's even crazier, but even in February and March, the pace is pretty brisk … in fact, this year, we're up over 40% to date, even with limited product that makes stocking everything at Amazon problematic. On some days, Laura comes in to help, but usually she's a remote employee, silently taking care of returns,

exchanges, and other order-related stuff, so Alex and Amy can focus on shipping.

“So, we have a new possible deal,” Alex tells me, before I can say anything. “There’s a new FedEx aggregator, part of the SCV economic development corp, that says they have great rates, better than ours.”

“And?” I ask.

“And they are better.” Alex says, looking a little uncomfortable.

“But?”

“But I don’t know what we give up if we go with them. Like, what happens with lost packages? Returns?”

“Do we have to ship from their warehouse?”

“No, but—”

“—But you gotta look into it a bit more,” I finish for him.

“Right.”

“No worries,” I say. “If it makes sense, do it. If it doesn’t, don’t.”

Alex sighs and looks relieved. I don't know why. He's the Director of Operations. Without Alex, things won't happen. It would be a very bad day. He has carte blanche to do what's right. And I have total confidence in his decisions.

Aside: David Ogilvy, the ad magnate, used to give his management a gift of Russian nesting dolls, to make the point that if you hire people who aren't as capable as yourself, your prospects will only get smaller and smaller—but if you hire people more capable than yourself, the company grows and grows. Alex is a lot better at I am at the things that keep the company running.

"Alex," Amy interrupts. "Shipping."

"Yeah, yeah," Alex says, looking at me apologetically. "It's Monday." I grin. "Go do what you gotta do."

Again, the theme remains: everyone wears a lot of hats. And everyone at Schiit I have complete confidence in. This is the only formula that makes sense, when you're not a 20 000-employee juggernaut.

Hell, I'd argue it's the only formula that makes sense, period … it's just that in big organizations,

it's a lot easier to hide in the team. I can't tell you how many times I've seen situations where 2 people do all the work of a 10-person department. I've seen it so many times that I believe it's the norm. Hopefully someone will prove me wrong.

Mo is sitting silently, listening to music as he's packing Modi 2s. He doesn't say much, but he's a critical part of the whole … silently making sure things go in boxes, a necessary part of a company, well, making things.

This is how it works.

Back To Broken Schiit

One round downstairs complete, I go back upstairs and cross my arms, looking at the broken (redacted.) Mike has gone on an errand somewhere with the production qualifier, and might not be back. It doesn't matter. He does a lot of his work out of his house, where he won't be bothered by distractions like high-speed internet access. And he and Dave might be back for one of their night-owl sessions working on firmware. Schiit is really a round-the-clock kinda deal.

Now, my broken (redacted) had been working just fine … until I'd put it in the first article chassis.

Had I shorted something in the process? Or had I just jostled something out of place? It was entirely possible I had the same kind of solder problem as the production qualifier.

But no. A quick disassembly and some prods with a screwdriver proved that everything was sound.

And yet … it would sit and bounce from half a volt of offset to minus a half a volt, and go on like that for several minutes. When it was warm, it was fine.

It didn't make any sense. I resoldered a few things, kinda at random, hoping for some change. No dice. I went over the whole thing with a 10× magnifier. Nothing. I ran it on the FLIR. Nothing strange. No dead devices, no crazy temps.

What had changed? Nothing but the case …

… but had I used the same (redacted?)

No. That had changed. I swapped the (redacted) out, and boom, it worked! I fired up the Stanford and checked performance against the stored values. It was running a bit better than before.

Cool. I now had a working (redacted.) Mike still had the latest one, but I had one with a working

transformer … other than the hum. Ah well. I added a note to my list:

> send an email to the transfo guys and let them know we need another prototype.

And So It Goes

Another trip downstairs. Some weird repairs that Tony can't figure out. Tweaking of a production step on the Ragnarok. A quick change to a new chassis. That's what most of what I do at Schiit …

… well, that is, in addition to chassis drawings, silkscreen artwork, manuals, product descriptions (fun fact: they're usually written before the product is real, and frequently before there's a working prototype—and then revised to reality, of course, if the product comes to light), sending stuff for photography, laying out ads and brochures, and the general marketing-y stuff like that.

It sounds simple. Something needs done, you do it. Or find someone who can do it better. And if they're not available, you pick it up. Or if you're not available, they pick it up. This is the rhythm of a small business, totally different than the structured world of a large corporation. I was once speaking at a marketing forum in a fancy

venue, where they put you up in a fancy room and feed you fancy food and have fancy things like formal nights and stuff like that, and met a woman who worked for a large pharmaceutical company. When she heard that I was a speaker and learned I had an agency, her eyes lit up.

"That's what I've always wanted to do," she said, breathily. "I've always wanted to have my own company." I laughed. "It's not so much about having your own company, as your company having you."

"What do you mean?" she asked.

"I mean, well, let's put it this way. You're the head of marketing, right?"

"Right."

"So if you need to send a logo to a trade show company—"

She shook her head. "The staff does that."

"Or if you need to print up a bunch of brochures—"

"Staff."

"Changes to the website?"

She crossed her arms. "I'm mostly strategy, really." I nodded. I knew how this would go. I could be polite and say, yeah, you get the right people, you'll be fine. But I'd already started going down the honest path, so why not keep on going?

"The thing is, when you're in a small company, there's nobody to fall back on. Everything lands on your shoulders."

"But if I set it up right … "

"No. That doesn't work. Do it. Learn it. Then hand it off. Maybe. Maybe you still want to keep some of it."

She frowned, clearly thinking I was crazy. "And that's what you did?"

"Yes."

"But … what if I don't want to do all the work?" I smiled. "You have an excellent career. It sounds like you have a great support network. Why would you want to give all that up?"

"But … " she stopped herself, looking nonplussed. "The freedom … ?" I laughed, long and hard. "Yes. And that freedom is hard, hard work."

To this day, I don't know if she ever started her own business. If she did, I hope it worked out …

and that she was very successful. But there are really no shortcuts. And I wouldn't have it any other way.

Now, of course this isn't a complete story. I didn't go into the rest of the day, or into the geek audio engineering in-jokes and good-natured ribbing that accompany Mike and Dave's time at the shop, nor into my evening at home, which usually has me sitting in front of the computer, working on one of the latest boards, or tweaking ones that aren't quite done yet, or showing off the latest prototypes for Rina when we're home—and having her shake her head or give me the thumbs-up, or waking up at 5 AM with a great idea for a new ad direction for Schiit, or answering some more emails and posts before I go to bed (and may have been drinking …), or the sketches of new ideas I make, or the long talks I have with Mike about future product plans and company direction, or the ongoing discussions I've had with some industry guys about some ideas to really shake things up. But I think you get the picture. I hope you enjoyed this little tale … of a life in the day of Schiit.

2015, Chapter 6

So Ya Wanna Get Into the Biz?

Prompted by a recent hire or two at Schiit, the ongoing resumes we receive on a weekly basis, and the more open-ended inquiries we get from time to time about

> I really really want to work in audio, what body parts amputations/small animal sacrifices/alien incantations do I need to get into this biz?"

I figured it might be time to talk about just what it takes to make a career out of audio.

Now, a disclaimer: no, this isn't a solicitation for resumes, nor can I predict what Schiit's hiring needs will be with any certainty.

No, instead, think of this as a general guide to getting into audio—whether you're looking to start your own company, or work for any company that makes great audio its mission.

Maybe this will help. And maybe it won't. Because first, let's start with the warnings.

Audio Ain't For Everyone

Let's get one thing out there up front: audio—as in great audio, not the consumer-flavored, Bluetooth-enabled, lotsa-features, convenience-with-big-bass-for-mass-box-stores stuff—probably isn't the easiest field to turn into big bux, sexy titles, massive power, or things that tend to impress the ladies.

Why? Let's run through the realities:

1. **Most companies who pursue great audio aren't large companies.** This means that they're rarely hiring … if at all. Some great audio companies don't have, or need, any engineering or marketing resources beyond their founders. So the chance of "getting in" is relatively small.
2. **Starting your own company has its own challenges.** Similarly, if you're looking at starting your own audio company, it's not likely to attract enough interest to net you a ticket into the venture capital lottery—unless it's tied into some big-buzzword, future-looking deal like distributed audio in the Internet of

Things. Not saying this can't sound great, but I bet that's fairly far down on the list of importance for these emerging companies.

3. **If you're on the engineering path, you're gonna have to endure the sneers and jibes of your fellows,** who consider audio largely a trivial field with completely known (and solved) problems. Saying that there may be something beyond −112 dB THD and IMD and inaudible noise makes you a bit of a shaman or High Voodoo Priest to much of the engineering establishment … well, unless you're working on one of those fancy "3D audio" or "object based audio" standards, where yep, there are shades of gray in the algorithm and implementation.
4. **Outside of engineering, if you're looking to be hired by a great audio company, the opportunities increase**—but may be increasingly frightening. Many audio companies need help in marketing and sales … but companies making great audio who are dependent on heavy marketing investment and an aggressive sales team may be the first to cut staff if results aren't what they'd expect. Many audio companies absolutely need operations help … but:
   a) They may underestimate the importance

of operations, preferring to work on the latest new sexy device they are cooking up—so the hires are never made
   b) Operations isn't a field that comes to mind for audio companies, so many people will simply ignore it
5. **Outside of engineering, if you're looking to start your own company, well … please, just no.** Unless your partner is an engineer. There's a pervasive idea that anyone can come up with a cool idea for a product and have it turned into a marketable reality using the Magic Capabilities of The Dudes in China. But that's an absolute fantasy. If you don't know why your product is better, and you don't know how to make it better, you will have nothing more than the audio equivalent of one of thousands of no-name rebranded Android tablets that sell for $44 on Aliexpress.

Aside on #3: oddly enough, despite the current climate regarding "no audible differences if it measures well enough" in the outside engineering world, I have yet to meet an engineer working in audio—even in very consumer-focused companies, or in the pro space—who doesn't believe that there are differences. They just don't feel comfortable discussing them. In the

same way you were "never fired for buying IBM" in the old days, you're not gonna get fired for delivering a device that is completely up to spec … but may not perform well sonically.

"Okay, you're a downer," you might be saying right now. "Why even bother? Sell real estate or get a ice cream franchise and call it done."

Why? Because great audio *is* for some people.

You can see it in the eyes of some show attendees. That wonder at being carried away by music … coupled with a deep-seated questioning stare that says, *why does this sound this way? How could it be better? Why can't I make this my life?*

If you're one of those people, audio may be for you. Because, you know what? Even if you're sneered at by other engineers working on satellite communications or smart home devices, even if it might not be the easiest field to get into, even if it's not going to net you the funds to buy a small island in the Bahamas, *you love it.*

And doing what you love, I think, is really what matters.

Getting Started: The Engineering Path

Yes, I'm gonna be unfair and start with engineering. Engineering really is the cornerstone of great audio. It's the easiest way to get in, whether you're starting your own company, or want to work for one.

> **Note:** easiest ≠ easy.

"So now you're gonna tell me to go out and get an engineering degree, right?"

Nope. Not at all. There are plenty of great audio engineers with no degrees. However, a degree will be helpful if you're shooting at a medium-to-large company. Even if they have small "great audio" enclaves, you're still probably gonna go through the churn-n-burn of a corporate HR department. That means: *no degree, no job.*

And, in reality, I have used little of the higher math that I learned in school. S and jw domain control stuff, sure, a little bit. But much of it isn't much more than basic algebra, backed by measurements, and underpinned by deep knowledge of how basic devices (transistors, tubes, etc) work.

I'll illustrate. At school, one of the classes I had to take was engineering thermodynamics. Now, thermodynamics are very important. You're not going to be able to design a reliable power amp without understanding thermo. However, the way professors go about it is absolutely retarded. They'll show you a "heatsink" with a weird cross-section that's circular and triangular and just plain wrong, and ask you to calculate its heat dissipation with differential equations.

When first confronted with this, I sat back in my chair and said, "That's dumb. Just look at the surface area and ambient temp, and you're close enough for any practical application."

Yeah, that professor didn't like me.

Bottom line, school complicates a lot of the basics of engineering with complex math that you'll probably never use. That math may be useful if you're looking to get a paper published in an IEEE journal, or if you're working on new DSP algorithms, or if you want to be a Ph.D in residence at a large company, but in day-to-day work, it's overkill. I nearly dropped out of engineering before they got to the control-system shorthand stuff … and even that I've only used a handful of times.

Plus, schools concentrate on simple circuits that are fundamentally unlike what you'll encounter in audio. After doing an "audio amplifier" on a breadboard in an engineering lab, seeing the schematic of an actual working audio amplifier will be overload. You'll wonder why the heck it's so complicated—even if it's a relatively simple design.

What I'm saying is, if you're expecting to come out of school and immediately be useful to a small-to-midsize great audio company, think again … you're gonna be fundamentally unprepared for the reality.

So what do you do?

First: start hanging out on diyaudio.com. This site can be thought of as "the leading edge of audio, mixed in with a thousand other crappy designs, dickish opinions, and complete drivel." But note "the leading edge" stuff. Some really good design work goes on there. Nelson Pass hangs out there. And the bad stuff is quickly dissected and dismissed. Joining DIYAudio costs nothing, and it's arguably better than any education or book—if you can keep up, and if you can start sorting the laughable from the laudable.

Bonus course: tubecad.com. Don't judge this book by its cover—it's not just tubes, and John Broskie has probably forgotten more cool topologies than I've ever imagined. Just know that this isn't a cookbook site. His stuff isn't usually completely worked out … but it's a great starting point to get you thinking.

Second: buy a book. Cordell's amp book is very good. No, I will not link to it. I am fully confident you are strong with the Google. Yes, I said "buy a book." Here's why: Cordell's book covers the basics of amplifier design from the ground up, including the whys and wherefores of different devices, strengths, weaknesses, why you use them, why those damn amps look so complicated, etc. It doesn't really get into small signal design, nor does it get into digital at all, but it's worth reading, re-reading, and taking notes, even if you already have your engineering degree. It'll remind you what matters for audio.

Third: start burning stuff to the ground—er, I mean, building stuff.
This is absolutely key. Don't invest in simulation software, don't create air-castles of circuits in fantasy-land. Get some breadboards and heatsinks (and, these days, surface mount break-

out boards and Bluetooth modules and ready-to-go DACs) and start hacking stuff together. Tube or transistor, it doesn't matter. Discrete is better than op-amp for learning, plus it shows better if you're asked what you've done in audio. But the point is this: you need to build things. Building is the best way to understand what's going on … and it can either become the foundation of your own company, or serve as an example of what you've done for a prospective employer.

Wait, one point of sanity: actually, tube or transistor does matter in one way: tube circuits frequently use voltages that can actually kill you. There's a great thread on diyaudio.com about tube safety. Repeat: you can kill yourself. You can also burn yourself on overheating power transistors, watch capacitors fly off the board because they were installed backwards, crack chips in half with the wrong power supply, and watch an entire design self-immolate (complete with flames) because it's oscillating uncontrollably. Be careful.

Fourth: ask lots of questions about stuff you don't understand.

Amp sits there oscillating? That's a problem. Ask about it on diyaudio. Take your criticism and

lumps without whinging or passive-aggressive buttholery. Come back with an improved design. And move forward. Same if that new DAC you have isn't working per the datasheet. Someone—maybe many someones—will have had the same problem.

And once you've done some cool stuff that works, it's time for a decision.

Decision: I want a company.

Congratulations and welcome to the insanity!

If you haven't read this book from the start, go back and do so. Figure that you're gonna need to have a second source of income for a couple of years, and you're gonna be spending some out-of-pocket money.

Or you can take your case to the crowdfunding sites. If the marketing is strong in you, you may be successful there. But beware … you'll also now have a large number of very vocal backers, all demanding your time. For a brand-new, untested company, this can be a very scary place to be. Your call.

Decision: I want to work for a company.

Congratulations, you're welcome in this great conclave of obsessives! I hope you have a lot of persistence. Like I said, most great audio companies are not exactly on a hiring rampage 24/7. Hell, a hiring rampage may be, "We need a new engineer!"

Remember, Sumo was a relatively large and dominant company when I got hired. We had two engineers. Theta was a juggernaut by high-end audio standards when Mike decided to go a different direction—and it had nine people … total.

So what do you do? You show work that is clearly relevant to what the company has already done … and keep reminding them that you exist, you're interested, and that you can help them do more, better …

And keep reminding them …

And eventually, something will happen.

Aside: I used to give talks on how to get published in science fiction at writing cons. Science fiction is pretty much the equivalent of audio in engineering. And my point always came down to: run the numbers.

- Of everyone who writes, maybe 10 % get the courage to submit something for publication.
- For everyone who submits a story (and gets it rejected), maybe 10 % send another one. That's 1 %.
- Maybe 10 % of that 10 % aren't completely crushed by the second rejection, and send another story. Now you're looking at 0.1 %.
- Maybe 10 % of that 10 % aren't completely crushed by the third rejection—and they might even get a note from the editor asking for a rewrite. Now we're down to 0.01 % … and your odds are now looking a lot better.

Everyone kinda blinked at me, until I took out a chart of all the stories I wrote and submitted, color-coded as red for rejected and green for accepted. I printed it out large-format so everyone could see the great interstellar cluster of red, punctuated by a few lonely green dots. Then they nodded.

A career in great audio engineering is kinda like that.

Bottom line: keep at it. Politely.

Getting In: The Not-Engineering Path

I'll be the first to admit that I don't know as much about the myriad of other professions that could lead to a career in audio … but I can say one thing for certain: great audio companies need people that aren't in engineering. Engineers can be great at designing a product … but suck at getting it out the door in a timely manner. Engineers can be great at pushing the limits of their designs … but have no idea how to communicate why it's special. Engineers may be wonderful at getting the most out of a limited budget … but have no idea how to sell it.

> **Aside:** in my marketing career, it wasn't unusual to find the key selling points of a new product buried as bullet point #57 and #73 of 89 points the engineer thought were important.

Operations is probably the best path in, but the least sexy. And you may need to do some selling of your own capabilities to the company. They may be unaware of all their own internal inefficiencies, problems, and challenges. Talk about ability to deliver products in a timely manner. Talk about stocking challenges. Ask them how they deal with incoming and outgoing inventory. Ask them if they've really had a sit-down with

all the shipping companies and compared rates. Find out how they're doing things, and discuss ways to make them better. One thing's for sure: both engineering-led companies and marketing-led companies would both be thrilled if the more "mechanical" parts of running a business "just happened."

Marketing is another. But is the company large enough to really need a full-time marketing person? Many aren't. Many just need some graphics, packaging, and an ad or three. Which means if you're a freelancer, you may be able to build a respectable business by being the go-to guy for a number of smaller firms (though you may have to be careful about working in competitive situations.) This also means the pool of companies with full-time positions is smaller. And those positions may be the dreaded catch-all of "sales and marketing." Which—make no mistake—means "sell my stuff, or else, and if you have some free time, some graphics and social and crap like that would be nice." Be careful. Find out what you're getting into.

Sales? Maybe. If it's a more traditional company selling through retailers, sure, you need someone to interface with them, help introduce them to new products, introduce new perks and spiffs and

stuff like that (and, given it's a small company, maybe run them down when they don't pay.) If it's a direct sale company, you're talking a sales associate position … ask what opportunities there are to move up, and don't be surprised if there isn't a clear answer. Many smaller audio companies may not be able to predict where they'll be in a few years.

Everything else. Customer service or tech support can be a great gateway to a future technical or engineering position, if you're (a) just starting out, and (b) have a high tolerance for some very buttheaded customers. Technician? Sure, again a good stepping-stone to engineering. Bookkeeping/financial? Yes, absolutely—especially if the company is selling to dealers. Even direct sale with international distribution is a barrage of accounting, wire transfers, etc. Office/admin stuff? Maybe, depends on how low you want to go to get into a company. Legal? OhLOLOL! Come on, these aren't multinational entities—if they need a lawyer, they'll hire one as necessary. Business management/MBA stuffs? See legal. Maybe there are a handful of firms that could use these kinds of services, but most are gonna blink at the corp-speak and buzzwords and wonder what the heck it has to do with them. Business

school, like engineering, rarely teaches the kind of case studies that are relevant to small companies operating in a passionately connected niche market.

And—as above with engineering—the same notes on persistence hold. You're fishing in a very small pond, and there ain't a lot of hungry fish. Make your case, keep reminding them you really would love to work in audio, and try to keep a dialogue up with the people who make the hiring decisions.

And, eventually, you may get the chance to join the insanity.

Why Go For Audio?

If I have to convince you that audio is a wonderful field to be in, it's not for you.

Yes. It's that simple.

If audio is for you—great audio, not just the next Beats copy—you already know it. You sit there, working on your latest design, or plugging numbers into your newest spreadsheet, or coding like hell on what's gonna be the next Fer-Shure-App-Success, and the music carries you away. It goes beyond the headphones you're using and the

amp they're attached to … but those matter, too, because you've gone up the chain chasing better sound … and wondering if it could be better.

If that describes you … audio may be for you. And in that case, don't ignore its call. Because being in audio can be a really amazing place.

Hope to see you here!

2015, Chapter 7

The Yggy Circus, Tidal Uprisings, and the Unknown Future

April 21st, 2015 is the day I am free once again.

Free of expectations, free of questions, free from the ramifications and repercussions of the idiotic leaks and promises I began making, almost 4 years ago. Because Yggy is released, and nobody knows what comes next.

I've said it before, and I'll say it here again: *never, ever, ever, ever, for any reason, no matter how much you've had to drink, no matter how much it will change the world forever, no matter how cool and wonderful and groundbreaking you've convinced yourself it will be, never ever for any reason talk about future products.*

No, really. NEVER.

Bottom line, I should have never mentioned Yggdrasil or Ragnarok. But I did. And, in doing so, I probably wasted half a man-year in my

time, Mike's time, and Nick's time, responding to questions about when they would launch, (honest answer: when they're damn good and ready), what features they would have (honest answer: no clue until the spec's finalized, and even then … well, development changes everything), how they'd sound (honest answer: really? We don't hear the same things you do, and how do you answer this when the product doesn't exist?), what they'd compete with (honest answer: see the Yggy FAQ, first question, but hey, that was the target, and targets aren't realities.)

Go back and read that bit about wasted time.

Then ask yourself: *what else could we have been doing, other than managing the spin on our own information leaks?*

Yep, there could have been new products.

Yep, we could have gotten a better handle on production.

Yep, we could have explored some exciting new tech we're only now getting into.

But we didn't. Am I bitter? Not at all. It was my choice to blab about upcoming products, secure in my ignorance about how difficult they'd be to get to market. It was my choice to dig that hole

deeper, when we first thought we had a handle on them.

But now Yggy's done. And we're heading into a grand new future.

Last Minute Tricks, or, The Inevitability of Problems

Yggy, like any product conceived by humans, had both its share of hole-in-one triumphs and irritatingly stupid things that held it back. I know some of you conceive of the product development process as a wonderfully smooth and seamless flow from a group of shining white-coated scientist-types, eyes glowing with the reflections of data from advanced touchscreen test equipment, to spinning 3D models on giant displays, to touchless advanced robotic manufacturing producing a flawless product from the predicted start date.

AHHAHAHHAHHAHAAAHAAA! Right. I'll eat my hat if the current iPhone didn't have at least half a dozen tweaks/kludges/unexpected workarounds in its first run.

> And, ah, one rather large—and I bet unexpected—failure mode.

Bottom line, the path to production is never linear, never smooth, and never completely predictable. Especially for a one-of-a-kind product like Yggy, using DSP code and D/A converters never used anywhere else.

So what tripped us up? Lots of things:

1. Rewind to more than 2 years ago, when we found that our planned DSP platform was end-of-life, or "not recommended for new designs." This meant that we had to re-do literally everything on a new, unfamiliar DSP platform. The upside was much higher processing power and greater bit depth for the calculations (resulting in a truly insane 18 000+ tap digital filter). The downside was that it was an entirely new platform that none of us had used before, so the learning curve was very steep.
2. Rewind to even before that, when Mike got the nutty idea to use the AD5791 DAC in the first place (he first brought the idea to me in 2010, shortly after the part was introduced). The AD5791 is NOT a friendly part to work with, for many reasons, including a data format that was fundamentally incompatible with audio, and an output that can include significant glitch energy. Early experiments in formatting

audio to run on Analog Devices' demo boards were not promising—in fact, so un-promising that we explored two other alternatives to the AD5791 before finally deciding that yes, they were usable for the Yggy. We came to this conclusion (to use the AD5791) less than 10 months ago.

3. Follow the AD5791 through the first working motherboards and analog boards. It took quite a few iterations before we were happy with our solution to its need for reference voltages, error-correction amplifiers, the glitch problem, and so forth.
4. Check out those first working analog boards, with a shiny new discrete balanced topology gain stage on them … a balanced topology not yet ready for prime time, as it turns out. We ended up using "simple" JFET buffers, rather than a complex topology.
5. Note the scare quotes around "simple" in Point 4 above. Simple is relative. Managing the performance of these buffers with complex, very-low-noise voltage references to set operational points took some iteration.
6. And let's go back to the DSP code, shall we? The early versions required a computer tether to reboot the processor on the fly, when it had a little oopsie or two when changing sample

rates. Not exactly something you'd want to bring to a show, let alone ship.

7. For fun, how about a shunt power supply just this side of self-immolation? That went through some iterations as well.
8. And let's not forget the USB Gen 3 input. It performs quite a bit better than the Gen 2 … but it also requires different drivers for Windows. This meant many rounds of qualification before we were comfortable releasing them (the good news is the Gen 3 drivers work for both Gen 2 and Gen 3 USB inputs.) This is something that wasn't fully worked out until last month.
9. And then there were the usual stupid problems—metal that needed revised to fit, the wrong connectors on the front panel boards, etc. But this is relatively minor—these kind of things happen with all new products.

It feels good to finally talk about some of these problems, because it wasn't like I could discuss them when we were going through the pain.

All these problems help explain why we're so late. But it also goes to show what you're really getting yourself into, when you say, "Hey, yeah, this is what we're planning to do … " whether that plan is next month, next quarter, or next year.

Seriously. Just shut up. You're aiming for "product launch," not "product escape." Talking about what you have planned in the future may feel good now, but it's usually not a great strategy for the long term.

Why? Because suddenly:

1. Everyone knows what you're planning. Can you say, "Lost competitive advantage?" Especially if they can move faster than you.
2. You've just invited a bunch of questions from everyone. You now have to manage implied release dates, even if they are only fuzzy suggestions.
3. You've opened yourself up for evisceration if things change. Different specs? Missed features? Slipped dates? Doesn't matter how much you said, "You know, we're only kinda-sorta thinking about doing it like this and that," you're gonna be held to it.

But we do blab. Lots of us. Including me. Because it does feel great! And sometimes you really, really think you're right around the corner from revolutionizing the industry.

Hint: it's not you who calls the revolution. It's your customers.

The Utterly Predictable Press

Which is a great segue to the mainstream press' audiophile-immolation-du-jour. Now done with Pono, they've moved on to Tidal.

Yes. Groan. There's tons of great opinions on the idiocy of the mainstream press in the audiophile press out there, so I won't belabor the points again (and, yeah, you already know them anyway, like *nobody can hear the difference, its snake oil anyway, it's elitist, the artists are greedy,* etc, etc.)

And … here's what gets me:

All these histrionics are over a company that's nothing more than another competitor to Spotify and Pandora, offering essentially the same thing at the same price.

On any other news day, *snore. Nobody cares.*

But … this competitor ALSO dares to offer a higher-priced tier that promises higher quality.

Now, it's *get out the pitchforks!*

Does this make sense at all? Yes and no. Yes, in today's moment-driven mainstream online press world. It's a lot easier to start a war than to get into a nuanced discussion. And if it's a class war, even better. Because everything has to be black

and white. It has to be compared. If you're in this camp, you can't be in the other. There's no possibility for a continuum.

Idiotic, yes. But this is what so many things have devolved to: people shouting at each other online over the most trivial stuff, with the press egging them on.

Here's what I say to Tidal: it's your customers who call the revolution, not the press. Stay focused on real quality, and you have a chance.

> But really, a compressed tier the same price as everyone else's? This is not the way to differentiate yourselves.

But I'll shut up now. We're not their marketing agency.

The Unknown Future

Okay. Back to Schiit. We're in a funny place now, aren't we? Because the future is unknown. Completely open. Full of limitless possibilities. We could do anything. We could start making toasters or car door handles tomorrow.

Or, much more boringly, we could just keep doing what we're doing.

The reality is somewhere in the middle. I expect we'll surprise some people this year. I also expect that some of what we announce will be utterly predictable—especially if you know Mike's history.

But I'm not going to promise what kinds of products, or even how many. Even some of our predicted products have seen schedules slip … so you may not even see the same number of products I predicted in December.

"Wait, what, are you getting lazy?" you might ask.

Nope. It's just that most of the stuff we're working on is pretty far off the "safe and easy" template. It's more unpredictable. It's harder to develop. It takes more time.

And—I'll say it—we're *taking* more time, too. Deliberately. Not because we're lazy, but because to keep things moving forward, we need to get stuff right … and we need to explore some surprising new things, from new manufacturing methods to different and unexpected technology.

"Oh, no, are you gonna change completely on us/go crazy and do Beats-like headphones/get acquired by Kanye?" you might be asking?

No. Relax.

We're still going to have fun and make great, affordable, and unique products. That won't change. Ever. It's what we are. Hell, our trip down the path of great and affordable will probably accelerate. A lot of the stuff we're looking into right now gives us virtual superpowers in terms of cost and efficiency.

"But if you don't tell us what you're gonna do, if everything is behind a cloak of darkness, how're you gonna deliver what we want?" some others may be asking.

It's simple. Because we listen to you. We do our research. We understand the market. And we're confident that we, ourselves, bring enough intelligence and insight to the mix to create products you want to have … and that may even be on an "out of the ballpark, category-redefining" level from time to time.

But there will be no more teasing.

No more pre-announcements.

No more blabbing.

That is, until the new stuff is ready. Then, you can tell us if we're insane … or not. Thank you again for your patience, your patronage, your requests, your suggestions, and your insights.

Here's to the last 5 years … and to a great 2015.

2015, Chapter 8

Reacting to Now vs. Planning for the Future

When I was recently on vacation, I had an "engineering moment."

Here's how it went. A friend (Jetse) and I went up to Campbeltown, Scotland, to visit the Springbank distillery. This may be an obscure distillery even to ardent Scotch aficionados, since it is part of the smallest recognized whisky distilling region in Scotland.

We went there for three reasons:

1. They make some great whiskeys (if you can find the original 21 year, you're in for a treat.)
2. They are the last vertically integrated distiller—as in, they do everything in-house.
3. We like to drink. Yeah. There you go.

During the Springbank tour, they took us through the malting barn, past the oven, then on to the grist mill, which was a curiously Rube Goldberg-looking contraption with multiple leather belts

attached to cast iron wheels, with axles all terminating in a fully-certified-for-steampunk cast iron box with a rococo manufacturers' badge stuck prominently on front.

"This grist mill grinds all of the grain for our whisky," the guide told is. "And, funny story, this is the same grist mill we've been using for 75 years. The company made them so well, they never break. They're so good, the company went out of business … they never had any service revenue."

Jetse and I shared a horrified look. Jetse is an engineer as well, working on ship propulsion systems.

"Uhhh … " I said, then trailed off. Best to leave them to it.

Jetse, being Dutch, is more direct. "Errr," Jetse said to the guide. "What happens when it *does* break?"

"But they don't break," the guide said, smiling with all the confidence of someone who knows nothing about complex mechanical things.

"I understand that's the story they give you," Jetse said, again being Dutch. "But when it *does* break, and it *will* break, what do you do then?"

"Uh … well … " the guide said, clearly not used to being challenged.

"Wouldn't it be a good idea to find a spare? Or have someone reverse-engineer this one and have another one 3D-printed?"

"I don't know anything about that," the guide said, looking truly uncomfortable for the first time.

Jetse looked at me, his eyes bugging out in incredulity. I knew at that moment exactly what he was thinking: *How can you run a business without a backup plan? Especially when literally* 100 % *of what you make runs through a single bottleneck made by a long-dead company?*

"Let's keep going," I told Jetse, nodding at the corridor leading to the rest of the tour.

Jetse shook his head, but came along. The rest of the tour went uneventfully, and we had (and bought) some very good Scotch.

But later that evening, over a plate of World Famous Mach Dunes Haggis Nachos (no, you *really* cannot make this stuff up), Jetse went back to that moment.

"Those guys," Jetse said, shaking his head.

"What guys?"

"The Springbank distillery guys. They're betting the company on a 75-year-old device with no support. When it breaks, they're screwed."

"It's a different world," I told Jetse. Campbeltown was, figuratively, a million light-years away from Jetse's shipyards and Schiit … just a tiny town way out on a finger of land in Scotland, with a handful of crazy distillers carrying on very, very old traditions.

"Oh, don't give me this 'we can all sing songs and get along'," Jetse said. "I saw the look on your face. You were as horrified as I was." I nodded. "Yeah. I guess we both had an engineering moment."

"Engineering moment." Jetse laughed. "No, just common sense. Would you run Schiit like that?"

"With one point of failure and no backup? Nope."

Jetse nodded. "You always have to have a secondary plan. If that grist mill goes out, they may lose months—or a year—of production. What happens then?"

"Then their market gets really interesting in 12 years or so." Because they would have a hard enough time judging demand and pricing for

their product, over the dozen or so years it would take to mature it. A busted grist mill might not affect their immediate results, but depending on what the future looked like, it could be a make-or-break thing.

Jetse sighed. "That's the Chinese definition of interesting. No, thank you."

The Pressure of Now

The above anecdote isn't intended to bash the good folks at Springbank—it's really just to illustrate how the "now" can come to eclipse everything in business. In Springbank's case, the "now" is that their grist mill *is* working, and there's no reason to think it might break soon. Plus, they have tons of other "nows" to deal with, whether it's the latest barley crop, or a supply chain issue with bottles or barrels, or dealing with the normal logistics and financial concerns of running an ongoing business …

… *plus* trying to guess what their market might be wanting to buy, 12 to 21 years in the future.

In that case, planning for the future is especially tough. With the pace of change, would you want to bet on, say, whether you're going to be driving

an electric car in 12 years … or if you're going to be driving at all in 6? Or whether NASA will have a functioning reactionless drive (or even warp drive) in 21 years? Or whether or not Google's life extension research will start paying dividends in a decade or two? Or a thousand other smaller things that are changing the world, right now?

Yes, I know, this is all crazy talk. But do some Googling, and you'll see that none of the above is out of the realm of possibility. And then reflect back on this … less than 20 years ago, Google simply *did not exist*.

With the future in flux, it's easy to keep your head down and focus on the "now." But that's not the only pressure keeping our focus short-term. Because, in today's corporate world, there's virtually no downside to reacting to the present, and no upside to planning for the future.

Read that again: *the vast pressure is for results now, not tomorrow.*

CEOs get rewarded on how well the company did that year—or even every quarter. Deliver a couple of quarters of bad results in chasing future fortunes, and you may not be there long enough to see that future unfold.

And it continues down the line. The C-suite may not get the attention they need at their banks if they don't deliver consistent results. Or they may be so tied in with the bank that they can't exist without them at all. The financial pressure to perform—*now*—is huge.

Worse, companies are endlessly compared to the competition. What are they doing relative to Gorblesnort Inc? How do they stack up features-wise to the new Arglebargle XYZ? Are they fully buzzword-compliant to the latest Megacrapdoodle standard?

Oh, they aren't? Well, what's wrong, we're falling behind! *Holy hell, gotta catch up now! Get those new features larded on! Get that Retina touchscreen pasted in!*

Yep. All acting in the *now*. Rather than asking:

1. Are our competitors even doing the right thing?
2. Do the features even matter?
3. Are the standards themselves killing us?

Bottom line, however you look at it, there's almost no reason to plan for the future. And plenty of things that keep you in the *now*. A company has to actively resist the extreme pressure to react to the *now*—in fact, you can make a case that it has

to act against its very nature—in order to think long-term and plan for bigger things.

Resisting the Call of the Now

I suspect that most of the companies who successfully resist the "call of the now" are smaller companies working in niche industries. You'll be able to find them by looking at the companies who are always there, at or near the top of some small industry, year after year, decade after decade. And most of these companies continue to prosper without changing much—resisting the fads and fashions of the moment is their most important calling.

But some biggies operate a bit differently than the typical "grab all you can as soon as you can" companies out there. Let's look at 3 of them.

Quick disclaimer: these are my personal opinions. I could be very wrong. Lots of you probably think I'm very wrong. That's cool.

The Musk Empire. I use this title for Tesla and SpaceX and SolarCity and all the other stuff Elon is up to. He's the poster child for planning for the future, set on upending industries and changing the fate of humankind.

Grandiose? Yes. Sustainable? We'll see. But when one dude can say, "Yeah, I'll answer your questions about my home batteries soon, but I'm a little busy delivering stuff to the space station, and, you know, making rockets land on their tails, like all fans of 1950s sci-fi know they should, and, you know, planning for Mars and stuff like that," it's a bit surreal.

And it's clear he has *absolutely zero interest* in *what's happening now.* He's blown off investors, deflected questions about current revenue and solvency, and done what he's needed to do to keep the innovations coming.

Again, whether or not the empire grows or collapses (or is the right thing to do) is something yet to be determined. But he's not worrying about next quarter—because he has a much bigger vision.

Google. Quick. What makes up almost 90 % of Google's revenue? Guess what. It's advertising. Yes. Google is an advertising company. At the same time, they're playing with self-driving cars and inexpensive smartphones and life extension and artificial intelligence and their CEO is involved in an asteroid mining startup and probably

a dozen other things I don't remember off the top of my head.

In this case, it looks like Google is actively looking to the future with a "throw a bunch of stuff on the wall and see what sticks." They've been criticized by investors (not exactly a long-term thinking bunch, ha) for excessive spending on research, and their approach can seem scattershot, but …

… if even one of their crazy initiatives pays off, hell, Apple's almost trillion-dollar valuation is gonna seem like chump change.

And, that, I suspect, is why they do it: because they have a stable base to build from. The challenge is finding the next exponential leap.

Apple (well, kinda). If you're expecting me to put in a good word about Apple's forests and sustainability initiatives, sure, yeah, okay, but let's see if those are more than point programs reacting to the disaster du jour.

What I'm really speaking about is Apple's (now-fading) ability to put aside the present and look around corners, as well as their (also now-fading) ability to say, "No, it doesn't need that … not until it's right."

The iPod was the beginning. I've gone on about what a breakthrough the iPhone was, and how it broke the model. I was lukewarm on the iPad, but Apple had enough vision and sense to turn it into a new computing phenomenon. I'm lukewarm on the Watch, as well, but we'll see how that turns out … it may be much more successful than I expect. The important thing to note is that:

1. None of these things were new.
2. None of them had the most features or capabilities.

Go back and read that again: these were all *things that companies had done before,* and *current competitors ticked more options on the features list.*

But because Apple was more focused on their own vision, they were able to resist adding long lists of stuff that would slow the product down, or make it buggy, or half-baked (er, um, Maps, and yeah, well, there are some problems here).

Why can Apple do this? In the past, a singular vision. Today, if it can keep it up, it's due to having the comfort of an amazing cash hoard. They have plenty of time to sit back, observe, and get things right. They don't need to make the next quarter.

But they could also be looking forward a bit more … we'll see how that goes.

Okay, Enough Already. What Does This Mean To Me?

Yeah, I know, some of you are tired reading about Scotch or Elon's Martian dreams, but there's are a couple of points I wanted to highlight in those above examples before I moved on to the practical stuff.

1. If you know you have a potentially company-destroying bottleneck, having a backup plan is the first thing you should be doing to ensure your future.
2. For companies that go against the grain and think long-term, it's important to ask why they do it—and how they can continue to get away with it.

Because … no matter how much I try to distill this to a set of rules, learning to plan for the future (or decide to act for now) isn't cut and dried. It will be different for every company.

And with that, let's break the rules … and create a rule.

Rule 1: As Cash(free) → 0, The Future → Now

Huh? Wut?

Okay. Read it like this: *"as your free cash approaches zero, your future becomes your now."*

Does that help? I hope so. Because if your company doesn't have any free cash, you should absolutely be thinking in the now. The now determines the next sale you make, the next prospect you serve, the next customer you make—and keep—happy. In this case, making sales, and keeping customers thrilled, right now, *are the most important things you can do to ensure your future.*

Corollary 1: If you're a startup company, your future *is* your now. Plan for grandiose dreams of world conquest later. Concentrate fully on selling and making customers happy. Period.

Corollary 2: Not all cash is created equal. Note the Cash(free). This means cash you put in yourself, or created through the normal operations of the company. Free cash does not include cash you got from investors, from crowdfunders, or from the bank. This kind of cash is Cash(encumbered). It has expectations attached to it. Investors expect to have an opinion … or even to drive the bus. Crowdfunders

need to be actively communicated with … and may sway your course with their own opinions. Banks want their percentage, or may want it all back, after a time. The value of encumbered cash might be much less than its face value to you … or, in other words:
Real Value(Cash(encumbered)) ≪ Face Value.

Note that two of the three examples of companies looking to the future outlined above have big hoards of cash. The third has proven access to cash, and an overriding vision.

Which brings us to the second rule:

Rule 2: $(\text{Vision} \times \text{Belief})^{\text{Compelling}} = \text{Success}$

Okay, to take this out of the equation realm, let's say, *"Decide on, and believe in, your your own compelling vision."*

Once you've gotten past the hand-to-mouth startup years and have plenty of Cash(free) in the bank, it's time to decide what you stand for, and why it matters. No company will be truly successful in simply bending to the consensus opinion of the now, or shifting with the current sands.

As an example, we once worked with Memorex as their digital agency of record. They were playing

with a whole slew of new products (thumb drives, music players, etc—remember, this was 10+ years ago). They asked for our opinion on their shiny new line.

In a then-uncharacteristic honest answer, I told them, "Well, basically, these are copies of other stuff from industry leaders. Some are fancier, some are cheaper, but they're the same basic thing."

"And?" the VP of Marketing asked, looking a little nonplussed.

"And it depends on how far you want to go. If you look forward a few years, you'll see that a lot of this stuff is just part of a shift to an overall integrated home entertainment streaming/repository system that is more computer-based than individual component-based."

"And this helps us now, how?" their VP of Marketing asked.

"It helps by knowing where things are going, so these products can be tweaked to fit more into that future—and bring you ahead of the competition."

"Tweaked how?"

"Like this thumb drive," I said. "Why isn't it also a music player? Why not have a WiFi chip in it so you can stream to it on a home network."

Marketing VP shifted in his seat, glancing nervously at his CEO. "That's a little, uh, advanced, for us."

And that was that. In the end, Memorex got sold to Imation for chump change, and continues to limp along.

So, *believe in your own compelling vision.*

Deconstruct that:

- Vision: something that looks forward past the state of today and all the demands of the now, and sets some guideposts for your company in the future.
- Believe in: you need to believe it, first. You can't use a Dilbertoid Random Vision Statement Generator to do this one.
- Compelling: Why would anyone else find value in your own vision? It has to matter to your customers as well.

When you have a compelling vision you believe in, you're head and shoulders above all the other companies living for the now, waiting breathlessly

for next month's results, and in general only serving as the factory floor for a hastily-constructed rendition of this year's buzzword-compliant device.

What's Schiit's compelling vision? How about this: *to completely upend the value and values of high-end audio by delivering unequalled bang for the buck and creating a fun environment by doing so.*

And yeah, I know, I said this can't really be broken down to a list of rules. However, rules are a signpost, so here's another:

Rule 3: As Entanglements ↑ , Control → 0

Or, in English, *avoid entanglements and obligations that affect your ability to control your company's destiny.*

Have free cash? Have a compelling vision? That's great!

- Except your dealers don't like the way your product roadmap is looking, so they start pushing something else.
- Except you have one gorilla client who dictates exactly what you do, even if you don't completely agree.

- Except you've signed a multi-year deal with a standards body that promised you the world … but only delivered a neverending nightmare of revisions.
- Except your single-source supplier just went belly-up.

You get it. When you get yourself into situations where you lose control of your own destiny, you're putting yourself in danger.

That dealer network may look tempting now, when you can stuff it with product to claim big sales numbers (but hey, what about getting paid, huh). That gorilla client may look like a wonderful path to stability, but what happens when the current management is ousted in a merger. That standard you just subscribed to might promise never-ending upgrade revenue from your customers as the standards change … but what about your own upkeep, and what about when the customers decide to opt out. That one supplier may seem like the only company that can do what you need them to … but what happens when they get a bigger and more demanding customer, and your stuff slips?

But there's more. Those are external entanglements. There are also plenty of internal ones, too. Like:

- Except your ego got the better of you, and you decided to "really show them now … " which led to an bottomless pit of development cost for a product that may never appear.
- Except you decided you really, really needed that Porsche 917 and the house on the hill, and suddenly your Cash(free) = 0.
- Except you hired the wrong people and checked out of the day-to-day workings too fast, and the business is rapidly unraveling … without you noticing.

Bottom line: the more you can avoid entanglements (internal or external), the more control you have over your future, and the greater the chance you have of seeing your vision through.

Want another rule? Sure:

Rule 4: If Play = 0, Probability(breakout) = 0

Or, *Play as soon as you can.*

On the other hand, having a vision and avoiding entanglements isn't all there is to it. You need to start playing as soon as Cash(free) and time permit. Experiment. Mix things up. Play "What if?" Feel free to surprise yourself with new ideas.

Or yes, in other words, kinda like Google.

Your budget may not be as big as theirs, but the principle is the same … *you'll never stay ahead if you rely only on your steady-state business. Find what's next.*

Playing is part of planning for the future. Because no matter how grand your vision, no matter how much cash you have, no matter how many people support your company, *there's always something that'll come out of left field.*

For Yahoo, it was Google. For Nokia, it was Apple. For SpaceX, it may be new competitors on the rocket front, or entirely new and crazy paradigm-breaking things like the reactionless drive.

For you, what is it? What's next? Where is the market going?

These questions you can't answer with focus groups and crowdthink. These questions are answered by your play and the intersection of new discoveries, proven vision, and ability to look ahead.

But start playing, as soon as you can. If Mike hadn't played with Modi, we wouldn't have a $ 99 DAC and a $ 99 amp … and we'd be a very different company today.

And a final rule, to bring this full circle:

Rule 5: If Plan(B) = 0, Enscrewedness = ∞

Or, *for all the big things, have a plan B. Or else.*

Back to the guys at Springbank. They have no Plan B for when their grist mill breaks. They're hoping someone can fix it, and get it working again in a short enough time frame not to interrupt their production.

Remember … if you're panicking, in the moment, in the *now*, you aren't planning for the future. For really important stuff like that—potentially company-destroying stuff—have your Plan B ready.

And … that's about it.

Except the unwritten rule I started with: the fact that you can't break everything down into rules. Every company will be different. And the future is not always the future, when your beholden to others who control your money. And your vision may be very simple (make stuff cheaper) or very complex (make humanity a multi-planet species). And you may have some entanglements that are unavoidable. And not enough time to play. And not enough forethought to see every Plan B.

But you know what? I expect that companies that plan for the future do much, much better than those that simply react to the present.

Happy futuring!

2015, Chapter 9

Hints, Teases, and Solicitations

Ready for something completely different? How about some info on what we're working on (or at least thinking about)?

Wait. Relax. This won't be a laundry list of what's coming next, nor will it be all-inclusive. That would kinda defeat the whole purpose of me yelling, "Don't talk about stuff until it's done," over and over again, wouldn't it? Nor will it have estimated release dates, extensive and imaginary spec lists, or even any guarantees what I'm talking about here will ever be turned into a product.

Which is why this chapter is called, "Hints, Teases, *and Solicitations*."

Note the "and solicitations." If you have input on what we're considering, or thoughts that go in another direction, we'd love to hear what you think.

But before we go into the "what-ifs, maybes, and pipe dreams," let's do two things first:

1. Talk about where we are now.
2. Talk about stuff we *won't* be doing.

Where We Are Now

Schiit, today, has a wide range of headphone amp and DAC products, together with a single speaker amp, a phono preamp, a dongle DAC/amp, a simple passive preamp, and a kinda-unclassifiable weird USB decrapification device. Prices range from $49 to $2299. That's a hell of a line, for a company that started only 5 years ago and has been entirely self-funded through operations.

Arguably, the line may be a little too big and sprawling. But it does break down pretty well into small, medium, large, and XL products (Magni 2/Modi 2, Asgard 2/Bifrost, Mjolnir/Gungnir, Ragnarok/Yggdrasil.) Each serves a niche, and each has showed increasing sales, even in the absence of serious updates to some of them.

It would be perfectly good to sit back, cross our arms, and say, "Well, there you go, let's keep improving what we have, and call it good."

But …

- We have only a single speaker-capable amp.
- We don't have any traditional stereo preamps.

- We have no multichannel products.
- We have no portable products.
- We have no wireless products.
- Hell, we don't even have any remote frigging controls!
- And no, we don't have any app-driven, internet-of-things, self-aware, Skynet-making stuff silently blinking on a dev desk …

And we have a pretty neat technology portfolio that includes:

- Extensible relay-switched stepped attenuation that has glitch muting and very robust code.
- Robust microprocessor-controlled "intelligent" amplifier management.
- Unique closed-form digital filter algorithms based on non-Parks-McClellan mathematics, implemented on a cost-effective DSP platform.
- Unique and proven interface between the digital filter algorithm above and multi-bit D/A converters.
- Most advanced adaptive clock regeneration system around, no buzzword bingo required (ask atomicbob for the jitter plots of Yggy after it's warmed up if they're not posted here).
- Extensive experience with cross-shunt push-pull (Circlotron) amplifier design, a true balanced, differential stage.

- A unique, dynamically adaptive output stage (from Lyr.)
- And a couple of things I can't talk about yet … (there's the tease).

If you look at our current technology portfolio, it might provide some roadsigns to where we're going. Or not. Because for everything we're "missing" in the line, we have to ask ourselves:

1. Does it matter?
2. If it matters, can we do it better/cheaper?
3. Are we excited about working on it?

For example, we don't have a traditional 2-channel line of separate power amps and preamps. And our technology portfolio—specifically, intelligent amp management and relay-switched stepped attenuators—perfectly supports the development of such a line.

But does it matter? Well, maybe. There are a lot of amps and preamps out there. How do we do it better and cheaper? Amps are a difficult proposition. Preamps less so … but I'm getting ahead of myself. And ignoring the most important question.

Because, believe it or not, #3 is more and more the arbiter of what we do.

What We *Won't* Do

Okay, I've said this before, but for the sake of any new readers (and easy reference), there are a few things we have exactly zero interest in. Here they are:

1. **Transducers.** Also known as headphones, speakers, and "them things that actually vibrate and produce sound." We're not experts in transducers, and we don't see what value we can bring to the market without a long, extensive learning curve and mucho development time. And we're not going to simply go to another manufacturer and say, "Can you slap a Schiit logo on this bad boy so we can sell 'em as our own?" So, no headphones and no speakers.
2. **Surround/Multichannel Audio.** Mike went there, and metaphorically got his leg gnawed off by a grizzly. Today, it's an even more complex and ever-changing morass of proliferating standards (Dolby, DTS, Auro, blah blah) and connectivity (HDMI 2.0 at the moment, but who knows about the future?) So, no multichannel, no surround.
3. **DAPs.** When you're sitting on the subway or working in a noisy open-plan office, do you really need much more than a smartphone

and a pair of efficient headphones? This is not a rhetorical question … I'm serious. Are $99 to $499 DAPs really that much better than phones in real-world use? Are there good reasons for carrying $1000 to $3500 devices around in your backpack or briefcase? Are there use cases that I'm overlooking? I'm ready to listen, but at the moment, enthusiasm for a product like this is so low that I don't think we'll ever do one.

What We're Playing With, Thinking About, or Otherwise Dreaming Of

Okay, now that I've dashed the hopes of anyone wishing for Schiit headphones, speakers, surround processors, AVRs, or DAPs, take heart! Because there's a ton of stuff we still keep glancing at and wondering, "Hmm, should we do something here?"

Let's talk about some of them, starting with our traditional headphone market. Our line here is pretty solid, with some notable exceptions.

- **Portable amp or DAC/amp.** Hmm. This is one I spent a lot of time on last year and ended up abandoning, mainly because I couldn't answer Question #2 well enough—that is,

how do we do it better or cheaper? This is something I could still get behind, though, if we figure out a way to make it unique enough. And we have some ideas … but those ideas point towards a relatively large device with very high power output and some unique takes on power supply and topology (discrete.) So we'll keep playing. Which is where I'd ask … would you be interested in a larger-than-normal portable amp/DAC that provides truly stellar performance? I'd assume that iOS and Android connectivity (self-powered USB input) would be a must, but let me know what you think.

- **DAC upgrades.** Mike already said that customers who've purchased our upgradable DACs will be rewarded with technology from Yggdrasil at some point in the future. As will new purchasers, of course. I can't set a timetable for these product/upgrade introductions, nor can I set prices … but I can say that Yggy isn't the last superburrito product you'll see from us.
- **Combined desktop DAC/amp.** We still get asked about this all the time. We actually have tons of DAC/amps. They're called Magni 2/Modi 2 , Asgard 2/Bifrost, etc, etc. This separation of DAC and amp allow you to choose

a combination that best fits your needs, to avoid the obsolescence of an embedded DAC inside an amp. If we come up with a solution that sidesteps the obsolescence problem and doesn't have problems with digital noise in the analog amp section, sure, maybe there's something there. What do you think?

Now, let's look at the "2-channel" world. In other words, speaker stuff. This may seem like the most likely thing we'd get into. But it isn't. Why? Because we're not super excited about it. One of the biggest things keeping us out of 2-channel is deciding what we want to do with speaker amps (if anything). Because, like it or not, they're the foundation of a 2-channel system. Or does it matter? Could we just do preamps and call it a day? You tell me.

- **Speaker amps.** The hardest thing here is doing something meaningful … for a non-eye-bleeding price. Do we shoot above Emotiva (in price) with big, heavy, impressive boxes filled nifty circlotron designs and intelligent amp management? Or do we try to go head-to-head with this behemoth price-wise, but with smaller and more efficient amplifier designs (think small Class AB, not D, think fan)? Neither approach is particularly appealing.

Above Emotiva, there's about ten billion small manufacturers, all with their own spin … and some with very compelling products. At Emotiva, well, there's Emotiva. Would something that doesn't look big and impressive sell well against their iron? I don't know. If you have thoughts, I'll be glad to listen.

- **Preamps.** Okay, now we're talking. I can see opportunities for remote-controlled, relay-switched stepped attenuator preamps—passive, active, tube, etc. There's a shortage of good affordable products at this price point, and I think we can really make some cool stuff here. But do they fly without amps? Is the remote a BLE module and smartphone app, or old-skool proprietary RF? (No thanks on IR, not huge fans here.) Still a lot of stuff to work out, but there's something here. What would you like to see?
- **Music server.** Argh. Yeah, I know, some of you want a non-computer-style server that still takes USB drives and SD cards and Wifi audio and stuff. But do you know what you call that? A computer. No matter how you slice it, a modern music server is a computer. It may have a lot of computer-y stuff hidden, but it's still either an embedded Linux or Windows box. This means software de-

velopment, support, maintenance releases, software/firmware upgrades, etc. That's a big undertaking to do it right. We're currently glancing in the direction of these things, but that's about it. Would it be better to simply do a super-uber-mega-USB-reformatter that "perfects" the USB packets and timing (which vary significantly from platform to platform, and even computer to computer) that could be used with any computer?

- **Wireless audio.** Bluetooth, nope. Not until it's capable of uncompressed streaming. WiFi, hmm … but we'd need RFID to allow easy pairing … lots of questions here. Let's just say we're playing at the moment. But is it really so hard running some cables?

And, a wildcard.

Should we be looking at an analog-to-digital converter using the same megaburrito filter as Yggdrasil? Mike did it once, for insane prices (see Mobile Fidelity's GAIN system (not GAIN 2). Yes, there are actually a bunch of GAIN CDs out there that were mastered with a complementary algorithm to Yggdrasil. They're, um, pretty insane. Should we do that again? But again, that gets us into an entirely different market … the pro market … which we know very little about.

And A Final Question

Running down the same old tracks and trying to keep up with the Joneses will get us nowhere fast. There's a ton of stuff I'm missing (and a couple of imminent things I can't talk about). But, to help us try to figure out what we're missing, can I ask you one question?

What's the one thing you dislike most about your current audio system?

Yes, *dis*like.

If you have a few moments, indulge me. I'd appreciate it!

2015, Chapter 10

Knowing Our Place?

First of all, thanks to everyone who commented on last week's chapter. We're listening to all of your thoughts, input, and suggestions … and I can already say that a lot of you are going to be pretty thrilled with where we're going. Last week, before heading off to TheShow Newport, Mike and I had some long conversations and solidified many of our plans going forward.

But … Hmm … TheShow Newport.

Let's make that (and trade shows, meets, etc, in general) the subject of this week's chapter, while it's fresh in my mind.

If you've been reading along from the start, you know I'm not the biggest fan of trade shows, nor the best planning and logistics mastermind on the planet. To sum up quickly for those new to the book: audio trade shows are a great place to hear some of the best gear in the worst possible venues on the planet, while enduring bad hotel

food and perhaps picking up a nice cold or flu along the way.

And, if you're an exhibitor, knowing you'll *always* forget something, that something will *always* break, and that things won't go the way you expect, *always*.

Like TheShow. When I found out we were in a glorified tent hidden well outside the hotel, rather than in a ballroom, there you go. Not what I expected. As *always*.

Yes, you read that right. Headphonium exhibitors got stiffed. The traffic was crap. If a headphone exhibitor was counting on exposure, this was a real kick in the nuts.

And, it was also a reminder of where we stand in the whole scheme of things, audio-wise. For those drooling over gold-plated audio jewelry, lovingly hewn out of solid blocks of unobtanium, headphone audio is still little more than a curiosity.

Maybe we should do shows at colleges. But I'm getting ahead of myself …

So Why'd You Go, Dummy?

Okay, yeah, I hear the snarky voice in the back of the crowd, saying, *Well, if shows suck so much and you got boned, why'd you go?*

Two reasons:

1. We thought we'd be introducing some new stuff. But it was late. Ah, well. Look to mid-August (but not at a show.)
2. We didn't know the degree of bonage we'd be getting. TheShow was at a new venue this year, and they changed the location of the Headphonium exhibits.

It could have been worse. We could have been cooking in a literal tent. The AC in the hard-sided contraption they had set up was so powerful that we had to wear sweaters. And they tried to get traffic to the Headphonium with signs and some relatively scary-looking vehicles (which have *what* to do with audio, but hey, who the heck knows what people think?)*

* **Disclaimer:** I am a car guy. But over-the-top black and gold Mustangs that recall very bad images of Trans Ams and dudes with chains and chest hair are, well, not my cup of tea.

But in the end, the proof is in the results. And even though the crowds came out to see the gilded two-channel rooms in "the tower," they largely ignored the poor little guys out in the Headphonium. I doubt if we saw more than a handful of people who weren't already at CanJam just a couple of months ago.

So, lesson learned: if we're going to be going to a two-channel show, bring two-channel stuff.

We have Ragnarok and Yggdrasil. We could have done exactly that. I just forgot about that before the show. We should have picked some great direct-sale speakers and showed up in the tower. That probably would have been worthwhile. Headphones, not so much.

The Business Lesson

Here it is, in one sentence: don't conflate two seemingly similar industries, or else.

Or else you'll get bitten.
Or else you'll lose money.
Or else you'll waste time.

We made the mistake of lumping two-channel stuff in with headphone stuff. Not a good idea.

Headphone stuff hasn't gone (as far) down the road to gold-plated Bentley audiophilia that nobody can afford.

> Though current trends are worrying … continual price inflation pushing the top end constantly higher, manufacturers now asking buyers "how much do you think we can sell this for?" rather than pricing at a simple multiplier of BOM cost, obsessive focus on aesthetics, cosmetics, machined-from-a-single-block silliness,* ad infinitum.
>
> > * Yes, I know, Apple does this. They also make a hundred million or so of their widgets a year. Not a hundred or a thousand or ten thousand. If we made a hundred million widgets a year and had a foundry sitting right next to the machining center so we could immediately recycle the Pacific Ocean of chips we'd be creating, maybe we'd do that too. Or maybe not.

I should have known. I've seen this same thing happen before in many industries. I've even seen it happen in our own industry, for the clients that Centric serves. Every once in a while, we get an audio company that wants to reach outside the traditional high-end audio realm. This is always:

1. Very costly.
2. Not very effective.
3. Abandoned after spending a ton of money.

Repeat after me:

- A Beats customer isn't necessarily a Schiit customer.
- A Levinson customer isn't necessarily a Mr. Speakers customer.
- An AVR customer is not necessarily a Levinson customer, nor a Beats customer, nor a Schiit customer.
- A sound bar customer, again, isn't necessarily any of the above.
- And yet all are "audio" customers.

Sounds simple, yes. But it's sometimes even easier to forget.

So, if you're starting your own business, the first thing is to know where you fit. Or else you might end up spending a lot courting customers that simply don't get what you do.

And this is especially important for shows and meets. Why? Because even though headphone audio is still the gnat on the fly's ass of high-end audio, the expectations of a decent show presence go beyond just plunking your stuff down on a

table. It takes quite a bit of money and energy to make a trade show work. Consider all of this:

Pre-show stuff:

1. Getting all your products together.
2. Getting all your headphones together.
3. Making sure you have every single connector/cable/etc you need for them.
4. Deciding on sources (we have gone through three different tablet sources before finally settling on one we like—this was not cheap in terms of money—or, especially, time. It takes a ton of time to set up sources, do updates, transfer all music, arrange it logically in libraries, etc.
5. Have something to identify the company made—backwall, or go crazier.
6. Have something to identify the products made—tabletop graphics or more.
7. Have some brochures or handouts or something.
8. Get hotel rooms (preferably at the show hotel, and not the show holding the furry convention down the street, because you were too lazy to get the rooms in time—ask me how I know this.)
9. Decide who's going to the show, and get them there, maybe involving airfare.

10. Send a press release if you have any new products and invite press.
11. Consider doing a pre-show or at-show event, and plan for that.
12. Pay for your show space.
13. Do an ad for the show guide (did you see ours?)
14. And about fifteen other things I'm probably forgetting. Ask Denise. She's very, very good at that stuff.

At-the-show stuff:

1. Setting up and breaking down the booth.
2. Fixing anything that goes wrong.
3. Finding replacements for the stuff you forgot.
4. Making sure all staff is there when you want them to be (staff scheduling).
5. Deciding on a policy for expenses and incidentals for staff (are you paying for drinks, for example?)
6. For really crazy corpocrat people, deciding on a policy for what staff can and cannot say.
7. Answering questions and helping people who stop by.
8. Resetting the players after demos so people know what they're listening to.
9. Dealing with any computer interface problems.

10. Dealing with any operational problems, including intermittent power from the hotel, etc.
11. Talking to press, doing demos, etc.
12. Again, about a dozen things I'm missing. Denise?

And post-show stuff:

1. Following up on everyone you promised to get back to.
2. Descrewifying the pile of equipment, headphones, etc you brought back from the show so it's ready for the next show.
3. Fixing or replacing anything that got damaged.
4. Planning for the next show.

Bottom lines, shows take a ton of time and energy—and, in true Murphy fashion, they usually happen exactly when you'd rather be spending that time on new product development or something more, well, critical to your business.

But shows can be valuable, too. If you're looking for dealers or distributors, absolutely. If you're introducing new product and know the press you want to see are gonna be there, you bet. Unfortunately for us, we were in neither mode for TheShow, and it was too far down the path to just pull out.

And if we'd had any presence of mind, we'd have remembered that we're in the middle of moving out of the "headphone only" mode. Although we're very well-known in the headphone community, two-channel guys are only beginning to discover us.

Which means we should have gone to TheShow with Ragnarok and Yggdrasil and speakers. Nothing more.

The New Show Strategy

In talking to other exhibitors at TheShow (hell, we had to do something, it was pretty dead), we've come to the following "brilliant" conclusions:

1. **Bring headphone gear to CanJams, meets, and other shows that are truly dedicated to headphones.** CanJams are excellent places to be in the headphone world. And there are plenty of them to go to, all around the world. A company could have a full show schedule by attending only CanJams.
2. **Decide on some speakers, or partner up with some speaker guys, and start showing speaker stuff at two-channel shows.** That includes shows like TheShow Newport and AXPONA and RMAF. Of course, that makes

RMAF very interesting, because of the added logistics of having both a CanJam exhibit and a speaker room. Especially since I don't think I'll be there this year (I'm having a big birthday, and yeah, sorry, guys, I am not going to sacrifice another birthday to RMAF, though I really do love you all … well, in a non-creepy way, of course.)

Yeah, I know. Not rocket science. But I'm not done.

1. **Do some of our own shows.** Those new products? They're probably going to be too late for CanJam NorCal, and too early for RMAF (which we'll have some other new stuff at). So, why not do something ourselves around mid-August? Details TBD, of course, but that's what we're thinking about.
2. **Start looking at having a presence at smaller, regional meets.**
   Beyond CanJam, there are plenty of meets. We need to make sure our newest and hottest stuff is there, so you can have a listen to it. How we'll do this is currently open to discussion, but I'm sure we can be more present than we've been.

And let's be clear—we're not going to be abandoning the headphone community for two-channel

high end. However, now that we do have (some) two-channel products, it's time to start showing up to the party. Which will be fun in itself… fending off the dealers and distributors, and watching people's eyes go crossed when they hear our name.

But let's take it a bit farther. What might happen if we decided to re-mix the traditional audio show entirely? After all, all audio shows pretty much use the same format: invite audio companies to buy as much space as they can and bring whatever they want. In that format, of course the guys with more money, time, and staff are gonna be the most impressive, every single time.

So, what about some new show ideas? I'm blue-skying here, feel free to add.

1. **A show where everyone gets the same amount of space.** In the headphone world, call it just one table. What would you do with that, if you couldn't bring your 30-foot backwall and light show? Who would look most impressive then?
2. **Better, a show where everyone is allowed to bring only a single system.** Regardless of how big or small they are. It's easy for us to look impressive and big with 8 separate

systems from $79 to $3998, but what if we had to bring just a single system? What would we bring? And why? What would everyone else bring? Would it make sense to go all-out, or impress with budget-constrained products?

3. **Or how about a show with "no system over $1000?"** How would that go? What would we find if the top-end is lopped off of the show, and we all have to bring budget gear?
4. **Or maybe better, a show segregated by budget?** This would definitely help prevent heart failure when people start asking about prices.
5. **Meets that challenge you.** You know, I've been thinking a lot about the whole objective/subjective dichotomy we have in audio lately, and I think it may be time to create a new kind of test. Not ABX with unfamiliar music, but something much more interesting—letting people choose their own music and switch between two different signal chains when they wanted (both level-matched and with distortion lower than audibility, you know, typical "good test results." The goal would be to see if (a) a consistent majority preferred one signal chain over another, and (b) to see how many people could consistently tell the

> difference. Yes, I know, this kind of test will never satisfy the ABX folks, but I think it might get us a little closer to the truth about whether or not there are some people who consistently hear audible differences in systems that should sound the same. I suspect the answer is neither 100 % subjective nor 100 % objective, but a continuum … some people can hear the difference, some can't, some care, some don't … and all of that is fine. Of course, this is a big undertaking, but … hmmm … I need to think about this some more, I kinda just write stuff as it comes to me.

So, am I crazy or what? (Yes, I know, you're probably laughing and nodding your head.) But I will say one thing: the past years I've spent in headphone audio has been more fun than all the years of two-channel, home theater, and video servers put together. There's still a fresh, raw edge to the headphone side that I really don't want to see plowed under in the ongoing rush to unobtanium pricing.

I'm hoping we can help keep the headphone side on its toes.

> I'm sure you'll let me know if we don't.

And with that, I'll sign off … with one last re-

minder … that as of June 15, we'll be celebrating exactly 5 years of being "shock over substance," "a flash in the pan" and "flavor of the month."

Here's to the next 5 years!

2015, Chapter 11

Food Scientists vs. Michelin Stars

On the advertising agency side of things, we have the "parable of the ice cream man" that we use to illustrate the old adage of "if all you have is a hammer, everything looks like a nail."

Huh? Ice cream men? Food scientists? Hammers? *What the heck is Jason rambling on about now,* you ask. Is he hungry? Demented? Just too damn hot in the unusual muggy summer heat that makes wearing headphones and writing a misery?

Maybe a bit of all of the above. But bear with me.

First, let's talk about the advertising agency biz. Ad agencies rank down there with used car salesmen and Kirby door-to-door shills in many people's eyes, despite some lingering *Mad Men*-inspired-but-really-never-happened fun-times memes. And yeah, I understand. When many agencies make a living by charging ridiculous multiples for every intern's time and have

big headcounts that do nothing but justify their (oops, surprisingly) large bills, the reputation is deserved.

But one thing a (good) agency does well is *solve marketing problems,* rather than *sell a specific service.*

For example, there are any number of marketing specialists like "social media consultants" or "native content advertising specialists" or "pay-per-click advertising companies" out there. And they're defined by their single specific slice of the marketing pie … social media or native advertising or pay-per-click or whatever. That's what they do. And if you, as a company, invite them in the door, you know exactly what they're going to try to sell you … either social media or native advertising or PPC.

And that's how they're like the ice cream man. The ice cream man drives around and sells ice cream. On a hot day, he's a welcome relief. The perfect solution (well, that is if you aren't paleo, or on a ketosis diet, or, well, hell, food is weird these days). But on a cold day, who wants it? Unfortunately, ice cream is what he has to sell, even on the cold days, the rainy days, and the

times you really have exactly zero interest in ice cream.

Just like those marketing specialists. If you want (and need) what they're selling, you're in luck. If you don't want or need it, you may end up buying something that does nothing for you.

Compare this to a full-service advertising agency or marketing agency. A (good) agency will sit down and listen to what you want to do, then offer ways to accomplish your goals. These ways may not be what you expected. They may not even be what you asked for. But a (good) agency will be able to pull from a variety of strategies and tactics that it has experience with, and recommend some that offer the highest probability of success.

So, you could say they sell ice cream. And tacos. And organic nuts. And fresh fruit. And lots of other things. But they're not just selling ice cream. They're selling what you need, rather than only what they have.

And that's the parable of the ice cream man. I've used it tons of times … when a company comes asking for a social media program when they really need direct mail (yes, still relevant—depends on the audience) or when they want to

change a perfectly good website that really just needs more traffic.

So what does this have to do with food scientists vs the cordon bleu? Read on.

The Insult That Backfired

Early on, soon after Schiit launched, the objectivist movement sniffed us out and began throwing various insults our way. "They must not use any test equipment, they just talk about how things sound." "Those designs are huge, inefficient, and overpriced … probably unreliable, too."

And my personal favorite, which went something like, "How can a tiny audiophile company imagine it could do better than the latest I op-amp, they have a worldwide staff of scientists and engineers working on those designs?"

When I first read that, I was taken aback. Because, yep, it sounds completely insane. One or two engineers at a tiny company that started in a garage with miniscule resources, pitted against the combined might of, say, I or Analog Devices? Most people would bet on the big guys.

But then I started thinking about it.

And I realized: *Frito-Lay probably has more food scientists working on the newest formula for Doritos than there are Michelin-starred chefs in the world …* but which would you rather eat?

And … *Darden Restaurants and Taco Bell, Inc and McDonald's probably have tons of food scientists stapled to their beancounters …* but wouldn't you rather have something fresh from a Cordon Bleu chef's kitchen?

Both of those are perfect examples of where *single individuals* trump *teams of scientists* armed with the latest technology all the time.

Why?

Two reasons, I think:

1. They actually give a crap about raising the bar on how something tastes (and have the knowledge to at least point in the direction of a solution.)
2. They don't have to worry about packagability, shelf life, merchantability, etc. (and yeah, I know, this is a big one.)

So, happy with turning the insult around (though I never said anything about it), I let it go.

Until today, when I realized this is a great way to describe what we're doing with digital audio.

Digital, Michelin Stars, and Flavor Technicians

I've gone over this before, so I won't go into insane depth again, but here's the gist of what we're trying to do with multibit digital audio playback: *preserve the original samples of the vast majority of music out there, in order to offer the best possible playback fidelity.*

And yeah, I know, some people think we're full of Schiit when we start talking about closed-form digital filters and how most DACs simply replace everything with a reduced-bit-depth approximation of the original. That's fine. I'll let them argue with Baldr (Mike Moffat) about that one.

But here it is: we *do* have a stated purpose. And that stated purpose is to *make the most out of the ingredients we have.*

Kinda like a top chef, huh?

> Or, well, maybe at least a good one. Throw me a bone here.

Contrast this to the approach taken by delta-sigma DACs. The vast majority of music in this world (about 99.99 %—do the math on the tracks available yourself) goes in at 16/44 and gets

transformed into a multi-order noise-shaped 2 bit to 5 bit approximation of the original. Yes, even in some very exotic designs.

Which, if you think about it, isn't that a lot like the food scientists at Frito-Lay and Taco Bell trying to preserve some shade of the original ingredients through latest-tech additives, preservatives, flavor enhancers and modifiers?

"So you're saying that everyone else is wrong?" you ask, arms crossed.

Not at all.

What matters in food is in the tasting. And there are some true food artists working in the molecular gastronomy space. And then there's personal preference, too. Sometimes it's time for a bag of Doritos. And sometimes it's time for yellowtail sashimi with lime and burdock.

If someone achieves the perfect result for your ears, don't worry about this analogy. Sit back, listen, and enjoy. Hell, we make delta-sigma DACs too, and they sound plenty fine. And, especially in the case of $99 to $149 products, there won't ever be a multibit option at that price range.

But, as Mike has said, we're gonna be bringing multibit technology to our upgradable products … and then you have a choice. Stay delta-sigma, or go multibit. It's entirely up to you.

More than DACs

Almost every day, someone asks why don't we do a Bluetooth audio interface, or a Class D amp, or digital room correction, or any other number of nifty audio-scientist-style tricks that are as easy as picking up a chipset and (in some cases) doing some RF certifications.

> **Aside:** which goes to show how simple and ubiquitous these massive technological solutions are. If you ride the teams-of-audio-scientists train, you do get quite an amazing toolkit. Easy wireless transmission of compressed audio, $5, 100 W amp chips, and automatic room correction for the price of a chipset, a license, and a microphone.

But every one of these solutions runs counter to our goals. Bluetooth audio is lossy and variable in terms of quality, depending on your source and how far away you are from the receiver. It's taking those original samples, applying perceptual

encoding-based compression, and decompressing on the other side to an approximation of the original. Yes, it can sound good, but it's not us. Class D amps? Same thing. They take the analog waveform (or, in some cases, direct digital input) and use noise-shaping and filtering to approximate it at the output. Digital room correction? Again, everything's being run through a number of very complex mathematical transformations—in this case, where the actual goal is to change the output in order to compensate for the room. Used correctly, this is impressive technology … but, like stability control on a car, sometimes it's hard to tell where it actually goes off the rails until something bad happens.

"So you guys are grumpy old buttheads, forcing the hard way on everyone?"

No. Not at all. We're explaining what we do, and why we do it. If what we're doing sounds good to you, then there you go. If not, not. Depending on how many people think we're crazy, we'll either continue to succeed, or die on the vine.

But … there *is* a method to our madness. There are *reasons* for what we do, besides "let's make a product that hits the current hype du jour and makes big profit."

If this parable helps you understand what we're shooting for in digital, then it's done its job. Beyond that, it's up to you to decide what your own personal preference is.

The Real Moral: Here's to the Misfits

Okay, I've explained what we're trying to do. But it's important to remember that there are plenty of other companies out there that are pushing the limits of what we can do in audio … and that don't have globe-spanning teams of scientists and engineers.

So here's to everyone who tilt at the monoliths. Who don't think a homogeneous, Lightning-cable-connected, DSP-corrected-to-the-eardrums solution based on big science and lowest-common-denominator parts assembled in an inexpensive corner of the globe is the One True Future Path. Who work with parts long obsolete … or create new technologies that elude even the giant engineering teams (remember, I licensed Nelson Pass' patent, not the other way around.) Who explore improbable ways to do things … and sometimes succeed. Who ignore the datasheets and reference designs and tweak it until they

have something better. Who listen, and measure, and come back to do it again and again.

Because, otherwise, all we have is Doritos. And that's a very, very sad place to be.

2015, Chapter 12

On Measurements (With a Side Order of Sanity)

Okay. Big subject. Big chapter. But it's a great subject: *what do we measure, and why do we measure it?*

When Schiit first launched, we downplayed the instrumentation aspect of our designs—providing only a handful of measurements, without attribution or explanation of how we arrived at them. This led to early speculation (in some corners of the web) that we were doddering mouthbreathers, rubbing magic stones and praying to the audio pantheon for help with more sound-good magic.

So, why didn't we provide complete measurements from the start? Well, plenty of reasons, none particularly intelligent in hindsight. From the top:

1. Our experience has been that *many* measurements, at least the commonly given ones, don't correlate highly with perceived sound quality,

even when those measurements differ by several orders of magnitude. Current example: compare a Magni 2 to a Vali, and tell me you can hear that the THD is 100× higher on Vali than Magni 2 at 1 V.

2. Our experience has *also* been that measurements vary from product to product, and cherry-picking one product with stellar measurements and using it as a reference is disingenuous at best, and once fudge-factors are applied to actual measurement ranges, measurements look a whole lot less impressive.
3. And, finally, our experience has always been that measurements can be spun as well as any fancy marketing-speak; how something is measured, at what point, with what filters and weighting, with what equipment … if you work at it hard enough, you can produce impressive numbers from virtually any product, numbers which then encourage people to argue about the relative superiority of a −103 dB A-weighted SNR vs. a 98 dB unweighted SNR, without knowing the technical reality behind it.

"Okay, so now you said that measurements are meaningless," you might be asking. "Why bother at all?"

Because measurements are absolutely meaningful—it's just that what measurements are meaningful changes at which part of the product development and production process you're in. Without measurements, it would essentially be impossible to develop our products … and without deep, high-quality measurement capability, the products wouldn't be as good as they are.

And, with that, let's start this introduction to what we measure and why in an unconventional place: by introducing you to the equipment we use, and why we use it. Relax, there won't be any test on this at the end.

The Metrology Players

"Metrology" is just a fancy word for the science of measurement. It's sounds more appropriate (and more costly) when used to describe a $ 3.5M atomic force microscopy system, but it applies just as well to the gear and techniques we use.

So let's take a look at what's on my desk, starting with the measuring stuff.

Stanford Research SR1 Audio Analyzer. This is the big boy that does most of the heavy lifting. The SR1 outputs both analog and digital signals,

and inputs both analog and digital signals, so you can look at distortion, noise, gain, max output, frequency response, jitter, and a whole lot more … no matter if you're measuring a preamp (analog in, analog out), a power amp (analog in, BIG analog out), DAC (digital in, analog out), or ADC (analog in, digital out.) Mike has his own, kitted out with all of the high-precision options in the digital domain. Mine is a more basic model, but it still does some very good digital analysis.

- **Why is this is on my desk?** I use it all the time. It's a lot easier to hook up a design, even in breadboard form, and know exactly what it's doing, then to squint and guess (well, at least after the basic power-up stuff happens … but I'll get to that.) It takes much of the mystery out of how a design is performing (at least when there isn't a cable/grounding/instrument glitch, but I'll get to that, too.)
- **How much will this set me back if I want one?** Ouch. Starts at $ 10k, can get $ 13k or so with options.
  More info at: thinksrs.com/products/SR1
- **What alternatives are there?** Audio Precision is the leader in this space. They cost more (the one model they have that can beat

Stanford in all aspects is about 2.5× the price). There's also the Prizm DScope, which is about the same price as the Stanford, but with not-as-good performance in the digital domain especially. There's also Rhode and Schwarz, which cost a mint and we have no experience with them, so I can't comment.

- **No, really, what if I want to do measurements like this for cheap?** Get a QuantAsylum QA400 for $199. It won't do everything the Stanford does, and it will absolutely blow up if you overload the inputs (the Stanford will not), but it's a great way to get started. We use one in production. QuantAsylum also has a much more interesting product coming, the QA405, that looks like it might provide a significant portion of the capabilities of the Stanford/AP/dScope universe for a lot less … but at the moment, it's vapor. Beyond that, it's soundcards and open-source software.
- **Hey, can I use an old Sound Tech or HP distortion analyzer to do the same thing?** Uh, no. Get a QA400 and invest in resistor padding/protection box and get much more comprehensive measurements.
- **Hey, my scope does FFTs, can I use that?** Uh, no. If that scope goes below about −80 dB on its FFT scale, we'll eat a hat … er, I mean,

be very surprised. The Stanford can resolve down about −150 dB.

Agilent (now Keysight) MSO-X 202 A Mixed Signal Oscilloscope. This is the device that they play with in the intro to the old Outer Limits series. In general terms, it takes an analog waveform input and displays it on a screen. Like most scopes these days, the Agilent is a digital scope, which also means it can easily look at one-shot and transient events. It also has a ton of other capabilities like built-in function generators, 4 channels of analog, and 8 channels of digital, so you can do mixed-signal stuff and logic analysis. Mine is one with all the options, so I can do all of that. And, best of all, what it displays is very high-res and analog-looking, which is a welcome change from early digital scopes.

- **Why is it on my desk?** For quick "Does it work? Is it oscillating? How's the clipping look?" type measurements, a scope and function generator will keep you a lot more sane than being stapled to an audio analyzer all the time. Plus, it lets you look at noise on the grounds, in power supplies, etc in a more natural manner.
- **How much will it set me back?** New, about $3500. Check eBay for refurbs. Mine was an

apparently unused refurb with all cables, etc for $ 2k.

- **What alternatives are there?** Tons. Agilent/Keysight is just what we use. If you're a Tek aficionado, don't get up in arms. There are also a ton of good Chinese oscilloscopes out there … if you're willing to put up with less analog-looking waveforms and funky UIs, there are some amazing values. The Siglent SDS1102 can be had for $ 350-ish and does a good job, if you're just starting out.
- **Danger Will Robinson!** Do NOT get a USB oscilloscope or pocket oscilloscope for your only scope, unless you have endless patience with software UIs or can decipher and control what is happening on a scope with cryptic numbers and four buttons.

Tektronix 7603 Oscilloscope with 7A22 and 7A18 Modules. In addition to the fancy new digital Agilent scope, my desk also supports a giant old 1970s-era Tektronix mainframe oscilloscope. Yes, a fully analog beast. Yes, we're talking hand-assembled modules made from discrete components here. While this may seem like an anachronism, it really isn't. See below.

- **Why is this heavy, ugly, giant-ass thing on my desk?** Because it offers some capabilities

that modern digital scopes don't. First and foremost, the 7A22 plug-in goes down to 10 µV per vertical scope division. Compare that to the new Agilent, which does 1 mV. When you're looking at ground noise or power supply noise, that's 100× more resolution on the old scope.

- **How much will it set me back?** Get one on eBay for $ 150 to $ 300. Modules are $ 50 to $ 100. Prepare to reinforce your desk, especially if it's from IKEA.
- **What alternatives are there?** Want the same capability on a new digital scope? Then you're buying a high-$ Tek and a blindingly expensive differential probe amplifier.

Fluke 8842 5.5 Digit Multimeter. This is a good desk multimeter that allows you to measure things like AC and DC voltage, 2-wire and 4-wire resistance. It has 5.5 digits, which means it has enough accuracy to chase down things like shorts on a board. It ain't very modern, though … good modern stuff goes to 6.5 digits easy. Look at the Fluke 884 A or something like that.

- **Why is it on my desk?** First reason: for checking if your power supplies are working, or if you put the wrong resistor in somewhere, or if you have DC offset, or a board short …

you need a good DMM, period. Second reason: battery-powered multimeters are always out of battery when you need them most … and you usually forgot to stock up on the 9 V batteries.

- **What will it set me back?** Used, $150 to $250. A new fancy 6.5 digit Fluke will be about $1000 to $1200.
- **Alternatives?** Come on, get a Fluke. You know you want one. If for the name alone. A test and measurement company being called "Fluke" must rank up there with Schiit.

Fluke 179 Multimeter. At the same time, I also have a battery-powered Fluke.

- **Why do I have it?** Because sometimes you need two multimeters. At least if the batteries are good.
- **What will it set me back?** $300 new, $150 to $200 used.
- **Alternatives?** Come on, get a Fluke. Seriously.

DER EE LCR Meter. One thing that most multimeters don't do is measure inductance and capacitance. Hence, I have a cheap LCR meter. Cheap because I don't use it much.

- **Why?** This is invaluable for checking transformer designs, especially picky quadfilar stuff for circlotrons.

- **What will it set me back?** $100-ish.
- **Alternatives?** Tons. LCR meters can get crazy expensive. If you need it, cool. We don't.

Seek Thermal Camera. This is a widget that plugs into an Android phone (or iPhone, if you get a different model) and shows graphically how hot something is getting. It's similar to the Flir One, and all of the standalone Flir units, but cheaper. Mike has a real Flir. Dave has a Flir One. All are good options.

- **Why is it on my desk?** Because it beats the hell out of losing thumbprints on hot components. Or, more seriously, because it shows you exactly what's getting hot on a big PC board, and how hot it's getting. If you have a couple of SOT-23s running at 150 °C, they won't be running for long.
- **Why Android?** Because my Schiit phone is Android. My Centric phone is iOS. I also use a Nexus 9 and an iPad 3. And a MacBook Pro and an Ultrabook. Haters, hate equally.

Siglent SDG1025 Arbitrary Function Generator. Technically not a piece of test equipment, since it only outputs various waveforms, but I'm including it because I usually use it with the Agilent scope.

- **Why is it on my desk?** Because if you want to look at square wave performance, or open-loop frequency response out to 10 MHz, you won't be doing it on a Stanford. Plus, it has knobs rather than a keyboard, so it works a lot more "analog-y" than punching in a number on the Stanford, or running an automated test routine.
- **How much will it set me back?** About $ 350 on eBay.
- **Alternatives?** Wow, cheap much? Get one. You'll thank me.

Of course, there's a ton of other stuff on and around my desk that is used as ancillary to measurement, like:

Powerstat Variac. This is an ancient device that allows you to bring the power up on a new design slowly (variable-AC, get it?) Powerstats last forever, so you only need to buy one. Per tech, that is. Very useful if you don't want things to instantly smoke when you botched something on the board. Hook up with an AC ammeter in series, slowly turn up … if it spikes, well … you messed something up.

Mastech HY3003F-3 Power Supply. This is a dual-channel 30 V power supply. If I want to put

something together fast on a breadboard to see if it works, I usually don't go through the trouble of building a power supply. That comes later. For tube stuff, I usually use the supply out of an old Lyr or something like that.

Weller WES51 Soldering Station, Yihua 898D Rework Station. The Weller is a fancy soldering iron with good control over temperature and interchangeable tips. Necessary for working with PC boards, duh. The Yihua is an inexpensive hot air blowing device that allows you to work with surface-mount more easily. Nothing fancy.

Whew. That's a hell of a lot of stuff. But now that you're introduced, let's talk about actual measurements.

The (Expected) Measurements (That We Publish)

"Huh? What? Caveats already?" you ask. "Why not publish both the expected and unexpected stuff?"

Hold on. We'll get to that. What I want to do first, though, is break down the typical measurements you expect to see on a piece of gear, discuss what meaning they have for us, what we expect

to see, and if there are correlations to audible differences.

So, here we go:

Frequency Response. This is the range of frequencies that the device can reproduce, usually expressed with a variance. For example, "20 Hz–20 kHz, ±0.1 dB." A device with this specification can reproduce all frequencies from 20 Hz to 20 kHz with no more than a 0.2 dB variance. Yes, 0.2 dB, not 0.1 dB. You saw the ± there, right? Kinda makes speakers with 45 Hz to 18 000 Hz ±3.5 dB not look so hot, hmm? That means they could easily be down 7 dB at 45 kHz and 18 kHz.

- **How hard is it to get great numbers?** For electronic equipment, not hard. A flat frequency response from 20 Hz to 20 kHz is cake. For example, most of our equipment is −6 dB at 0.16 Hz and 300 kHz to 500 kHz (without input filtering at 150 kHz or so, so it doesn't reproduce AM radio). That makes 20 Hz to 20 kHz at 0.1 dB nooooo problem. It's also why we sometimes have two numbers, one at ±0.1 dB and one at ±3 dB. The only struggle would be with capacitor-coupled tube amps like Vali and Valhalla 2. In those cases, the bass

−6 dB frequency is defined by the RC filter of the coupling cap and the load, or about 7 Hz to 8 Hz for a 32 Ω load and 7 Hz to 0.8 Hz for a 300 Ω load.

- **What do we shoot for?** Flat. Duh. It's actually very hard to make something that's not flat, unless you are deliberately including tone controls. Otherwise, something is monumentally amiss.
- **Does frequency response correlate to audible differences?** Sure it does, if it is grossly off. But "grossly off," these days, should be the exclusive domain of transducers—that is, headphones and speakers. It is very easy to get ruler-flat response from electronics.

THD. This is the total harmonic distortion, usually expressed in a percentage at a certain output and frequency, or over a certain frequency range for a specified output. For example, "<0.003 % at 1 V RMS, 20 Hz to 20 kHz," or "<0.01 % at 150 W RMS into 8 Ω." There are a ton of gotchas with this measurement, including the "total" word. Total harmonic distortion includes all harmonic components—2nd, 3rd, 5th-order stuff, and more. That makes it a less illuminating measurement than looking at an FFT and saying, "Well, it's all

2nd order at 80 dB down, you can't even see the 3rd in the noise."

- **How hard is it to get great numbers?** It depends on the topology. It is dead-easy to get great numbers for a super-high-feedback, IC-based design. At that point, THD will be dominated by layout quality (especially in power amplifiers) and power supply design (ditto for power amplifiers.) For a tube run at low plate voltages with no overall feedback, well, you may be looking at several percent distortion—and at that point, it could be audible. Numbers can be made to look better with weighting, especially if you're talking THD+N, but these designs will usually be exposed by a high noise floor.
- **What do we shoot for?** It depends on the topology and the design. In general, we shoot for −60 dB minimum, or 0.1%, at typical outputs. The only amp we make that doesn't make this number is Vali. Most everything is much, much better, though.
- **Does THD correlate to audible differences?** Like frequency response, only if it's grossly off. While 0.1% THD is high for an electronic design, it is VERY good for a transducer. Again, the transducer dominates.

IMD. This is intermodulation distortion, an important measurement because it can give us an idea of how a design performs with respect to non-harmonic distortion, which is usually more objectionable than harmonic distortion. This measurement is usually expressed as a percentage at a certain output, and type of IMD measurement (CCIF, SMPTE.) For example, "<0.006 % at 2 V RMS full-scale output, CCIF." The type of measurement is important, because CCIF usually measures 19 kHz and 20 kHz tones (or another pair of tones 1 kHz apart—we use 19 kHz and 20 kHz) and looks for a resulting 1 kHz tone. SMPTE uses 60 Hz and 7 kHz (or two other non-harmonically related tones.) We've found both of these tests to better illuminate what's going on with a marginal design, which has led to some specific tests that we do that go a bit farther … but I'll get to that later.

- **How hard is it to get great numbers?** Like THD, it depends on the topology, and it's pretty easy to get good numbrs from a high-feedback design. For poorly designed tube amps, it can be high. Numbers aren't usually going to look better with weighting, though, so IMD is less easy to game than THD or THD+N.
- **What do we shoot for?** Same as THD. As

with THD, only Vali doesn't meet the −60 dB number.

- **Does IMD correlate to audible differences?** Again, only if it's grossly bad … but since it's not harmonically related, it may be more audible than harmonic distortion … and transducers usually swamp the electronics numbers, again.*

* It's not like we hate transducers or anything, they're just harder to get right from a measurement standpoint. That's why it's always best to start with a headphone or speaker you love, before going crazy on amps and DACs.

SNR. The Signal to Noise Ratio of a product is easily the most highly correlated to actual audible differences … simply because a noisy product is, well, noisy. Plug sensitive headphones into it and listen to it hiss. Or hum. Or both. Typically expressed as a dB number below a reference, with weighting, like: "−102 dB, referenced to 2 V RMS, A-weighted." Pay attention to the reference and the weighting, because that's where the first number can be gamed. Big time. For example, if we measured Ragnarok's low gain output in the same way you normally rate a power amplifier—that is, referenced to maximum output and A-weighted—it would have a staggering spec like

−135 dB. However, this is because it's referenced to 29 V RMS and A-weighting rolls off where most of the power supply noise will be.

- **How hard is it to get great numbers?** Again (you'll get tired of hearing this), it depends. If you're gaming the numbers, rating A-weighted referenced to max output, it'll be a lot easier. If you're using a simple, single-ended design with excellent power supply rejection and/or a low-noise power se easy as long as you understand thermal noise … that is, the fact that high-value resistors just sitting there in the signal path (or used in gain stages) create thermal noise. Pick a 30 kΩ/3 kΩ pair for a diff amp drain load and source degeneration respectively, and your noise levels will be a lot higher than 3 kΩ/47 Ω, for example. Tubes are more variable from device to device and manufacturer to manufacturer, so you may get noisy tubes and quiet tubes—hence bigger fudge factors being applied to tube amps. However, all in all, it's not tough to get good numbers. Just remember to watch the reference level and weighting when comparing products.
- **What do we shoot for?** At least −100 dB at low gain, or at full scale output, unweighted.

We get there most of the time. One big exception: Mani. You try delivering 1000× gain and see how it works out, noise-wise. That's just the reality of phono preamps. Now, −100 dB may not sound particularly impressive, but −100 dB unweighted really translates to inaudibility in practical applications. And, as I mentioned before, we're notoriously conservative on our numbers. We usually add a 6 dB fudge factor to the published specs, based on worst-case measurements. And, it really matters where the noise is. Noise at 60 Hz with no harmonic components (like from tube heaters run with AC) is surprisingly hard to hear. Noise at 120 Hz from glitch rectifiers and unsnubbed transformers with tons of harmonic components is buzzy and irritating.

- **Does SNR correlate to audible differences?** Absolutely. As stated above, a noisy component will be noisy … you'll hear it.

Crosstalk. This is how much of one channel bleeds over into another. You'll usually see these numbers expressed as "−XX dB over Y to Z range," like: "−68 dB from 20 Hz to 20 kHz." In the old days, crosstalk was an important spec, because phono cartridges were so limited that they could only deliver about 20 dB to 30 dB of channel separation. Today, you'd think it would be easy to have

infinite channel separation, but the reality is that everything inside the box influences everything else … PCB traces "talk" to each other electromagnetically, poor power supplies can bleed channel content across from one to each other, even output jack resistance and load impedance comes into play.

- **How hard is it to get great numbers?** Harder than you think, especially when you can see some numbers based on insufficient information or complete fantasy. Crosstalk numbers of −80 dB, −90 dB, and −100 dB and more should be looked askance at. Is that only at 1 kHz? Then maybe. It's a lot easier to isolate PCB traces at 1 kHz than 10 kHz or 20 kHz. Is it into no load? Again, then maybe. But real numbers based on an actual 20 Hz to 20 kHz measurement and with actual physical output jacks that have actual physical resistance are usually a lot lower.
- **What do we shoot for?** At least −60 dB from 20 Hz to 20 kHz, and at a 32 Ω load if the product is a headphone amp, or 8 Ω load for a speaker amp.
- **Does crosstalk correlate to audible differences?** Even in this day and age of potentially infinite channel separation from a digital

source, probably not … even at −60 dB and higher.

Output Impedance. This is the impedance of the output stage of a product. It will be more important in power amps, where a lower output impedance is better (unless you're of the "current output" frame of mind, which I won't discuss—but if it works for you, have fun!). This is usually expressed in terms of ohms, or, in the old days, in terms of "damping factor" into a specified load. For example: "0.05 Ω, or a damping factor of 160 into an 8 Ω load." These are the same measurements. For preamps, you may be looking at 75 Ω or 600 Ω (or even higher for some tube pres). It's important that power amps be low output impedance so they do not affect the frequency response of the headphones or speakers they are powering (transducer impedance varies with frequency, whereas for amps and preamps it should not—or vary predictably due to a Zobel network … like I said, this gets deep really fast …) It's important for preamps and source components to have a stated output impedance that is not too high, or long cable runs may be problematic.

- **How hard is it to get great numbers?** One more time: it depends. For high-feedback de-

signs, it's dead-easy to get low numbers. If you throw 100 dB of feedback at something, you can have essentially a 0 Ω output impedance. (Or, if you wanna be tricky, use an error-correction output stage for true 0 Ω output impedance, or even negative impedance … and watch out—we were messing around with an error-correction stage that welded a 1/4 inch headphone plug to a Neutrik jack when I pulled it out while the amp was playing.) For tube amps, low output impedance will be much harder, since the intrinsic output impedance is much higher and there is less gain to play with. Hence the 14 Ω output impedance of Valhalla 2 in high gain mode.

- **What do we shoot for?** As low as possible on amps, and 75 Ω for SE output sources and 600 Ω for balanced output sources. The SE and balanced output impedance is just set with resistors in sources.
- **Does output impedance correlate to audible differences?**
  Yes, it can. If you use a high-output-impedance headphone amp with multi-driver IEMs, you're going to run into gross frequency response problems (because they use crossovers). The same thing happens to a lesser extent with dynamic headphones that have impedance

variations at certain frequencies. And the same thing happens with speakers. In addition, source components with too high of an output impedance and long, highly capacitive cables can roll off high frequencies.

Maximum Power. This is how much power an amplifier can put out, usually expressed in watts RMS into a specified load, or multiple specified loads. For example, "1.5 W into 32 Ω," or "60 W into 8 Ω and 100 W into 4 Ω." These numbers are very important for understanding how an amp will work with your specific transducers. If you're trying to use an amp rated for 110 mW into 32 Ω for HiFiMan HE-6 headphones, you aren't going to be very happy. But also, if you're using an amp rated for 1000 W into 99 dB efficient speakers, you also may not be happy if the amp is noisy.

- **How hard is it to get great numbers?** Heh. We should turn this around and say, "Which numbers?" At one point, the Federal Trade Commission regulated how power output was to be measured, since the amp companies were gaming the numbers so much. They don't pay very much attention to that today. But here's what you're looking for: power in RMS watts per channel, all channels driven, into a specified load. Not RMS? That doubles

your stated power. Not per channel? Hey, you can add the numbers! Now Ragnarok is a 400 W amp (both channels into 4 Ω, not RMS.) Pretty tricky, right?

- **What do we shoot for?** Depends on the application and the amp. Typically 1 W or higher into a typical load for headphone amps, and 50 W or higher into 8 Ω for speaker amps.
- **Does maximum power correlate to audible differences?** Absolutely. If your product can't drive your transducer without clipping (distortion), you'll hear it. To be fair, though, most headphones don't need much power to run.

Full-Scale Output. This is a simple measurement for source components, typically DACs, specifying how much voltage it puts out for a full-scale signal (0 dB). You'll usually see this in terms of RMS voltage, like "2.0 V RMS."

- **How hard is it to get great numbers?** This really ain't about great numbers, it's about knowing what kind of output you have.
- **What do we shoot for?** The consumer standards: 2.0 V RMS for single-ended sources, and 4.0 V RMS from balanced sources.
- **Does full-scale output correlate to audible differences?**

Believe it or not, it can. If you're trying to compare a DAC that has 1.5 V RMS output to one that has 2.5 V RMS output without level matching, well, the higher output one is usually going to sound better.

Gains. This is a simple measurement of how much a product amplifies a signal, usually expressed as a simple ratio or in dB. For example, "Gain=20 (26 dB)." This is a much more important measurement for complex signal chains, such as vinyl … where you have to start with 100 to 1000× gain, then factor in the 400 mV output spec into the gain of your amp, and decide if you still need an active preamp. For simpler systems, it's important to know that the higher the gain, the noisier the product.

- **How hard is it to get great numbers?** Again, this isn't really about good or bad numbers, it's about appropriate numbers.
- **What do we shoot for?** The ability to drive an amplifier to its full output with a standard consumer-level source, and gain switching for transducers that need lower noise.
- **Does gain correlate to audible differences?** No, except for the level-matching caveat above … stuff that's louder tends to sound better.

Power Consumption. This is a simple measurement of how much power a product uses. Well, usually simple. In the case of speaker power amps, the FTC mandates measurement at full output into the lowest specified load. So that leads to big numbers for Ragnarok that don't really reflect its idle power. It's rated in watts, like "400 W."

- **How hard is it to get great numbers?** Depends on what you mean by great. Low power consumption either means low-standing-current design or "standby" modes with keep-alive transformers and high complexity.
- **What do we shoot for?** We don't. Every product uses the power it needs to.
- **Does power consumption correlate to audible differences?**
  LOL. No.

Size and Weight. Yep, these are specs too. We take them. We provide them. No, they don't correlate with sound.

Whew! This is getting deep. But hold on … there's more. I still haven't talked about what we don't publish … and how we use measurements throughout the process … and one of our not-so-standard measurements …

The Measurements (We Don't Publish)

Wait. Stuff we don't publish? Why not? Are we embarrassed? Are we hiding something?

Nope. No embarrassment, no hidden agendas. The handful of measurements we don't publish fall into really only two categories:

1. Too easy to misunderstand.
2. Too technical to matter to anyone other than us.

Let's have a look at some of these measurements.

Jitter. Oh gawd, people love jitter. They just love, love, love it. They love to throw numbers like 0.2 ps at the screen and talk about "femto" clocks and how their jitter is just the lowest possible number it can be. But there are several problems with this:

- The number of individuals who can actually measure sub-ps jitter is very, very, very low—it requires a $30k instrument that does exactly one thing: measure jitter
- "Femto" clocks—clocks with femtosecond level jitter—frequently only have femto performance at very high frequencies … and are dependent on layout, logic, power supplies, PCB noise, etc, etc … so they may not really mean anything

- Most don't measure jitter where it matters—at the word clock to the DAC—because the number will be higher than those sub-ps readings

We measure jitter on all our designs at the word clock with an interval analyzer. On the best designs, we can get down around single-digit ps numbers. On other designs, it may be >50 ps. In either case, they're not impressive when you're comparing to 0.2 ps. So we don't publish those measurements. Nor do we publish our eye diagrams, jitter impairment tests, and jitter spectrums from the Stanfords … but a quick look at Yggy's jitter specs, tested here independently, confirms exceptional performance.

- **So, how hard is it to get great jitter numbers?** I could be snarky and say, "Not hard if you make them up." But, bottom line, very hard, especially when measured where it matters.
- **What do we shoot for?** 2-digit ps in non-Adapticlock products, single-digit 50 ps in Adapticlock products, both measured at the word clock with an interval analyzer, and correlated with a benign jitter frequency distribution as measured on the Stanford.

- **Does jitter correlate to audible differences?** It shouldn't. Modern numbers, even on some fairly terrible interfaces, really should be below the limit of audibility. We'll leave it at that.

Open Loop Gain, THD, Frequency Response, Slew Rate. These are the same as the gain, THD, and frequency response measurements outlined in the "what we publish" section above, but applied to the gain stage without feedback, or "open loop." Since we do primarily discrete designs, and primarily designs that have open-loop bandwidth greater than 20 Hz to 20 kHz, it's important for us to characterize the designs in an open-loop state. This helps us optimize them for their particular usage. Note that these are early-stage measurements.

- **So, how hard is it to get great open-loop numbers?** Not too hard with inherently linear stages—from complex multistage solid-state amps to simple tube designs that are run with proper voltages.
- **What do we shoot for?** Great numbers for both linearity and bandwidth. And by "great," we mean less than 0.1 % THD, greater than audio bandwidth, etc.

- **Do open-loop numbers correlate to audible differences?**
  According to some of the Pundits That Be, unless the slew rate is insanely low, no. However, we have noted sonic correlations between a constant −6 dB per octave falloff outside the flat passband to infinity (no lumps, humps, bumps, or other weirdness going on.) Of course, pure objectivists will say we're fooling ourselves on that one. But hey, maybe if we can convince ourselves that there *are* differences when there *aren't,* maybe they can convince themselves that there *aren't* differences when there *are*. Neener.

Gain/Phase Margin. Ah. If you're not familiar with Bode plots, this one will be relatively boring. Briefly, these are measurements that determine how stable an amplifier design will be, if you're going to use feedback. There's always a time delay (phase difference) between the output and input signals of an amplifier. This time delay varies by frequency. And if the time delay means that the output signal is in-phase with the input with an open-loop gain greater than one, boom! You have an oscillator. An amplifier doesn't do anything until you give it an input, and then it amplifies it by a specified amount. An oscillator just starts

singing by itself. Which you absolutely don't want in audio … especially when some circuits can oscillate at frequencies that take out FM radio!

- **So, how hard is it to get good gain/phase margin numbers?**
  It depends on the design … more complexity and greater bandwidth (as with current feedback) make it more challenging. However, the main thing to note is that measurement of, and compensation for, an amp design should be done at the PC board level, not in simulation or on a breadboard, because parasitic coupling comes into play here.
- **What do we shoot for?** An amplifier that isn't an oscillator into any sane load.
- **Do gain and phase margin numbers correlate to audible differences?** Again, they shouldn't, at least if the amp is actually stable into all loads.

Measurements, Beginning to End

Measurement isn't something you do once to a single golden sample and call it done. It's an ongoing process, from the first breadboards to "naked" prototype PC boards to production

qualifiers in chassis to production itself. Here's a run-down of what we do at various stages in the game. It may not be entirely complete, but hey, this is a looooonnggg chapter, and I'm doing this from memory.

Early Stage Measurements. We're talking breadboards and naked PC board prototype stage here. This is when we're either just playing around (on breadboards) or trying something that may or may not end up being a product (on a prototype PC board, usually without a chassis, usually green rather than red … we play with quite a few things that never become products.

At this point, what do we look for?

- **Gross instability and oscillation.** This is the biggest problem in both analog and digital designs. This measurement usually starts with a scope and a waveform generator, because it's easy to just take a look and see what's going on … as well as see the fundamental of the oscillation if it is oscillating. Is the sine wave "fat" on the scope? Then it's oscillating. Is a square wave reproduced with a ton of overshoot and ringing? Then it's only conditionally stable. If it's apparently stable on the scope, we move to the Stanford … and that's where it

becomes interesting, because sometimes stuff that looks stable might have a noise floor of only −100 dB, when they should be −130 dB to −140 dB. This means it's oscillating at a low level or only conditionally stable … which means more compensation is needed.

- **THD up to maximum output.** While we're on the Stanford, we'll look at THD performance, but usually in terms of what specific harmonic components are present. Typical notes in my engineering docs might be: "−95 dB 2nd, −115 dB 3rd, 4th and up in noise at 1 V RMS into 32 Ω, −75 dB 2nd, −90 dB 3rd, −75 dB 4th, rest in noise at 15 V RMS into 32 Ω. Of course, these numbers will change based on PCB layout, so they have to be re-measured every time.
- **Maximum power output.** Where does it clip? We look at this visually on the scope and at 5 % THD on the Stanford.
- **Clipping behavior.** What does it look like when it clips? Does it have sticking problems (nasty waveform irregularities that sound bad and can damage transducers)? If so, that needs fixed before moving further.
- **Noise.** This is another important one, and is 100 % the domain of the Stanford. I noted that a giveaway to unstable behavior is a high noise

floor, but that's not the beginning and end of it. Is the noise floor where you'd expect it to be, given the overall circuit design? Where are the power-supply-induced peaks at 60 Hz, 120 Hz, 180 Hz? How big are they? Running down where the noise is coming from is usually the domain of the antique Tek scope—with that, you can see exactly what grounds are being contaminated, and by what.

- **Heat.** What's getting warm on the board? This is where the thermal camera comes in. Calculation only gets you so far … if you have a couple of 5 W resistors each dissipating 2 W … next to a hot transformer … better measure it. Though, in our case, it's usually miscalculating the power dissipation of an SOT-23 transistor or something like that. If it doesn't fly off the board before the thermal camera comes out, that is.
- **Input performance (digital).** In addition to the above tests, digital devices get verified for different input frequencies and levels, as well as different levels of degradation of the input signal, to make sure they lock to the signals we specify.
- **Jitter performance (digital).** Also on the digital side, we'll typically take a look at jitter performance once everything is on a PC board.

Qualifying Measurements

Once we get past the naked board measurements, we get into stuff that's intended for production. So that means we're now measuring boards in chassis … hopefully final boards, but the measurements let us know just how final they are.

This is where we repeat everything above, and pay more attention to the actual numbers across the board (including IMD, output impedance, and other stuff I didn't mention above.) We'll measure 4 to 6 prototypes and see if they are all similar, or if they're all over the map. If they are similar, and the numbers are what we expect, then it's then on the path to production.

Well, except for a few more measurements …

- **Multiple qualification.** Yeah, we like our Stanfords, but Dave has an Audio Precision SYS-2722 as well. And yeah, it's not the latest Audio Precision, but if the AP numbers and the Stanford numbers agree, that's a good sign that everything is right with the world.*

  * Getting consistent measurements isn't as easy as you might think. A bad cable, a power cable running over an input cable, noisy AC

(our building is very bad), RF interference, ground loops (like from forgetting the scope is still connected when you're running the analyzer) and even broken equipment all affects your results. We actually blew up the analog generator section of one of the Stanfords … not enough to get it to quit working, but enough to make it behave oddly (wrong output at some ranges) and poorly (high noise and distortion.) Once the problem was confirmed with a loopback test, the instrument had to go back to NorCal for repair. And this is on professional-grade measurement gear. On a QA400 or sound-card based system, it can be much worse.

- **Stress tests.** This is also where we get to short the outputs, stress the protection system, and see if the products survive. It is also where we see if static will cause problems with digital and analog inputs, and where we'll run tube designs for a few weeks (or months) and re-measure tube characteristics to make sure there's nothing unexpected going on.
- **Custom tests.** And sometimes we come up with custom tests. For Ragnarok, we were having problems optimizing the algorithm that controlled its bias and DC offset (read, they

were blowing up unexpectedly), so Dave wrote a special version of the Ragnarok firmware to push out the numbers applied to the DACs that set bias level about 1× a second, together with the instantaneous readings for bias current and offset. This allowed us to see in real time what Ragnarok was doing, and make changes to the algorithm that eliminated the problem.

Production Tests. And here's where we do it all again—on most of the first run of products, as well as first articles of each run and spot checks throughout. Plus additional testing that I'll get to.

Now, some of you are saying, "Wait a sec! Does that mean that you don't run everything through the Stanfords?"

In short, no. Why? Two reasons:

1. Because the variation in production (and failure rate) is exceptionally low. In a typical run of 2000 Magni 2s, for example, 2 to 4 of them will fail on first power-up.
2. These failures are usually gross and easily seen on a scope.
3. While the Stanfords are excellent instruments, they are cumbersome to use in production. Yes, we could spend the time setting up automated

tests, but automated tests can be easily as fallible (false positives, false negatives) as other instrumented tests, and they provide no benefit for gross failures.

That said, should we be measuring everything in production with automated test equipment? It's getting to that point, yes. Testing all the input sample rates on a multiple-input DAC and all the I/O on a Ragnarok is pretty time-consuming. So we may be looking into adding some Audio Precision gear specifically for that (they have some stuff specifically targeted at pass/fail testing … and those instruments would live at our PCB assembly house, so they have to be simple to use.

Beyond that, production has its own set of measurements we need to do … measurements that cross over into the actual production process.

- **Programming and verification.** Lots of our stuff needs to have firmware installed, and the firmware needs to be verified. No testing can be done before the firmware is installed and verified.
- **Biasing.** Some of our products need manual biasing of the output stage, like Lyr and Mjolnir. This means we actually have to adjust potentiometers and measure voltages across

an output resistor to set the bias … and then verify the rest of the measurements.

- **Tube Matching.** Some products use tubes that need to be matched. We do this in-circuit for the specific product the tubes are used for. And, for even more detailed work (such as looking at a new design or qualifying a new batch of tubes), we have a computerized curve tracer as well.

And then, of course, after all of this, there's a final listening test. Yes, our final test is subjective. You'd be amazed at what kind of things it can pick up—including tons of stuff that would sail right through a scope or a Stanford. Stuff like scratchy volume pots, operational noises or glitches, or interference only at specific output levels is something that automated or instrumented testing isn't going to find easily.

Measuring the Unexpected

Okay. Is that enough? No. Let's go deeper, and talk about one of the measurements we do that is off the beaten path. This measurement appears to correlate at least loosely to subjective impressions, and it unearths some surprising problems in gear that otherwise measures very well.

"So why not release it for the world?" you ask. "If this is such a breakthrough, everyone should be using it!"

Well, we're not sure it's a breakthrough. Our sample size is very small. And any correlation it has with sonics is loose at best. And it won't matter for the hard-core objectivists who have decided that there are no sonic differences between competently designed components, no way, no how, nuh-uh.

And it's not a breakthrough test. It's a simple extension of the old IMD idea, but this time with three, four, or five sines—a multitone distortion test. The theory is the same as IMD—can we reveal non-harmonically related stuff with a more complex signal? Usually we use four tones, ranging from 50 Hz to 15 000 Hz, but we've run more and less. We've used different tones. We're still playing with this, so don't take it as gospel.

Aside: it is dead-easy to set up a multitone test on the Stanfords, but I'm not sure how easy it is to do on other products.

So what do we see when we do multitone tests? More non-harmonically related stuff in designs

that sound not so hot. Sometimes in very surprising places … not even related to the beat frequencies themselves. Like the example of the Perfect DAC.

The Perfect DAC was not one of ours. It was sent to us by a friend who wanted to get some measurements for it. This was a delta-sigma DAC, manufacturer and chipset redacted, with a very fancy power supply and all the buzzword-compliant stuff people like to hear about these days. We said, "Sure, why not." And ran it through its paces.

And … in terms of standard measurements, this DAC blew everything we've ever measured away. I mean, vanishingly low noise floor, virtually undetectable power supply harmonics, insanely low THD, flat frequency response …

… until you looked at the IMD, which gave numbers a bit higher than you'd expect, given the THD results. And the numbers weren't related to the 1 kHz spike … they appeared down low, below 100 Hz.

What? We ran through our multitone test (it's easy to do digital multitones on a Stanford as well, not sure about other analyzers) and the low-frequency numbers went bonkers. As in, there

was a broad range of non-harmonically related distortion components from 10 Hz to 90 Hz, at a fairly high level (−50 dB or so). −50 dB is potentially audible. And it was up nearly 90 dB from the baseline measurement.

So what happened? I don't know. With digital, there are more variables, and noise-shaping and decimation are math-intensive, algorithmically based operations. Perhaps there's a glitch in their algorithm. I don't know. It's not our DAC, and it's not something we were going to spend the time to dive into.

So … while we putter around confidently with all of the accepted measurements, maybe there are still realms out there where "here be there monsters."

That's why we still listen. And measure. And come up with new measurements. And listen again.

And I'll leave it at that.

2015, Chapter 13

Detours in Balanced-Land, or "Improving on 'Perfection'"

"Perfection?" some of you are already sneering. "If you're talking about Mjolnir, that ain't no perfect amp."

Yep. No kidding. Nor is anything else. Hence the scare quotes on "perfection." The fact is, though, Mjolnir was our first statement amp, our first cost-no-object design. I know this sounds crazy in these days of $15 000 headphone amps and $3500 portables, but at the time, Mjolnir was our first shot at an amp that didn't need to hit a price point.

"It'll cost what it costs," I remember telling Mike, way back in 2012.

And that's was our mindset as we went about developing Mjolnir. Exotic transformers, sure. Premium parts, no problem. A whole new chassis design, absolutely. The only two things we didn't pay much attention to?

1. **Aesthetics.** Mjolnir was supposed to be a simple amp—while it wasn't cost-constrained, it was supposed to be the simplest expression of a balanced amp we could come up with. So super-custom chassis designs weren't really what we were focused on.
2. **Convenience.** Also, as a simple, no-frills balanced design, we didn't think much about convenience. Right down to making it only a balanced design—which couldn't be unbalanced.

And balanced-only was Mjolnir's Achilles heel. There weren't a ton of people who wanted to commit to having only balanced headphones, no matter how high a performance level Mjolnir delivered. That really held it back. If I could go back and whisper in my ear (oh, the days a TARDIS would come in handy), I would have told myself:
"Big dummy, even if people are super hardcore into balanced headphones, what about their friends who bring their single-ended stuff by? Drop a single-ended output in there, FFS!"

But even then, I might not have listened to myself. Space was at a premium in the Mjolnir design, since it was based on through-hole parts. There probably wasn't space for summers. And even

if there was, we didn't have the trick summer design we ended up using in Ragnarok.

Between the Current and the Future

Besides the realization of how big of idiots we were for not including single-ended outputs on Mjolnir, several other things happened on the road to Mjolnir 2.

1. **Ragnarok happened.** Once conceived as a much bigger Mjolnir that could run speakers—and could be configured as either a solid-state or tube-hybrid amp—Ragnarok morphed into a real technology statement, doing stuff that no other amp does.
2. **Ragnarok tube *didn't* happen.** And, in the process, it became very apparent that a tube hybrid version was in nevergonnahappen-land.com. At least without fans. And different transformers. And maybe not even then, since I refuse to do tubes that are merely "cosmetic." And Ragnarok's topology doesn't lend itself to much more than cosmetic tubeination. And Ragnarok's topology doesn't work without the intelligent microprocessor control we have wrapped around it (more on this later.)
3. **Yggy happened.** Yeah. Finally.

4. **Yggy trickled-down happened.** Even before we started shipping Yggdrasils, Mike handed me the upgrade board for what would become Gungnir Multibit. Although this happened well after Mjolnir 2 was in development (I've had protos running on my end table for almost a year now), it definitely provided the impetus to see how far we could take the "midrange" line.

So, in light of all of this, maybe a Mjolnir 2 was inevitable. But that didn't mean the first prototype did much … except catch on fire.

A False Start

The first Mjolnir 2 prototype I did made a lot of sense. To me, anyway. It was nothing more than a simplified version of Ragnarok—and 100 % solid state. It did away with the microprocessor management system, moved to surface-mount design, and didn't use the crazy relay attenuators like Ragnarok. It couldn't drive speakers, of course, so the parts were a lot smaller, as well.

And it did exactly one thing: blow itself up.

The problem with Ragnarok's topology is that it really needs the 24/7 oversight of a microprocessor

control system, constantly checking and setting bias and DC offset. Without this in place, using only a differential servo to set the bias, the first Mjolnir 2 prototype would thermally run away and self-immolate.

Yeah, I could have added a bias servo as well, but that would be even more complication. I didn't want to end up with a design that used 200 parts to keep the 40 active ones in order. Yeah, philosophical nuttery. And perhaps I would have gotten past that eventually, but …

… but this was also about the time that it was sinking in that there would never be a tube Ragnarok. I kept trying to figure out ways to make it work within the thermal envelope of the current amp, but that kept coming out gobiteme.com.

So I had a great solid-state amp with Ragnarok (now beginning to ship), but it would never be solid state. And I had an unlistenable, unmanageable Mjolnir 2 prototype that would only work with:

a) A billion more parts.
b) Full microprocessor control, à la Ragnarok.

Frustrating. I put the whole mess on a shelf and tried to forget about it for a while.

Until one day, Mike asked me about it. And everything changed.

Don't Overlook the Obvious

"Mjolnir 2?" I answered Mike, groaning inwardly. "How's Mjolnir 2 going? Pain, fire, death."

"That good?" Mike said, sounding amused.

"Yeah. It doesn't work without the bias control. And the Ragnarok tube ain't gonna happen. And we really have to think about how all this fits together."

"Why?" Mike asked. "Does everything have to line up perfectly? Each product in its own little box?" I sighed. No, it didn't. But, "It has to have some kind of consistency," I told him.

"Sounds like a way to get stuck in a rut." I frowned. He was right. We could internally focus-group ourselves to death—without ever doing a focus group. All we needed to do was to talk ourselves into doing the same old thing, every single time.

"There's only one thing that makes sense," I blurted out, completely frustrated. A friggin Tube Mjolnir. A Tubenir. Hell with solid state, we have that solved with Ragnarok."

"Then do it," Mike said.

"But … " I began to protest, and trailed off. Because in the process of saying it, I realized, *This is exactly what we need to do.* Not another microprocessor-controlled technological tour de force, but something simpler, something more visceral, something that would really make Mjolnir 2 stand out.

Aside: I didn't know how much a tube hybrid balanced headphone amp would stand out at the time. As far as we can tell, Mjolnir 2 is the ONLY tube hybrid balanced amp out there, period. Of course, I could be wrong. Please enlighten me if I am.

"But … switching from tube to solid state, we're gonna get complaints," I told Mike, finally.

"If we do anything, we'll get complaints," Mike said. "Just the act of making a decision means we'll have complaints. The only way not to have complaints is to do nothing at all." I nodded. I knew that. *To avoid strife, say nothing, do nothing …*

But switching to tubes … while it would be undeniably cool, and while it would move Mjolnir

away from every other balanced end-game headphone amp out there … what if someone wanted solid state? We always had people asking for Lyr as a solid state amp, too …

And then it hit me again. It was finally time to do the solid state tubes.

The Long History of the Schiit LISST

Solid state tubes, like I say on the product page, are not a new idea. There have been a number of companies that have tried to make them. Here's how they go about it.

1. They pick a type of tube and look at the response of the tube at different plate and grid voltages. These "tube curves" define the transfer function of the tube—a transfer function that is significantly different than most solid state devices (except a SIT, which is a story for another day.)
2. Then, they try to come up with a combination of solid state devices (and passives) that mimic the curve of the tube. Sometimes this combination is very complicated—some have used literally dozens of parts. Sometimes this combination is pretty simple. For a look at one scheme, Google "trioderizer," for a simple

way to make a JFET work a lot like a tube (and also for an example of why you never let engineers name anything.)

3. Finally, when they've gotten as close to the tube curves as possible, they announce to the world, "We have the perfect copy of the gold-grid, pinch-waisted, JimmeeJoeBob 12RU78 from 1959, come and get perfect tube sound forever!"

The reaction to #3 above is predictable, of course: tube die-hards cross their arms, squinch up their faces, and prepare to be supremely unimpressed by what the solid-state tube sounds like. Even if it sounded better, the solid-state tube company has created a perfect environment where *nobody will ever admit it.*

And that usually sinks the solid-state tubes.

That's why we've taken the 100 % opposite approach—putting a definitively solid-state device in a can and not tweaking it to sound like a tube, and promising only, "Ya wantcher solid-state heah, we gotcher solid-state heah!"

What's interesting is how long we've been playing with this idea. From literally the first days of the Lyr (early 2011), I wondered if we could replace the tubes with a solid-state device. I even did

some experiments that showed, yep, you could, as long as the solid-state device could stand the tube rail voltage (about 200 V.)

From there, I shelved the idea. Mainly because I wondered if something that would only work in our amps would be broad enough. I mean, Lyr, Lyr 2, and Mjolnir 2 all use current sources to bias the tubes, so plugging in a depletion mode MOSFET will bias up just fine. But there's no guarantee it would work in every circuit.

So, one part of Mjolnir 2 was taking out the solid state tube idea, dusting it off, and making it work. Because it's one thing to hang some parts out of a tube socket to see if it works, and a whole different thing to have something that mechanically replaces a tube.

And LISST was actually fairly complex to get working right.

First, we had the mechanical design. Stick a couple of boards in a can, sure. But how? With what connectors? In what can? How do you stick it in the can, period?

Second, we had the topology. Sure, we wanted something simple, but what if a "triodized" version sounded better? Best to do both.

Third, we had the production. Our assembly house in Simi Valley has done a bunch of complicated stuff for us, but nothing that would require assembly like this.

For mechanical design, we actually ended up doing two PC boards in a "T" configuration, with slots for alignment of pads that are soldered directly to each other (no connectors at all). The vertical part of the T holds the depletion FETs (which dissipate about a watt or so of heat), and the horizontal part of the T holds the pins that go into the tube socket. All of this is then slid into a painted steel can full of epoxy—yes, LISST are fully potted, and therefore pretty much unserviceable. If they fail, we swap them. But, if they're used in our equipment, it's unlikely they'll fail. Lifespan should be in the many hundreds of thousands of hours.

For the topology, we actually built and listened to both the standard and "triodized" versions—and ended up liking the ones that didn't have any transfer function trickery applied. Simple as that. Could we make different LISST versions? Sure. Will we? Not unless we find a way to do it better than the current model.

For production, we ended up working more closely

with the assembly house than we ever have before. They were the ones who specified many of the production steps, including full potting. LISST is the first product we essential get, final and tested, from the assembly house. Kinda like an iPhone, but in a really ugly box.

> **Aside:** I am sorely tempted to skewer the super-over-the-top-packaging-as-art corner we've gotten ourselves into by offering a future "luxury" product in both "Ugly packaging" or "Fancy packaging," that you choose when you order. The cost of the fancy packaging would be accurately reflected in your choice. So, if this "luxury" product normally was $499 in the ugly packaging (that is, standard Schiit packaging), it might be $680 in the fancy packaging. The difference would be that we'd be giving you the choice of spending money on something you only look at once and then hide on a shelf—or not.
>
> > **Aside to the aside:** This future "luxury product" is entirely fictional. No points will be awarded for guessing what it might be. This has only been a thought exercise, prompted by some really over-the-top headphone packaging I'm looking at right now.

Okay, Finally, On to Mjolnir 2

Once I was past my tube phobia, Mjolnir 2 was a comparative breeze. And by "comparative," I mean, "only three prototype board versions, after we threw away the self-immolator."

But let's back up a bit. Because Mjolnir 2 isn't just "Mjolnir + Tubes." From the start, it was intended to be much more flexible and friendly than the original Mjolnir. And, at the same time, I gave myself permission to spend some more money to improve performance even more. Most of this ended up in the board and transformers, but I'm getting ahead of myself.

Instead, let's start with the "flexible and friendly" part. Flexibility, to us, meant at least one thing: single-ended headphone output. We actually ended up with a lot more than that, so let's take a look at the whole picture.

- **Single-ended output.** In a circlotron, single-ended output means one thing: summers. Circlotrons are inherently balanced topologies that cannot be deconvolved to provide single-ended output. In fact, some circlotrons have both outputs sitting at a rail (say, 40 V). Mjolnir 2 and Ragnarok aren't like that, but this serves to show how big of a surprise you

can get if you just try to use one output to ground. So, we added a version of Ragnarok's summer to Mjolnir to provide both single-ended headphone output and single-ended preamp output.

- **Gain switching.** No surprise here, considering it's in virtually all of our other amps. But good news if you're running IEMs or high-sensitivity headphones. Although the original Mjolnir was very quiet, Mjolnir 2 is easily 20 dB quieter in low gain mode. One trick thing, though: gain switching is via two relays, one placed in the exact center of each channel's circuit.
- **Input switching.** Mjolnir didn't have input switching at all—you chose one input, and grounded the single-ended input if it needed it with a rear-mounted switch. That's it. For Mjolnir 2, we decided to add real input switching, so you can select either the balanced or single-ended input—again with a relay that switches it right at the back panel, just like Ragnarok.
- **Front panel switches.** And, in a controversial move, we decided to move the switches for gain and input to the front panel. A result of switchgate? Nope. The prototypes had front-panel switches while switchgate raged.

We're not categorically opposed to front-panel switches; it's just that many of our products do better with rear switches. Take power, for example. There's no way the Mjolnir 2 power switch would fit in the front without compromising the placement of transformers.

But I think what's most interesting is some of the stuff that isn't seen at all.

- **4-layer board.** Look at Mjolnir 2's PCB photo on our site, and you can't see that it's a 4-layer board. Yes, I know, 4-layer boards are kind of a thing right now (or even more layers), but Mjolnir 2's board is 15.5 inch by 7.7 inch—this is a HUGE 4-layer board, especially considering the 2 ounce copper. What this buys us is much more flexibility for optimal routing of power supply and output traces—which, in circlotrons, are pretty much the same thing. Going to a 4-layer board gave us the ability to create a truly optimal layout with 16× the current carrying capability of the original Mjolnir. The board also costs about 2× as much the board in the original Mjolnir.
- **Crazy power supply.** Note that Mjolnir 2 now has not just one, but two of the "billion pin" C-core transformers we used on Mjolnir. These complex transformers give us the capability to

support the amazing range of power supplies needed in Mjolnir—2×18 V, 2×6 V, 4×25 V, and 1×200 V—without capacitive coupling, voltage doubling, or other trickery. And the billion-pin transformers, like the 4-layer board, cost over double what the standard 48 V A c-core transformers cost.

- **Super-trick servo.** You may have noticed how proud I was to get rid of the DC servo in Ragnarok. Now, it's not that I hate DC servos—far from it, I'd rather use a (good) DC servo than a (good) coupling cap—but the best servo is no servo at all, just like the best coupling cap. Without the microprocessor bias, I couldn't eliminate the servo in Mjolnir 2 … but I could make it act like it wasn't there at all. Huh? Well, with a differential topology, common-mode noise disappears. So, the DC servo feeds one side of the output directly … and the other side via capacitive coupling. The result? The servo content that isn't DC is cancelled to zero. Boom. (Almost) no servo.

The Most Boring Development in the World

Mjolnir 2 tube, or Tubenir, or whatever you want to call it, was almost comically boring in development. The first prototype started up and

ran, and I've had versions running on my side table for almost a year.

So why so long to get this to market? Simple. Because nothing is ever simple.

Although Mjolnir 2 V0.90* started up and ran, that didn't mean it was good to go. It had a bunch of problems. The biggest one was that the new high-voltage/SE/heater transformer was, well, rather low-voltage. We didn't have enough volts to regulate the 200 V rail. We didn't have enough volts to regulate the 18 V rail. The heaters worked fine, yeah, OK, but that doesn't get you far.

* You can tell how confident we are in a prototype design by the numbering of the board—V0.90 means we're pretty confident. I've started stuff at 0.20. 1.0 is release. Mjolnir 2 went through 0.90, 0.95, and 0.97 before 1.0. I've seen us do 0.97, 0.98, 0.99, 0.995, 0.998 … so sometimes we're a little overconfident.

It was, however, enough for us to get an idea of what it would sound like. And from that day, it's sat atop a Gungnir (first standard, then Multibit) in my main listening system.

"Wait a sec, are you saying this is better than Ragnarok?" somebody is certainly asking now.

And no, I wouldn't say that. However, for my preferences (a slightly "wetter," tonally rich presentation), it's ideal. Wonderful thing, personal preferences.

Aside: and before you ask, my Mjolnir 2 uses tubes—nothing super fancy, just some JAN 6ZB7s. The LISST I like just fine, and I think it's a big step up from the original Mjolnir, but the tubes are just a bit better (again, for me.) Paradoxically, I like Lyr 2 with LISST better than the real tubes (stock tubes, anyway.)

In addition to the transformer problem, my trick front-end current source was running a bit too hot … and I'd screwed up some part outlines … which meant I had a bunch of through-hole parts sticking up in the air, each leg twisted around to the right surface-mount pad.

But really, the development of Mjolnir 2 is largely a story of transformers. We went through 4 prototypes before they got it right. The next version, with higher voltages, wasn't high enough. The version after that hit the voltages … but hummed physically. The 4th version finally hit the right voltages and didn't hum.

Aside: Mike wasn't super thrilled with Mjolnir 2 until the last transformer. Getting the 200 V rail in regulation really changed the character of the amp, from one that was almost a little soft, to one with great dynamics—without being hard. Now you can't keep him away from the amp.

And, you know what? I just realized that I still need to put a chassis on the Mjolnir 2 I have—a 0.95 board on the first prototype bottom chassis.

Ah, well. Back to development. In addition to transformer stuff, Mjolnir 2 needed a completely new chassis—and it remained to be seen if our "hiding the toggle switches in a slot" would work. And yeah, the first chassis prototype showed that we had a bit of work to do. The switches didn't line up vertically. But really not bad, for the last chassis we'll do with 2D CAD. The tweaks were small enough to order the chassis without a second chassis prototype.

And despite everything going smoothly, we weren't ready for either TheShow or CanJam. In this case, blame Yggdrasil. Yggdrasil and Ragnarok have changed how we make products pretty dramatically—from a single guy putting single boards in 2-piece chassis to, well, pretty much everyone working on subassemblies on a

production line basis. Those complex products really were a lot to digest, but I think we're getting a better handle on them now.

So yeah. Beyond that, how do I create drama with development that's that easy? Short answer: you don't. And maybe you start looking over your shoulder. I've said this before, but easy projects are scarce on the ground. If you have one, the next development might not be so easy.

Or maybe we really are getting good?

Nah, not possible …

2015, Chapter 14

We Launched a DAC and Got a Movement

Or, the *Commodification of Off-the-Shelf Audio D/A Converters.*

"Huh, what?" you're probably asking. "What's Jason on about now? Yeah, you have a couple of multibit DACs, but hell, you can't even keep them in stock."

Yeah. True. And you've asked me about some other things, like what our subjective listening tests are like, and stories from the deep past. But let's face it. The subjective listening tests are nowhere near as interesting (or detailed) as the instrumented tests. After all, they're subjective.

Hell, I can summarize our subjective listening test regime in a few bullet points:

1. We build a prototype.
2. If I built it, I listen to it. If Mike or Dave built it, they listen to it.

3. The person who built it tweaks it a bit, with the help of instrumented measurements too.
4. When we're happy, we pass it off to each other for a listen. We don't say anything about how we think it sounds. (Yes, I know, the subjectivists are sighing, thinking, 'nonverbal cues', blah blah, woof woof.)
    a) As an alternative, we will pass off several different versions of the prototype to each other, with each product simply marked with "X,", "Y", or "Z," or with a switch on the board that engages different modes.
5. We'll see what each other thinks, make a few tweaks, and repeat.
6. When we've all agreed it's ready to be heard, we have a small panel of trusted listeners that let us know what they think. This may or may not result in more tweaks.
7. When our small circle of trusted listeners think it's good to go, it goes into production.

And yes, it's that simple. No ABX, no automated/robotized tests, no blindfolds. Remember, this is *subjective* testing. *Not* objective.

So now I expect you have a few questions, which I'll attempt to anticipate here:

Wait a minute! You don't do this blind? You don't even try to level-match? You don't even

make a minimal attempt to use ABX? No, no, and no. See above. This is subjective, not objective.

But this invalidates everything you're doing, don't you see? No, we don't see. We don't believe it invalidates anything at all. Especially when multiple listeners with no information or prompting come up with the same notes. Especially when multiple product variations are involved.

Well, you could subtly be influencing them— This is why we hate objective/subjective discussions. Because it always comes down to this. We use objective tests, which we go to great lengths in equipment setup and troubleshooting to make sure they're objective. We also use subjective tests, which we conduct exactly the same way we have for over 20 years (in Mike's case, over 30 years). If the combined results of this don't jive with your own preferences, you have my apologies. But this is how we do things. And you *did* ask.

Ferget that schiit, how do I get on your trusted listeners panel? Unfortunately, we're not looking for anyone else at the moment. If you are interested, though, send us links to stuff you've

reviewed and let us know why you'd like to be part of the group. I can't guarantee a response, though, since, like I said, we're not really looking for anyone now. And you'd have to demonstrate that you know what an NDA is, in any case.

So there you go. For objective testing, we use a very wide range of excellent test gear. For subjective testing, well, *we listen to it*. That simple.

On to DACs and Movements

Okay, so let's talk about the real subject of this chapter, which will surely raise some eyebrows. Here's the gist: *if your D/A converter is off the shelf, it will become increasingly harder to differentiate yourself, and therefore charge premium prices.*

Read that again, and break it down:

1. **If your D/A converter is off the shelf,** (translation, a typical consumer D/A converter chip from the handful of manufacturers surviving after consolidation—and, by "typical consumer D/A," we de facto mean "delta sigma.")
2. **it will become increasingly harder to differentiate yourself,** (how many times have people sighed recently when a new product was introduced, and said, "ah, just another XYZ

D/A converter." *Hint: much more frequently these days.* It doesn't help when formerly pricey D/A converter ICs now have 200 variants with very slightly different names, some at prices that make them usable in $300 designs.)

3. **and therefore charge premium prices,** (who really buys the $260 hamburger on the menu? Come on. Be serious.)

Or, of course, I could be 100% wrong. Maybe you'll still be able to use $2.15 DAC ICs and charge $3000+ 5 years from now. Sometimes high-end is counter-intuitive.

But I don't think I'm wrong.

Here's why: the awareness of D/A converter technology is at an all-time high right now. At least part of this is because of the relatively affordable Yggdrasil and Gungnir Multibit (and eventually an even more affordable Bifrost Multibit, as Mike already mentioned.)

In reviving a technology that everyone thought was comfortably dead and buried, we've stirred up a lot of contention. Some people think we're 100% right, some people think we're 100% full of crap, some people think we're somewhere in the middle. But the fact is, we've started a lot of

conversation, in a market that was largely heading towards a Delta-Sigma/DSD monoculture with a handful of off-the-shelf chips.

> **Aside:** and, of course, it's not just D/A technology, it's digital filter tech as well. With an off-the-shelf delta-sigma chip, you're stuck with their digital filter, or filters, most of which aren't, well, super exciting.
>
> > **Aside to the aside:** if you want to see the difference visually in the Yggdrasil's digital filter, look at its stopband rejection plots. As Mike stated, we have wider bandwidth than Parks-McClellan filters, plus an amazingly fast transition band.

"So multibit or die, right," some will scoff.

Actually, not at all. Although Mike is 100 % committed to multbit technology with defined linearity specs and no missing codes (as well as to the supercomboburrito filter), there are a ton of different approaches that a company can take, if it isn't going the datasheet/app-note/E-Z off-the-shelf route. All of them demand much more investment in time and effort than just going "off the shelf."

In fact, let's take a look at some of the other

companies doing stuff that isn't off the shelf. Here are just a few. Apologies if I missed your favorite ones:

- **TotalDAC.** They use discrete resistor ladder DACs to do multibit conversion. Very expensive to implement well, but it's definitely not off the shelf.
- **MSB.** Same here. Discrete resistor ladders for multibit. See above comment.
- **Soekris Module.** This is a DIY discrete resistor ladder DAC. You can find out more about it on DIYAudio.
- **PS Audio.** The DirectStream converts everything to 10× DSD and does not use an off-the-shelf chip.
- **Chord.** They're using their own digital filter and D/A implementation on an FPGA. Not off the shelf.
- **Metrum.** Does multibit with their own modules, which may use industrial DAC chips. Again, not off the shelf.
- **DCS.** Uses their own "ring DAC" topology that isn't off the shelf.
- **Meitner.** Decodes DSD without an off-the-shelf chip.

I think there are some Lampizators that also do DSD without an off-the-shelf chip, and I know

there's an open-source hardware design for decoding DSD without an off-the-shelf chip as well. But let's face it. I don't know the whole market, so I may have missed a few.

But when you look at the list above, you'll immediately notice a few things:

1. Not all of the "non-off-the-shelf" DAC makers are doing multibit. In fact, it may be that the majority of them are going to single-bit or their own delta-sigma format rather than multibit. So "non-off-the-shelf" doesn't have to be synonymous with "multibit."
2. Most of the products from these companies range from "wow, that's expensive," to "holy hell, that's a car." I have nothing against anyone who can drop car-money on a DAC and not flinch, but I personally would have some very deep soul-searching to do. This is a good topic for an upcoming chapter—*"when to say when."*

And I think that's why, when we launched a DAC, we got a movement. Suddenly, non-off-the-shelf technology on both the D/A and digital filter side were (relatively) affordable—in fact, undercutting some popular off-the-shelf implementations. When there was a clear-cut price distinction between the custom implementations

and off-the-shelf boxes, then there was a ton of room to position off-the-shelf converters at premium price points. Now that a Gungnir Multibit is $1249 with both the supercomboburrito filter and $120 worth of fully-spec'd-for-linearity-and-no-missing-codes D/A converters, suddenly the value proposition of $3000 off-the-shelf D/As is looking pretty grim.

> And when Bifrost Multibit is around half that price … well, there you go.

It doesn't take much of a trip around the web to see the effects, either. Multibit vs delta-sigma is a heated topic of debate in the Cavalli DAC thread here at head-fi. Multibit wouldn't even have been mentioned last year. Multibit is starting to be discussed at sites that were previously militantly pro-DSD. Multibit versus delta-sigma arguments are happening everywhere from Audio Asylum to Reddit. Again, none of this would have happened last year.

But now that there are (relatively) affordable multibit DACs based on unique tech … the discussions are happening.

A Response to the Arms-Crossed Crowd

Now, there are certainly some of you who are sitting back, frowning, arms crossed, and saying, "Well, you know, it's really the implementation that counts! Just because there are app notes and reference designs for off the shelf D/As, the tweaks and special sauce that a true audio aficionado applies to the design is what matters."

And yes, to an extent, you're correct. We have long gone past the datasheet and app notes with off the shelf chips. Yes, it does make a difference.

But be honest—does it make a $2.15 part worthy of a $3000+ player?

Be *really* honest.

"Oh, well, there's R&D cost, that has to be amortized," you'll respond.

Then perhaps they need to make more of them, we say. But that would take lowering the price.

And, let's not forget—the R&D/tweaking/etc that led up to Yggdrasil spanned 5 years. We had to literally write the book on how to use the D/As we chose. For 4 of the 4 years, we weren't sure it would work. Our solutions go so far beyond the datasheet it's not even funny. Remember, this is

arguably the first new multibit design (not using discrete resistor ladders) in the last 20 years. We had no help from the manufacturer. This was all 100 % in-house.

So, if you want to talk R&D cost of tweaking an off-the-shelf, designed-for-audio product and something that's 100 % ground up, I think we may have a bit of a better case for R&D needing to be amortized.

But that's not how we price products.

And that's why you'll soon see a multibit design that's even less expensive.

Alternately, some of you are sitting back and saying, "Well, what really matters is how it sounds."

To which we agree wholeheartedly. Even after all of our blathering, you may find you like someone else's product better. You may even like an off-the-shelf implementation better.

That's why the world is so wonderful and varied … because we can agree to disagree.

Why It Matters

"Well, if I like something else, and it's off the shelf, why does this even matter?" you may ask.

It's simple: because monocultures are dangerous. Monocultures don't change, and monocultures are vulnerable. If everyone is doing the XYZ delta-sigma DAC that does 4× DSD, then it's a pretty fundamentally boring universe, isn't it? What incentive is there to move forward?

But when manufacturers choose to step off the shelf … then amazing things can happen. Suddenly, there are much more interesting things to talk about … it's not just about whether Manufacturer Y's implementation of the XYZ chip is better than Manufacturer Z's, or if the tweaks applied by Manufacturer X really give it the edge. Now we can discuss multibit vs delta-sigma, and multibit delta-sigma vs 1 bit, the different kinds of digital filter algorithm, and much more.

And in those discussions, in those clashes of ideas, is where real progress is made.

That's why it matters.

2015, Chapter 15

Tales of Odeon, or How *Not* To Start An Audio Company

I've mentioned before that I had a speaker company right after I got out of college, and I've shared a couple of anecdotes from it, perhaps the most famous being the tale of Eddie and the Zagnut Bars. I think that story, in itself, says everything you need to know about Odeon.

But you've asked for more, so here we go. Hopefully at the end of this chapter, you will:

1. Have better understanding of what it takes to start your first business from literally nothing.
2. Gain a little respect for the need for some *tiny* amount of capitalization in a company.
3. Not think I am brain-damaged, crazy, or certifiable.

Note: I am busy digging up some old Odeon photos, which are from a pre-digital era. For the moment, you'll have to imagine what I'm talking about with a few crappy photos of some of the

survivors, shown here. Yes, I still use Odeon speakers.


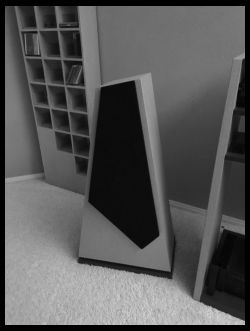


Buckle up. This'll be a long chapter.

Paleolithic Odeon

So how'd I get started with speakers?

Well, to answer that, you have to go all the way back to college. I'd always been into audio, which in those days meant things like Sony receivers and 10-band equalizers and eventually SAE surplus equipment being closed out through DAK (a wonderfully odd wholesaler in the San Fernando Valley from those days when you actually had to get off your ass and drive somewhere if you wanted to buy things—well, or do mail order—there was no internet.)

And that's what I was listening to at the time I decided to make my own speakers: an SAE integrated, a Mitsubishi CD player, some cassette deck I forgot (if you were into car stereo, and I was, you had to have cassettes), and Technics speakers. Hardly high-end.

At that time, I'd also started dismantling and reverse-engineering various bits of car audio gear, but my junior-level engineering really wasn't up to the challenge of switching power supplies, so I quickly realized that building stuff like that was beyond me. I could do simple stuff like parametric EQs and crossovers (small signal stuff), but you couldn't build a car audio company on that—you'd need amps. Plus, all the metalwork and tooling … I was 20, I didn't even know where to get started on that.

But I was already working audio into my college curriculum. I was planning on doing a novel single-ended noise reduction system* based on spectrum analysis and frequency-variable switched capacitor filters (and when I was a senior, that's what I did, complete with a near-production-worthy case.)

* Remember noise reduction? I designed and built something that was essentially a noise

gate—that would sense when there was enough high frequency content in the music to mask noise (tape hiss, this is still when cassettes were viable, remember) and allow it to pass through, but then would reduce the bandwidth to mask noise when there wasn't enough high-frequency content in the program material. In real time. Without encoding. I eventually got something that worked pretty good, but the era of noise reduction was ending … it wasn't viable (though my experience with switched-capacitor filters did get me my first job, so hey, not all bad.)

But anyway, back to speakers. The speakers I was using were perfectly OK, if you liked boring square wood-veneer boxes with three drivers up front. In other words, they were like most speakers of the day.

Being young and idealistic, I figured I could do better in two different ways: one, using more drivers so that each could be more optimized for its frequency band (please stop laughing) and two, in an aesthetically pleasing truncated-pyramid design that was both functional (less internal reflections) and cool.

And it wouldn't have fake bark on it. Yeah. I hated wood.

So I started researching drivers and sketching up ideas. There was a place called Speaker City nearby in Burbank, which is where I ended up getting most of the drivers—mostly Vifa stuff, except the 15" driver for the bottom woofer, which I forgot what it was.

Yes, you read that correctly. My first speakers used:

- A 3/4 inch tweeter
- A 1 inch tweeter
- A 3 inch dome midrange
- A 6 inch midwoofer
- A 10 inch woofer
- A 15 inch subwoofer

Yes, in one speaker. Not only that, in one speaker that was about 12 cubic feet in a truncated pyramid arrangement. This translated into a speaker that was about 50 inch tall with a 32 inch×24 inch base. In other words, friggin HUGE.

Now, I didn't have a shop to make these speakers in. Nor did I have good woodworking tools. Nor did I have any experience woodworking.

Did it matter? No. I bought some 3/4 inch particleboard (not MDF), got my dad's ancient Craftsman circular saw and Black and Decker jigsaw, bought a 6 foot long rule, and started cutting away.

On my parent's patio.

Yeah.

They were very understanding of their son's mental disorder. They became less understanding when the weather turned bad and I had to bring the speakers into the living room, as-yet-unfinished.

And finishing was one thing I hadn't figured out very well. I ended up using tons and tons of Bondo, a long finish sander, and a gallon of house paint in bright white, plus black spray paint for near the drivers.

But eventually they did look a lot like professionally-made speakers. I put together the crossovers (based on calculations from Speaker Builder magazine, much like the calculations that determined cabinet size and port size) and lined the boxes with fiberglass.

Then I fired them up.

How did they sound? Glorious. Pure perfection. Like nothing I've ever heard before. And that's why you should be very careful with the stuff you build. You lose perspective. In retrospect, they didn't sound as bad as they could have, but they were very very far from good. The

upward-pointing 3/4 inch tweeter was relatively innocuous, but the complex crossover, bizarre mix of drivers, and big boomy sub were not ideal.

But they did sound, well, BIG. Eddie loves them. In fact, he has them to this day (a second-generation "improved" pair that were much the same.)

From there, I went on to build a small bookshelf speaker (6 inch woofer, 3/4 inch dome tweeter on front, 3/4 inch dome on top—I had a weird thing for top-firing tweeters, thinking they would sound more "airy.") Again, in a truncated pyramid box. These actually sounded pretty good. I still have them.

In fact, they sounded good enough that I sold a few pairs of them during my senior year in college. And they looked good enough to impress my notoriously hard-to-impress artsy friend Jose. Enough that we starting wondering if we could actually make a business out of it.

Which led us to driving to our first CES in Vegas, to see what the business was like.

Yeah, I know. Stop laughing.

The Dreams and Reality of CES

I remember the drive out to CES, because I thought I had finally found the path to fortune (and away from the work-for-a-defense-contractor path that seemed to be the inevitable fate of engineers from my school.)

Jose and I had grandiose dreams: if we could capture only 1% of the total market (as estimated by TWICE, a trade magazine), we'd be a $75 million dollar a year company. We talked about fantastic ideas of building cities on the hilltops and cliffsides in the surreal desert landscape on the way to Vegas. We talked about using the money to start a car company with design that didn't suck.

In short, we were insane. Learn two things from this:

1. When someone comes to you with a grandiose business idea and says, "I only need to get 1% of the market (or 10%, or whatever), RUN. Fast. They have no idea what business really is.
2. Don't count your money before:
   a) You have a product
   b) You are selling it
   c) You are selling it at a profit

d) You have paid for all the parts
e) You have paid all your people
f) You've been doing this long enough that there is money left over after (e) and (f)
g) And even then, know there will be gotchas.

Yeah, we were young and naive. Now I have Schiit and an agency and Jose has his own company doing specialty costume for giant movie franchises. So, if you want a positive message from this, it's this: *don't give up*.

At CES, though, we found only more fodder for our grandiose fantasies. We didn't spend any time in the high end ghetto (we didn't even know it existed back then), but went straight to the main floor. In those days, the main floor was dominated by big audio and video names—from Pioneer to Rockford-Fosgate. They all had big, glitzy booths and impressive displays.

The one thing they didn't have: interesting-looking speakers. Square wood boxes as far as the eye could see. Jose and I rubbed our hands together in glee: how could they miss the nascent demand for stylish loudspeakers?

After the excitement of CES, we knew what we'd do. We would make speakers. And we would change the face of the speaker industry.

Yeah. Dumb. I know, I know, I know.

Building Our Shop

While we may have been idiots in terms of basic business, we were smart enough to realize one thing: my parent's patio was not a factory. We needed something bigger. Something, preferably, enclosed. And something cheap, because I was just starting my first engineering job out of school and I had no money.*

* Let me define "no money." No money means, "everything I was making went to either paying bills or chipping in to the speaker company, the speaker company had no cash flow, and Jose was living with his parents as well with only sporadic work. This is no money. This is not like starting Schiit, where I could put in some money and had income from other sources so the company could continue to reinvest in growth.

So, what did we do? We built our first shop in back of Joses' parents' house. This setting is worth a few words. Jose's backyard was big, which was good. It was also a blasted apocalyptic landscape dominated by the dead stump of a

giant olive tree. Some years past, the family had tried to unearth it, but had never completed the job, so it just sat there like an arboreal meteorite in a crater several times its size. This somewhat limited the size of shop we could build, as did a number of eucalyptus and smaller olive trees in the yard.

But we found a space, roughed out a plan for a 20 foot×20 foot shop with a slanted tin roof, and an actual working door and window. Since we didn't have the time or money to do a real foundation, we decided to make it a raised floor on 4×4s pounded into the ground. The exterior and interior were unfinished plywood. The whole thing cost only around $1100 in wood and materials and only involved minor injuries in its construction. Very palacial, huh?

And permits? Ha! The neighbors were far enough away in general, and probably had unpermitted stuff on their property. We didn't bother.

After construction was complete, we found a couple of little glitches in our plans, most notably, well, lack of power. We hadn't even thought about electricity, and we needed decent power to run a table saw, a jointer, a radial arm saw, a sander, a drill press, and an array of hand tools. And yes,

you read that right—we were actually setting up as a wood shop, to actually make speaker enclosures. Pure, unadulterated craziness, yes.

A few extension cords later, we had our power. One other little problem was lack of heat, though. To fix that, we picked up a big kerosene heater. And yes, you read that right as well—we were running a kerosene heater in a wood shop … with no dust collection system, and with a reasonable amount of high-VOC painting going on. It's really amazing we didn't end up blasted to the moon. I did almost lose another chunk of finger in the jointer one day, though …

Aside: funny, that half-baked shop lasted through the 1994 earthquake, 6 years later.

Now that we had our home, Jose turned his attention to improving our aesthetic designs. While the truncated-pyramid look was different and distinctive, he was shooting for designs that were much more radical.

And they did look amazing. Jose took the basic idea and ran much farther with it than I had ever imagined. His stuff was absolutely amazing. After Jose was done designing, we had arguably the most distinctive line of speakers around (see the photos, and consider this is 1989.) Radically

angled and futuristic, they looked like nothing else out there.

There was only one small problem …

Designs Near-Impossible to Make

If the earlier speakers were hard to make, these reached into the heights of near-impossible. The top speaker, a 6 foot tall monster, featured a seamless wraparound grille on its central spike, which was inset into a stepped base. It was also slanted both front and back, making the cabinet work a complete nightmare. It also needed close-tolerance, high-precision construction that was very, very difficult in a backyard shop. The only answer in the short term was lots of bondo and finish work (we eventually set up a set of pin router jigs that made things a tiny bit easier).

And note that I said "speaker line." As in, Jose designed a complete line of 4 different speakers, from the monstrous "Point Ones" to the bookshelf-sized "Point Fours." You know, so we could look like a serious speaker company.

Hint: this is a very, very astoundingly bad idea.

When you're starting up, you should start with a single product. Get it right. Sell some. Prove

you can make it consistently and that there are no gotchas. Then move on to the next. DO NOT start with a line. Lines are about ego. And ego gets you in big trouble.

We did our best with the line, though, using our anechoic chamber (the backyard) and some primitive measurement equipment to run frequency response plots and tweak crossovers, drivers, etc. Were they stellar designs? No, but they weren't completely crap. We tried.

So were 4 speakers enough? Not apparently. I also designed two subwoofers for the bookshelf speakers (a 10 inch and a 12 inch). Both of these were very cool looking, but pretty bad as subs go.

And more … this was also the era of the Bose Acoustimass sub/sat systems. Everyone wanted Acoustimass. These were horrendously overpriced plastic speakers with a bandpass subwoofer. But they were small, and they had reasonably impressive bass for their size.

So, I decided we needed a competitor to them—the Point Five. This used small 4 inch woofer and 3/4 inch dome tweeter satellites with a bandpass subwoofer featuring two 5 1/4 inch woofers. In a rare burst of sanity, the subwoofer was a sim-

ple rectangle, but the satellites were our neat truncated-pyramid design.

I mention this design for several reasons:

1. I still remember putting together the first prototype, holding the tiny woofers, and thinking, "There is no possible way these will sound any good."
2. I still remember taking them into the living room and firing them up—and standing there with my jaw on the floor. These little, tiny, cheap speakers played better than virtually everything else we made—they were accurate, they played loud, and the bass was amazing from a half-cubic-foot subwoofer. Eddie and Jose came in and stood there, floored. As Eddie said, "Screw everything else, let's just build these."
3. However, they were far from perfect. Their impedance was pretty brutal. Run it from a Sumo Andromeda 2, and they sounded like giant-killers (top executives from JBL and Infinity came to listen multiple times the first time we showed them at CES.) Played from a cheap receiver, not so much.
4. Despite being handmade, and sold through a dealer model with dealer margin, they sold for only about ½ of what the Acoustimass did.

Yeah.

5. They were arguably the first speakers sold in 5.1 sets for home theater.

In retrospect, I should have listened to Eddie and built only the Point 5s. We could have refined the design and made it truly stellar. But we were young … and had egos to feed.

So, with such a huge, sprawling line of speakers that were very, very hard to build, the result was that we had to make them pretty much 100% ourselves. Sheets of MDF came in one side, and finished speakers came out the other. We experimented with how to best finish the speakers, starting with Formica laminate, then abandoning it when the laminate peeled right off the speakers in a hot car. Plus, the glue really sucked to use. Plus, it had visible seams.

What we finally ended up using for finish was both a blessing and a curse. It was called Zolatone.

Some of you are groaning. For those of you who aren't, Zolatone is a textured, multicolored spray finish which can charitably be said to look a bit like granite (in some of the more sober colors) and is very, very good at hiding small blemishes. Paint a fairly rough box with Zolatone, and it came out looking perfect.

Plus, it wasn't wood. Remember, we hated wood. Plus, it was the 80s. Zolatone was one of those 80s abberations, like putting saxophones in bands that had absolutely no need for them.

Zolatone was also pretty easy to use and to spray. It didn't require expensive equipment. It coated well. It wasn't horrifically expensive.

But … one little detail … the VOC content of Zolatone would probably set off alarms in downtown Los Angeles if sprayed anywhere within 10 miles. This stuff was nasty, smelly, and had more weird solvents in it than I imagined could ever be in something sold over the counter.

Painting inside with Zolatone was suicide. Getting a properly ventilated paint booth was out of the question. Getting real respirators wasn't gonna fly, either.

So how did we deal with it? Simple. We painted outdoors.

And yes, I know, this is not exactly OSHA-compliant painting procedure. Nor is it particularly environmentally friendly. But we had no money. We were young. We got away with it. And that's why we did it.

Around this time, we decided it was time to do our first brochure. To do a brochure, we needed pictures. And we wanted to do cool pictures. Sounds pretty normal so far, right?

Well, our idea of cool pictures involved taking the whole line out to Vasquez Rocks—this is the place where Kirk battles the horrible plastic lizard in the original Star Trek—and using the bizarre-looking rocks as a backdrop.

Did it involve a professional photographer? No.

Did it involve picking a great day for photography, like, a day that wasn't 40°F and windy enough to knock speakers over? No.

Did it involve getting insurance? No. Hell, we didn't even know you needed something like insurance.

So what did we do? We threw the speakers in Jose's brother's van, drove out to Vasquez Rocks, hauled the speakers out to a clearing in front of the rocks, set them up, and began taking pictures.

After a while, a park ranger came by on a horse and asked what we were doing.

"Taking pictures," I said, thinking, Well, isn't that kinda obvious.

"Do you have a permit?" he asked.

My brain melted. *You needed a permit? To take a few pictures?*

"Do you have insurance?" he continued.

"Uh, well … " I began, and trailed off. Because there was no way we could fake our way out of that.

So, a few minutes later, we started hauling the speakers back to the van.

Now, the photo shoot wasn't a complete disaster. We actually got some good shots for the brochure. As usual in those days, I did the darkroom work in my bathroom, set up a brochure in Pagemaker, and then took the photo to be married to the layout and printed (back then, scanners were a little too insanely expensive—hell, we were considered pretty high-rent because we actually had a fax machine!)

A few hundred bucks later, and we had a brochure. Which meant we could start marketing. And by "marketing," I mean, "Driving around to

stores and showing them our speakers, while they looked very, very confused."

Because this was the pre-internet era, remember. There were only three ways to get your stuff in front of prospective customers:

1. Get individual stores to carry it. Which is what we tried to do.
2. Sign up distributors, stores, or chains at trade shows. I'll get to shows.
3. Get a rep and have them do both of the above. We didn't know about reps.

In the end, the "driving around to stores and showing them the product" strategy got us exactly zero dealers. The dealers seemed very confused by our new, stylish, futuristic, non-square, non-wood product.

It seemed like, maybe, just maybe, the market for adventurously-styled speakers was not as big as we thought.

Our Savior: Trade Shows

So how did we sell anything? Simple. We went to trade shows. Specifically, CES. This is when we discovered the "high end ghetto." High end audio wasn't on the main floor of CES. It wasn't

in the ballrooms off the main floor. It wasn't on one of the other specialty main floors, like they had for car audio. It wasn't even in the rat-mazes of smaller booths that flanked the main floor.

No, it was in the oldest, ugliest, and smelliest hotel we'd ever been to: the Flamingo. And, as an added bonus, it was grouped together with the adult video exhibitors.

Yes, you can see where high end rated in the EIA (now CEA) world.

And the first show was pretty exciting. We'd gone out of our way to bring the whole line, but we ended up playing the Point Fives most of the time. They were impressive enough that most people thought the 6 foot tall Point Ones were playing. Eddie had to walk over to the tiny subwoofer and cover the bass port to kill all the bass to make them believe it was those tiny boxes. That's the show the JBL and Infinity engineers came by, then brought their engineer friends, then brought some upper management. Very exciting! If they thought we had something, maybe we did.

Of course, we did all we could to get attention, which meant we didn't fit in very well. To put it mildly.

You see, Eddie was in charge of the music. And Eddie had less than 10 nanoseconds of patience with boring audiophile music—you know, simple female vocals, string quartets, smooth jazz. Eddie called it, "Music that would sound good on a walkie-talkie." So Eddie played pretty much whatever the hell he wanted. This ranged from Kraftwerk to 2 Live Crew to Kiss. Usually at earbleeding volume.

And everytime he cranked it up, people streamed into the room, blinking and grinning, as if they just coming off being extras in Apple's 1984 commercial and were trying to acclimate to the real world.

Aside: I seriously wonder if we were the first people to play stuff other than jazz, classical, and vocals at a CES. Because people certainly acted like it

Now, while Eddie was blasting 2 Live Crew, the room next door was a little less than thrilled. The room next door was actually occupied by Avalon Acoustics, one of the first companies to jump in on the ultra-price bandwagon with some $15 000 speakers. Of course, they were trying to play their jazz trios and string quartets at 80 dB

and look all serious and audiophilic for the press that was coming by.

After several complaints (including one that said we'd knocked the pictures off the wall of their room—with our two 5 1/4 inch woofer sub, ha ha), we finally reached a truce: we'd turn it down when they were doing serious demos or when they had press. Ironically, we didn't have to turn it down much—all the traffic was coming to the Odeon room. They didn't like that, either.

And that first show got us enough interest, and enough orders, to think we really could make something of Odeon. We actually took several cash-in-hand orders from international companies. Suddenly we had distribution! And people who seemed to understand what our cool-looking speakers were all about. We even picked up a couple of dealers. Maybe this would go somewhere.

But the problem was, in-between shows, not much happened. We were 100 % dependent on shows to keep the orders flowing … and between our sprawling product line, late deliveries, some quality problems, and other growing pains, we never really had steady growth.

But Eddie continued to give great demo, and in general be a character. I linked earlier to the story about the Zagnut bars (and why we always, always win the Stupid CES Story of the Century), but I didn't cover all of his escapades, which included fun stuff like using the megadollar tube amps we'd borrowed for one show as burger warmers (they fit perfectly between the tubes) and asking one prominent press personality if he'd jumped out of a helicopter and landed on his head, after he "grounded" his shirt to the carpet with nonconductive plastic wires and roach clips. Eddie was, in reality, the original heart and soul of the communication style behind Schiit—he unknowingly shaped a lot of what we did.

A Real Shop?

During the course of Odeon's growth, we did eventually move out of Jose's backyard. We found a relatively palacial cinderblock building with a roll-up door in Sylmar for cheap, next to the National Guard Armory and a meat-packing plant. One thousand whole square feet of wonderful space! Of course, the heat and AC didn't work, but (as Mike said) it was cheap.

Now, did we finally put in a dust collection system?

No.

Did we fix the heat? No. We kept using the kerosene heaters.

Did we fix the AC? No. We sweated.

Did we buy better tools? Not really. Though we did Frankenstein our table saw into a 5-horsepower monstrosity with an 8-foot fence and no safety guards at all. Again, it's amazing people didn't lose limbs at Odeon.

And we did get better at production, largely due to Jose's brother setting us up with some pin router jigs for the bigger pieces, and Jose building his own set of jigs for some of the smaller and more complicated parts. We got to the point where we were building a good product, and were able to deliver it on time.

And we did some very cool stuff. The later Super 5s and Super 3s (which I still have to this day) edged into the "seriously good" category, largely due to heroically thick cabinets, minimal crossovers, time-aligment, and lots of measurement and tweaking. We only made one pair of each, though. If they had ever sold, they would have been very expensive (in Odeon terms—$2000 to $3000 a pair.)

But … bottom line, the orders didn't take off and multiply. We made enough to keep the doors kinda open, in that we could pay for parts, and pay Eddie and Bob (Jose's brother) a pittance, but Jose and I never took a salary. In fact, most of my disposable salary from my job at Sumo went into Odeon, never to be seen again.

And … the biggest kicker for me … I was holding down a more-than-full-time job, in addition to running Odeon—and "running Odeon" meant I did a lot of hands-on production. A typical day had me getting up at 6 AM to go into Sumo, do my engineering work, get out and get to Odeon around 6 PM, then work to midnight, painting boxes, assembling crossovers, bolting speakers together, stretching speaker grillecloth, whatever needed to be done, until about midnight. Then to bed and start again. Weekends were just more Odeon and less Sumo, but there was tons of Sumo stuff to do on weekends. And to further complicate things, I was beginning to design stuff for Mike Moffat at Theta.

After 18 months of 100+ hour weeks, I was well and profoundly burned out. Jose was tired and broke. Eddie and Bob weren't making enough money to live. And Jose was starting to see interest in his sculpting abilities, starting with a

commission from Treasure Island hotel in Vegas.

Jose and I agreed. It was time to stop fighting the good fight, and wind it down.

Prophetic Words (Ha!)

Once we decided to close up shop, I felt two very profound things:

1. **Freedom.** I wasn't chained to Odeon anymore. All my money didn't need to go to Odeon. I could buy things for myself. It was truly an amazing feeling.
2. **Loathing.** Specifically for self-employment. "I'll never start another company," I said, confidently, as money started coming in from the Theta work and things began looking up. I told Sumo's metal vendor that he was a fool for starting his own thing, working for someone was way better, and paid a lot more.

But things change. Sumo stumbled in the post-Asian Financial Crisis market, and Mike and his partner had a falling-out at Theta. Suddenly, my gravy train was looking, well, pretty thin.

And that's what led to me eating my words less than a year later, and starting my second company: Centric. Luckily, that one went somewhere.

But it would keep me out of engineering—and out of audio—for another 15 years.

2015, Chapter 16

When to Say “When”

Since the introduction of Yggdrasil, we’ve been sucked into more and more discussions that include what Mike likes to call “the Audio 1%.” At the same time, I’ve watched some listeners go all the way up the chain to car-priced DACs to try to find some incremental upgrade to Yggdrasil’s capabilities.

So, I gotta ask, “When do you say, ‘when?’” I think it’s a question we should ask more often. I’ve seen people go down some very deep rabbit-holes (as in, spending $100k+ on audio gear, while living in a mobile home). I read jokes about it—*welcome to the hobby, sorry about your wallet.* I have had long conversations with enthusiasts angsting over whether or not they should go for the latest megadollar product intro … *am I missing out? What if it’s as good as they say? I could get that and be done with buying forever …*

Except you know it won’t be done. Ever.

Go back and read that again. The current latest and greatest will be superseded by something that is even later and greater, and the gnawing little voice will come back … *what if it's dramatically better? Can I afford to miss this? Just one more purchase …*

Except it's not just one more. Not the last. Not done. Not ever.

Knowing when to say "when" starts with accepting this fact. There will always be something better. Or at least something that seems better. Or claims better. Always. Forever and ever. Even if you buy $ 20 000 monoblock amplifiers. Even if you commission Nelson Pass to personally hand-build bespoke 1000 W Class A 4-chassis monoblocks out of iridium and diamond at a cost that rivals the GNP of a small nation. Even if you find that lost crystal transmitter in grandma's attic that calls down the space aliens, who give you an antimatter-powered direct-brain-stimulation with 100 % guaranteed PerfectPerception™ synapse certification, containing every concert in Earth's history, from a 10 000-year-old bone flute to the latest EDM.

Accept that there will *always* be something better.

It. Will. Never. Be. Over.

Lessons from Other Hobbies

"Okay, so how can you sit back and pontificate about all of this, all cool and Spock-like?" some will be asking. "Are you perfect, while all of us are flawed?"

Ho ho. Nothing could be so far from the truth. I've gotten sucked into the black hole. Big time. I just took my header in a different hobby: cars.

I've owned exotics. I've owned high-dollar perfect restorations. I've owned several of both at the same time. Because the corollary to *There's always something better* is *there's no dose like an overdose,* or, "If one is good, two is better, and many is best."

Except … what the hell do you do once you have your exotic?

There'll always be a better one next year. And even if you can ignore the one next year, the body style will change in another two or three years. And then you're back on the lot, with a trade that cost you $ 30/mile to drive it.

And what do you do once you have your perfect restoration?

Every time you drive it, it's less perfect … and the more you look at it, the less perfect it seems. Surely someone could have done a better job on the frame, you think. And you begin wondering. And you start looking at cars for sale again.

And, when you're into the racing side of cars, you rapidly become familiar with this expression: *Anything can happen on the track.*

Your new Ferrari 458 was just beat by a modded STI on a road course? Sure, seen that. New Viper killed by an electric 510 at a drag race? Sure, seen that. A whole bunch of people at an illegal street race deciding that the old 50s iron that a couple of old-timers brought out was way cooler than racing, transforming the whole thing into an impromptu car show? Sure, been there too. Sometimes it's not about being the fastest, or being the most exotic, or being the most perfect.

And those are usually the best times.

And then you start realizing, *You know, I don't really enjoy any of these damn things, all I do is clean them and wrench on them and then worry every second I drive them … so what the hell am I doing?*

Well, either that or you get enough money to go full-on crazy and get a warehouse and a staff to maintain them … but there's no real danger of that around here.

And you realize even more, as you clean and wrench and angst and shop for even more cars:

1. This crap is eating my life. There's a whole lot more interesting things than cars out there.
2. Nobody really gives a crap (except for me) how new and exotic or old and perfect my stuff is.
3. Buying the ultimate is never the ultimate, never the end.
4. The buying high doesn't last very long, and frequently comes with an expensive hangover.
5. The more I try to buy my way out of my dissatisfaction with my current stable, the less satisfied I am.

But that's the reason you see me driving much less crazy iron these days, even though I have even greater capability to spend on cars.

Because I realized, *There's always something better.*

And decided (after wasting tons of money, time, and angst) *it didn't matter.* In fact, super-exotics and perfect restos now amuse me more than

anything. I sit back and nod, knowing how they really feel … the paranoia, the insecurity, the whole crazy merry-go-round. Someday they may realize just how crazy they are, and say "when." Or maybe they won't.

And that's how I feel about car-priced DACs and amps that cost like home remodels: amused. Because I know those ultimate systems will soon be up for sale on Audiogon, as the owner chases the next "last and best purchase."

So, How Do You Let Go?

Okay, wonderful, but how do I achieve this same level of don't-caredness? you might be thinking.

Good question. I'm not sure I have all the answers, other than really burning yourself out and learning a big painful lession, but I hope I can give you a few signposts.

1. **When someone says,** "The new Snortledorfer III DAC is the closest thing to live music, and clearly justifies its $ 28 463 price tag," or "Of course the $ 2000 NambyPamby 1000 isn't as good as the $ 11 500 Pflugternik Amazonian S.VII, because it is brighter and more etched,"

silently append the following statements in your mind*:

a) In his or her opinion.
b) According to their biases.
c) Among the handful of products they've heard.

2. **Think of the amazing world of experiences** you're shutting yourself out of, when you OCD-out on one tiny thing (like audio.) Instead of buying that car-priced DAC, imagine:
   a) Taking a trip to Antarctica (or another exotic destination where you have enough time to get out beyond the tourist traps and broaden your horizons.)
   b) Or go to a few dozen (or a few hundred) live concerts
   c) Or, closer to home, touring vineyards in France
   d) Or, more practically, learning French. Or Chinese. Or English.
   e) Or, wackier, learning to fly a plane and getting your pilot's license.
   f) Or getting the gear you need to brew your own beer
   g) Or transforming that back room into your perfect workspace, just like you always wanted.
   h) Or getting the tools you need to take up

woodworking
   i) Or a 3D printer
   j) Or hell, just buying a car
3. **Keep reminding yourself,** *It's never over.* No matter how much you think your audio journey will end with your big-ticket purchase, know that it won't be. Remind yourself of when you've said, "This is the last big buy," or "That should do it … forever."

* I think the continual regurgitation of **opinions** as de-facto, settled facts is the biggest source of angst, flames, and in-fighting in audio—and, by extension, one of the biggest sources of "not saying 'when.'" If every reviewer … hell, if every *person* on a forum was required to repeat, *"This is my opinion, I am biased like everyone else, and I haven't heard everything,"* until they got this *viscerally*, in their gut, the world would be a lot better place.

News frigging flash:

- Not everyone likes the same sonic profile
- You may be 100 % opposite the preferences of a self-professed expert
- The person may be totally wrong or mistaken
- These de-facto pronouncements are sometimes (er, usually) made based on short au-

ditions in different locations, and sometimes made as regurgitations of other pronouncements, with no knowledge of the product at all
- And … you don't know if the person doing the pronouncing has an agenda

Bottom line, there's a lot of great gear out there. There's terrific stuff for everyone. But you, and only you, can decide *what's for you.*

So How Do You Decide When To Say When?

Aha. This is a harder question.

First, it's a hell of an egotistic thing for me to say, "Yes, I can advise you on how best to spend your money." Bottom line: I can't. I'm not you.

Plus, I bet there's a big variance in how much you *can* spend, and how much you *feel comfortable* spending. If I was to guess, the "how much you feel comfortable spending" would be a better way to judge when to say when, especially if you append, "for a product that isn't critical to my life," to the phrase.

Or at least that's how I'd play it. If I can't easily and comfortably spend the money, I'm already at the "when." And even if I can easily and spend

the money, I might still pass on spending it. If I already have a system that meets my needs, or is close to what I'm contemplating, I'd most likely pass.

But I'm very fiscally conservative. You may be more adventurous.

And you may even be the kind of person who never has to say, "when." You have the wherewithal to simply keep buying and expanding your audio gear without regard to price. You have multiple large homes, expensive cars, and deep liquid cash reserves. You are *true* Audio 1% types.

But I think a lot of the people buying 1% gear don't necessarily fall into this type. I think that many of the customers for 1% gear find buying something that costs five figures is a significant expenditure, especially when there are no easy financing or leasing options.

And this is when you have to decide: audio or lifestyle? Sound or travel? Gear or personal enrichment? That's a decision only you can make.

And, of course, there are tons of people who aren't looking at 1% gear at all, and for whom a $1000 or $100 system purchase requires a

profound sit-down-and-think-on-it session before making the plunge.

And I understand. I've been there. For much of my life, a $100 system was about all I could afford, unless I made it myself. For another big part of my life, $1000 was crazy, unless that was just the retail price and I got it via a gear swap or designed it for a company I worked for. And I went through both the three-figure and four-figure stages thinking, "Man oh man, it's gonna be great when I can really put together a cost-no-object system!"

But a couple of things happened on the way to that. One was the audio market veering out of control into full gold-plated-Bentleyland (when I started working for Mike at Theta, I thought a $3000 DAC was beyond the pale … when $16k DACs started appearing, it was a step so far up the ladder, I knew I'd never be comfortable there.) The other was that I realized that I enjoyed the cheap Cobalt about as much as the Thetas, when it came right down to it. Yes, it wasn't quite as good, and no, it wasn't going to light the world on fire, but it was fun, musical, and enjoyable. And that was enough.

So I learned when to say "when" in audio pretty

early. When I say "Magni 2 and Modi 2 may be the only amp and DAC you'll ever need," I mean it sincerely. When I tell someone, "Don't bother with the amp and DAC at all, spend more on your headphones," I'm 100 % serious.

Can you spend more, and get more? Sure.

But is it *fun?* Maybe that's a better metric. Is it as thrilling as that first taste of exceptional sound? Is it so much better than going out and experiencing live music, or traveling, or learning to cook or brew or fly? Is it so much fun that it eclipses everything else in your life?

Maybe it is. *But maybe it isn't.*

Aaaaand … *if you aren't having fun with your gear, your gear is having fun with you.*

2015, Chapter 17

The Multibit Revolution Gets Cheap

Author's precis: Mike Moffat will certainly cover this in his own blog, but he's still got a few decades of catching up with the present to do … so, until then, you have my perspective on the Bifrost Multibit.

Well, maybe not *absolutely* cheap, because $599 is still a decent chunk of money, even when counted in 2015's inflation-decimated dollars.

But $599 is what we were charging for the Cobalt 307 back in 1993, and that was considered a breakthrough in terms of value-for-dollar. When you consider that $599 in 2015 is $364 in 1993, the Bifrost Multibit is pretty amazing.

Don't believe me? Consider:

- Bifrost Multibit uses a 100% multibit DAC, where the Cobalt 307 used a hybrid multibit+delta-sigma DAC—and it's a multibit DAC we would have swapped a testicle for in 1993 (weeeellll, maybe not literally, but you know what I mean)

- Bifrost Multibit uses our proprietary time- and frequency-domain optimized, closed-form filter (AKA the "supercomboburrito" filter), where the Cobalt 307 used an off-the-shelf digital filter
- Bifrost Multibit has 3 inputs, including asynchronous USB 2.0, where Cobalt had only 2—optical and coax

And, the big one:

- Bifrost Multibit is built as an upgradable platform, so future changes in USB input or DAC technology don't mean instant obsolescence, where the Cobalt 307 was never intended to be upgraded—it was a "disposable" product

So, for about 40 % less cost, the Bifrost Multibit delivers a hell of a lot more value—and a hell of a lot less obsolescence—than a DAC Mike and I designed 22 years ago. So, even if it isn't "cheap" in absolute terms, it's certainly moving in the right direction!

Okay. Pause for breath.

I kinda jumped in here, getting all bean-countery on you right up front. I probably should have made it clear that this chapter is about Bifrost Multibit. And yes, it'll have all the back-story and design details you expect.

But leading with the financial aspect is interesting, isn't it? It shows that we *can* push costs down—and deliver more—in high-end audio.

And that's something both Mike and I are very proud of.

Multibit on the Desktop

From long before Yggy, Mike had been planning for multibit upgrades all the way up and down the upgradable DAC line. In fact, he'd already made special provisions for the current requirements and the low-voltage power supplies of the Analog Devices SHARC DSP that we were using.

And, as soon as we introduced Yggdrasil, people began wondering how far down the line it would go. Some guessed that Gungnir would be a relatively easy target, but less of you were convinced it would hit the Bifrost.

And for Gungnir, transforming it into a multibit DAC was relatively straightforward, as I've covered before. In fact, it was so straightforward that Gungnir Multibit is very similar to Yggdrasil, in terms of number of DACs, type of DACs, and discrete output buffers.

But Bifrost was another matter. Although the electronic provisions were in place for the DSP engine, there was still a lot of debate about which DAC we'd end up using. After all, Bifrost has less than 1/4 the total analog board space as the Gungnir, so the packaging would have to be extremely efficient.

In fact, early development had Mike and Dave wondering if we'd have to go in a different direction than the Analog Devices DSP:

"What about the filter?" Mike asked, referring to our supercomboburrito filter (that runs on the DSP that needs quite a bit of board space for it and its accouterments).

"Weeellll, there are some really good sample rate converters we could use," Dave said.

"Sample rate converters?" I asked. "For what?"

"Instead of the filter," Mike said.

"Instead of the burrito?" I asked, mouth hanging open. "And *sample rate converters?*"

"Not *asynchronous* sample rate converters," Mike said, waving a hand. "*Synchronous.* Use it to bring up the input sample rate 4×—synchronously—

so you get, say, 176.4 out for 44.1 in, then it gets interesting. If they sound good."

"There are some good ones out there," Dave said.

"A sample rate converter?" I repeated, wondering inwardly, *How am I gonna explain this to everyone … hey, we got your shiny new multibit Bifrost here, but it uses a sample rate converter, and yeah, I know we talked some Schiit about sample rate converters, but those were asynchronous and these are synchronous …*

"A *good* sample rate converter," Mike corrected.

Still, I shook my head. Copy-wise, it would take some fancy footwork. But I remained silent. Because the Bifrost Multibit was still vapor. There were no prototype boards, no detailed schematics, no decision on a D/A. There was an AD5761 16 bit DAC, similar to the AD5791 and AD5781, but that family is processing-intense, needing headers written in before each sample … which would use up even more board space.

There was also the geriatric AD1851 16 bit audio DAC from Analog Devices, which was paradoxically still in production, and not super expensive. I mentioned that one to Mike and Dave, but they just muttered and left it at that.

And it wasn't long before Mike found the DAC he really wanted to use.

The Bizarre DAC

"I've got it," Mike said, one day when we were both at Schiit. "The AD5547."

I frowned. "What the hell is that?"

"It's a DAC. Look it up."

I did. "16 bit?" I asked, doubtfully.

"Look harder," Mike said.

I did. And I realized that the specs on this DAC were really amazing. So amazing, I should call out the relevant points with numbers.

1. **It features an integral nonlinearity (INL) plot that's better than ±0.5LSB.** This is a spec they never provide for audio DACs, because (a) it would be terrifying in the case of audio multibit DACs and (b) it is not possible to measure delta-sigma DACs in this way, since the output depends on the preceding and following samples.

2. **The THD performance is actually scary good for a 16 bit DAC—far better than 16 bit.** 16 bit level THD, from a theoretically perfect 16 bit DAC, is −96 dB. Most 16 bit DACs from the Jurassic Age of Digital didn't hit this number. The AD5547 is −104 dB, much better than 16 bit.
3. **Its noise level is down at the 22 bit level.** Think about that for a bit. Why would you need 22 bit equivalent noise for 16 bit audio? The answer is, you don't. And in the old days, you couldn't. Many, many DACs from the early years of the PCM era couldn't do 16 bit noise levels.

And, it was a dual DAC. As in, we'd only need one chip per stereo Bifrost. This meant less board space, and an easier implementation … except for one thing: the AD5547 is a parallel input DAC.

"Huh?" you may be asking. "What's that, and who cares?"

Well, most DACs take their data in serial form—data goes in on one single pin. You send it that stream of data, plus a few clocks, and it's happy. Most audio DACs are built so they easily interface with the output of typical USB and SPDIF receiver chips, so it's kinda like building blocks. Hook up a few pins and go. Even I can do it.

But the AD5547 takes the data in parallel, on 16 separate pins, plus other clocks to ensure it converts samples at the right time, and both at once. Bottom line, the clock logic and data input to this DAC are fundamentally different than audio DACs—so there's no easy paint-by-numbers implementation. Mike and Dave would be figuring out how to interface a sample rate converter (or, I still hoped, the full burrito filter) to another bizarre DAC—one completely different than the AD5791 and AD5781.

There are many days I'm happy to be an analog designer.

Hard Design Decisions

From there, Mike and Dave seemed to drop the development of a Multibit Bifrost. It made sense—we were getting deep into Yggdrasil production and ramping up Gungnir Multibit, so they had plenty to do on those fronts without any other distractions.

I didn't realize how close we were to having a Bifrost Multibit until Mike asked, "You know, if Bifrost Multibit had an op-amp analog stage, how upset would you be?"

“Why?” I asked. “Are we really that space-limited? I know a real discrete I/V is a nightmare, but if we do parts on both sides of the board—“

“You need two gain stages per channel,” Mike said. “It’s a multiplying DAC.”

“Oh,” I said, letting that sink in.

“It comes down to the burrito or a discrete I/V,” Mike said. “And even then, I don’t know—“

“Burrito,” I said.

“You sure?”

I nodded. “No question. The burrito is the DAC.”

Mike rubbed his chin and considered for a moment. “How’re you gonna spin this, after all the words about discrete?” Mike asked.

I laughed. “It’s a lot easier to explain this than sample rate converters.”

“So you wouldn’t be upset?”

“No.”

“Good,” Mike grinned, and handed me the finished, working Bifrost Multibit board.

Yep. That's right. Mike and Dave had just gone ahead and done it. And, as far as I'm concerned, they made exactly the right decision to retain the filter and go to op-amps for the I/V stage.

Why? Because discrete I/Vs are (a) complex, (b) have complex power supplies, or (c) both of the above. Any simple discrete I/V was out of the question because we didn't have the four independent, non-ground-referenced floating supplies they would require. So that meant it would have to be complex.

Complex, as in Theta Gen V level complex. And the discrete I/V I did for the Theta Gen V had something like 250 parts on a 4 inch×6 inch Teflon PCB—for a single channel. The total size of the Bifrost Multibit analog board is 4 inch×5 inch.

Plus, going discrete in the Bifrost Multibit was completely different than our other DACs:

1. The discrete stage in the Yggdrasil and GMB are both just buffers. The DACs used there have voltage output. So, they are very simple discrete stages, just four active devices per channel. However, as measurements clearly

show, this simplicity does not compromise distortion performance (this is usually the penalty paid … simple discrete amps typically have high THD.) That's why when you see some "2-PPM wonder amp" it usually has about 80 active devices per channel. We can argue till the cows come home which sounds better.

2. The discrete stage in the standard Bifrost and Gungnir is actually a small amp stage—not exactly a discrete op-amp, due to its very specific gain structure and open-loop bandwidth beyond the audio band—but it also takes a voltage output of the DAC, amplifies it a bit, and passes it on. No I/V necessary.
3. The AD5547 is a multiplying DAC, which means it needs two very well-matched gain stages for exceptional accuracy, not just a single I/V converter. Which means it would need twice the circuitry. Fully discrete? As I mentioned, aha, no.

So … even with surface mount components, the math was simple. It all pointed to "nope, sorry, divide by zero, sums to nothin, system is unstable, ain't gonna work around here." Anyone looking for a comparison to a mythical discrete-output board for Bifrost Multibit will be disappointed … because *we simply never tried to do it.*

> And no, there is no conspiracy where we're holding back a better analog board that will magically appear in 6 to 8 months, when sales start slowing down in the spring.
>
> > Now, that doesn't say we'll never have a better analog board for Bifrost. Hell, it's possible that even better multibit DACs will be available in the future. Or a new delta-sigma technology will force Mike to recant and go 100 % pro-DSD. Though the third sentence is, I think, about as likely as a meteor hitting me on the head before I finish this current sentence.
> >
> > > And, sentence finished. There you go.

So, was that it? Did Mike just hand me a board and call it a day? Is that all the development drama I could come up with?

Nope. Though Bifrost Multibit, like many of our latest products, was remarkably drama-free. The most fun we had was taking the finished Bifrost Multibit to a meet and the Schiit Show, 100 % incognito, to see if anyone noticed.

And, you know what? Some *did*.

"That's not a stock Bifrost," one of our early listeners told me at the Schiit Show.

“Sure it is,” I told him. “It says so right on the product mat.” (We’d had mats made up with product specs and pricing, so people new to the gear knew what they were listening to—one of the smartest things we’ve ever done.)

“No it’s not!” he insisted. “Is there, a … an upgrade, ah, coming?”

“You know we don’t talk about products before they’re released,” I teased, grinning.

“Come on! Tell me.”

And I did, after asking him not to repeat it on the forums. He smiled even wider. “October, huh?”

“We hope,” I reminded him. “Lots of things can change.”

And lots of things could change. We could possibly have introduced Bifrost Multibit earlier—well, maybe—but, more likely, it could have been pushed later. That’s why we don’t talk about new products. Because they aren’t products until they’re on the shelf.

Now, for the arms-crossed brigade out there saying, “Well, you heard it from the top, Schiit punks us at the shows, there’s no telling if what they’re showing is what it actually sounds like,”

let me clarify:

1. We took the prototype to at least one meet, where pretty much everyone knew what it was (we told them and asked them not to talk.)
2. We took two prototypes to the Schiit Show, because, hey, we thought it would be fun to see if anyone notices.

Beyond that, here's our super-secret method for choosing the best-sounding products to display at shows:

1. Denise asks for what we need.
2. Alex pulls it off a shelf and sends it.

Yes, that's right. No special tweaks, no magic burn-in protocol, no hand-picked-for-best-measurements, no take-it-out-and-see-if-it-works. Just off the shelf and on to the table. As if you purchased it.

And … if you're still miffed at us having a bit of fun at a meet and our own show, please accept my apologies. We were just having a little fun.

Now, while we were having fun at the shows and beating up on the Bifrost Multibit prototypes, lots of other wheels were turning. Because Bifrost Multibit wasn't just a product launch, it was a product launch, a product change, an upgrade

launch … and also the launch of a new AKM chip in the delta-sigma Bifrost.

> Yes, we did a delta-sigma update too with the AK4490. And that's all that needs to be said about that. Not that it's bad. Just not super interesting. Except for the price drop.

Why is this a big deal? Because we couldn't simply set up a single product page and make it live when the time came. We had to:

- Verify specs for the Bifrost Multibit and the Bifrost "4490"
- Write copy for the Bifrost Multibit, the Bifrost "4490", the Bifrost Multibit Upgrade, and the Bifrost 4490 upgrade
- Get photos taken of both boards, both on the motherboard and off
- Set up the upgrade system to handle the coming Bifrost upgrades
- Buy tonnage of parts and make about a billion Bifrost Multibits and Multibit boards, so we'd be ready for the onslaught of upgrades—while we'd had some practice with Gungnir, Bifrost had about 12× the number of products in the field, so the potential volumes were daunting
- Argue about, and set, pricing—increases in our run size and lower parts cost meant that

we could reduce the price of the Bifrost and have a very aggressive price set on the Bifrost Multibit*
- Get the tonnage of boards through the PCB assemblers
- Produce finished Bifrost Multibits to have on shelves for the launch

* Mike correctly predicted the consternation from the price drop. Yes, it's an unusual thing to do in audio, but I'm still convinced it's the *right* thing. Even if the bitchfest that ensued made my cynical side say, "Ya shoulda taken tha money!"

Now, of course we can't get into production without at least one little surprise. Considering the complexity and scale of this product, *one* little surprise would be a minor miracle. So we ended up with two.

- **Surprise the First.** The prototypes (that we had listened to for hours, taken to shows, etc.) were one sample off between channels. This is what happens when you're working on interfacing a DSP to never-before-used DACs. What this translates to is phase shift between channels that varies with frequency. No, we hadn't heard it. No, nobody else had heard

it, though some listeners had commented on "there being something discontinuous between the lows and highs" and "woolly bass." Given the price point and musical performance, we didn't discover the error until we were into the final qualification phase. This also goes to show why you shouldn't bring prototypes to shows … the prototypes weren't all they could be. Har har. In punking the showgoers, we'd punked *ourselves.*

- **Surprise the Second.** The production first articles had a singing voltage reference. As in, it was oscillating. It was at such a low level you'd probably never hear it, but it showed up on the Stanford right quick. We made a couple of component changes and fixed it up before the full run.

Which meant that the Bifrost Multibits that went to RMAF were not only the first production Bifrosts (with both Surprises addressed) but were also correctly labeled as Bifrost Multibits. On Friday, October 2nd, Can-Jam attendees were able to hear Bifrost Multibit in its full glory for the first time.

And, I'm thrilled to say that not only did we announce on time, but we actually shipped on the day of the announcement. Yes. As in, the

first Bifrost Multibits shipped within minutes of their orders.

Almost like we're getting good at this, I sometimes think, then worry about the next intro …

Because you can always get caught off guard.

2015, Chapter 18

Death (and DNR) of a Product

Okay, I've covered the death of a product before, way back in the Asgard 2 chapter. But, while Asgard died, it didn't stay down. It rose again, an even more capable amplifier, in Asgard 2 form.

Now, let's talk about products that die and stay dead. Do not resuscitate, DNR, finis, dirt nap time, however you want to call it. Because in the last few weeks, we've had a couple of DNRs at Schiit.

Oh, you didn't notice?

Well, that's a hint as to why they died. Because, let's face it, if you aren't going to bring back a product, there's a reason. So let's talk about the products, and the whys and wherefores of why we chose to euthanize them … and why they aren't coming back.

"Wait a sec, what products are you talking about?" you ask.

#1: Modi Optical
#2: Loki

Yep, there you go. Two DACs down, five still standing. We think that's plenty. And that's a great place to start on the whys and wherefores.

Philosophical Reasons to Kill a Product

There are tons of reasons to kill a product. Many of them are simply practical … but I'd rather start with the more philosophical reasons. Because the philosophical reasons really are the most important. The philosophical reasons keep you on target—delivering products that only you can make, ensuring they are relevant, and helping your customers choose easily between them.

> And, really, no kidding … don't discount the philosophical reasons. Agree or disagree, love them or hate them, it's one of the major reasons Apple is one of the most valuable companies in the world.

So, philosophically, why did we kill Modi Optical and Loki? There are two:

- **Managing Product Sprawl.** As I've noted in the past, our line is pretty big. Maybe a little too big. It can be intimidating for someone

who isn't familiar with us to wade into a list of 7 DACs with prices ranging from $ 99 to $ 2299. Which one do they need? What are the differences between them? Why move up the ladder to a more pricey model? Questions like this resonate in people's minds, and can actually lead to someone choosing to go elsewhere. If a company only offers a single DAC, it's easy to understand what they do and what it does. When you have a huge line, it's not so hot. So, philosophically, we were looking for ways to reduce the size of the line. And that led us to Modi Optical and Loki. Without Modi Optical in the line, Modi became a simple two-tier choice of Modi or Uber. Without Loki, we didn't have the distraction of DSD (people constantly asked if Loki was a better DAC than Modi, not knowing it was a completely different DAC that likely wouldn't play their music *at all*.)

- **The Multibit Revolution.** After Yggdrasil, Gungnir Multibit, and Bifrost Multibit, it should be very clear where our heart is at in digital—in bringing back exceptional multibit DAC performance based on modern platforms and a unique digital filter, so you can enjoy your current digital library to its fullest. Or, "DACs for the music you *have,* not the mu-

sic you *have to buy.*" (Compare the current number of tracks available in 16/44 and those available higher, and you'll see that high res is currently about 0.1% of total music—and that's *all* high res, not just DSD). And, at the same time, making multibit DACs that are also ready for sane high-res formats, not just 16/44.1. Sure, delta-sigma has a place, and that place is in exceptionally affordable products like Modi and Fulla. But with our hearts in multibit, Loki was a philosophical sour note, an odd man out—it championed DSD, which is the complete antithesis of multibit.

Now, of course, philosophical reasons will only get you so far down the path to murdering a product. If Loki or Modi Optical had been selling by the cartload, I'm sure we would have found a reason to have a stay of execution.

But they weren't. So let's talk practical reasons for axing a product.

Practical Reasons to Kill a Product

Okay, these reasons are a lot simpler than the philosophical ones … but you'd be amazed how many companies ignore them. I've seen it a dozen times before, from Sumo to the present day.

Why?

Well, I'll admit it. It's really hard to kill a product. You spend a bunch of time developing something, refining it, and running across the lines to get it to production, and you've got some real emotion invested in it. You don't want to see it die. You want to see it do well.

And, it's very easy to rationalize and say, "Well, sales may pick up soon," or, "Well, with one or two tweaks, it might really take off," or "Well, it's not really taking that much effort to keep in the line … "

But, you know what? Sales won't pick up unless it's put on sale … and then you're dependent on sales until the end of your natural days. *Don't. Ever. Do. Sales.*

And you know, those "one or two tweaks" may not be so easy, or you would have done them to start. And those one or two tweaks may take the product out of budget, or have it colliding with your other products. So, you're taking a chance of going down a dead-end road. Your call.

Aaaaaannd, "not that much effort to keep in the line," only seems like not very much effort when you're staring at racks and racks of products you

can't move. As soon as you have to make a second run of a slug-selling product, you have to devote time and money to it—and know that your money will be tied up for a looooong time.

Do I sound like a beancounter? Yep, you bet. Because this is called "practical reasons," not "feel-good excuses."

And, with that, let's take a look at some practical reasons to Kevork a product:

- **Slow Sales.** It doesn't take a Ph.D to figure out that if a product isn't selling, it's a candidate for culling. And with Loki being the least popular product we'd ever introduced, well, there you go. It took us three years to sell the first and only run of Lokis. Never again. Second least popular product: Modi Optical. We did a few runs of this one, but it never lit the world on fire.
- **Pain to Support.** Loki also generated far more support inquiries than any other product, because DSD is a supreme pain in the ass to get working on many software players, and also because, since it was a low-cost product, most buyers didn't want to spend more money on a software player and used Foobar, which is not exactly the most intuitive thing to configure. Hell, people had so much trouble with Loki

that we had to take it off of Amazon and recall all stock—it was generating too many one-star reviews. Coupled with slow sales, this made Loki a helluva target. Modi Optical wasn't as bad, but it generated more than the usual number of inquiries from people who were confused about what format to send to it, and about problems with very long optical cables, so it wasn't exactly a winner in this department, either.

- **Not Unique in the Line.** Loki was certainly unique, but with the proliferation of music players that transcode DSD on the fly, and the general reduction in interest in DSD, it became less interesting. Modi Optical was just a subset of Modi 2 Uber, so it wasn't unique enough to keep in the line.
- **Costs Too Much.** Okay, now we're starting to get into factors that didn't, well, factor into the decision to kill Loki or Modi Optical. But it bears mentioning, because this is another important practical reason to kill a product: it costs too much to build, relative to what you can sell it for. In business-speak, this is a "low margin" product. Most businesses have a multiplier that they would like to see on the build cost, and most businesses allow some flexibility on it … but some allow too

much. Sumo was very good at rationalizing low-margin products … and it didn't help the company at all. Fun fact: Schiit's lowest-margin product is Asgard 2—but that's not going anywhere, because it sells very well, and it pretty much never breaks. Seriously, I think their failure rate is less than 0.1%.

- **Hard to Build.** Again, this is another factor that didn't really count towards Modi Optical and Loki meeting their own personal Valhalla. But again, I have seen this eat companies' profits alive. Sure, maybe the bill of materials cost is low, but if it takes hours of tweaking to make each product, then there's a problem—especially if your other products are relatively easier to make. I don't think it's a mystery that our hardest-to-build products are Ragnarok and Yggdrasil. But again, they aren't going anywhere.

So, when it was time to make more Lokis and Modi Opticals, we simply … didn't. We let them sell out, then we turned off the products. They're gone, done, fini.

Aside: When it was clear we'd finally sell out the first run of Lokis—almost 3 years after the introduction—I went to Mike and said, jokingly,

"It's time for you to get working on Loki 2." His response was unprintable in a family magazine.

> **Aside to the aside:** No, seriously, Mike wants nothing to do with DSD anymore. The Loki was an interesting experiment during peak DSD hysteria, but now that the DSD advocates are realizing that Sony's vaults of DSD aren't going to swing wide anytime soon, the interest in native DSD decoding is at an all-time low.

So What About All Those Lonely Orphan Products?

What about all those thousands of Lokis and Modi Opticals out there? Are they suddenly stuck out in the cold, lonely orphans doomed to live out their lives in sorrow?

Of course not. They simply keep doing what they're doing, usually until long after the warranty period. If they break, we fully support them (hint to manufacturers just getting started: just because you cancelled a product doesn't mean you can stop supporting it—plan on going well past the warranty expiration date on the last one you sold—this means having any unique parts in stock). And that's about the end of the story—

unless you come back to the philosophical aspect again.

Because here's the bigger question: what about the potential customers who wanted an affordable optical-input-only DAC, or an affordable DSD-only DAC?

It's easy to say, "well, they move on," but let's take a deeper look at this—at what is driving the demand, and how big the demand really is.

What was driving the demand, on the optical side, was largely Apple devices, and to a lesser extent, CD players with optical outputs. And we've seen the future with Apple … most of their newer devices (Macbook Air, Macbook, new Apple TV) are losing the optical connection. So the market is shrinking. And if you think there are tons of CD spinners being sold these days, think again.

And AVRs and TVs don't count on the optical side, even though they have tons of those connectors, because with the format confusion between surround and 2-channel PCM, it's best to simply stay out of that mess. It creates more work for support to have to explain, over and over, to a neophyte with a flat panel or AVR, that optical won't necessarily pass 2-channel PCM from all

sources due to the idiotic copyright restrictions built into the HDMI inputs.

So, the market is changing—the biggest driver of optical uptake (for us) is moving away from it, and the complexity of AVR/flat-panel devices makes that segment unpalatable. We're about simplicity and fun. So there you go.

With DSD, it's a longer discussion. Because I completely understand that if you have several hundred legal HDCD rips, or if you've gone all-in on the selection of new classical available in DSD, maybe you want to have a DAC that does DSD.

Aside: but "does DSD" is something worth discussing, too, since most DACs convert DSD into an intermediary multilevel delta-sigma format, so is it really DSD? Another philosophical question—one we'll leave to the manufacturers of DSD DACs. It is worth noting that there are some DSD DACs that maintain DSD's single-bit nature all the way through—sometimes to the point of not even using a DAC chip, but simply a switch, to decode the signal. Those manufacturers clearly believe in DSD on a philosophical level, and you can probably bet they've done a lot of work to optimize their products for DSD … so if you're all-in on DSD, it's probably best

to go all-in with a manufacturer that isn't just including DSD in their buzzword compliance database.

However, our experience seems to indicate that most audiophiles are kinda like us, with thousands of physical CDs, or thousands of (legal) lossless rips of CDs, and maybe with a Tidal subscription that gets them access to 25 mm tracks—all of which are 16/44, and as a total universe, vastly dwarf the amount of music available in high-res, let alone DSD.

So DSD is not in our future, unless the uptake changes in some dramatic way. To us, basing our design decisions on music that is far less than 0.1 % of the market (and has the built-in presupposition that you need to purchase, or re-purchase, the tracks) makes little sense. The same way making our design decisions based on formats that are not purchasable in any way, shape or form (8× DSD, 32/768, we're looking at you) makes little sense.

But that's us. Other manufacturers will disagree. They'll argue that we can only move the industry forward by supporting bold new formats. And, you know what? Maybe they'll be right in the end. But we think, for the foreseeable future,

that making the most of the music we have—16/44—makes the most sense.

Products Should Stay Dead, or, The Paradox of Cancellation

I don't know why it always works out like this, but whenever you cancel a product, suddenly three things happen:

1. A whole bunch of people call and email you asking where the product is—and, by "a whole bunch," I mean, "10× to 50× more than ever gave a crap when they could just click "buy it now."
2. If you sell through dealers and distribution, you'll get five large emailed orders for the cancelled product—not enough to make a run worthwhile, but sitting there, tempting.
3. You'll say to yourself, "Hell, I never should have cancelled the stupid thing, look at all this interest."

And then, if you're silly, you'll un-cancel the product and do one last run. And that's when you find out three other things:

1. The whole bunch of people who came out of the woodwork are tire-kickers, and won't buy the product.

2. The dealer/distributor orders got chopped in half, or cancelled, in the time it took to produce the product.
3. You'll say to yourself, "Crap, I was right the first time, should have kept it dead."

Sumo un-cancelled products. Every time we did, it worked out exactly like the above. Every time, we kicked ourselves. Every time, we got stuck with a bunch of slow-selling stuff.

And then the rationalizations began again … *it's not really all that much effort to keep it in the line* …

Sometimes it's time to kill your babies. Do it. And don't look back.

2015, Chapter 19

The Most Difficult Design Brief

Okay, so let's imagine for a moment that amplifier design briefs are conducted kinda like hiring a sketchy spaceship on a certain desert planet—as in, at a table in a seedy bar, with both the marketing wonk and the engineer keeping their cards as close to their chests as possible. Such a design brief might go something like this:

"I need an amp," says the marketing wonk, casually, as if the company can survive on nice thoughts and well-wishes, as if they don't really need, well, like, *a product*—a damn good product—that people actually want to buy.

The engineer sighs. "What kinda amp?" he answers, after a time, knowing that there are a thousand ways to give the marketing wonk what he wants.

"A good amp, and cheap," replies the marketing wonk.

“How many watts, how many channels, powered by what, what application, how big, what kinda numbers you want, THD, noise, you know … details?” the engineer rattles off, crossing his arms.

“It’s a headphone amp, maybe a watt, and small, desk-friendly,” the marketing wonk says.

“That’s doable,” the engineer says, thinking, this is a cakewalk, throw one of the big TI headphone drivers in a box with an op-amp and wall-wart and I’m done in time for dinner.

“But I was thinking, well, you know, about something special, like fully discrete,” adds the marketing wonk, not looking up from his drink.

The engineer clears his throat and says nothing, thinking: *Now you’re talking real amplifier design. Thermal stability. Bode plots. Testing with reactive loads. Tons of parts. Only a company like Schiit would be crazy enough to do something like that on an inexpensive desktop amp.*

“Price?” the engineer asks, finally.

“Hundred and fifty, hundred and seventy, tops,” the marketing wonk says, looking hopeful.

The engineer keeps his poker face, quickly calculating the cost of a complex discrete product with a ton of parts. "Anything else? You want a DAC in there too? Navigation system? Slow-cooker and microwave? Remote control?"

"No, no," the marketing wonk waves his hand, as if to cut through the engineer's rapid-fire onslaught of ideas. "Except … "

"Except what?" the engineer prompts.

"Except, we were also thinking, well, like, maybe a tube … " the marketing wonk trails off, smiling hopefully at the engineer.

"A tube?" the engineer asks, speaking very slowly, as if to a child. "As in, a vacuum tube? As in, you want this to be a tube amp?"

The marketing wonk nods.

The engineer, laughing uproariously, pulls out his blaster and blows the marketing wonk to atoms.

The Challenge of Inexpensive Tube Amps

Okay. I had my fun. Let's now get down to it. And that "it" is this: inexpensive tube amps are a huge design challenge. That is, if you want to do them right.

And they're an even bigger challenge if you want to do them right, make them flexible, and make them act like, well, a modern solid-state amp.

But I'm getting ahead of myself. First, let's talk about what makes a tube amp so expensive. Tube amps are fundamentally different than solid state amps, and many of those fundamental differences equate to higher cost. Things like:

Output Transformers. Many tube amps use output transformers to drive low-impedance headphone loads, or speakers. Tubes are fundamentally high-output-impedance, high-voltage, low-current devices (in the context of 32 Ω headphones or 8 Ω speakers), and output transformers translate high voltage and low current to lower voltage and higher current. One problem: output transformers are not cheap—in fact, good output transformers usually cost much more than a power transformer. Which, in itself, is one of the most costly things in any product. So, add the cost of two good output transformers to an amp design, and watch the MSRP skyrocket. That's why many inexpensive tube amps are OTL, or Output Transformer-Less designs. Like Valhalla 2. Through careful selection of tubes and design of the output stage, Valhalla 2 can drive many low-impedance, high-sensitivity

headphones credibly … but it will never be able to drive them like a tube amp with output transformers can. Luckily, Valhalla 2 is absolutely stellar with high-impedance headphones, which need less transformer magic.

Tube Outputs. And, let's face it, simply using output tubes (rather than solid-state outputs in a hybrid amp) is costly. The tubes can be expensive themselves, plus you'll need a lot more heater current to run them (see below). That's why Step 2 in cutting tube amp cost is usually losing the output tubes and making the design a hybrid. There's no shame in this—a well-designed hybrid can be very, very good. But it is, in general, a cost-saving measure.

High Voltage Supplies. To run tubes right, you need high voltages—think hundreds of volts. Tubes need a lot of voltage to run linearly and give their best performance. This means high-voltage capacitors, custom transformers, high-voltage-rated parts, maybe even through-hole parts exclusively, depending on the voltages you're shooting for. When you're running 200 V rails to get 100 V on the plate (like we do in Valhalla 2, Lyr 2, and Mjolnir 2), the whole design is going to be more costly than an amp running 30 V or so (a typical solid state design). Unfortunately, many inex-

pensive tube amps are "starved plate" designs, running at 12 V to 24 V … which is a recipe for high distortion and a very colored amp design.

Tube Heaters. And the power supply keeps getting more complex when you talk tubes. You also have to take into account the tube heaters, and the current you need to run them. Hey, those glowing tubes glow because they are actually hot. It's not cosmetic. Tube heaters do exactly what you'd think: they heat up the cathode so there's a convenient cloud of electrons to make them work. No heater, no worky. The problem is, tube heaters can draw 300 mA to 600 mA of current at 6.3 V (AC or DC)—each. That's automatically 2 W to 4 W of power dissipated just in the heater of a single tube if you're running AC heaters … and even more if you're running DC heaters. This is a pretty beefy power supply in itself. As far as we know, no inexpensive tube amp has ever just thrown a tube in there and lit it up with an LED … which would make the tube literally just for show.

So now you know a little bit about why an engineer might moan and groan when asked to design a good, inexpensive tube amp. Sometimes, the "good" and "inexpensive" can seem to be mutually exclusive.

“And yet you did a pretty good job with Vali,” some of you are saying. “That’s a helluva good sounding little tube hybrid, for hella cheap.”

Yes. And it did sound very, very good.

That is, if your cans happened to fit a fairly narrow profile (relatively inefficient, full-sized headphones that did not reveal the high noise floor), and if you were tolerant of a big design foible with the subminiature tubes (high microphonics), and if you didn’t need a preamp out, or gain switching, or wanted to roll tubes …

Yes, it sounded great. But it was a very limited amp. It wasn’t for everyone. It was no substitute for the do-all Magni 2 and Uber.

Which is why I started to ask myself: *What if we could make an inexpensive tube amp … without the limitations of Vali?*

Yes, I know. Get out the blaster.

Adventures in Seconding the Vali

As I mentioned, Vali had plenty of obvious things we could do to improve it. Adding a gain switch and preamp outputs were the two most top-of-mind. But simply adding a gain switch and

preamp outputs wouldn't alter the basic character of Vali … and it wouldn't address the real limitation of the design, which was the microphonic tubes.*

* Although we got very good at sorting out the microphonic Valis and keeping them out of customer hands, they were a huge liability in the shop. First, the tubes had to be tested, sorted, and matched, then installed in a product, then burned-in for a few days (an unusual amount of time for an inexpensive product), and then, when they failed—usually by continuous ringing—they had to be manually desoldered, replaced, and the whole process started all over again. When a run of 1000 Valis yielded 150 that had to be reworked—after pre-sorting the tubes—you can begin to understand why this was such a disruption.

So, it was apparent, very early on, that the best way to improve Vali would be to change the tube.

But what tube? And what would be the ramifications of changing tubes?

The original Vali used an unusual tube: the 6088 subminiature pentode, in triode-strapped mode. These submini pentodes were odd in many ways:

- They didn't require a very high plate voltage—30 V to 40 V would be fine for them. You could probably even get away with a 24 V supply and 12 V plate, as some inexpensive amps do.
- They didn't require a lot of power for the heaters—only 20 mA at 1.25 V (yeah, only 25 mW to the heaters … big difference from 2 W for a 6DJ8-style tube.)
- They were direct heated, so the cathode and the heater were inextricably intertwined, which led to some limitations when it came to looking at a way to add a gain switch to the amp.
- They were inherently noisy (hiss), which limited their application as a preamp stage.
- They were very microphonic (diiiinnnnnnnnnnnggggggggggg …) which, as above, limited their application as a preamp and for sensitive headphones.

For the purpose of choosing a new tube, the first two oddities are the most important. Could we find a tube that worked at Vali's low plate voltage (about 45 V) and wouldn't cause the power supply to self-immolate with the amount of heater current it needed?

Furthermore, could we maybe run the heaters via AC? We had a 14 V AC wall-wart. A small drop

would get that into the range where two series 6.3 V AC heaters would work, or 12.6 V heaters in parallel.

We started out by looking at some very similar tubes—different subminiature pentodes, specifically Russian rod pentodes. Rod pentodes were close enough in terms of heater and plate voltage to drop in and tweak. These devices have one huge advantage over Vali's old tubes: they are pretty much non-microphonic. However, they were also noisier. And their operating point wasn't super-conducive to a headphone amp (higher gain, hard to get rid of the gain and noise because of the direct heated cathode.) And, they were very high distortion.

So we moved on. We looked at some very cool subminiature Russian triodes … but those needed 400 mA of heater current at 6.3 V. This would be a helluva load on our power supply, and, with the tube inside the chassis, the Vali could substitute as a portable grille (seriously, you're looking at 60 °C or so).

So we moved on again. This time, I started thinking, "You know, what we really need to do is hang the tubes outside the chassis, then the heat from the heater won't be a real problem. So,

we looked at some 7-pin pentodes that only used about 150 mA of heater current each. That could be viable, but I was nervous about the supply of those NOS tubes.

So we moved on again, looking at different tubes (12AU7, etc) that worked well at low volts and had 12 V heaters. I didn't like the performance of those tubes, and their new-production nature meant those tubes would be costly.

Deep into the project, I was stuck. The real problem was the power supply. Using a single 14 V AC wall-wart, yes, I could step it down to run 12.6 V heaters, and use a linear voltage multiplier for a high-voltage rail, or I could use a switching supply to derive a regulated 6.3 V DC for a 6.3 V tube, but neither of those options appealed. What I really needed was a power supply like I used in the higher-end amplifiers, with multiple AC voltages for multiple purposes. Making one single AC voltage try to do everything was an exercise in compromise.

But when you're using a wall-wart, you really could only get one AC voltage on its barrel jack. That's the rub. Two connection terminals = one AC voltage. Done and done.

At that point, I almost decided to leave well enough alone. After all, Vali was selling well. It was a good amp for many headphones. It just wasn't really, well, universal or flexible.

Except …

What if our wall-wart guys could do a wall-wart with two AC voltages? And have something like a DIN plug connector, rather than a barrel jack?

If they could do that, then I could run the tube heaters off AC at 6.3 V, so there wouldn't be any power dissipated by DC regulators … and *also* get an AC voltage high enough to rectify or multiply into a proper tube plate voltage.

Now, I'd never seen a wall-wart like that before, but it was worth a query. I fired off an email, expecting to receive a "sorry, we can't do that," response. Instead, I got a picture showing a wall-wart with a DIN plug connector. *"Like this?"* the email asked.

Aha, now we were talking! I arranged to get a sample of the wall-wart made up.

A Vali With 2 Wall-Warts … and a Bizarre Gain Stage

While I was waiting for the wall-wart, I realized that I could get an idea of the performance of the power supply simply by using one of our 6 V AC wall-warts and one of our 16 V AC wall-warts with a voltage multiplier … and so that's how I started hacking up a Vali.

And, while I was at it, I figured, *Why not just use a 6BZ7 tube, like Lyr 2?*

We had an excellent supply of them, and while it might look silly sticking out of a small amp like Vali 2, it would be a good test case to see what we could do.

And it's a good thing I started there, because it quickly showed me one of the limitations of even a two-AC-voltage design: there would have to be some compromise on how high a plate voltage we could run. 100 V like the Lyr 2 or 120 V like the Valhalla 2 just wasn't going to work … because we also needed that same AC voltage to supply Vali 2's solid-state output stage.

So, I started looking on how we could maximize voltage on the plate. What we ended up with is a truly bizarre combobulation of a tube and a PNP

transistor (this is a hybrid, remember). But the transistor is NOT the plate load … the plate load is still a resistor, though a much smaller one than you might expect. I haven't seen this topology elsewhere, but that doesn't mean that someone hasn't done this before. Most good analog ideas have.

So what does this weird gain stage get us? It gets us a full 60 V on the plate, which is plenty to get good performance out of 6DJ8 tubes and tubes in the same family. Now, 60 V may not sound like a lot, but in the context of inexpensive amps that sometimes use 6 V to 12 V on the plate, it's a huge, huge difference. It also means we can fully swing the discrete output stage into clipping, so the output stage is highly efficient. Also, with a new biasing arrangement, the output emitter resistors are almost 5× smaller than some of our other amps, for lower output impedance.

Sounds like a ton of design and optimization work for a very simple stage, right? That's because it was.

And, it wasn't over even after I'd gotten the first dual-voltage wall-wart and PC board prototypes for Vali 2. Because, in the process of testing which tubes we could use in the design (besides the stock

6BZ7), I found that the new combobulated tube-PNP hybrid gain stage ran at different operating points for each tube we used. Some were just a little bit off … but some were different enough that the amp's power output would be affected.

So, I added a bias-centering servo. Which makes the little Vali 2 a fairly complex example of a modern hybrid: a tube for voltage gain, discrete bipolar transistors for output stage, an integrated circuit op-amp for bias control … on top of a dual-AC-voltage, high-voltage+heater power supply with regulated rails for both the tube and bipolar output stage.

Sounds complex? Yes. To achive simplicity and flexibility in use, internal complexity is usually required. Not that this is a bad thing, nor is Vali 2 overly complex (in fact, it uses only about 10 more parts than the original Vali.

A Thoroughly Flexible, High-Performance, Affordable Tube Hybrid

So what's the end result of all of this design work? An amp that is, by all measures, a significant step up from Vali. It has:

- Higher output power

- Gain switching
- Preamp outputs
- Tube rolling
- Easily replaceable tube
- A single tube to keep rolling costs down
- By far the most sophisticated power supply in its class
- A unique tube gain stage
- A unique discrete output stage
- No starved-plate design or integrated buffers
- And, like the Ubers, a brushed aluminum top and knob

So how does it sound? Well, Mike, I, and the rest of the (formerly skeptical) shop think it leaves Vali in the dust. But the final judgment, of course, is yours.

And when you get right down to it, sound or not, flexibility or not, the main thing I hope to communicate with this chapter is simple: amp design isn't simple or easy—even if it's "just" a revision of an existing product. When you go off the reservation and start talking about discrete design, or tubes, it's a much bigger undertaking than you might think.

Perhaps even more interestingly is that in Vali 2's case, there was no design brief. No discussions in seedy off world bars. Hell, we didn't even have

the amp in the overall "2015 plan." It was always a, "Well, let's see what we can do with this." As I mentioned, I came close to simply letting Vali soldier on for a few more years. But in the end, the relatively slow pace of development put me in a "no pressure" mode that allowed me to explore some really different designs … from the weirdo wall-wart to the combobulated gain stage.

And in the end, I think Vali 2 is really a big step up for inexpensive tube amps. But hey, I'm biased.

2015, Chapter 20

Looking Forward to 2016

So it's the end of another year. A good time to look back at what went right, what went wrong … and what might be coming next year.

Aside: And yes, I know, Mike has been teasing you about The Gadget, but that's really early dev stuff right now—think Theta Frankenstein stage, if you remember back to the earliest days of high-end digital. I can't say much more about that, other than to reiterate it is like nothing else ever done in digital … and that I'm sure it'll generate its fair bit of controversy when it arrives (is it right, is it wrong, is it fair, is it really "better," etc.)

But I'm getting ahead of myself, with an aside only two sentences in to the year's recap. So let's turn around and look at what we did … and what was good and bad with that.

Most obviously, we introduced several new or updated products. Much less than last year,

though, as I promised. Mike and I met after-the-fact of 2014 and asked ourselves, "Why did we do so many products? Why so many updates?" Hell, there seemed to be products we updated for no good reason, when we looked back on them. And, like I said last year, it was just too much. It strained our capacity for growth, and caused any number of production glitches, hitches, delays, and challenges.

Or, in numbers, last year we had 12 product intros. This year, we had 6:

1. Yggdrasil (final-friggin-ly!)
2. Gungnir Multibit
3. Bifrost Multibit
4. Mjolnir 2
5. Vali 2
6. LISST

One new DAC, two significant DAC updates, two significant amp updates, and a new idea: solid-state tubes. Not too bad, pacing-wise. If we averaged that, it would be a new or updated product every two months. Still maybe a little faster than we'd like, but is LISST really a product unto itself (it seems to be—they are selling like crazy) and are updates really new products? The answer to both is "well, probably not." LISST are built for us start to finish by our production

partner in Simi Valley. Updates—especially DAC updates—don't change the chassis significantly, so a big part of the "oh craps" that we run into in production are eliminated.

So, if you toss the updates and the LISST the only new-new product was Yggdrasil. By that metric, we had only one product introduction. Positively lazy!

Except … maybe not so much.

Surprise the First: Multibit Profusion

As I mentioned before, Mike handing me the Gungnir Multibit before the Yggy was actually shipping was a huge surprise. I'd figured we'd be able to do something with Gungnir … you know, sometime far in the future. Instead, we were staring at the ability to do something with it, like, *right now.*

But *should* we? We actually talked it through. Like most of our talks about product strategy, it went something like this:

"You know, most companies would milk the Yggy for at least a year before introducing a downmarket variant," I told Mike.

“Yeah, but we aren’t most companies,” he told me.

And that was that.

Gungnir Multibit followed Yggy by only about 3 months. Bifrost Multibit followed that by 5 months. While the world was still struggling to digest the bizarre, non-DSD, 21 bit Yggdrasil (see next section below), we’d sucker-punched them with an entire line of multibit DACs, starting at about 1/4 the Yggy’s price.

Why is this significant? Two reasons:

1. As mentioned, we didn’t drag this out two years, spacing the products at “audiophiley appropriate” times. Instead, we said, “Hey, let’s get the maximum multibit benefit out there at all feasible price points, as rapidly as practical.” Suddenly, a whole lot of stuff out there starts looking overpriced.
2. All the multibit platforms we use are entirely new, and not based on moribund or NOS DAC chips. This means we aren’t limited on how many we can make, and it means we’re bringing entirely new things to the market. Suddenly, off-the-shelf audio DACs don’t seem so interesting.

And well, actually, let's add a third reason: despite the huge R&D effort needed for Yggdrasil, we didn't start with a hogged-outta-solid-aluminum, gold-plated monstrosity that cost as much as a car … even though we certainly could have. That set the bar for all of our multibit DACs, on down the chain.

No, wait, let's add a fourth: because our DACs are upgradable, 2/3 of the multibit DAC line is available as upgrades to people who already own our DACs. This is a huge win for everyone—saving even more cost.

So yeah, DAC-wise, the end of 2015 is dramatically different than the end of 2014. In 2014, we had zero multibit DACs. Now we have three.

Holy moly, what a change.

Surprise the Second: Multibit Reception

When we introduced Yggdrasil, I wondered what the world, currently drunk on dreams of 32/384 and 8× DSD, would think about this bizarre, 21 bit, no-DSD DAC. I knew what it sounded like, and I knew that 99.9 % of digital music is 16/44, but, you know, what would everyone *think?*

Turns out I shouldn't have worried. The biggest problem was actually keeping up with demand. We broke the world supply chain for AD5791s (or at least it seemed like we did). We had to move up runs. We had to plead with metal suppliers to deliver early. We had to set up an entirely new, separate line to keep making Yggys.

What did people think about it being 21 bit? Except for a very few questions, nobody seemed to care or notice. All they cared about was the sound.

How about the "crippling" lack of DSD (as some opined)? A non-issue. Even the most ardent DSD fans are beginning to realize that the Sony vaults are not opening wide, and the format will remain a tiny niche.

Aside: yes, for some who have large investments in SACDs that they then have ripped, maybe multibit is not ideal. We get that. It's cool. We can't please everyone.

And beyond popular response (which is really what matters), what happened? Amazing stuff:

- A blockbuster review of Yggy in Computer Audiophile
- Product of the Year from Computer Audiophile

- Product of the Year from Headphone.Guru
- Gungnir Multibit in the Rolling Stone Holiday Gift Guide
- A big interview of Mike Moffat at Audiostream

So yeah. Reception. I think it's been fine.

Surprise the Third: Big Changes to Analog

Another surprise, for those who have been watching our prior "second gen" products, is how much the analog side of things changed. Both Mjolnir 2 and Vali 2 are significantly different than their predecessors—far more than a few new features, tweaked cosmetics, or fractionally higher power output.

Mjolnir 2 went from a solid-state amp to a tube amp. It also became the first amp you could choose to run solid-state or with tubes. It was literally nothing like its predecessor, save for one transformer, and the use of a circlotron topology.

And, introduced with Mjolnir 2, LISST also allowed Lyr 2 and Lyr owners to switch to solid-state if they felt like it. It's a small product, but it allows us to offer something that nobody else does—the ability to run several different amps as either tube hybrids or all solid-state.

And let's not forget that Mjolnir 2 added something I bet you thought you'd never see: front-mounted switches. (Yeah, I'm being a bit jocular here, so shoot me.)

Vali 2 saw similar huge changes—literally, the output device size and the overall size of the chassis are almost literally the only things unchanged. It's a monumentally more flexible amp than the original Vali. Heck, it even allows tube rolling. Hell, it even takes LISST!

The point is: these are huge changes, not a quick feature-add or incremental update. Expect to see this trend set the tone for future next-gen products—or perhaps the next-gen products will be even more radically changed.

Bottom line, on the analog side, we're not standing still, either.

Right and Wrong in 2015

2015 was boring. And boring years are very, very good years. Excitement, in manufacturing, is rarely excitement in the positive sense. Excitement usually means things are breaking, stuff is late, someone messed up, and you are boned.

Boring is what happens when most things go right.

So let's start with right:

In-stock intros. A couple of years ago, you would have fallen out of your chair if we introduced a new product and it was orderable. In stock? Inconceivable! And yet, this year, 5/6 of the products we introduced were in stock when we announced them. As in, we announce, you order, you get shipping notice 15 minutes later. You know, like a real company. Maybe even better than a real company (I'm looking at you, Apple … come on, "Available November?" Sheesh.

> **Aside:** and yes, I know, we were 15 000 years late on Yggdrasil, but hey, the 5 products that followed it were all in stock on intro. And not for microseconds.

Significantly less problems with new products. No, we were not perfect, but there were no embarrassing "total recall" moments like we had with Asgard 2 a few years ago and Fulla last year. Other than a few transformers coming off boards on early Yggdrasils (which could not have been foreseen, since they were the same size and mounting as we've used in literally tens of thousands of other products—and quickly addressed

in production, and everyone affected getting free return/exchanges) and a very bizarre glitch with the Bifrost Multibit in a small percentage of systems (so small, we had to bring back some units because we couldn't replicate the problem—and we cannot get the returns to replicate it all the time, meaning it's system-dependent), the products are solid.

Better stocking. Although we're still not perfect, we're much better at keeping products in stock. Yes, we were caught out a bit by Mani demand, and we've had some struggles keeping up with the multibit DACs, but long, grueling out-of-stock statuses are significantly reduced.

Lower cost parts. Thanks to Alex bringing on new distribution partners—ironically enough, found when we were breaking the world supply chain for AD5791s—we've been able to reduce costs, while using the same high-quality parts.

Higher quality parts. A new metal vendor, and a couple of intense meetings with the current one, and our chassis are now looking better than ever—even the inexpensive chassis like Vali 2.

Increased operational efficiency. A new hire, Tyler, has been putting into place the systems we need to operate as a real—er, I mean, *much larger,*

company. At the same time, Alex has been pulling out the stops getting everyone cross-trained and finding what they are best at doing. As a result, we cruised through this holiday season without having to add staff.

What about wrong?

Sure, we do lots of wrong things. Every company does. This year, it wasn't too bad, though. Most of it was functional stuff:

Overbuying. We went a little crazy at the end of summer when it came time to get stocked up for the end of the year. As a result, our overall parts cost went through the roof. It looks like it's a gamble that has largely paid off, this holiday season, but it just as easily could have been very bad. As an entirely self-funded company with no bank lines or receivables financing, we are extraordinarily free of pressure to perform at some arbitrary metric, but we also need to keep an eye on cash flow.

Bad inventory visibility. With so many parts spread between us and our production partners (the assembly house in Simi Valley and vendors in the San Fernando Valley), keeping up with *what we got* and *what we need* started outstripping the ad-hoc systems we were using. It also contributed

to the overbuying spree. As a result, we've been putting in a real inventory management system, and recently started using it. Yeah, I know, basic stuff.

More metal problems. What's a year without screwy metal? Well, so far, we wouldn't know, because every year has had screwy metal. This time, it was a new vendor literally making thousands of bad parts due to no oversight (at their firm). They got to eat it … but we had to rush through several replacement runs. The good news is that we are working much more closely with the vendors now, and I believe they understand, more than ever, what we need.

But this is nit-picking. It was a great, great year on virtually all fronts.

What About 2016?

Yeah, yeah. I know, 2015 is old news. So what will you see from us in 2016 … that's what you really want to know, right?

And that's tough. I can't exactly outline our product plan in detail. Because, well, you know, that's talking before you actually have something to sell, which we're trying very hard not to do.

And, frankly, because we're not 100% sure of what we'll introduce.

So let's go broad-brush, and let me see what I can talk to you about.

Advertising. Yeah, big deal. But it's a fact. You will be seeing more of us—in Stereophile, at Computer Audiophile, and maybe a few other places. Not a lot, nothing over the top, but it's time for the 2-channel world to know a bit more about the Multibit DACS. Also because …

2-channel. I'm certain you'll also see the first dedicated 2-channel products from us in 2016. As in preamps and amps. How many products, and when remain to be seen. But you will see some things from us. I hope. And they'll be very, very cool. I think. At least in part because …

New topology. It's not often you see radically different stuff on the analog side of things, but I think we're ready to unveil something I've been working on for a couple of years … a new, inherently balanced, current-feedback topology that is not a circlotron, and is not supersymmetry. This is by far the best-performing topology we have ever come up with, both in terms of measurements and (in our opinion) sonics. This topology can stretch

from the 2-channel world to the headphone world pretty easily.

Manhattan Project. I've said all I can about that, but I do believe it's a 2016 product. However, it might not be.

Another secret. We have another shocker for 2016, which I can say even less about … but I can confirm it is NOT digital.

Improved distribution. We're hoping to extend our "direct partner" model—as with Electromod—next year. Also, I hope to have some 100 V product for Japan, as we have some resources to work natively there. These are grand experiments … we'll see how they do.

More book. I'll keep writing, if you'll keep reading. I may be doing some fiction again as well.

Portable? Still not real excited on this one. I know it's in vogue, and a lot of people want them, but to do it right, it will be huge, hot, and heavy. Wow, that sounded weird. I'm still trying to convince myself we can do something different and meaningful.

And, of course, we reserve the right to come up with a surprise or two … like we did with Vali 2 at

the end of the year. And, after all, the Manhattan Project wasn't even a glimmer in Mike's eye at this time last year—it was totally out of left field this year.

Thank you again for all the support, comments, and suggestions! We're listening … and we'll continue to try making the most fun, high-value stuff out there!

2016

2016, Chapter 1

A Longer Discourse On Marketing

Okay, first let's clarify: this *isn't* the marketing book I promised. This is marketing as a primer—how to set it up for success, what's an agency and when to use one, blah blah—and marketing as done by Schiit—what's right, what's wrong, where we're going, etc. The hope is that this chapter can do just two things:

1. Illuminate the sometimes-murky waters of marketing for people who are interested in the, ahem, "art."
2. Provide some guidance to companies looking to market in the audio neck of the woods.

And, well, yeah, some of this content may appear in the marketing book, but this is more laser-focused, and more about us. You see, I do listen to you …

So, Why Do You Need Those Marketing Pukes At All?

This is what I'm asked the most, at least by people who aren't in marketing. (Unsurprisingly, marketing people think they are absolutely indispensible, but, ah, we'll get to that …) I think what they're really asking is this: "After we've seen some of the highest-paid people in the company cavort with trendy agency types, lounge in meetings over lattes, travel to exotic locations for trade shows, hang out with cool directors and photographers, and maybe have an affair with the actors and models, how the hell do they get off with the easiest and coolest job, when I have an ******* boss breathing down by neck for impossible coding goals with yesterday deadlines?"

Weeeeellllll … where to start?

First, marketing really isn't as cool as it looks, whether you're talking Mad Men or real life. Marketing has its own pressures, not least of which is being blamed for crap sales when the product is really the crapifier—and by "blamed," I mean "fired."

Second, yeah, I know, your salespeople already know all their prospects, they don't need market-

ing's help with that, marketing gets to play while they do all the tough work.

Third, yeah, I know, your engineers are second to none and the products should sell themselves, just look at the 167 new features they added to your new thermometer (and in their list, only #74 and #151 are differentiators from the competition, but I digess.)

And this shouldn't be an apologia for marketing. Let's get down to the core question, which I believe) is:

What does marketing *really* do for a company?

Okay. Deep breath. Done right, marketing *can* mean more new sales, more repeat sales, happier customers, and faster growth.

Note the "done right." Done wrong, marketing can piss people off, help kill crappy products faster, lose any sense of purpose or measurability and just burn money, or even take the company to the edge of bankruptcy.

So how do you do it right? I'll get to that.

For now, let's talk about the necessity of marketing. Marketing, in its most basic form, is about informing potential customers about your product or service. You need a new car cheap, cool, we make the most reliable and inexpensive cars. You need a new smartphone, great, we have one that runs on atomic power and never needs charged. You need something to clean your windows, we have a robot servant that also makes sure your kids are safe.

Okay, so I'm getting a little wacky there. But you get where I'm going: *all* companies need some form of marketing.

With no marketing, nobody will know about your product or service, and by definition, you'll have zero sales. This is not the recipe for a successful business venture.

Now, the arms-crossed crowd is getting grumpy. "Ah, come on, I know plenty of "best kept secret" products that really sell themselves and don't need marketing."

Oh really?

So you know about a product that comes in a white box with no text, nothing at all to identify

it, no way for someone to know what it is or where to buy it … and it sells?

So, when you go to the store, all those shiny packages there yelling about how their brand of overly-processed GMO corn+flavor enhancers is better than any other—those are not marketing?

And so, when you drive by a roadside stand, that hand-lettered sign saying, "Straberry's $ 3" is not marketing?

Aaaaand so, when you read an article in the Wall Street Journal about new trends in smart thermostats that lists products by name and URL, that's not marketing?

Marketing is far more than annoying ads. It's can be as innocuous as identifying what you do … and it can also be the most brazen, blatant, and irritating thing out there.

Where marketing gets its evil reputation from is when companies go too far. A little packaging and PR? Sure. Running some text ads on Google Adwords? Sure. Taking out a spread ad in a magazine? Sure. Blaring a constant barrage of the same commercial every 5 minutes over every radio and TV station in sight? Annoying.

Spending $100 million to have your name supplant the beloved name of a major stadium? Boo. Laser-etching your logo on the earth-facing side of the moon? Truly nasty, that.

Again, I'm getting silly. But you get the point. Companies need marketing, if for no other reason than to let people know what they do.

More sophisticated marketing builds on that, to target the kind of people who are looking for what you sell … and to tell them why it's better than other products. And if your marketing pukes are good, that's what they do—and that's why you need them.

When To Market, and When To Shut Up and Fix Things

Let's start with something very simple. You should NOT be marketing unless you have three things:

1. A clear definition of what your product or service is.
2. An identified audience who is, or should be, interested in your product or service.
3. At least one unique advantage your product or service has over the competition.

"But—" you begin.

NO. There are no "buts" here. If you don't know what you're making, who wants it, and why it's better, you're done. Gone. Outta here. Don't even bother marketing, because it will kill your product dead even faster.

Okay, I'll be nice and allow three questions.

"But my product or service is entirely new, I don't know how to explain what it is," You ask. Well, then you should figure out how to communicate what it is and why someone would want it. Because if you can't do that, how do you expect to ever sell anything?

"But my audience is everyone," you say. Or, I don't know who my audience is." Two comments: your audience is *in no way* everybody. Unless you have ten-figure marketing budgets, in which case you can claim, "Well, probably pretty much everyone who has enough money to buy one of these after seeing our ad." And, if you don't know who your audience is, where they hang out, and what they like, figure it out now—before ever embarking on a single ad. No … before doing a *business card.*

"But my product looks like everyone else's, it's not really unique," you say. Well, you better find some way it's unique, or you're going to be in a world of hurt. Is it less expensive? More functional? Prettier? Made from diamond-hard alloy reverse-engineered from UFOs? Designed by hermit monks living on a 58 square meter island 1000 nautical miles away from any major landmass? There are always differentiators. Just make sure they really are different, and *that they matter in a good way.* The fact you have a high-res display ain't gonna sell your smartphone, and the fact that your smartphone display is 8.7 inch diagonal may be different, but it is NOT good.

So, do you have those three things? Then you can start marketing—or, in other words, creating the words, visuals, and outreach plan to communicate what you do and why it's better to the target audience.

"No, wait a sec, I have those three things, but I don't know how to create the words, visuals, or the plan," you say.

And I grin. Because that's where marketing can—and should—earn its keep: helping companies with a clear vision of what they do and why it

matters connect with people who just may care enough to spend money on them.

And that's where the "marketing village" comes in. Because many companies don't have all the talent necessary to go from that three-item list above to an effective, measurable marketing campaign.

So let's talk about the players.

All the Pukes in a Room: The Marketing Village

Before we go any farther, let's talk about the kinds of pukes—er, I mean, *people*—we have in marketing … and who they work for, because lots of them may not be company employees.

First, the organizations. In addition to the company marketing their own products (let's call them the Arglebargle Widget Co, or just Arglebargle, for convenience), there can be many other organizations involved in their marketing:

- **Marketing Agency.** This is a "catchall" kind of marketing company, one that does many different things to help a client like Arglebargle get the word out, and drive sales. They may do websites, ads, PR, branding, direct, and collateral design under one roof. A marketing

agency is a good choice for a smaller company with limited budgets, because they're usually more focused on results.

- **Advertising Agency.** You'd go to them for the "serious" ads—television ads, big glossy print and outdoor work, etc. This is kinda a throwback in today's world—an agency that only produces, and places, advertising. Sure, there are some agencies out there that strictly do ads, but the line has blurred between ads and experiences—especially online—so expect an ad agency to do websites, branding, and other things these days.
- **Design Agency.** These agencies focus on the visual side of things. You'd go to them for packaging, collateral, business systems (fancy-speak for cards/letterhead/envelopes), display graphics, presentation design, info graphics, etc. Sometimes put down as the "pretty picture guys." Many are not very strategic, but this isn't always the case.
- **PR Agency.** These are the guys who are connected with the press—they have the contacts with editors that help them get your news in print (or online) and your products reviewed. Of course, PR now has a strong online and social component, so the lines get blurred with what's next.

- **Social Media Agency.** This is an agency that creates, manages, and sometimes provides strategy for your Facebook, Twitter, Snapchat, Pinterest, ad nauseum, if you are silly enough to waste a lot of time on this. They are usually pretty useless for micro-social strategy and implementation, like activity on Head-Fi, for example.
- **Web/Mobile Agency.** These are the guys who build your website or create your apps. Of course, your marketing agency, ad agency, or design agency may also be able to do that too.
- **Branding Agency.** These guys are, at least theoretically, the strategists. The ones who figure out what you should say to differentiate yourself from the competition. The ones who come up with powerful, emotional ways to say it. The ones who wrap it all in a compelling visual experience that sets you apart. And sometimes it works that way. Sometimes it goes off the rails, and you get delivered a brand that doesn't look or sound anything like what you are. This is the component that most companies should keep in-house as much as possible … because nobody knows who you really are as well as, well, you.

Then, the contractors. Because sometimes you don't need all the agency baggage to get some-

thing done, especially if you have the core of your branding in-house. And believe me, there's tons of agency baggage. But I'll get to that.

- **Video Producer.** Because you need to have some video shot, edited, and a voiceover and effects added. Whether it's an educational video or for a commercial, there are plenty of producers who are available to help you out.
- **Graphic Designer.** Like a design agency, but cheaper and with less pretention and infighting. Providing you can do the strategy and words, much less expensive than the full agency treatment.
- **Photographer.** Because you need pictures. Good pictures. And believe me, you don't want to do them yourself. That's its own rabbit hole.
- **Illustrator.** Because you need something illustrated—a product diagram, an infographic, a cartoon. And it's not like you usually have an illustrator sitting around the office.
- **Copywriter.** Because you need good copy. And writing is hard.
- **Various contractors.** Because you can get people who will write PR, help you with shows, put content on your website, etc, etc—all the various little marketing tasks.

So now you're probably sweating a bit, thinking, "Holy crap, I have to hire all of these?" and "Why doesn't any large company simply do this internally?"

Well, to be short, "no" and "holy crap," respectively.

The longer answer as to how many of these dudes (and dudettes) you have to hire is, "Depending on how much you can do internally, and how great your needs are, you may not need to hire any of them at all."

However, the likelihood that Arglebargle Widget Co, a small start-up, has the staff or time to take everything on, well, that's pretty low. So you may end up hiring a graphic designer to execute an ad you came up with, or even a contractor to help set up a trade show.

And, assuming that Arglebargle knows its product, audience, and differentiators, knows where they hang out, and can put together copy for an ad or brochure, that may be all they need.

Aside: the more you can do in-house (without going broke or overburdening your staff), the better. Yes, I know, the agencies will howl, but there is nobody who knows Arglebargle better

than Arglebargle, especially when the company is just starting up. A key component to Schiit's ultra-efficient marketing is simply the fact that I do all of the strategy, copy, and visuals, do some of the press interface, and we use a contractor to make sure I don't screw up the trade shows. That's it.

And Schiit isn't the only company that acts this way. One of Centric's former clients had a CEO who was intimately involved with their marketing. Any marketing item that could affect revenue—from a trade show graphic to an ad or the website—was under his direct review, and he gave specific input. Nobody doubted who was running marketing, even though he had an in-house staff to manage much of the day-to-day aspects of it, and an agency to advise and execute. The success of this approach was evidenced by the company growth—from $80 million to $450 million while we were working with them.

Now, as far as the question about why a big company doesn't just hire everyone they need in-house, well, that's both simple and complex. The short, verbal form would be, "Well, because they're expensive … and they're also fickle."

Let me explain. Hiring a full in-house art department, brand management department, copywriting, video production, ad placement, PR, etc, etc would be extremely expensive. And, unless the company was exceedingly large and in need of literally boatloads of content (think Target—they use an in-house agency), those people would usually spend most of their time sitting around.

And, at the same time, creative talent is unique, in that many of them want the prestige of working for an agency (usually thought to be better than working for an in-house department), or they would rather freelance … and they tend to move around pretty often, to where job opportunities seem better, so they may not be with the company long, even if they take the job. This, by the way, applies on both the "art side" and the "programming side." Both should be considered creative.

Because of the cost and uncertainty, then, most companies like Arglebargle prefer to contract with the agencies and freelancers on an as-needed basis, to produce the specific marketing items they need.

Which leads to an interesting dichotomy: Arglebargle and its agencies or contractors have

different goals.

For Arglebargle, it's simple:

- They want to have a project done
- They want something to be effective
- They want it to fit within a specified budget
- They want to have it work with what they've done before
- They want it on time

For an agency, it's also simple, but different:

- They want ongoing work, not projects (think retainers)
- They want to do something cool that wins awards
- They want more budget than the cheap-ass client provides
- They want to radically redefine your direction
- They want the time to do it right, even if that doesn't fit your schedule

And yes, I know, not all agencies are like this, but the fact is: most agencies don't want project work, most agencies are looking to do cool stuff, and most agencies would rather radically redefine your look and message, rather than build on what came before.

So be careful … and know that you may be working at cross-purposes. While what the agency

proposes may seem cool, make sure you're not funding a Shiny Marketing Object—a "cool" new tactic with no guarantee of success.

Schiit's Marketing and Herding Cats

I mentioned above that a lot of Schiit's "secret sauce" is simple: I do most of the marketing from start to finish. What I didn't get into, though, is why that's good, what we're doing, what happens when I can't do it, or what we're doing this year.

So let's talk about that, and then let's talk about a general prescription for doing marketing right (or at least right-er—herding all the marketing pukes … er, I mean cats, in the same direction.)

First, let's talk about what Schiit has done, and what we're doing this year.

In the past, we've done a website, a few online and print ads, and trade shows. Oh, and we post what we're doing on Facebook, for some reason.

This year, we're adding a few print and online ads, and a show or two.

"Whoa, big change, huh," you might say. "Don't go wild or nothin."

To which we say: *Exactly.*

Getting distracted by Shiny Marketing Objects can quickly eat your budget. Sure, we could do an all-out social program, or create an augmented reality app, or produce videos, or do lenticular animated tip-ins, or heat-activated urinal billboards … but these are highly speculative, expensive tactics that may or may not work.

See, the thing is, success in marketing is usually pretty boring. In fact, it can be summed up in just two principles:

1. Stick to the stuff that's measurable.
2. Make sure it pays off before doing more of it.

This means that our primary areas of marketing are:

1. **The website.** It's a direct-sale site, so it's where the money comes in. This is the most important thing for us, with a bullet.
2. **Micro-social.** Like this book. We're present where our customers are. This is very important.
3. **Online ads.** Because they are 100 % measurable and trackable down to purchase.

So why are we engaging in trade shows and print ads? Aren't they less measurable? Yes, they are. But with the success of #1 to 3 above, we can now speculate a little. And by "a little," I mean, A *real* little."

- **Shows** are where you are going to see the press, and get coverage on new (or existing) products. They're important.
- **Print ads,** in the past, we've done to support the shows. Now, we're venturing out into Stereophile very deliberately—to let them know about our multibit DACs, and to prime the audience for upcoming 2-channel products. In this case, we're building a foundation.

Note that nowhere in the above are there any highly speculative or expensive tactics. This is based on my experience of 20+ years in marketing. While cool, they are more risky. So they should be left until you're literally bursting at the seams with extra money. Until then, focus on your product, and do the simple marketing stuff.

You'll also note that nowhere in the above are any focus groups or research. We're close enough to the market that we don't believe we need it. We also have enough confidence in what we're doing that we don't believe we need to second-guess, delay, or get a different opinion on what's coming. We could be wrong. We'll see.

So what does this all mean to a company wanting to do marketing right? Well, much of the above. Let's sum it up:

1. **Start by defining your "first 3."** Don't move an inch until you have this down. You can go crazy and flesh this out into a full message platform, but that's a bit crazy, especially for a small or start-up company … and beyond the scope of this chapter, anyway.
   a) A clear definition of what your product or service is.
   b) An identified audience who is, or should be, interested in your product or service.
   c) At least one unique advantage your product or service has over the competition.
2. **Determine with what you can do yourself.** Can you do copy? Layout? Video? Programming? None of the above? All are fine answers. Just be honest.
3. **Determine what conservative, measurable tactics will be most critical to your success.** If you're selling affordable products direct, a great e-commerce website shoots to #1 immediately. If you're selling a very sophisticated and expensive service, your sales presentation and collateral may be as important as your site. General rules include:
   a) E-commerce is always #1 if you're selling direct
   b) Online ads are easy to measure and tweak, offline not so much.

c) Mass social is almost always a loser, first-thing; however, engaging in micro-social communities that are predisposed to be interested in your product can be extremely important—and it's measurable
d) Don't discount "old" tactics like direct mail—it still works, especially high-end direct mail—and it's easy to measure
e) Print ads may be worth doing, especially in niche industries
f) Don't get wrapped around an axle on "brand look" or how something "feels,"—these are just excuses not to do something

4. **Hire in-house, contractors, or agencies as needed to help you execute these tactics if you can't.** Just know that their goals may not be the same as yours:
   a) Always have a firm grasp on the "first 3." Go back to that as a reference when they present something to you.
   b) Reject anything that doesn't support your "first 3."
   c) Pay attention to fine print on contracts—number of revisions, who owns the art-work, etc.
5. **Measure and refine before you do anything else.** Okay, so you've gotten a site in place and you're running online ads. Great. That

doesn't mean it's time to go out and embark on your augmented reality app. Instead:

a) Make sure Google Analytics is installed on your site
b) Make sure it's tracking sales, or, if you're a leads-based business, that you've set up goals to track
c) See who's coming from what referral sites, keywords, and ads
d) Do more of what works, and less of what doesn't

6. **Go back to #3.** And do the next thing on your "conservative, measurable tactics" list.

And yeah, I know, Google Analytics is a whole marketing course in it self, as is Adwords (online/mobile search ads). That's beyond the current scope of this chapter. However, if you don't have the time or inclination to learn about these things, know that you can always hire a contractor or an agency to help you—and DO IT. *Listen to them. Learn what works.*

Because marketing that doesn't work is nothing more than ego—and can be, eventually, the downfall of the company.

2016, Chapter 2

What To Do While Waiting for Ragnarok (or, *Alternate Strategies for Creating Your Own Well-Capitalized Company*)

One of the things I'm asked the most is, "How the hell do I start a company like Schiit, if I'm not able to put a whole bunch of money into it?"

This question is usually accompanied by a pained, bewildered expression—a pure distillation of all the obstacles, slights, screw-ups, and stacked decks that keep everyone from simply telling their boss to get a bag of sand and a little hammer, they ain't comin in the next day, or anymore at all.

And I understand that expression. That was me at 19, still in college but dreaming about starting a multi-million dollar speaker company while making cabinets in my parents' living room.

That was me at 23, out of college and immensely broke, plowing every paycheck into Odeon, the failing speaker company that was having trouble breaking into 6 figures in sales, let alone 9. That was me at 26, speaker company dead, when I was working for Mike doing Cobalt, and smugly sworn off ever trying to start another company. And that was me at 28, thrown out of the warm embrace of Theta and Sumo into the harsh reality of having to get a job, or start another company, or *something … anything …* to keep from being broke in very short order.

And, when I was honest with myself, it came down to this: I always wanted to have a company like Schiit … but the investment necessary, and the obstacles involved, simply seemed too large.

And yes, sure, you can read rah-rah stories about entrepreneurs who found venture capital or angel funding or whatever, but the reality was that I didn't have the network or the connections to even start thinking about that. The pre-Centric days were also the pre-internet days, and the chance of convincing a VC to put a ton of money into an audio venture (low return) helmed by a very young (and even younger-looking) kid with one dead company under his belt was, well, notgonna-happen.com. It took another half-decade or so

before they started throwing stupid money at web development firms headed by twentysomethings.

All of this is a long-winded way of answering the above question in the shortest and most brutal way:

"Many times you can't."

But What Can You Do?

In the negative above, note the positive outcome. At 28, I wasn't able to start a company like Schiit with my own money, nor was I able to find capital to make it happen … *but I did start a company.*

And that's something to keep in mind. It's important to think about what kind of company it was, too. But I'm getting ahead of myself.

First, let's tear down a bunch of the stupid crap you may have heard about starting a company:

1. **You can get money from investors.** Ahhahahhahaahahaa! No. Unless you have an in with a VC, plus an experienced board, a roster of impressive advisors, and you hit the exact right time with the exact right concept that's the hotness du jour, no. Forget it. Put it away. Do not pass go, do not collect $200 million.

2. **You can get money from banks and creditors.** Not really. Not in the sense of them bankrolling a pure start-up from an average person without a trust fund or demonstrated cash flow. Once you're running, they can provide enough rope to hang yourself with … er, I mean, things like receivables financing and other financial equivalents of push-up bras. Enter at your own risk, and beware of addiction.
3. **You can get money from credit cards and such.** Partially true, but how many cards can you really get, and at what limits? What interest rates? Yes, this can happen, but it can be a very deep hole, and one that is very difficult to climb out of. I financed some of the fast-growth part of Centric on credit cards, and it took many years to dig out from it.
4. **You can get money from crowdfunding.** Ah, yes, the panacea of the twenty-tens. Now not smelling so good. Some high-profile failures and endless delays will do that. Also, with the crowdfunding environment more competitive than ever, expect to drop some fairly serious money producing the marketing for it—and promoting it. And remember, that money comes with gotchas—big gotchas in the form of thousands of backers, all baying for their

product, right now.

And also let's review some of the true stuff:

1. **It's not for everyone.** Running a company means there's no place to hide. You can't say, "Well, that's the strategic directive for our division," or "I told Bob to take care of it, I don't know why it's not done," or even "Joe, can you copy these things for me?" Everything rolls uphill to you. Not comfortable making decisions? Not good at handling stress? Maybe having your own company isn't the best idea.
2. **It's difficult.** Yes, it is. Not only are there tons of complexities and gotchas in whatever your company does, there's also the miasma of "stuff" that surrounds a firm. Business regulations. State and federal statutes. Quarterly taxes. Payroll and HR stuff, when you get there. Accounts payable. Bookkeeping. The list is very, very long. Are you ready for that?
3. **It can eat your life.** Want to spend endless time on the beach, enjoying every day as your children grow up at your side? Weeeellllll … that probably isn't gonna work, at least not during the start-up phase. And that "start-up phase" can easily stretch from 2 years, to 5, to a decade. It was a full 10 years after I left

college, and 4 years after I started Centric, before I took my first one-week vacation.

"Okay, fine," you're saying. "I get it. And I'm cool with all that. And I still don't see how you start something like Schiit without being made of money."

Well, in actuality, Schiit didn't consume that much cash at all. As I noted in an early chapter, Schiit started, all-in, at about $10 000. The first run paid for the second runs, which paid for the next, and so-on. And it has continued to this day, entirely self-funded, with zero bank loans, zero investment, and zero crowdfunding.

The sting? No salaries for 2 years. Yes, we paid Eddie and Tony, but for Mike and I, nothing.

Which brings me to the first "thing to do while waiting for Ragnarok:" *Trade time for money.*

But let's go into more detail about that (and other) strategies.

Strategy the First: Trade Time for Money

Yep. You can do exactly what we did. Invest lots of time, not money. It works like this:

1. **Identify a compelling product you can feasibly build with a small investment.** For us, this was Asgard. Yes, we announced both Asgard and Valhalla when we started, but it was vapor. Well, almost. We had a working board.
2. **Determine if there really seems to be a demand for this product.** Yes, I know, this is the most hang-your-butt-out part of it. I had cold sweats about whether or not the Asgard would actually go … but based on the market at the time, it certainly seemed like it would.
3. **Make sure you can talk to people who might buy the product without going broke on marketing.** If the plan involves huge advertising investment, that's big money—which negates the advantage of trading time for money.
4. **Build the product.** Note this isn't a slam-dunk. But without a product, you don't have a business—at least not in this scenario. See below.
5. **Sell it.** It has to sell.
6. **Reinvest the money to make more.** More of the first product, or additional products.
7. **Repeat 4 to 7 until you can pay yourself.** Note that you may be hiring people—no, wait, you almost certainly *will be* hiring people—

before you pay yourself. Don't get a chip. Don't huff and pout. This is the way it works. The people who work for you benefit first. If you can't defer satisfaction, you probably shouldn't get into business.

8. **Repeat 8 until you're bored.** Then sell the company and go play golf or make spaceships or something. Or not. You know what I mean. I hope.

Pitfalls of this approach? Sure, there are plenty.

First and foremost, I'm sure you may be thinking, "I can't go two years without a salary, no way, no how."

But how true is this?

What if you started the company in your spare time, while you maintained your regular job? That's what I did when I started Schiit. I stayed at Centric while Schiit was ramping up.

Yes, I know, that's a ton of time. During the Centric/Schiit start-up phase, I was easily working 80 hours a week. Work eats your life. You won't see your kids. That's a pitfall—but it's also an inherent part of trading time for money. How bad do you want it? That'll always be a key question, no matter how you end up going about starting your own business.

Other pitfalls include:

1. Your first product doesn't take off at all.

At that point, the whole scheme collapses. This is what happened with Odeon loudspeakers. Our speakers didn't take off. So we spent years tweaking the line, coming up with more exotic designs, and trying to figure out ways to get it to run. But it didn't. Chalk this up to a product that probably didn't sound as good as it should, looked strange, and didn't have the advertising support that the dealer base needed in order to take us seriously. Remember, direct sale wasn't an option back then.

So what do you do if your first product doesn't take off? Remember my joke about having "Christmas presents until the end of time" if Asgard didn't sell? Well, that's one option.

Another is to come up with another product and try again. It really depends on your appetite for iteration, your budget, and your stamina.

But make no mistake: success from the start isn't a given. Be realistic. Have a Plan B. And know when to cut your losses.

2. The product doesn't take off as fast as you hoped.

This is bad, too … but maybe not as bad as you think. I'd intended Schiit to be a "hobby business," something that we could run out of the garage, or out of a 400 square foot shed in the backyard, or something like that. If Schiit had been slow-growth, that's where it probably would have ended up.

The problem, of course, is that with a slow-growth business, you're going to be a lot harder-pressed to start hiring the people you need to get you working sane hours. So you may end up looking at many, many years of having a main job and a hobby business … never seeing your wife or kids … working 16/7 and hating life.

If you end up with a slow-growth business, you really have three options:

a) If you can make enough money at the hobby business to live on, and you really love it, quit the day job. That'll preserve your sanity and let you focus more on the biz—which may make it grow faster.
b) You can also decide to go ultra-high-end and bespoke—that is, into the semicustom realm, where sales prices are higher. This may get

you to a point where you can lose the day job. Of course, this also presupposes that you have a compelling product worthy of ultra-high-end, bespoke status.

c) Finally, you can throw in the towel. Maybe it's simply not worth it. Maybe you don't have the right product for the time. There's no shame in that. I've done it several times. Schiit and Centric are only my successful businesses. There have been others.

3. You fall prey to the unexpected.

A new competitor comes out of nowhere, with a more compelling product at a lower cost, killing your sales momentum. Your first product proves to be spectacularly unreliable, causing havoc among owners and diverting all your time to service. The market shifts to a new connector/standard/software and your product becomes obsolete.

In short, you never know what you're going to happen, and the cards may come down against you. There's nothing certain about starting a business, and there's no "sure bet" product—not even after market research, focus groups, persona analysis, scattering goat entrails, positive

thinking, EST, praying to your copy of How To Win Friends and Influence People, going with your gut, etc, etc …

So, in this case, what are your choices? Go back to #2 and decide to buckle down, step sideways, or cut your losses.

Yeah. I know. Even trading time for money is risky. You can chew up years of your life and end up with nothing. Or you can win big. However, in either case, you're not jeopardizing your retirement. That's the benefit of trading time for money.

Bottom line, the odds are better than Vegas or the lottery. What are you gonna do?

Strategy the Second: Start Low

"So what happens if I don't have ten grand laying around to start a business, and can't wait two years to start taking a small salary?" you ask.

Well, you can simply wait. Or you can *start low.*

What is "starting low?" It's choosing, deliberately, to start a less capital-intensive business. When Theta was imploding and Sumo folding up, I

didn't have the money to put into product development, nor the time to wait on a salary (nor was it really viable in a pre-internet direct sales era.)

So what did I do? I started a marketing agency. (AKA "design agency, advertising agency, interactive agency, or whatever-marketing-y-stuff-we-could-do-for-money agency).

An agency is much like a consulting company. It doesn't need a ton of money to get started. You don't need to invest in a first run of products. You don't need huge investments in capital equipment. You don't carry inventory. Your margin for stuff you do in-house is 100 %. It's part of the wonderful "service economy" that was going to save America's GNP, back a decade or so ago.

Aside: What's truly hilarious is how many manufacturing companies tried to run themselves like a service business in the wake of offshoring. They weren't actually building anything, so products became abstract, an entry on the balance sheet. Their only job was to market them … and rake in the money. Or that was the theory, anyway. It has worked for some … and not so hot for others.

Sounds wonderful, right? Why would anyone in their right minds go into business making a product when something like that was available?

Lower start-up costs, lower risk, higher margin … hell, we should all be doing it!

Four words: *low barrier to entry.*

When you're getting into a service business, you're now up against anyone who can print a business card, throw up a $500 Wordpress site, and claim some expertise in your chosen business.

Service businesses are easy to set up, so there are a ton of them. Go into most companies with the pitch of, "We're a design agency," and watch them roll their eyes. They get pitched every day. They know that there will always be agencies lined up behind their current one ready to take their business.

But … in the case of Centric, we had a unique pitch. We only went after technology companies. And we led with, "I'm sure you have plenty of agencies calling on you, but I bet you also don't have any that are headed by an engineer, who can really understand what you're doing, who can talk to your tech team and translate their blatherings into real benefits, etc … "

We also went directly after consumer electronics, because, like, well duh. My background at Sumo and Theta gave us a real edge.

Because of this unique approach, we got a lot of business. And Centric did very well. Four months after we started up, we booked 10× my highest previous salary of business in a month.

Sounds great, right? But even with a unique pitch that gets you above the noise, starting a low-capitalization service business has plenty of pitfalls. Like:

1. **Big customer bosses.** When you're in the service business, your life really isn't your own. It's a lot like working for a big company, with a big boss who runs everything.

Except, in this case, it's usually a lot of bosses. Because a successful service business will have 10+ active accounts. And each of those accounts will have a boss who controls your interaction with them.

Sometimes this is great. Sometimes your bosses are sane. Or at least benevolent. Sometimes they will even fight for your right to charge them more, if the big boss on top of them changes your designs at the last moment. Sometimes, they can seem more like a friend than a boss, and the relationship can last for years.

Sometimes it's not so great, though. Sometimes they'll stomp in and demand free changes … on a website that's 95% done. Sometimes they'll blame you for an ad that flopped … after they rejected your ideas and dictated their own. Sometimes they'll have you run another round of ideas because their gardener/wife/dog didn't like it, after you sweated blood for 18-hour days to deliver the first ideas.

Bottom line, if you really want to be your own boss, service isn't where it's at. Remember, it's called "service."

2. Big customer losses. It's always a scary day when your biggest customer gets acquired. Because you have no idea if they'll be working with you anymore.

No matter the platitudes, the acquiring company has their own ideas … and you may suddenly see a big piece of your business simply disappear. If you have a dozen accounts and everything is spread around equally, this may not be a big deal. If that one client was 60% of your business (not uncommon), you may be, well, done.

3. The dreaded "wave." When you're not working on your business, you're working in your business. It's natural for many service companies

to scramble for sales when times are thin, and sit back when times are fat. This leads to the dreaded "wave" of feast and famine.

It took Centric over a decade to figure out how to get off of this wave. Before, we'd have revenue swings of 20 % to 30 % from year to year—both up and down. When you're talking about maintaining an in-house staff, losing 30 % of your revenue from one year to the next is not a great way to do it.

Bottom line: in a service business? Always be selling. Always. No matter how bright the future looks.

But for me, the biggest pitfall of Centric is something I really can't put into a generic list. Because the biggest pitfall was that Centric kept me from doing something in audio for a decade and a half.

When Centric was super-successful, it was easy to scoff at audio as something that I flirted with in the past, but really wasn't applicable today. But with every "big boss" or "big loss," I looked wistfully at audio, and wondered if I couldn't do something with it again. And it took until 2009 to act.

Strategy the Third: Go Oblique

"Okay, that's cool, but I don't really want to do anything but audio—no agencies, no service businesses, etc—so what do I do if I don't have a ton of money and time to come up with a product.?"

In this case, you can *go oblique.*

This uses the same idea as starting low—low capitalization—but in a directly related field. Going oblique lets you start in audio, and stay in audio, until you have enough capital to create crazy products like Yggdrasil.

Huh?

Here are some examples of oblique businesses:

- Headphone cables (or other cables)
- Audio accessories and tweaks
- Audio retail

"Wait, wait, wait!" you might be saying. "Are you saying to start a cable company? Aren't there already plenty of them?"

Well, yes and no. Here's the thing. Headphone cables—and audio accessories and tweaks—are relatively low-capitalization, low barrier to entry types of businesses. You can start very small,

much smaller than a full-blown electronics company that needs to worry about hundreds or thousands of parts and safety/emissions compliance.

And for audio retail, it's possible to get terms (usually), which means that capitalization requirements are much less. And even if you can't get terms, it's easier to get a bank interested in *Sonny's Sony Shack* than *Sonny's Maywerk Electronix*.

Let's look a bit harder at these options:

Headphone cables, general cables, audio accessories and tweaks. Yes, there are a ton of cable and accessory companies out there. Some of them sell exotic stuff at eyewatering prices, and some are affordable.

But no matter how many companies there are out there selling cables and accessories, I bet there's still room in the market for a savvy competitor. Where's the well-known option for affordable headphone cables made in the USA? Yes, I know there are a ton of them, but these are usually smaller companies without a lot of brand recognition. Where are specialty ultra-shielded products for phono use that are easily accessible and don't cost a bazillion dollars? Where are the sensible

tweaks that have real, measurable results and don't need a second mortgage? For the right individual, starting a strongly branded company with great prices in those realms may be a path to success.

But note the "may be." Like everything else, this is a gamble. I may be talking out my butt. Proceed with caution. YMMV.

Audio retail. In recent years, we've seen the rise of several strong competitors in the audio direct sale environment, including companies like Sonic Electronix. These are directly taking on the more established audio online sales companies like Audio Advisor and Crutchfield—and they're growing fast.

Yes, these are big names. But starting an online retail presence doesn't mean you have to sell everything from Sennheiser and Sony. It's possible to be a lot more focused and selective. (For a brick and mortar take on that, check out Cloth and Metal, who sells Schiit—amongst many other things.)

Or, you could even take a hybrid approach—selling your own cables and a select offering of other audio products. Or you can look into a whole new retail model. The point is that it may

be easier to start in retail, build a customer base, and then look into making your own complex, high-risk products.

Strategy the Fourth: Intrapreneurship

Whew. This is turning into a long chapter. Don't worry, the last two strategies will be shorter.

So what is intrapreneurship? Well, it's like entrepreneurship, but within a company that you work for.

Think of it like this: If you work for a company that's close enough to the audio biz, you may be able to convince them to fund an excursion into audio … with you at the helm. You could even end up with royalties based on performance, or even ownership, if the company is a separate entity?

"Never gonna happen," you grump, crossing your arms.

Well, it does happen, and it has happened many times in Mike's companies. Theta worked that way. Dave was a critical part of Theta … but his contributions were paid in large part via royalty. I did Cobalt for the same reason. Other employees made other products and were paid the same

way. At least one major audio company got its start through Theta employees (and that's not counting Schiit.)

And Schiit works the same way. We're funding the efforts of a couple of employees who want to do things that are off the beaten path. If they turn into products, everyone benefits.

So yes, it happens. If you're close enough to your chosen product, and you have an idea, it doesn't hurt to ask.

Strategy the Fifth: Buy the Dream

And to wrap up, here's the big one. Go all-in. Mortgage the house. Sell the cars. Cash in the 401k. Because, let's face it, a lot of established families who say, "We have no money to start a business," really mean, "We have no money we're comfortable using."

And yes, I totally understand. I never had the courage to hang my butt out that far. Because the pitfall is obvious: complete financial annihilation. I'm not comfortable staring bankruptcy in the face.

But maybe you are. If it works for you, kudos.

Bottom Line: If You Want To Do Something, *Do Something*

And there it is. You can sit back and sigh wistfully, you can make up excuses to put off work, or you can angrily pound the keyboard, spreading your message board wisdom about how current audio products are fubar.

But … every sigh, every excuse, every comment does only one thing—they take away the time you need to make a successful business.

You in?

2016, Chapter 3

The Importance of Service

Okay, so let's go to a place I haven't talked a lot about: service. Or, more specifically, what happens when something breaks.

Yes, I know, this may seem like a topic that's about as interesting as a plumbing manual, or discussing the finer points of automotive wrecking, but in actuality, service is really fascinating. In my opinion, doing service well is the primary thing that separates a mature company from a hobbyist venture or early-stage enterprise—and it is absolutely key to providing the highest possible value.

"Wow, that sounds pretty over-the-top," some of you might say. "Isn't a great product more important?"

Well, exceptional service starts with exceptional products—products that are designed to break as little as possible. And before you go into full "Well, duh," mode, this isn't as easy as it seems. Even

the simplest product is a complex assemblage of parts … and it's usually the first time all of those parts have been put together as a whole. This means there are complex interactions that you may not see (at all) until you've built a few hundred—or a few thousand—of them.

And those "complex interactions" will come in first through service. *The amp is making a funny noise. A channel is down. The DAC won't lock to a digital input. Radio is playing through the phono preamp. A volume control sounds scratchy.* Get too many of any one of the above, and you're looking at a complex interaction … which means you need to make a change.

Which means: *if you aren't paying close attention to service, you may miss critical data that will help you make your products better.*

Again, I know, this sounds like a "well, duh," moment. But it really isn't. Many organizations—even very small ones—have a service department that is kept as far as possible from the lofty towers of the CXO suite. (Yes, even when a "CEO" title is laughable … I personally wouldn't think about taking such a grandiose title until I was looking at an organization of several hundred employees. At least. So let's just refer to the

upper management as the "HMFWICs"—head MFers Who's In Charge.)

Why hide service from the HMFWICs? Tons of reasons:

- The HMFWICs think it is beneath them
- The HMFWICs pontificate about how their time is better spent in processes, logistics, research, engineering, deal-making, golf-playing, drinking and schmoozing, going to fancy retreats on the company dime, etc, the latter three usually being the most honest
- The staff truly wants the HMFWICs to spend their time better at processes, logistics, research, engineering, etc, so they divert service concerns away from them
- The staff below the HMFWICs want to hide a perceived "dirty" part of the organization from the upper management

Well, here it is, for all HMFWICs: service is *not* beneath you. And spending your time better at processes, logistics, research, engineering, etc *includes* service as a *primary* component.

Or, in other words: ***HMFWICs, pay attention to service, or else.***

And please note these are not empty words. Every day at Schiit, I spend time in service. Sometimes

even doing repairs directly—stuff that's puzzling, off the beaten track, or that occurs in clusters comes to me.

Yes, you read that correctly: it's entirely possible your broken amp was repaired by yours truly. Yes, even today, with all of our growth.

"Well, that's an entire bag of crazy-snacks," some CEOs will say. "You could be designing new products … working on products … making deals."

Yes, I could. And our products would suffer as a result.

Because I take some of the service burden directly, I feel it keenly. I'm highly incentivized to make sure our products are as dead reliable as possible. And this incentivization results in a feedback effect, so things get better and better.

This is why our service load is so light. Between Tony, Jesse, and myself, we're spending only about 12 to 18 hours per week (total, not each of us) for well over 100k products in the field.

So is everything rosy and perfect? Not at all. Things break. Weird things happen. But through service, we see them quickly … and we get them fixed.

But service isn't just about stuff that breaks after a while in the field. In fact, service really starts with the warranty, how you troubleshoot, and how you handle D/A products.

So let's talk a bit about all of those aspects …

Warranties: Not Just Rolling the Dice

You've probably noticed that our products have different warranty periods. The larger products, from Asgard 2 and up, have 5-year warranties. The smaller products like Magni 2 have 2-year warranties. Fulla has a 1-year warranty. And tubes are covered for only 90 days.

How did we arrive at these numbers, you might ask? Did we roll the dice and see where they fell? Did we use a giant Wheel of Fortune? Did we sacrifice small animals to read their entrails?

No. The warranty periods were chosen deliberately in every case. The decision was based, roughly, on two factors:

1. **How long we figured a product would last, given typical usage.** Note the "typical usage." This will be important later.

2. **How long the average competitors' warranties were.** We wanted to be longer than most of them, without going full-crazy like some outliers offering 20 years (or even more.) More on this later.

"So does that mean a Magni 2 is made like crap, so it has 2.5× less warranty than an Asgard 2?" you might be asking. "And, if some companies can do 20-year warranties, why can't you?"

Both good questions. Let's look at them in turn.

First, no, there's no reason a Magni 2 won't last as long as an Asgard 2, given the same usage. Or a Ragnarok. Or a Gungnir. They're all built to substantially similar standards, with the same gauge of sheetmetal used in the chassis (well, except Ragnarok and Yggdrasil, which are thicker), the same FR4 2- or 4-layer boards, the same premium-but-not-audiophile-insane components, etc.

But note the "given the same usage." Magni 2s (and the other small components) typically get used harder than the larger gear. Magni 2s are much, much more likely to get thrown in backpacks, carted around many different places, hooked up into different systems, etc. Because of this different usage, and to keep prices down, we

decided to limit the warranty on smaller products like Magni 2. Fullas are even more likely to get bashed around. So the warranty is even shorter. Tubes? 90 days covers DOAs, and that's basically all we need for NOS tubes.

And yes, we said it. *"And to keep prices down."* When you're doing $99 products, this is absolutely a valid concern when setting warranty period. Sorry, but that's reality.

So what about those 20-year warranties? Why not up the warranty term for everything in the line? Isn't that a great way of saying how confident you are in your products?

Well, yes … but in the era of lead-free solder, it also may be a little foolhardy.

"Wait, what?" you may be asking. "Isn't lead-free a good thing?"

Absolutely, when it comes to not growing third arms out of your forehead and having kids that can do basic math and be trusted not to eat the furniture. However, in terms of electronic products, lead-free solder still has some question marks around it. Although we haven't seen it in our products yet, lead-free solder can grow "tin whiskers," which can cause shorts on the

board—and all sorts of problems. Perhaps this won't be an issue down the road … but we don't know. Where we used to be comfortable with a 20+ year lifespan for leaded solder, lead-free may be more like 10 years max. Or maybe not.

But would you want to stake your company on a "maybe not?"

Yeah. There you go.

And, finally, although it sounds a bit morbid to say this, every product needs an expiration date. 20-year warranties expose you to the "what if lead-free is a big issue in 10 years" problem, but infinite warranties expose you to something even worse—infinite liability. And companies typically don't like infinite liabilities. They don't do good things for balance sheets.

Melodramatic? Maybe. But if 80% of your boards fail in 15 years due to tin whiskers and you have a 20-year warranty, weeeeellllllllll … that's a huge bag of bad stuff. And, if you have a boatload of products with infinite warranties, good luck explaining (to an investor, bank, or corporate suitor) your strategy for servicing them 30, 40, 50, or more years down the line … when parts may simply not be available.

And that's why you'll see 5 years from us, max. Excuse us for wanting to keep prices down … and excuse us for wanting to be around for the long haul.

"We've Got a Deader"

Okay, now let's move on to the first "service case," to use some corporatese: the D/A product.

Yes. It happens. You open up a great new shiny wonderful thing, plug it in, hook it up … and it doesn't work. Just plain dead. No workee.

And yes, it happens to everyone. Us included. No matter how much we test, no matter how long we burn something in, there are going to be D/A products. This may be due to shipping damage … or it might just be its time to expire.

So what do you do?

If you're Schiit, you do what's known as a "Rapid Return/Exchange." We send you a shipper so you can send the defective product back to us at zero cost, and we send you a replacement as soon as the defective one hits the mail.

Yes, I know. Nothing groundbreaking, at least in terms of large-company policy. However, the more

discerning may have noticed that we say nothing about “Rapid Return/Exchanges” on the website, just some stuff about “If you need warranty service, you’ll pay shipping one way, and we’ll pay the other.”

Yep, sure, that’s the letter of the law. And that’s usually what happens, when your Asgard expires three years after you bought it.

But if it’s new and dead, we will do a Rapid Return/Exchange every time we possibly can. Which is about 99.9 % of the time. But if you happen to buy the very last black-chassis Asgard 2 and it shows up DOA, weeellllll … then it might have to come back to us. But that would be on our dime.

The point is, we can’t guarantee a Rapid Return/Exchange every single time, 100 % of the time … which is why it ends up being undocumented on the site.

The Art of Troubleshooting

Okay. Let’s say the shiny new product shows up … but it doesn’t work quite right. It’s noisy, or it glitches when hooked up to your laptop’s USB port. Or it works happily for a year or two,

then drops a channel. Or suddenly refuses to show up in your computer sound output panel.

In these cases, it's time to troubleshoot.

And this is where it sometimes gets fun. Because there are some people out there who absolutely, positively refuse to believe there can be anything wrong with their system other than their Schiit product, and will spend much time and effort defending their system, rather than doing some simple troubleshooting.

Or so it sometimes seems. I see a lot of the support emails, and it's amazing how some like to argue that their 100 % electroplated yak-hair and virgin rubber cables couldn't possibly be the problem because they cost several times the retail price of the gear they're using it with, when a simple cable swap would show that one of the cables had expired. Or that there is absolutely no way that Microsoft could update their USB drivers without them noticing it, despite the fact they're running Windows 10 on a computer that's constantly connected to the internet.

In fact, running through the numbers, about 7 out of 10 "problems" with our gear are not related to the gear at all, but is usually one of three things:

1. **USB port power management problems.** Number one with a bullet. We have very few problems with drivers these days … but getting a USB port not to throttle power sometimes seems to take an act of God.
2. **Bad cables.** No kidding. You would not believe how many bad cables are out there.
3. **Operational error.** You would not believe how many Bifrost problems can be resolved by asking if the right input is selected.

Now, we know that the USB port power management problems are real problems, and we continue working with C-Media (and through them, Microsoft) on minimizing them. But the reality is that companies like Seagate and Logitech also have problems with their hard drives and input devices being powered by USB ports (do a quick Google search if you don't believe me). Plus, El Capitan has been nicknamed "El Crapitan" by many USB DAC users (not just ours.) Again, don't believe us, Google it.

And if companies like Seagate and Logitech can't get 100 % compatibility … aaaaand if Apple has worked for 8 months to undo the problems with El Crapitan and are only finally getting it worked out in 10.11.4 (beta), weeeeeeellllllll … there are *gonna* be problems.

Those top three issues are why we provide a number of troubleshooting guides on the site (USB Problems, DAC Problems, and System Problems). These guides are quite detailed, and will almost always allow us to determine if there's a problem with our product or not. If there's something wrong, we'll get the product back, fix it, and get it back to the customer as fast as possible.

But sometimes, people don't want to go through the troubleshooting. If they don't, or can't, we have only one choice—to take the product back and have a look.

And many times, there's nothing wrong with it.

And that's arguably the worst thing that can happen. "Fault not found," is not cool. The customer doesn't want to hear this. And we don't want to send it back without really trying to get the product to fail. That's why a "fault not found" repair gets a whole lot more scrutiny:

1. It's run through two separate technicians, to ensure that one isn't missing something (one of them is usually me.)
2. It's burned in for a while to see if it only happens when warm or after a period of time.
3. If it's a digital product, it's checked on several computers.

And yes, sometimes there really is no fault to be found. Those are painful to send back, but I'm absolutely certain that our current process will find something wrong, if there is something to find.

Of course, there are always problems that don't slot neatly into a "type." Those will take some more back and forth, maybe even with the help of some photos or screen-captures. In every case, it is much better to do this via email, since we have a complete record of what we've said, what we've recommended, and what the customer has tried.

Hence our insistence on email support rather than phone support. I can't tell you how many times I've seen phone support end up in endless loops due to lack of information … or completely inefficient as they have to enter all information into the ticketing system in real time.

And, since email support is so close-coupled to service, we know what typical problems are real … and which may be phantoms. Have a Lyr of a certain vintage with noise in one channel? You're gonna get one email from us asking you to swap the tubes from front to back, and if the noise doesn't change channels, one more email to send

it in. (Bad batch of gain resistors, believe it or not.) Have an Asgard 2 that's wonky? We're going to be pretty insistent on you going through the entire System Problems troubleshooting guide. Because they soooooooo rarely break.

Beyond that, what breaks is mainly what you'd expect. In order:

1. **Switches.** Things that move are usually what break. See the top three items on this list. Note that this doesn't mean the switches are crappy … they're just the most used, and the most likely to be snapped off if the product is mishandled in shipping.
2. **Connectors.** By this, I'm including everything—RCA jacks, tube sockets, AC connectors, etc. Another mechanical part. Not surprising they sometimes have issues. However, the Neutrik headphone jacks have been stunningly reliable. I've only seen a single case of a broken jack … ever.
3. **Potentiometers.** Yes, sometimes potentiometers break. Usually in shipping. This is despite using very nice Alps pots.
4. **Wall-warts.** Despite pre-testing wall-warts, sometimes they arrive DOA, buzzy, or they expire after a few months.

5. **Transformers.** Every once in a while, we'll get a transformer that expires before the warranty. Note we're getting into the "Every once in a while" phase here. Once in a looonggg while.
6. **Everything else.** Output transistors sometimes show up dead, but it's so rare that the occurrence is probably in the low double digits. Sometimes input transistors show up dead. How that happens, I'm still guessing. Sometimes we get a noisy capacitor or resistor (yes, it happens). And sometimes a relay driver dies. However, we have yet to have a truly dead relay.

But if you have problems, take a deep breath. We'll get it fixed up … whether that means getting you to swap a bad cable, or if we have to bring it back and replace the whole thing.

In the Shop

Once a product comes back to us for service, it gets just that: service. We do very little "board-level" or "module level" service, and essentially no "swap for refurb" service.

These euphemisms mean, in order:

1. **Swapping a whole board without troubleshooting the existing one.** Fast, and very

inefficient. If you're running a typical high-margin niche product (read, overpriced at over 8× the BOM cost), sure, you can just swap boards and scrap the deaders. But this really isn't an efficient way to run.

2. **Swapping a module without troubleshooting the existing one.** At times we'll do this on complex boards, like Yggdrasil's USB input or DSP board. But in either case, those boards go into a queue for troubleshooting later.
3. **Swapping the whole product for a refurbished one.** When a product is essentially unrepairable, or only repairable at great expense (or with low-cost labor), this is a common strategy. None of our stuff is designed like this, so it doesn't apply.

So what do we do? We actually go through the board, troubleshoot it, and fix what's wrong. This requires a much higher level of technician than one who is simply swapping boards and modules, but it's absolutely worth it. It's worth it, because the cost of troubleshooting an existing product is lower. It's worth it, because it promotes greater understanding of what actually went wrong. And it's worth it, because it encourages us to keep our products easily serviceable. All of this contribute to a feedback loop that further reduces service problems and service costs.

And now you see why I do some of the service. It's very useful to dive into a board and see what's really happening. And once you've seen a few with the same problem, it's easy to go upstairs and make a change to improve the next run. That's why we've seen less and less problems over time … and less problems with new product introductions. We learn from service … and that learning is instantly applied to improving the product.

And I've seen some pretty bizarre stuff. Some of it's our bag, some of it is the customer's bag, and some are mysteries.

Our bag:

- The aforementioned Lyrs and Valhallas that got bad gain resistors. Yep. No way to predict it. And no way to know it would be a problem, until they started failing a couple of years later.
- Some Bifrost Multibits with a distortion problem on sample rate change. It took us many weeks to replicate this and find a firmware fix, but we have one now. If you have a funky Bifrost Multibit, let us know and we'll take care of it.
- Old Asgard 2s with humming transformers. These still show up from time to time.

Their bag:

- An Asgard that a cat threw up in. Cleaned it up. It worked fine.
- An Asgard that got orange juice dropped in it. Again, cleaned up and worked fine.
- A Ragnarok with a near full-grown cockroach in it. No clue if it caused the problem … still wondering how it got in there.
- A Mjolnir with about 5 pounds of dust inside of it, nicotine stains on the outside, and various other stains I don't want to know about. When cleaned, the noise it had been making went away. Not a huge shock.
- A Mjolnir that had been disassembled … including having holes drilled in the bottom (why?), a broken potentiometer, screws snapped off and chassis gouged to heck and back. This was supposedly an attempt to install an audiophile fuse. That repair bill was pretty big.

Mysteries:

- A Magni 2 Uber that literally smoked the output stage—as in, it burned the output resistors and charred the board. That one got a whole new board. Why? Unless it was overdriven into a shorted load for a long time, I have no clue.

- A Valhalla 2 that hummed (through the headphones) for no reason at all. Checked all the grounds. Replaced a bunch of parts. No dice. Nothing fixed it. That one got scrapped.
- A Mjolnir 2 that came back, "eating" expensive NOS tubes. I took that one, figuring it would be easy—a bad heater regulator, overbiased, oscillating, something like that. Nope. Heaters were perfect. Biased perfectly. Not oscillating. Sounded great. Stuffed it on the burn-in rack for days with stock tubes. No problems. That had to go back to the customer with a "no fault found" diagnosis.

The Penalties of Inattention

I've seen companies brought to their knees by their service load, and I've heard stories about many more. Ed Miller's description of the last days of Great American Sound—killed by unreliable amplifiers, desperately trying to get ahead of a tidal wave of broken gear—really resonates.

But those are big, grandiose kinds of ends. It takes a lot of bad product, for a very, very long time, to really kill a company. Allowing your service problems to get to that level is probably pretty rare. Because, at some point, the troops

have to go to the HMFWICs and say, "You know, we really have a problem."

But I've also seen companies not much larger than Schiit labor with 7 to 10 full-time service technicians … and consider that 100 % normal.

And that's probably the real problem with inattention to service … it can seem to be relatively benign. It can seem to be "normal" and "all right," as long as the company is operating, producing a reasonable profit, and the team of service techs are keeping ahead of the burden.

The problem is, "normal" and "all right" continue to get redefined, until they aren't. Until the company isn't doing so well. Until the service techs can't keep up.

Doing service well is one of the primary differences between an established, viable company and a hobby business or start-up. And keeping service close to the top is, I'll argue, one of the primary differences between a company delivering top value … and all the rest.

I'll say it one last time: *service is worth your attention.*

2016, Chapter 4

How We (Didn't) Move Again

If you've been reading this book, you've already heard about our first move (out of the garage) and the second move (out of the tiny Newhall office to our current industrial box) … plus, I know I've mentioned we expanded our current industrial box to make space for Ragnarok and Yggdrasil.

So, for those who are keeping score, that's two moves and an expansion. And that's plenty enough for me, like, ever.

Why? Moving is hugely disruptive. When we're 1/3 of a garage in size, the disruption may be only a day or so. When we're 1800 square feet in Newhall, you're talking more like a week total downtime. Now that we're 8300 square feet, the prospect of moving is really scary. We could be talking several weeks of disruption until we really got ramped back up again. Because now it's not just about moving boxes—it's about moving all the racking, fixtures, test benches, test gear, listening room equipment, office infrastructure,

shipping systems, finished goods, ad infinitum. The racking in itself is a huge project, because we're talking serious racks, 20 foot tall and bolted to the slab, designed to hold tons of sheetmetal and transformers.

So it was with quite a large amount of dread that I heard Alex talking about needing more space, shortly after the beginning of the year.

"We can push back on the suppliers," I told Alex. "Schedule things out. Have them ship only when we need the parts. You know, the old 'just in time' thing."

Alex sighed and shook his head. "A lot of them don't have the warehouse space."

"What about finding new suppliers?" I asked.

"For boxes?" Alex said, gesturing up at the racks. Packing material—boxes, box inserts, and foam—take up an insane amount of space.

I frowned. Our box guy was really the best in the business. He'd been with us since the start. He'd engineered a dozen really cool boxes for us, to the price point we needed. He wasn't late.

But he didn't have warehouse space. I knew that. And I wasn't about to throw him under the bus.

“Plus, there’s transformers,” Alex said. “Those take up a ton of space.”

By “transformers,” Alex meant “wall warts.” The only thing we buy from China and ship via ocean freight is wall-warts. We’ve tried and tried to find a US supplier, but everyone we’ve talked to has proposed prices that would make Magnis, Modis, Manis, Wyrds, and a bunch of other products unsustainable. So we order, like, a metric schiitton of wall-warts about 3 to 4 times a year. Tyler does the import dance with the powers that be (I imagine it is something like the dance a bee does when it finds a field of promising flowers, but with a whole lot less pollen) and pallets and pallets of wall-warts eventually show up at our back door. Due to the length of time involved in customs clearance and ocean shipment, we couldn’t do those as a just-in-time kind of thing. And they take up a ton of space.

“And we’re running out of burn-in space, in any case,” Alex said. “The new big products (meaning Yggdrasil and Ragnarok) are really killing us.”

“Can we put in more racks?” I asked, wandering back to the expanded space we’d leased a while back. It had looked huge at the time. Now it looked really cramped.

“We can put in bigger racks,” Alex said.

“Do it.”

“But I still don’t think it’ll take us past summer. And Mike has his Manhattan Project, and you have the (redacted) and the (redacted) and Mike says he’s gonna get the (redacted) back on track—”

“But things will slow down coming into summer,” I said. I really, really didn’t want to think about moving.

Alex was silent for a while. Then he laid out the single statement I really needed to hear, in order to think about this seriously:

“I look at it this way: we can move in 6 months, or we can move in 18 months,” Alex said.

“What?” I said, not immediately grasping what he was saying.

“We can move this summer, or the summer after. There’s no way we’re moving in winter. Ever.”

I shuddered, imagining a weeks-long business disruption, right in the middle of the December–January peak season. It wasn’t a pretty picture.

And Alex was right. If we were going to move into a larger space, it would have to be in the summer. Anytime from April to August would probably be fine (assuming no big product launches, and, well, that might not be such a good bet.) Beyond that small window, we'd be going crazy just trying to keep up.

So, did I think we could go 18 months in our current space?

Three words: *No. Friggin. Way.*

Why So Spacious?

Now, some of you are probably stifling yawns and wondering, *Why the hell is this even an issue? Get more space and quit whining about it!*

Well, leaving aside the business disruption aspect, space (or facilities, ifn you wanna be fancy) is one of the most critical aspects of a business. Get too little, and you'll end up moving too soon. Get too much, and you can hang yourself out to dry.

The lease—or mortgage—on your space is one important component of your overall cost of doing business. You don't want an oversized lease on space you'll never use. That's just money down the drain.

And—maybe even more important—it can have a negative psychological impact on your staff if you never fill up that space. Centric once took an office that was much too large for our staff, in the heady days of the last years starting with "19." We expected the need for a much, much larger programming staff than was realistic, once the years began with "20" and ended with "01." Sure, we weathered the worst of the Web 1.0 bust pretty easily, but we never really filled up the area we had allocated. So it was a giant, empty, echoing space with a few desks that the guys used to fly gliders and for practicing their putting. And, even though nobody said much about it, I could tell it weighed on them. *When are we going to fill up the space? Are we doing as well as we thought we would? Are we doing OK at all?* And we ended up staring at that space for seven years.

No thanks, not again.

But, in this case, we really were outgrowing our space. The question was how much to look for—and to lease or buy.

Aside: Buying. Yeah. Buying. Kinda like a (much longer) lease, except (hopefully) you're building equity, plus you get to write off the interest and the depreciation on the building, mi-

nus you have to pay property taxes. If you have a conventional accountant run the numbers, and you have the money for a down payment, buying absolutely makes sense …

… well, that is, if your business continues to sustain its current sales and profitability (or expands), and if commercial real estate doesn't have a big oopsie (read, 'significant downturn') and if you're comfortable with the government looking down your shorts (most business property loans are SBA loans), and if everything goes well on the business and personal side (I've seen business partners get roasted by commercial real estate investments that went sideways).

Yes, that's a lot of "ifs." But still, we were thinking about buying when we started this process. You bet. It increases the pool of prospect spaces we could look at. And we thought we could manage the "ifs." But, if there's one thing I've learned, is never take business property investment lightly. There's a lot of value in being more nimble and adaptable … hell, if you can do month-to-month without much penalty, do it. Leases still boil down to, "Pay us this much every month … every month … even if you can't."

Alex didn't care if we leased, bought, stole, built,

whatever … as long as he had his space. The more the better. Twenty-five thousand square feet would do us just fine, in his opinion.

But 25 000 square feet is HUGE. It was bigger than any business I'd had, by a very long shot.

I figured we could get by with a doubling … to 15 000 or so. Especially if it was all floorspace. The upstairs area in our current building is pretty much useless—imagine carrying Ragnaroks and Yggdrasils up and down stairs for testing, or assembling static-sensitive components on carpet. Yeah.

But 15 000 to 25 000 was enough to start ballparking. Even if the small side made Alex nervous … and the larger side made me start biting my fingernails.

Looking for a Few Good Spaces

Funny about those "best laid plans" things. As soon as we started looking, we immediately came upon what seemed like a perfect candidate: a 21 000 square foot space about a mile from our current building. The space was:

- A lot newer and shinier than our current space

- Almost completely empty, with no second-floor buildout at all—it had been an indoor trampoline space (seriously, do people pay for this stuff?)
- Near one of two local breweries—close enough to walk to, but not as close as the brewery that just opened up behind our current space
- Directly to the side of the informal Paul Walker memorial

Yeah. As in, people would hang around with flowers and candles and stuff, right out to the side of our proposed new space.

"Perfect," I told Rina. "You can get a bunch of friends to come by and do a big loud séance or something, when we go see the space with the realtor."

"Which would do what, exactly?" she asked.

"Get us a discount, if we play our cards right."

Rina looked at me askance, as if to ask if I'd been out to the brewery already that day. But I was serious. The space was for sale, not for lease. And if it had been on the market for a while, maybe I could talk them down even more …

Alex drove by and looked in the windows, then came back and pronounced it promising. "They

have one big box office up front, but that's about it. And they have a *loading dock!"*

"Great," I thought, imagining someone driving the forklift off the loading dock one day. We already have a big gouge in one of the racks. But a loading dock would be handy for the big stupid wall-wart shipments.

So. Alex was thrilled. Rina was thrilled. I should be thrilled. But they didn't know that an empty box meant a lot of buildout, and that a loan at prevailing rates didn't give us any real month-to-month advantage over a lease, plus there'd be property taxes …

And, in the end, the whole thing ended up being moot. When we contacted the broker, they made sad noises and said, "Sorry, that's in escrow."

And the "perfect" building evaporated, just like that. Lesson learned: call the broker first, before making any grandiose plans.

Which still left the problem of where we'd end up going.

Not a big deal, I figured, and fired up LoopNet (a listing service for commercial properties.) I plugged in our min and max square feetage,

selected “lease or buy,” and sat back, preparing to be wowed.

And got … basically nothing.

Ah, schiit.

The Duh Moment

“Vacancy rates are only about 2 % in this valley right now, for the kind of space you’re looking for,” said the guy from the SCVEDC (Santa Clarita Valley Economic Development Corporation.) “And that’s factoring in at least one large building that isn’t completely ready for occupancy, and … ”

Well, that explains the Loopnet results, I thought, tuning out.

“But the brokers may have some pocket listings, you said,” Sue prompted him. Sue was my business partner at Centric for many years, until she came to her senses and took a job with the SCVEDC. She’s looking much happier now. Marketing is a killer industry that can eat your life. More on that in a future chapter or two.

“Well, yes, but we’d have to contact them and see what they have,” said the SCVEDC guy.

“I think that’s what they want,” Sue said.

Alex and I both nodded vigorously. If we couldn't find a space large enough to move into, we'd have to go to "Plan B."*

* Plan B was to find another warehouse, 5000 to 10 000 square feet worth, move all stock out of the current building and store it there, and set up the current building for nothing but production. We'd lease a truck to shuttle stuff back and forth as necessary. I don't think I have to go into why this wasn't a wonderful option—the ongoing business disruption, the need to have someone drive the truck, the inevitable Murphy's Law crap about never having the parts you need where you need them, etc. But if it was that or nothing, well …

"Well, we can certainly make some introductions," the guy from the SCVEDC said. "But have you talked to your current landlord about whether or not they have another building, something you can maybe swap into a lot easier? You still have time on your lease here, right?"

"Yeah," I said. Almost three years, in fact. It would be another pain in the ass. We'd have to sublease it, or find our the extortion needed to break the lease. Another problem with moving—there's always the old space to deal with.

"Well, that's what I'd do," he said. "Talk to them, first. Maybe you can get out of the lease easier, find something they already have."

I looked at Alex. "You were right. We should have taken that extra space here, last year." Last year, about 4200 square feet of space—the unit next to ours—was available. We weren't really large enough to use it at the time, though Alex had looked at it with a gleam in his eye. I'd nixed the idea, but if we had taken the space, it would have given us a total of about 12 500—and might keep us from having to move for another year, or even more.

Alex looked innocently at the ceiling, but I knew he was thinking, *Yep, I was right.*

"It's probably best if we talk to some brokers," I told the guy from SCVEDC. "I don't know if our landlord has any more buildings around here, and this building is all leased up."

"That's fine. I'll pass along some names."**

** The Santa Clarita Valley Economic Development Corporation is a very aggressive local agency bent on bringing business into our neck of the woods—and retaining existing businesses. This is very important in California—and the

US's—fundamentally anti-business climate. I won't go into a political screed, because this is not the place for it, but the reason companies do things like inversions and leave profits out of the country is that *the USA has the highest corporate tax rate in the developed world.* Yes. Seriously. Look it up. Plus California state income tax. Plus other "hidden" business taxes and fees. Without research tax credits, an IC-DISC, and other business tax credits available for hiring in California, it would be a significant drain on our competitiveness. Hell, it probably is anyway. But for some reason, it seems like it's more intelligent to have us (and every other business) spend tons of money on multiple individuals and agencies in order to try to reduce their effective tax rate, (and, when large enough, keep money out of the country) rather than set a logical tax rate that companies are more willing to pay—and that their compatriots will call them out as cheap-asses if they aren't willing to pay … but I'm starting to ramble here. Bottom line, if you're in business, see if you have a local Economic Development entity, and see if they have some free advice for you. The SCVEDC has significantly impacted the way we do business, in a very positive way.

And that, I thought, was that. We'd talk to some brokers, see what they had … and if they had nothing, we'd proceed with Plan B.

But first, I figured, having Alex talk to the landlord and see if they had anything else wouldn't hurt …

The Easy Out

And, surprise surprise: *that 4200 square feet we missed was going to be available as of Halloween, 2016.*

Yes, that space right next to us.

Yes, that space we could just punch a hole into, like we'd did for our first expansion.

Yes, that space that would mean essentially zero business interruption.

I am stupid. I actually had to think about it a bit. Here's why:

1. **It didn't get us to what I considered "end game," at 15 000 to 20 000 square feet.** It was only 12 500. But considering I still don't know what our "end game" is, well, that's kind of a crappy argument.

2. **I wanted to see if the brokers came up with anything.** Because, you know, the grass is always greener, or something.
3. **It would make us the biggest tenant in our building, by far.** Another 5000 square feet or so would make us the whole building. And the building is kinda old, and the landlord is a cheap-ass. But it's not like it's falling apart or anything. Well, except the air conditioning in Suite A. Which they (finally) just replaced after a catastrophic meltdown.

So yeah, I slept on it overnight. Then I sent an email to Alex and copied Mike, saying, "Well, it's kinda like, 'well, duh,' right?"

Mike agreed. Alex breathed a sigh of relief. And I began the process of signing the lease to expand our space, come this Halloween.

And that's how we *didn't* move.

At least not yet.

2016, Chapter 5

The Subjectivist/Objectivist Synthesis

Color me stupid, but I'm going to wade right into the subjectivist vs objectivist debate, and see if both groups can find a happy place.

In the process, I'm going to (attempt to) distill the objectivist/subjectivist debate into a few lines, give some examples and anecdotes (yes, I know, not data) from my experience, call out some interesting factoids to think about, outline (what I see to be) the stuff that both objectivists and subjectivists get right and wrong, and attempt to find a synthesis in the end.

Now, don't get me wrong. I don't know where this is going to end. Like most of the stuff I write for this book, this chapter is written more as a stream of consciousness, usually over a single evening, or maybe two. Don't expect a scientific paper. It won't be endlessly footnoted and cited. It will call on you to do some Googling if you

want to dig deeper into the stuff I bring up.

Spoiler alert: both subjectivists and objectivists are wrong on lots of things, at least in my opinion. Both camps have members that get way, way, way too into their dogma, both camps have advocates that really need to chill out and have a beer (because, let's face it, the best mod that improves musical enjoyment is probably alcohol … it will almost certainly beat the pants off of any car-priced DAC), and both camps need to realize that, even if one camp "wins," the world will spin on, despite any histrionics.

Yes, I know I am completely insane. And yes, I know this may not end well.

Ready? Let's get started.

The Subjectivist/Objectivist Debate

If you've never experienced a subjectivist/objectivist debate, you must be new to the internet. Sorry about your sanity. Here's how pretty much every objectivist/subjectivist debate goes, in 12 lines or so:

Subjectivist: "I think my new Arglebargle X1000 sounds way better than the Craphound PST-1."

Objectivist: "No, if they both measure 20 Hz to 20 kHz flat, have THD below 0.1 %, and have a low output impedance, they have to sound the same."
Subjectivist: "I think my experience trumps your measurements."
Objectivist: "No, humans can't perceive anything beyond that, see (insert links to tests here.)"
Subjectivist: "Well, I hear a difference and so does (insert anecdotes about friends, spouses, dogs, fish, etc)."
Objectivist: "Anecdotes aren't data! You're fooling yourself. (Insert words about scientific method and significant results here.)"
Subjectivist: (Sigh.) "Just leave me alone to enjoy my Arglebargle with the other folks I'm talking to here."
Objectivist: "No! Don't you see you're being taken advantage of by evil companies selling overpriced gear?"
Subjectivist: "You probably just can't afford good gear!"
Objectivist: "You're nothing but a shill for the man!"
Subjectivist: "Ad hominem!"
Objectivist: "Ad hominem!"

And then you repeat the last two lines, ad infini-

tum, ad nauseum.

Even in this oversimplified distillation, we end up with some really interesting questions. Why does this particular objectivist* feel he has to insert test data as the ultimate arbiter of human perception into a discussion that was simply among excited audio-gear owners? Why does this particular subjectivist* reject the idea that measurements can offer anything of value? Why can't the two groups simply ignore each other and both live happy, but separate, lives? Does the objectivist think they are "saving" the subjectivist from exploitation, or is the motive more inward-focused? Does the subjectivist really honest in his perceptions, or are they colored by the cost of the component? How did the objectivist come up with their measurement thresholds that certify transparency? What mood was the subjectivist in when they made their assertion? What experience does the objectivist have with audio test equipment? Did the subjectivist try the product in the same exact location as their own equipment and use the same exact music they are familiar with?

* Not all objectivists feel the need to insert themselves into subjective conversations, just

as not all subjectivists reject measurements. Groups are always a continuum of individuals. As I said, this is an oversimplification.

Yeah. It gets murky, real fast. And it gets even murkier when you really start looking at both sides. Here are some factoids to consider. (And yes, thank you, I know that the plural of "anecdote" is not "data"—I believe I've already mentioned this.)

Consider:

1. **Every engineer I personally know who is working in audio is a subjectivist**—they think that stuff that measures good sounds different, without exception. Just protecting their jobs? Perhaps. But many of these same engineers started like me, as an objectivist … until they had a subjectivist experience or two that changed their minds. And yes, I know, I don't know every engineer on the planet.
2. **It seems that many of the most fervent objectivists aren't working as engineers,** or aren't engineers working in audio, and most don't seem to have any experience with the measurement gear they cite measurements from. Again, I know there are exceptions to this rule. Maybe their fervor is simply

Robin Hood syndrome—trying to help the poor deluded subjectivists?

3. **Most of the subjectivists I've met are unwilling or unable to factor in their own rationalization** when proselytizing for gear, when the reality is we all rationalize w*hat we have chosen* as *the best*—and the more we've invested, the more powerful the urge to rationalize. Nobody wants to look like a fool who spent car money on audio gear that makes no difference.
4. **Some objectivist thought leaders—the kind that have written books—don't portray their findings in a very (ahem) objectivist manner.** Instead, they use emotionally charged language, actively demean subjectivists, and are dismissive of any approach that doesn't align with the path that they have codified. This shrill, "die-heretic-die!" tone is not exactly conducive to rational debate … and (in my opinion), it doesn't seem to indicate security in their belief systems.

"Ah, you're not being fair," some will say. "You're painting with a broad brush, constructing a straw man, I am gonna haul out all my captain-of-the-debate-team tricks on you."

Yes. Which is why I said this is *my* experience,

and that it's *not* absolute. But if it's fair for an objectivist to ask a subjectivist if they've done a blind comparison of their gear when they anoint the Arglebargle as the greatest and bestest thing on the planet, it's fair for a subjectivist to ask an objectivist how they determined the measured limits of human hearing, and their experience with making said measurements on an Audio Precision, Stanford, or dScope.

And—editorializing more heavily here—it's entirely fair for both sides to ask each other to lay off and have a drink instead, because these endless debates become tiresome in a big hurry, IMO.

Now, I can see the attraction of both sides. Subjectivity seems to be the more holistic, natural, touchy-feely way of finding the best gear for you—never mind understanding the measurements, how does it sound? Objectivity seems to be the most rational and comforting—no need to spend megabucks, just check the measurements.

But subjectivity can drive an unnatural mania for more and more gear, while objectivity is the best tool for companies trying to make a quick buck—if it measures well, get it done, push it out, don't worry about the sound.

So how do you reconcile the two? Is there a middle ground?

Well, for me, there's been nothing but middle ground, at least since the early days at Sumo. Once I realized things really did sound different, I accepted that measurements can't tell the whole story. But I didn't throw out measurements. Nor did I throw out the subjective experience. And, from time to time, I've tried to correlate measurements with sound quality, or come up with measurements that relate more closely with perceived quality.

"Enough about you," some might be saying. "Is there a middle ground for me?"

I don't know. Only you can answer that. But I hope both sides can learn something from the other. To get to this synthesis, let's look at what (I think) each gets right and wrong.

The Right and Wrong of the Subjective and Objective Approaches (IMO)

If you're closed-minded, you may want to skip the rest of this chapter. Because, as I've mentioned before, everything is a continuum. There is no black and white. Objectivists are not "good."

Subjectivists are not "evil." Both sides have some very good points—at least in my opinion.

What *most* objectivists get right:

- They provide a foundation to work from—objective measurements will tell you whether your gear is broken, performing to spec, or you're fooling yourself
- They provide a leveling effect, in that inexpensive gear can measure and perform similarly to very expensive gear in some cases—this keeps ALL of audio from ascending into gold-plated Bentleyphilia
- They are absolutely right in that some aspects of performance should be measurable, and that these measurements should be done
- They hold manufacturers to certain standards and help prevent true stinkers
- They provide some common-sense rules on how to get better performance from your products, such as impedance matching, power requirements for specific output levels, etc.

What *some* objectivists get wrong:

- Instantly dismissive of everything except measurements, leading to endless arguments over meaningless specsmanship, like whether or not 0.0007 % THD is worse than 0.0005 %

- A dogma-based assertion that performance above or below certain standards (frequency response, THD, noise, output impedance, etc) assures complete transparency—and that all transparent gear sounds the same
- A dogma-based assertion that all aspects of human hearing are known
- A dogma-based assertion that we are measuring everything we need to measure in order to fully characterize an audio system
- Blind faith in blind testing and the results of flawed studies using blind testing
- A seeming need to preach to audiophiles uninterested in engaging with them, coupled with an inability to realize when to say, "when"

What *most* subjectivists get right:

- They are open to new ideas, like the idea that there may be something more than what we can measure
- They (usually) don't feel the need to go out and convert audio objectivists into subjectivists
- They tend to talk more amongst themselves, compare more gear, experiment with different kinds of gear, tweak gear, and in general push the limits
- They encourage manufacturers to continue to improve their own products, both relative

to their current line and with respect to the competition
- They are usually more willing to take a chance on new and unproven companies and technologies, which can help drive innovation

What *some* subjectivists get wrong:

- They blow subjective differences way out of proportion, even when the actual differences are tiny
- They don't take into account their mood/feelings/physical condition/intoxication when passing judgment on gear (like, if you get something in a crushed box with a broken knob on a day when you've just gotten passed over for promotion, versus a day when you have a wonderful first date, discovered a nice bonus in your checking account, had a wonderful dinner, and a few scotches …)
- They frequently make snap judgments based on brief listening at meets/friends houses/with different gear/under different circumstances, and these snap judgments can endure for months or years
- They are subject to rationalizing their purchases, especially when those purchases are very costly
- They can get obsessive about gear and go crazy

spending, spending, and spending even more, trying to find that "last 1%" of performance

Again, if you're feeling a bit in a huff right now, go back and read those "some" and "most" disclaimers on the summary headings. And consider that there are some who place themselves firmly in the subjectivist *and* objectivist camps. I've said many times before that I would never want to design audio gear without extensive testing (see the test and measurement chapter for more info on that one—by far the longest chapter to date, by the way), yet at the same time, I wouldn't want to have the blind pursuit of measurements be the sole arbiter of my designs. I believe that the different stuff I make sounds different.

Yet, at the same time, I try not to blow these differences out of proportion. Will any of our amps fundamentally transform your headphones into something completely different? No. If you don't like your headphones, you should be seeking new headphones long before an amp or DAC.

And I live by my beliefs that our products sound different, as well. My "main stack" is Mjolnir 2/Gungnir Multibit, not Ragnarok/Yggdrasil. Mjolnir 2 is a warmer, wetter, "happier" amp than Yggdrasil, and Gungnir Multibit is also a bit

more euphonic than Yggy. The combo may not be the absolute ultimate word in resolution and transparency, but I like the way it sounds better than our top stack.

However, note this is a subjective choice … and also note that I'd still be happy with Ragnarok/Yggdrasil. The differences are, as I mentioned, relatively subtle.

So don't get upset if you don't fit the broad categories above. I'm just listing what I noticed off the top of my head. It isn't meant to be comprehensive, or definitive.

Bad Data All Around

"Wait a sec," some of the sharper-eyed readers are saying. "I noticed you cited 'flawed studies' up in the list of objective wrongs. What's up with that?"

Okay. Deep breath. Let's get one thing clear from the start: I don't think either the objectivists or subjectivists have definitive, clear-cut proof positive of the limits or capabilities of human hearing.

The reality is that there's not a lot of money in proving audibility, one way or another. Audio

isn't another $100B drug, nor a breakthrough in cheap and clean energy, or even a new cleaning product that could break $100M in sales every year. (And this is an important fact to keep in mind.)

As a result, the studies I've seen are relatively fringe, ad-hoc, limited, flawed, or all of the above. The most frequently cited, Meyer and Moran, used sources that were not actually high resolution, and is directly countered by Bob Stuart's paper on the audibility of digital filters. At the same time, there is physiological evidence that the brain responds differently to sounds containing inaudible supersonic components (google "Inaudible High-Frequency Sounds Affect Brain Activity: Hypersonic Effect") and there is significant evidence that experienced listeners can hear the differences in amplifiers (but not DACs) in Innerfidelity's recent top-end gear shootout.) Further confounding things are unproven assertions like our own Mike Moffat's statement that he believes hearing is integrative, rather than differential.

Like I said before, I'm not going to go into minute detail in citing every paper, nor in picking them apart. If you're more interested in the subject, I suggest you start Googling, reading papers, re-

freshing yourself on what a statistically significant result is, and coming to your own conclusions.

The point I want to make is: there isn't any study large enough, perfect enough, and significant enough to end all doubt on the subject of what is, and what isn't, audible. Nor will there likely be one in the near future. There's simply no payoff in it.

And, come on, admit it: even if there was an ideal study, a certain tinfoil-hat crowd would claim that there is some giant deceitful conspiracy behind it all!

However, from reading many papers, viewing the results of many tests, hearing many anecdotes, observing my own reactions and the reactions of others, I'm willing to bet on what a potential "perfect" test might find.

Here's my bet on a one-line synopsis:
There are some people who can discern subtle differences below the commonly accepted limits of human hearing, and some of this group of people find those differences meaningful enough to care about.

"So you're saying that if I can't hear the difference, I'm a tin-ear?" someone is shouting. "Well, I

never! What an insult! I'm gonna grump off into a corner!"

No, I didn't say that at all. All I'm saying is that different people have different capabilities. You wouldn't want me flying your plane, at least without glasses. Nor would I expect to pass a sommelier's deductive tasting evaluation.

And (cue us never selling anything ever again), I don't think my ears are the last word in goldenness. There are plenty of times I think I'm fooling myself. I doubt what I hear.

But (big sigh of relief), that's why I have other listeners I trust, both inside and outside of the company. Some of these listeners are much more discerning than I am … they have picked out tiny details that I've never heard, and called me out when I brought in "ringer" products for a listen.

"So why don't you guys do a study?" someone is surely asking. "You can have the perfect data, you can confirm or disprove your hypothesis, you can end the arguments forever."

Well, setting aside that I doubt we'd be able to do a study with the scope and rigor necessary to be considered definitive, you're talking about:

a) A heckuva lot of time

b) A heckuva lot of money
c) A heckuva diversion from what we're doing

And the payoff would be … what? Half the audio world disbelieving and (probably) reviling us, no matter which way the study ended up?

"And maybe killing your own business, if nobody can tell the difference," someone snidely remarks.

Actually, I don't believe that. I think that even if one great study proved complete audio objectivity, the audiophile world would continue spinning, completely unaffected.

So, sorry. We're not going to try to prove our hypothesis. We're simply going to keep acting "as if." Which seems to be exactly what most of you want.

So What Can Each Side Learn From the Other?

"Wow, you really went off into left field, didn't you?" some are probably saying. "Is there any point to all of this blather?"

Well, like I said, this exploration might not go anywhere.

But maybe there is something here. If you zoom way out, you can kinda look at the subjectivist position as being the exploratory one—the drive that keeps us trying to make increasingly better-sounding gear. At the same time, the objectivist position as being the foundational one—it keeps us from going off into cloud-cuckooland as we explore the fringes of audio perception.

I think both sides have valuable insight. The subjectivists remind us that thinking "all is known" is not a good bet. Look at the guy who wanted to shut down the US Patent Office in the 1800s because "everything had already been invented." The objectivists remind us that there are fundamental rules (like impedance matching) and that our quests into the unknown can be costly, frustrating diversions.

If I was a pure subjectivist, here's what I'd take from the subjectivist side:

- You have no reproduced audio without science and the scientific method, and it's worth learning more about this
- Those "huge" differences others are talking about may not be so huge at all, or they might even be cognitive bias

- Spending big is not always the answer when it comes to great sound—being more discerning may keep both your ears and wallet happy

If I was a pure objectivist, here's what I'd take from the subjectivist side:

- Science hasn't fully characterized everything, whether you're talking perception, medicine, physics—question your own hypotheses and be open to revising your position
- No difference in perception to you may be significant to someone else with different perception—and cognitive bias cuts both ways
- Spending big on audio is a largely harmless pastime that doesn't affect you—let them be, and concentrate on something that makes you happy

And to both sides:

In a thousand years, when godlike AIs are unearthing the data-foundations of the Human Internet, they'll be shaking their metaphorical heads at the silly stuff we get into arguments about. Audio isn't a cure for cancer, it's not a new physical particle, it's not a hyperloop transport system. It's a *fun* pastime that helps you *enjoy* music to its fullest.

Sit back. Relax. Buy each other beers. Laugh at yourselves. Try something new. Ask questions. Read and digest. Interact and learn. Because, let's face it, if we're really here to share, and if everyone really has something to contribute, we should all be reading a whole lot more than writing.

Happy objective subjectivizing, or subjective objectivizing …

… or simply sitting back, listening, and enjoying some great tunes.

2016, Chapter 6

Being Comfortable With What You Are

When I was away on vacation, we met up with an acquaintance who came from a financial background—specifically, he managed a hedge fund. He'd just declared himself retired, and very, very comfortable … at 38.

Among the things we talked about was, of course, Schiit.

His first question was probably typical of a non-audiophile reaction, when confronted with our company.

"High end audio? Well, I just bought a pair of Beats headphones—" Mr. Retired began, but quickly stopped himself when he saw the glance pass between me and Ken—another friend we were traveling with.

"Yeah, I know, I know, they're crap, that's what you're gonna say, that's what everybody says, and

I know they aren't all that, but I like them, and they're convenient. I was just mentioning them because, well, if you make stuff like that, you need to be able to tell me why your stuff is better."

"Actually, we don't," I said. "We make headphone amps and DACs, which might be used with your headphones … unless they have internal amps or are Bluetooth or something."

> **Aside:** shoot me, I'm not familiar with the Beats range, nor am I interested enough to look it up.

Mr. Retired looked irritated and confused at the same time. "Wait, you don't even know if I can use your products?"

"Nope," I admitted.

The furrows in Mr. Retired's brows deepened. "And it seems like you don't care, is that what I'm hearing?"

"Now you're getting it," Ken said, grinning. Ken knows a lot about Schiit—and will be helping us in the future as we expand international distribution.

"Why?" Mr. Retired said.

“I can probably count the number of Beats owners who have bought our products on all fingers and toes,” I told him. “So as soon as you said, ‘Beats,’ I figured, ‘Well, that’s it,’” I told him.

“So if I said the name of some brand—some brand your stuff works for—“

“Like Sennheiser,” Rina added.

Mr. Retired nodded. “Like those guys. Sennheiser. Then you’d care?”

I shrugged. “Maybe, maybe not. Sennheiser makes a ton of headphones. If you’re talking about one of their entry-level models, you probably don’t need our products.”

A sigh from Mr. Retired. “Okay, so what if it was a good Sennheiser? Then you’d start the sales pitch?”

“Nope.”

Mr. Retired shook his head. “But, okay, so … ” he stared at his drink for a while. “If I had these good Sennheisers, then you could tell me why someone would want to buy your products.”

“Yes,” I admitted.

“But he probably wouldn’t,” Ken added.

Mr. Retired sat way back in his chair and studied the three of us with an expression of deep unease, as if we were some kind of interesting mold he'd found on his Ferrari's leather seats.

Finally he shook his head and blew out a long breath. "How do you ever expect to get as big as Beats with an attitude like that? I mean, sure, your salespeople have to make the numbers, but if the top management doesn't care—"

"We don't have salespeople," I said.

"Well—" Mr. Retired blinked a few times, opened and closed his mouth, and sputtered: "Wait. What? No salespeople? How do you move your stuff?"

I grinned. "Most people buy direct online. We also have some international distribution."

"But you have to sell them, right?" Mr. Retired asked.

"Actually, most people sell themselves. Of course, we have Nick to answer questions, but he's prohibited from doing a hard sell, and our policy is never to comment on competitive products."

Mr. Retired's eyes goggled, and he raised his hands alongside his head, as if to cup his exploding brain.

"Are you kidding me?" he asked, addressing the room. "You don't try to sell, uh, anyone at all, and you don't care about selling to Beats owners, the biggest headphone market out there? What's wrong with you?"

I grinned. "Nothing at all. We're self-funded, have zero bank debt, no receivables, are profitable, and continue to be production-limited in terms of growth."

Another series of blinks, as Mr. Retired processed what I'd just said. "Production-limited? What does that mean? That you could make more if you had more money? Then why don't you just get more money? You could grow faster—"

"It's not that simple," I told him. "Many of our higher-end products are constrained by parts availability or lead time. Or simply by our ability to build them at the quality level we want. Our business has been evolving quite a bit over the last five and a half years, and we're still learning."

"Learning? Bring in experts! Hire the best in the business! Take those bank loans! You could

really blow your company up!"

In more ways than one, I thought. "Nope. I'm not interested in what I can make next quarter, I'm interested in building a company that lasts, a company that's fun to be a part of, and a company that some fans simply love."

"But … you could be so much *more*!" Mr. Retired exclaimed.

And that's where I had to shake my head sadly and say nothing. Because to me, his vision would be so much *less*.

Knowing What You Are

Now, let's take a step back. I'm not surprised Schiit is confusing, scary, and incomprehensible to someone who's come from the dog-eat-dog world of the financial markets. His hedge fund would be just as confusing, scary, and incomprehensible to me. He thinks I'm crazy because I'm not in it simply to get as big as we can and get acquired, and I think he's crazy because I can't imagine working for nothing more than the money, nor staring at a long, long stretch of ennui in retirement.

I intend to do fun things until I drop—it's just that some of them may look like work. I understand if a lot of people don't get it.

But, it all comes back to knowing what you are.

Schiit is a natural outgrowth of my personality, and of Mike's personality. Both of us like doing cool audio gear—especially affordable stuff. Neither of us likes doing sales or making deals. Or debt. Or finance. We're lucky because enough of you like our products, which allows us to do more cool, affordable stuff (which is our idea of fun) without getting into the sales and finance side of the business.

If you let us continue to do this forever, I pretty much expect you'll see us doing the same thing as long as we can. We may get bigger, we may get into different markets, but you can expect the same basic principles in everything we do: compelling value, great fun, and unique irreverence.

That's what we are. And we're comfortable with it.

What Are You?

"Okay, that's all well and good, but what can we learn from this?" some of you are asking.

And yeah, I know, some of you are also asking, "Hey, what about them new products, boy, this has been a quiet year so far, what's up with that?"

Aside: You know what, let me answer that second question right here. Yeah, I know. This is a messy way to do it. But it rewards those who read the whole thing. So, here goes, in convenient numbered format:

1. **Yes, this has been a quiet year in terms of product intros.** Largely by design. Hell, our distributors are absolutely ecstatic we're not introducing new stuff. "I can finally catch my breath. Maybe," as one of them said. And we needed a breather, as we digested the production process on some big new products. The line for Ragnarok and Yggdrasil is significantly different than last year, and moves much more efficiently. And yes, I know, we're still in backorder from time to time. Blame the Analog Devices DAC availability largely for that.
2. **At the same time, there's a LOT of stuff going on behind the scenes.** I think I commented that there are "many (redacteds) right now, just nothing to announce. Don't

worry, there'll be some cool and (I think) game-changing products coming in the summer, and before the end of the year. There's probably more than we can release effectively, but that's the way things work out sometime. Don't expect us to flood you with a dozen new products at once; instead expect the hits to keep coming, once they start coming, at regular intervals.

3. **Part of the delay in new products is lead times.** A lot of the new products we're working on use significantly more custom parts—hardware, molded parts, extruded parts, assemblies—and those custom parts have lead times. Sometimes those lead times are longer than we'd like. One very significant new product may slip because the lead times on a custom piece of hardware are 10 to 12 weeks—and we couldn't verify that this trick piece of hardware would work until the prototypes were assembled. Expect more of this in the future, especially as products get more complex.
4. **There are both desktop and "2 channel" products coming this year.** I'm very, very excited for both. Don't be surprised if the 2 channel products start with preamps first,

rather than power amps. Even then, I think you will be floored by what we have in store. Grr, heatsink extrusions are a pain in the rear end, especially when they're, well, *different*.

5. **None of the new products will have a "3" on them, however.** Just in case you're thinking of one of our "2" series products. Like I said before, it is very hard to improve on the "2" series products without radical changes—in effect, making them different products.
6. **I really can't list everything we're working on.** There are literally 14 things on the board right now. It's likely some of them won't become products. And there's no way we'd try to introduce 14 products in the remainder of the year, anyway.
7. **Some of the product intro chapters are already written.** And one is a real saga—a product with more twists and turns than Ragnarok. Seriously.

I think that's about all I can say. At least without getting into a product-announcement-before-the-product-is-ready-complete-with-renderings-and-prices. Which, as you know, we don't do. Anymore.

Okay. So now that that's over, let's talk about what you can learn from being comfortable with

who you are, from a business perspective. No, I haven't forgotten that some of you would like to start—or are starting—your own companies right now.

I personally think you can learn lots of things, but the primary point is: *you don't have to conform to succeed.*

Go back and read that again. Then read it a third time. Then sit back, have a glass of your favorite beverage, and really think about that. In an era where everyone wants to distill the recipe for a zero-hour workweek into a slim novel, and seemingly every business wants the secret to being everything to everyone (while making the founders billionaires when the sellout inevitably occurs), *you don't have to play those games.*

Not if you don't want to, anyway.

Sure, if a pressure-cooker atmosphere, cutthroat environment, making money from nothing more than other money, and the chance for big gains trips your trigger, go ahead and venture into the financial world. Have fun, and I hope you achieve your goals.

But if you're looking for a business where you can have some fun, where your employees actually

like coming to work, and where you can see yourself "playing" at forever, you can do that as well. Simply by being comfortable with what you are.

By being comfortable with what you are:

- You'll have more fun
- You'll want to do more
- Your employees will want to be part of the fun
- Your employees will want to do more
- Your products will stand out in the sea of me-too stuff
- Your products will sell themselves
- Your fans will love the products
- Your fans will spread the word

"Yeah, kum-by-yah and all that, but are there any examples of this other than Schiit?" someone is probably asking.

Yes. But they're thin on the ground. Everyone who starts a company instantly gets "advice" from people like Mr. Retired, and most people who starts a company think that appealing to the broadest possible audience is the best plan, and the majority of people who start a company don't want to give up on a sale, so the cards are stacked against companies like Schiit.

I do have great respect for Taction, which is a company with some unique tactile transducer technology that is ramping up to create their first headphone, the Taction Kannon.

"Ah, that's basshead stuff," somebody here grumbles. "Not *real* audiophile gear."

Yes, and bassheads have often been marginalized by audiophiles. To some audiophiles, bassheads are only about how hard a headphone can hit. They're not discerning of the finer points of treble response. They (gasp) might not even listen to serious, acoustic music!

Which, to me, is a perfect, perfect market.

You know what? Bassheads have fun. They embrace what they like. And they are a market that has been underserved in terms of quality. Original Beats had a lot of bass, but they didn't have a tactile transducer like Taction.

And what is Taction doing with this? Their show graphics prominently feature a shot of Earth from space—being shattered into a million flaming pieces by giant headphones. The headline? "Not everyone can handle a Kannon."

This, my friends, is *knowing what you are.*

And, it's also tapping into a marginalized and underserved market. You know, like the marginalized and underserved audiophile market … as seen from the perspective of a Beats buyer.

Want long-term success? Go ahead. Embrace what you are.

2016, Chapter 7

On Modding, Hubris and Reality

So the discussion has turned to modding, hmm?

Actually, that's a great topic, as we wait for new products to mod. Why do some people like to mod products? Why do we actively discourage it? Are Mike and I perfect, impossible to improve upon?

In short, *I don't know*, *lawyers*, and *no, of course not*, are the answers to the three preceeding questions. But that's a heckuva short chapter, so allow me to be a bit more verbose.

And, allow me to add a bit more structure, starting with the whys and wherefores of modding, at least as I see it.

Modding: The Good and Bad

Yeah. I get it. The temptation to find some easy tweak that unleashes an even more immersive, emotional audio experience can be great. Hell,

I've done it. The Sumo amps I use to this day are modded I removed the input coupling caps, adding 2× the filter capacitance on the main rails, and cranked up the bias on the output stage.

And now, someone thinks they already have me in a corner. "HA!" they cry. "Why didn't Sumo just ship the amps that way? See, I bet you do the same thing at Schiit!"

Not so fast. Sumo didn't ship the amps that way, because of the following reasons:

1. With no input coupling cap, the DC servo, as implemented, acting against the varying DC offset of a preamp (or, worse, a passive preamp) causes whooshing, flatulent noises when you turn the volume pot. I chose to accept this. Most customers wouldn't. Note: a different implementation of the DC servo would not exhibit this problem.
2. 2× the filter capacitance would put a serious dent in the amp cost, especially when they were thin margins to start. We didn't want to raise prices, and, even back then, I knew my big-filter-cap fetish probably didn't make much difference anyway.
3. Turning up the bias turns the amp into a big space heater. Again, I chose to accept this, whereas our customers probably wouldn't—

at least not if the amp was billed as Class AB, rather than Class A.

At Schiit, things are different—we have no amps with input coupling caps, and if we use servos, they are designed not to have the same problem as the Sumo amps. We go crazy with filter capacitance in general, because the amps are smaller, and we have more cost to work with, since we're operating as direct sale. And we're not shy about making amps that run hot, since we've done all the calculations on capacitor life (and device derating) and know the temperatures they run at are not a problem.

Fun fact: the vast majority of the bleating you hear about "heat is the enemy of electronics" comes from one component: capacitors. Electrolytic capacitors are rated at some modest lifetime, at some specified temperature. For example: *3000 hours at* 85 °C. This causes some people to freak out and go, "Whoa, only 3000 hours, hell, I run 3000 hours a year, it's gonna fall apart/explode/kill my cat!" Wrong. Read the rating again. Then go here: illinoiscapacitor.com/tech-center/life-calculators.aspx, plug in the datasheet numbers, and you'll get numbers like 65 456 hours at 45 °C

(or, a typical capacitor temperature in an Asgard 2). That's 7.5 years of being continuously on.

But I digress. To get back on point, I get it. A few small tweaks for much better sonic performance … what's not to like?

And there are entire communities which support modding, and encourage modders to go farther. I've mentioned diyaudio.com before, and I'll mention it again. Many people come there as modders, and end up making their own products from scratch—or even designing them.

And, honestly, I think communities like diyaudio are where some of the most interesting ideas first see the light of day. The rapid feedback, iteration, simulation, building, and testing of audio stuff makes the corporate world look pretty staid and safe in comparison. Things can happen faster there. Designs can be worked out faster there. Commercial products have come out of the community. Some really, really interesting ideas are being explored. Hell, Nelson Pass is an active member.

Bottom line: if you want to learn what makes stuff tick, and what you can do to make the products

better, I recommend you take up residence at diyaudio.com.

In the audience, someone is nodding knowingly, but wearing a cynical smile. "Well, if you're so cool with modding, how come you go out of your way to discourage it on your products?" they ask.

Good question. I already gave you the short answer: lawyers. But you deserve a longer answer. So here you go:

1. **Modding doesn't always make things better.** One of the most popular mods (mentioned on the Vali 2 thread) is op-amp swapping. This is not generally a good idea, unless you know what you're doing and have some basic test equipment to verify your exotic op-amp is not oscillating. (And, since most of the op-amps in Schiit products are simply DC servos, it won't really do anything, anyway.)
2. **Effective modding frequently requires engineering knowledge.** Beyond test equipment as mentioned above, engineering knowledge is important for many truly effective mods. Using the Vali 2 example above, the single bestest mod you could do for it would be to swap the PNP transistor that's convolved with the tube for one with higher gain and better linearity. Except, well, I already picked the

highest-gain transistor with the best linearity, at least among available devices. If something else shows up, well then, maybe it's time for a running change (see below.) Beyond that, did you know there's an optimal bias point for BJT outputs, but turning up the bias on MOSFETs only makes them more linear? Did you know that op-amp is just a DC servo or part of the protection system … or not even in the circuit? Did you know that capacitor wasn't even part of the signal path, so installing a fancy one will make no difference? Did you know why the designer chose the components he or she did? Hint: it's not always about cost.

3. **Modding can hurt you.** A decent amount of our gear has rail voltages that range from 200 V to 250 V. Poking around op-amp circuits that might have ±15 V on it is one thing. Prodding high rail voltages is not a great idea, because shocks of that magnitude can hurt you pretty bad.
4. **Modding will piss you off when we won't service the product.** Any modification immediately voids the warranty. And, any modified product we receive for service has to be returned to stock before servicing (and we do get a fair bit of it). Depending on the mods, this can be very costly.

5. **Encouraging modding increases the chances we get sued into the ground.** There's a phrase screened on the back of every one of our products. It's called "no user serviceable parts inside." This means, "don't open this box." If we ever say "hey, sure, you can open it up to change a fuse, swap caps, whatever," then we have invalidated that phrase. So, in the Litigious States of America, when someone decides to mod their Valhalla 2 … when it's plugged in … while they are taking a bath … (don't laugh) … and their estate throws a lawsuit our way, one of the things their data forensics guys will be looking for is if we have ever implied that opening up our product and working on it was a good idea. If they find one email that invalidates the "no user serviceable parts inside" phrase, we are much more likely to be screwed.

Aside: Some may see #5 above as a distasteful corporate CYA. I see it as part of our corporate responsibility to ensure we are around if you ever need service. A company killed by a lawsuit can't sell you matching product, give you advice on your current products, or fix your broken gear. So excuse me if I draw a hard line on that one.

So, Are Designers Perfect?

Of course not. If we were, there wouldn't be Series 2 products, or Ubers, or Multibits, or Gen 3. Simple as that.

Here's the reality: *everything* is a learning process.

There are a million ways to design a product, so the act of design is one of exploration. I've written several chapters on how products twist and turn and change on their way to production. We start down one path, find that it dead-ends, and try another. We change and tweak and measure and iterate. In the end, we hope to end up with a product you'll enjoy. But there ain't no possibility that it's perfect.

Producing a product teaches us even more. Sometimes catastrophically, as when a first run has to be scrapped. We've been there. But usually it's more gradual. There's a better way to put it together. Customers have commented on this feature, or that annoyance. Parts come and go—good parts reach end of life, and better parts appear. That's why it's not unusual to have a couple of small running changes in a product over the course of its lifetime. These small changes need to be done for a good reason (such as parts going obsolete), and they have to be documented.

Because, in the future, products will need to be serviced. Having a hundred variations isn't conducive to serviceability.

Going to a next-generation product should be more than just a collection of small, running changes—it should be a significant re-design that seriously augments the product. Like adding functionality (gain switching, preamp outputs) or a complete re-think (tube or solid-state modularity) or even radical new directions. The next generation should be informed by the current generation, by customer feedback, by market realities … but again, there ain't no way they'll be perfect.

So, yeah. Lots of lessons. Good designers take note of these lessons and use them to improve their products.

The smug dude is back there again, arms crossed. "But if they're not perfect, that means the products can be improved. Hence, modding."

Ahem. Not so fast. Add the word "easily" between "can" and "be" and the phrase doesn't necessarily ring true. *Anything* can be improved. But if it takes an entire redesign, that ain't *easy* … and that ain't modding.

Can Designs Be *Easily* Improved?

That's really the question, isn't it? And I'd like to say, "No, not with well-designed products."

But my opinion means nothing to someone who wants to swap a fuse or a socketed op-amp or attach magical damping dots to the components. They're going to do that anyway, because they're relatively easy. When you get down to swapping coupling caps, playing with bias, and adding bypassing, you're in the hot-soldering-iron crowd, which is much more hard-core—what they consider easy may seem very daunting to many modders.

And, in reality, such a quick, facile, and dismissive answer doesn't really serve you very well, does it? So I thought I'd do something kinda silly … go through our products and discuss modding possibilities in a stream-of-consciousness manner.

Now, this doesn't necessarily include all products, nor is it likely to be 100 % inclusive on what can be done to them. Nor is it an endorsement of hacking up our products. Remember, *No User Serviceable Parts Inside.*

Magni 2 and Uber. Ha. We've already done the mods on Magni 2 Uber, or at least the ones

that matter. The Uber has a better gain stage—specifically, a complementary-driven VAS—than the Magni 2. Yes, this is well beyond a mod you'd typically see. We also added more filter capacitance, preamp outs—a functional mod—and some cosmetics. Why didn't these mods make it into Magni 2? Because they increase cost. What about op-amps? Well, the only op-amp in Magni 2 and Uber is for the DC servo, so swapping it out should net you zero sonic gain. Also, it's surface-mount, so it's a pain to swap. How about coupling caps? There aren't any—not input, not output, not interstage. How about a linear supply? No, Magni 2 and Uber use a linear supply with an AC-input wall-wart. So unless you're going to reengineer the gain stage, there's not much to get excited about in terms of easy mods.

Asgard 2 does use an interstage coupling capacitor. Would swapping it out improve performance? Maybe. We chose the one we use for sonics, size, and price. Adding a pair of $85 capacitors to it might make it better, but then you're talking about a much more expensive amp. Turn up the bias? It's already Class A, and it runs very hot. Not a good idea—and not adjustable unless you swap components. You'd better know what

you're doing there. Oh, and the op-amp? It's just a DC servo.

Valhalla 2 uses both an interstage coupling capacitor and output capacitor. This is typical of tube amps. Good luck getting anything better in the space available, though. If you do, could it make things better? I don't know. This one is much more space-constrained. Re-engineer it to eliminate the output capacitors? Sure, that's possible, but that's a whole new power supply, down to the transformers—and a more expensive amp. Personally, I'd leave it alone and think about maybe some fancy input tubes (WE417s or 5670s with a pinout adapter?)

Lyr 2 also uses an interstage coupling capacitor. Same caveats as on Asgard 2 and Valhalla 2. Lyr is pretty packed. The op-amps? DC servo, again. Don't bother swapping it. Turning up the bias? Only if you know what you're doing, and what you're measuring … and you have a fan. Lyr 2 runs hot. I'd also look at those WE417s or 5670s with an adapter.

Modi 2 and Uber. Yes, these use op-amp gain stages, so you'd think most people would be all over that. However, you'd better know what you're doing: Modi 2 uses a rail-to-rail op-amp

(and requires one). Both use surface-mount components, so there are no easy swaps. Personally, I'd do the Scotch Mod and call it a day (a couple of glasses, things sound much better …)

Bifrost Multibit. Yeah, yeah, lots of people would like to have a discrete gain stage. I've explained why that isn't happening. Cliff's notes edition: space. So why not swap the op-amps? Well, because they are extremely specific to the design, and they're also surface-mount. Yes, I know, we are no fun at all. Personally, I'd be happy the Bifrost is modular and upgradable, so when we figure out some new things, I don't have to throw the product away. Now, don't go hoping for a discrete stage … again, notgonnahappenland.com. Because space.

So what do we have, from this partial survey of our products? A whole lot of stuff that doesn't have much in the way of easy mods. Sure, they could be made better … by re-engineering the product, and/or changing the retail price significantly. But those aren't mods. Or at least they aren't mods that will make people very happy.

"Wait a sec!" Someone is asking. "What about the expensive products? You skipped over them entirely!"

Yeah, and Wyrd, and Mani, and SYS (fun fact: someone recently asked if we could put an RK27112 pot in a SYS—those of you who know the measurements of this pot are laughing real hard right now). Don't panic. There's no big conspiracy. Like I said, this is about representative products, not everything …

… no, you know what? Okay. Let's talk mods on Ragnarok and Yggdrasil. First, get good at firmware for PIC 24F microcontrollers. In the case of Yggy, learn to program Analog Devices SHARC DSPs. Then consider: no coupling caps, no DC servo, insane amounts of filter capacitance, heroic power supplies (including a shunt-regulated, choke-input analog supply on Yggy). And that's only the start. That's why modding, in any meaningful way, is going to be very difficult on the pricey products. Sit back. Relax. Have a drink.

And remind yourself: Yggy is modular.

So What Do I Do If I Want To Mod?

If you want to mod, nothing we say is going to stop you, right? So I think, in general, my advice would be to spend some time trying to understand what you're modding, how it could

be made better, and if modding it has a good chance of success.

Yes, I know. This takes technical knowledge. It's not as simple as just swapping parts. But it does give you a much better chance of being happy with your mods.

Now, this is still gonna be problematic in the case of our products, since we don't provide schematics, don't discuss possible mods, don't provide technical support for them, don't get into the deep whys and wherefores of our designs, and don't provide warranty service for modified products.

So, here's another idea: build something yourself. DIYaudio has many well-documented community designs that have schematics, layouts, even PC boards and parts lists. There's also a bunch of people who have already modded the product and weighed in. You could build yourself a Pass project, or a CFA amplifier, or a bunch of other things. And you can modify it to your heart's content.

Or, choose a popular kit amplifier and build that. They're usually very well documented, and easy to modify.

In both cases, you'll be learning a whole lot about electronics … including what mods matter, and why. Which (call me crazy), I think is a whole lot more interesting than swapping parts randomly, and hoping for the best.

Happy learning … building … and modding!

2016, Chapter 8

A Perspective On This Moment in Digital

Or, what we have our eye on, as things keep changing.

Someone new to digital might ask, "Why this subject? Can't I get everything I need to know from the various forums and press out there?"

To be brief: *maybe*. If you're looking for something to buy this second, sure, there are a ton of hypemeisters … er, I mean, seasoned and rational press and bloggers out there. There are a ton of people who'll try to convince you that (insert buzzword here) is absolutely critical to enjoying digital audio, and there's no way you should be without it.

But if you're looking for some perspective on what might be a real change in digital audio, rather than just another flash in the pan, you may want to look a little deeper.

And so, since the world has changed since the last time I dedicated a chapter to digital audio (in those ancient days when DSD was the rage), I figured it might be time to look at what's happening, from streamers to formats, and give you our perspective on it.

And yeah, this is *our* perspective. It's a perspective that comes from being in digital audio literally from the start. Read Baldr's blog for some really, really early stuff. Mike's seen the rise of digital recording (prior to CD), the rise of CD, the rise and fall of DAT, DCC, HDCD, and SACD, the beginnings of high-res, the start of portable players, the first beginnings of computer audio, the ongoing domination of streaming, the rise and not-going-much-anywhere of DSD, and now, the emergence of MQA. And he's been involved in quite a bit of the above on the design level. I've seen pretty much all of the above except the rise of digital recording, since CDs were my first "serious" medium.

And, if that above paragraph isn't a great illustration of why many of our DACs are upgradable, I don't know what is.

In fact, let's recap:

- Digital Audio Pre-Physical Media Era (late 1970s-on)
  - Digital recording prior to LP pressing
- Digital Audio Physical Media Era (1982-on)
  - Compact Disc
  - Digital Audio Tape
  - HDCD
  - Digital Compact Cassette (compressed)
  - MiniDisc (compressed)
  - SACD
  - Blu-ray Audio
  - CD Computer Audio
- Digital Audio File Era (Late 1990s-on)
  - iPods/DAPs
  - MP3
  - Lossless
  - High Res (24/96, 24/192, etc)
  - USB Computer Audio
  - DSD
- Digital Audio Streaming Era (2005-ish on)
  - Spotify
  - Tidal
  - MQA

And yeah, it's a messy sum-up, because I've merged formats and delivery methods and stuff like that into a whole bunch of bullet points. But note the four eras. Most listeners are squarely in the streaming era, but many of us krazee

audiophiles are still using files.

> **Aside:** I certainly am, most of them 16/44.1 CD rips. And Tidal.

But yes, that's a lot of change. And so, I thought, why not take a look at the things that are changing, what we're paying attention to, what we're playing with, and where it might take us.

With a side order of perspective.

Change the First: Computer Audio, and the USB Decrapification Wars

In the 18 months or so that have passed since I last bemoaned the state of computer audio and the USB interface, how much has changed? Is it better or worse? What's going to shape computer audio going forward?

In short, not much and quite a bit, no and yes, and the USB interface.

To elaborate on change, things look about the same on the surface. You can still buy a computer and use it as a digital audio source to feed a DAC, connecting typically via USB.

But:

- **Operating systems have changed.** Whether you're using Mac, PC, or Linux, you're probably looking at a new OS. And the bad news is that the OSes are worse than ever when it comes to streaming audio.
  - On the Mac side, El Crapitan has earned the loathing of many audiophiles that found problems with audio dropouts via USB … all the way up to the .4 release.
  - On the PC side, Windows 10 tries much harder to make all this USB audio stuff automatic … while sometimes installing the wrong drivers, or having even more draconian power management that makes USB unusable. Some systems even re-write the registry to re-enable power management after you've edited it to turn off power management! And if you think you're safe because you're on Windows 7 or 8, be careful, because Windows gargles donkey balls: pcworld.com/article/3073457/windows/how-microsofts-nasty-new-windows-10-pop-up-tricks-you-into-upgrading.html
  - On the Linux side, things are a lot better, but even then, early "modern" Linux variants could sometimes cause fits.
- **USB itself has changed.** As mentioned, USB port power management has only gotten

worse on the Windows side. On the Mac side, things go in phases. Most new Macs don't seem to have a problem. On Linux, most users are sophisticated enough to deal with it. USB is, unfortunately, not 100 % plug and play, thanks to systems that throttle power to the USB port (to save battery life, or be "green.") We would have exactly zero problems if USB actually met its spec. And more:

- The USB decrapification revolution has become a flood. Holy heck, I don't know if we would have ever introduced the Wyrd if we knew it would cause such a flood of decrapifiers on the market. Now, there's at least a half-dozen devices (hilariously, all more expensive than Wyrd) vying for the crown of Decrapifier du Jour. Some people are using 2, 3, or even more of these devices in series … and reporting sonic differences. My only comment is that with a half-dozen devices, some used multiple times in a single system, how do you know when to stop? Or is it better never to start?
- USB 3.1 changes everything … eventually. USB 3.1 is making its way onto computers as we speak. It has the potential to change literally everything, with the ability to deliver up to 20 V at 5 A (100 W) of

power. Now, of course, that will probably be throttled, especially since a constant 100 W draw would deplete a laptop battery in short order. But, with that kind of power available, you could think about serious, serious headphone amps (and even speaker amps) being powered via USB 3.1.

- **Wireless connectivity is now a much more viable option.** No, not Bluetooth. Leave that off. It's not lossless, and until it's lossless, we will spend exactly zero time on it. But if you're talking WiFi, then yes, it's quite possible to send your computer's output to a streamer—either commercial or DIY—and lose the wires. Some of these, especially the DIY options, are pretty "beta-like," and some might have you wishing for the comfort of good old reliable USB.

So what do we think is important about all of this? Well, to sum up, "OSes will continue to suck, USB will continue to be throttled (on some systems) and this wireless thing is pretty interesting."

So what are we doing?

OSes? There's nothing we can do about OSes, other than to make sure Microsoft has the right driver for our devices, and hope it detects and

installs correctly. It seems to be better these days, but we still have problems from time to time.

USB? As far as USB goes, well, we have to look at ways to make sure that the maximum number of systems can simply "plug and go" with our DACs. Although it's convenient to blame the computer manufacturers for throttling USB power, the reality is, nobody cares when they just got their shiny new DAC and it doesn't work. So, expect to see internally-powered USB inputs in at least some of our products—that is, USB that does not draw power from the USB bus in order to work.

Aside: "So why didn't you dumbasses do this from the start?" some are asking. Well, for two reasons. One, not all of our DACs can have self-powered USB. Consider Modi and Fulla, which are USB-powered. Two, it has been our experience that keeping the internal power away from the USB input has sounded better. However, as we continue experimenting, this may change.

So what about an uber-Wyrd to crush the other decrapifiers? Well, uh, probably no. Because Wyrd is still the least expensive decrapifier, and we think it succeeds admirably at its primary job.

That's never to say never, of course, but it's not like we've felt a need to talk about such a thing.

WiFi Streaming? Well, don't expect to see a wireless device from us soon. We're playing with some of the DIY options, and some of them are very good. However, they're also all a bit buggy. They're certainly nothing we'd want to make the investment in to develop and maintain. But they are a good option for people who want to send music losslessly from a remote computer, or pull music from a network drive through a tablet interface.

Wireless streamers are definitely one of the hottest areas in audio right now … and I'm sure we'll end up seeing more of these, and less traditional computers, as time goes on.

And with that, let's devote a whole section to them …

Change the Second: WiFi Streamers, and the DIY Revolution

As mentioned above, it seems that almost every day, some new kind of small WiFi audio streaming interface is being written about, talked about, reviewed, or introduced.

And yes, there are some commercial WiFi streamers—products that connect to your DAC and allow you to send music from your computer wirelessly, or play your music from a network-connected drive. But I think it's much more interesting what's happening on the DIY side, specifically amongst the small Linux computers out there today.

"Computers?" you might ask. "I thought you said this was a streamer, not a computer."

Well, here's the thing. All WiFi streamers are computers, even the commercial ones. They pretty much have to be. After all, how else would you get it onto your wireless network? It's not like Bluetooth, with its relatively simple pairing.

So, here's a recipe for a typical WiFi streamer:

- One small Linux-capable computer (Raspberry Pi, CuBox, Beaglebone, etc)
- One audio OS (Volumio, Rune Audio, I'm sure there are others …)
- Some kinda storage (usually an SD card or micro SD card)
- Some kinda power supply
- If you want to get fancy, some kinda case

Get your credit-card computer, download the OS, stick it on an SD card, install it on the computer,

set it up on your WiFi network, plug it into your DAC, and control it via its web interface from your tablet or phone. Done.

Total cost? Under $100, easy.

And you can add fancy audio stuff, if you'd like. Add a SPDIF output. Or add a whole DAC, if you don't already have one.

The point is, WiFi streaming is really taking off, and if you don't feel like having a computer sitting atop your audio system, it's not difficult to have a setup where it doesn't matter where your music is, and you can control everything via a tablet or phone. With the low cost of tablets these days, you could easily have a dedicated system for your listening room.

So why aren't we doing our own streamer? Lots of reasons:

1. There are already commercial streamers out there for relatively reasonable prices.
2. There are already DIY streamers out there for insanely cheap prices.
3. Anything we do would have do add something to the space, not just be a Pi in a box with Volumio.
4. It would require too much of our resources to develop and support it.

So, sorry, guys … a streamer really isn't in the cards. Well, unless we can think of something very, very compelling. What could we do that isn't already being done? Can we make it reliable enough not to crush us in support? Can it be made at a price point lower than the other commercial offerings, but add enough value that it's worth more than DIY? Lots of questions here.

Change the Third: Phones and DAPs and Tablets Oh My (USB Power Strikes Again)

One thing that hasn't changed much is how many people want to use portable devices as sources. A decent percentage of our inquiries revolve around the question of "Hey, can I use my phone (or tablet) as a source?" Or, "Hey, is (insert name of DAP here) a good source?"

The answer, however, has changed a bit. Back when I last wrote, Android was just getting support for USB Audio 2.0. Now, it's pretty much there, if you're running Android 5 and above. (I say "pretty much there," because it seems there are some device-specific glitches.) So, Android has the potential to be a good, plug-and-play source for USB-input DACs.

But …

you knew this was coming

But … USB power strikes again. In this case, USB power management is more important, since your phone's battery life can be significantly affected by the 150 mA to 250 mA power draw from a typical USB audio receiver.

This is why Apple devices automatically pop up the "this device draws too much current" warning when you plug in most USB-powered DACs. And then, of course, in typical Apple fashion, it simply shuts them down and doesn't allow the DAC to work at all.

Some Android devices will supply that amount of power, and battery life will suffer. Some will not. We can't test every Android phone on the planet, so we only know that you may get lucky, or you may not, or you may get lucky depending on how much battery you have left.

"Wow, sounds like self-powered USB inputs are the ticket here, too," you say.

Well, maybe. Let's look at this a bit more.

Self-powered USB and non-portable devices. Yep, you bet. Great idea. Plug into your Apple

or Android device and everything's cool. Except, well, you're chained to your desk again. Which may be fine if you're planning on using a tablet as a source. But it's less than ideal if you want to be truly portable and use your phone.

Self-powered USB and portable devices. Typical USB audio receivers draw 150 mA to 250 mA. That means you're looking at a 1500 mA h to 2500 mA h battery … just to run the receiver for 10 hours. Well, more like 8 hours, when you start getting into regulator losses, battery variation, etc. That's a pretty dang big battery just to run the USB audio receiver … and you still haven't provided power for the DAC or headphone amp! This is why you're looking at big fat batteries … or long charge times … or both … on a lot of portable devices with self-powered USB ports.

But there's a better answer. This one just showed up, pioneered by Audioquest:

Low-power USB audio receiver. Audioquest didn't use a commercial chip from C-Media or XMOS. They developed their own code for a low-power general purpose microprocessor, and created a USB input receiver that (IIRC) draws only about 25 mA. This means that they can plug

and play into Apple stuff, and maximize battery life on all devices. To us, this seems like the way to go. But don't get all excited—this takes a lot of investigation. Bottom line, we're looking really hard at USB and how to make it better and more convenient across the line.

"Okay, so we know that USB is a pain, and that you're paying attention to it," someone asks. "What else are you working on? When do we get our portable DAC/amp? When do we see a DAP?"

I've answered these questions before, and I'm afraid not much has changed since we last talked about digital. A portable DAC/amp would require us coming up with something truly market-shaking, and something we really, really like. I'm not saying it won't happen, but we haven't come up with anything that doesn't have significant downsides (like being HUGE or being "just another portable DAC/amp"). A DAP is just out of the question. There's no way we're going to spend engineering resources recreating player software, music management, and storage that have been amply worked to death in phones.

But yeah, we get it. Not everyone listens at home. Sometimes you have to be portable. So we keep

looking at the problem … .

Change the Fourth: Formats, and the Definition of Insanity

Deep breath. Okay, here's the big change.

- First, the definition of insanity: doing the same thing over and over again and expecting a different result.
- Second, this is where the Digital Audio Eras outlined above come into play.
- Third, yes, it's still insane, even in a different era.

So here we go. Audiophilia is making the transition from the File Era to the Streaming Era. This is a huge change. It's the first era where you may end up owning, well … exactly ZERO music. After all, if it's on-demand, all-you-can-eat, and cheap-per-month, why buy anything at all?

Well, before Tidal, there was certainly a reason: if you wanted uncompressed music, you needed to buy files or physical media. Now, with Tidal, 16/44.1 uncompressed is on the RIAA Roundup's All-You-Can-Eat Menu. At $20 per month. About the cost of a CD or two. And, let me tell you, it's great! It's amazing to be able to browse through

a library bigger than anything I'd ever have at home, choose something completely uncanny, and click Play.

But I *still* have files.

Why? Because many hotel internet connections aren't fast enough to stream Tidal. Most in-air internet, same problem. In-car? Not a chance unless you're in the latest LTE environment … and lucky.

Hell, sometimes *at home* on a 100 Mbit/s line, Tidal chokes. This makes you reminiscent for the days of CDs and files really, really fast.

And let's talk Mike. Mike lives about 20 miles away from the Schiit office, and his internet options are exactly two:

- Dial-up
- Shared RF

So his data rates make Tidal a complete no-go.

But, you know what? Data rates will get faster. Hotels will pull their heads out of their butts. LTE coverage will be better and faster. And high-speed internet will eventually reach the hinterland. There's no question that streaming is the long-term end-game. Heck, we already see it in the mainstream.

So what does this mean to audiophiles? Well, if the story ended at "Tidal offers 16/44.1 uncompressed streaming of XX million tracks," then it would be pretty much all good.

Yes, you can say, "But I want high-res." And that's fine, you can buy tracks from HDTracks or Pono Music or whatever. But high-res streaming needs even more bandwidth, which means that it's not going to happen—today, anyway—on Tidal.

Or you can say, "But I want DSD." And that's fine, you can buy tracks from the guys who sell DSD. But DSD streaming needs even more bandwidth, which means it's not going to happen—today, anyway, on Tidal.

But you do end up with a solid win-win. A Tidal subscription in a high-bandwidth world would be enough for most audiophiles. Audiophiles that want high-res or DSD can simply purchase those files and enjoy them. And anyone who wants to own 16/44.1 could just buy CDs and rip them. Everyone's happy.

But no. Now we have to do this insanity all over again. The biggest change: enter MQA.

If you're to believe (most of) the press, MQA is the greatest thing since sliced bread, an amazing new format that promises an even-more-crystal-clear-window-on-the-artist's-intent. It fixes everything that's wrong with digital! It's a watershed moment for audiophiles!

The implication: *everyone needs to get on board, posthaste! All DACs need to support it! Everyone should be gearing up for the revolution!*

But wait. Didn't we hear this, not too long ago? Kinda sounds like DSD, right? Because thousands of DSD titles were just around the corner (I mean, hey, DSD was a Sony thing, and they had, what, how many recordings they could release in DSD?) But 3+ years into the DSD revolution, all we have are a handful of recordings. Inquiries about DSD have dropped to nil on our side of the fence. Stick a fork in it, it's done. Call it the last format of the File Era.

Now we're barreling straight into the Streaming Era … and MQA is hailed as the answer to shoving high-res audio down a 16/44.1-sized stream, heralding the musical revolution for all devoted audiophiles.

Hence the definition of insanity. Here we go again.

"Well, but, MQA is totally different from DSD, and SACD, and HDCD, and all the other formats that wanted you to re-buy all of your stuff, because it doesn't have a "D" in it at all, and it makes high-res streamable," some might say. "That means you just gotta subscribe to Tidal, which (has promised to have at some unspecified time) MQA."

Oh, okay. So you don't have to re-buy it. You just have to subscribe to it. Yeah, in a way, this makes sense for the streaming era.

But even the assertion that MQA is the "easiest" new format to acquire has a lot of questions wrapped around it:

1. When will Tidal offer MQA?
2. How much will MQA cost? Still $20 a month? Or something else?
3. How much of Tidal will be MQA? Just a small percentage, or all of it? A lot of the heavy breathing has orbited around the idea of "all of it."
4. How much are you going to spend on an MQA-enabled DAC? Should this be factored into the reaquisition cost? Hint: of course.

But wait! Now we hear that Warner Music has signed up with MQA. Oh my goodness, the vaults are gonna open wide! Surely this is a sign!

Well, no. This means there are even more questions:

1. How many titles will Warner Music release in the next year or so? Thousands, or a dozen? Remember, DSD was a Sony thing (as in, they paid NO royalties on it), and they had big vaults, and very little happened on the DSD front.
2. Will Warner Music be OK with Tidal streaming their MQA stuff on their all-you-can-eat menu?
3. Will the releases be remastered, and therefore not directly comparable to the old titles?

Remember, an LOI (letter of intent) doesn't cost much. Reissuing a big library … when it might be available on streaming … hmm …

So, to me, it still looks like insanity. Because isn't this what they're saying?

"Hey, just turn over the entire industry to us, and we'll make it all good."

A little extreme? Consider that MQA wants:

- Licensing fees from the recording industry
- Licensing fees from the digital audio product manufacturers
- Hardware access into the DAC or player of your choice (because no software player)

- Subscription fees from every listener via Tidal, and/or
- re-buying a bunch of stuff re-released by the recording industry

If this isn't turning over the whole industry, I don't know what is.

And, you know what? Turning over the whole industry to accomplish a worthy goal might not be so bad. But it seems that MQA's technical side generates more questions for each one that is answered (hat tip to Michael Lavorgna of Audiostream.) And, well, the measurements don't seem to be very, well, high res:

archimago.blogspot.com/2016/01/measurements-mqa-master-quality.html

And that's the big change in formats. Instead of everyone asking "when're you gonna support DSD," it's changed to "when're you gonna support MQA?" With the same level of hyperventilating and implied-"I ain't gonna buy unless you absolutely guarantee you'll support this format, even though there ain't really no content for it yet."

Well, excuse us if we (again) sit on our hands.

Let's wait a bit, and see how this shakes out. If the entire Tidal library is MQA'd (at a reasonable

price) within a year, and if Warner releases a few hundred good titles on MQA, then, hmm, maybe it's going somewhere. And you guys can dig up this chapter from last year and say, "ha, you were wrong!"

But if it's 3 % of Tidal's library at $ 40 per month and Warner has 7 great titles out within a year, then yeah, we'll continue …

… to best support the 99.9 % of music out there in 16/44.1 PCM.

2016, Chapter 9

The Elephant in the Room

Let's talk about the elephant in the room.

Yes, it's an old saying. But it's a good one. It conjures up images of snooty, tuxedo-clad partygoers at a high-class soiree, all of them discussing the latest cause celebre or achievement du jour—you know, like the Prodigal Son is Going to Harvard, or the Doctor Daughter was just recognized by the AMA for her contributions to Medicine, things like that—and all the while a big, stinky, gray, wrinkly pachyderm tramps over the beautiful carpet, knocks over the canapé table, slurps down the entirety of the organic punch from the heirloom crystal bowl, and takes a steamy dump in the middle of the dancefloor.

And all this time, the Beautiful People carefully manage to never quite look at the elephant, their gaze darting past as if it didn't exist. And their conversation never mentions the beast, ever, even as their voices have to rise above its periodic bellows. The elephant will not interrupt the Oh

So Perfect Party. They will simply will it away, unsee it, retconn it from existence.

Until someone new comes into the room (someone who probably is wearing jeans and a fart-joke t-shirt), points at the elephant, and says, "What the hell is that thing doing here?"

And that, my friends, is a pretty good description of where high-end audio is today. The Grand Old Guard throws an elegant party—but they don't want to talk about the beast that's stomping up their dance floor.

The name of the beast? *Price.*

As in, big dollars. As in, bring a giant wallet. As in, *too much for you, pal.* As in, *if you have to ask, you can't afford it.* As in, *if you think the prices are high this year, wait till next year.*

Okay, I'll say it.

High end audio pricing has gotten completely and utterly out of hand.

Yes, this is a response to my brief visit to the recent TheShow Newport. This is the show where several exhibitors commented, "Wow, you're room is always busy. What are you doing?" This is also the show where, down the hall, an exhibitor used

the phrases “stunning value,” and “a real bargain” when describing a $ 20 000 preamplifier—without any hint of irony. This is the show where many visitors to the Schiit/Salk room wanted to add a zero onto the end of all of our prices. Like this:

“How much is that integrated amp?” asks the visitor.

“The Ragnarok? $ 1699.”

The visitor nods. “Seventeen thousand … ”

“Nope. Seventeen hundred.”

A blink and a blank look from the visitor. “Ah, what? Seventeen thous—“

“Seventeen hundred.”

The visitor looks around the room, like they’re being snookered. “Wait a second. How much are the speakers?”

“Twenty-eight ninety-five.” We correct him.

“Twenty-eight, ah, twenty-nine thousand?”

“No, hundred.” We sigh.

The visitor’s brow draws down in deep furrows. “So how much is this whole system?”

“Including cables, about $ 8600.”

“Eighty-six, ah, hundred—as in, under $ 10 000?” Asks the visitor, his voice small.

“Righto!”

And that’s when we want to explain that the electronics should cost less than the speakers, because you really should put most of your money in the transducers, but the visitor has tuned out. Apparently, it’s quite a shock to find a whole system for less than half the price of that “stunning value, bargain” preamp.

Now, we weren’t the only people doing affordable systems, but we were one of the few. I walked around a bit. The cheapies were all pretty much packed. Elac was standing-room-only, you couldn’t even really get in for a demo. Some of the more esoteric (read: weird-looking) high-priced rooms had good traffic. But the ones with more normal-looking gear and $ 50 000 monkey coffins were pretty slow.

I left the show in a funk. The insane pricing was so pervasive that it colored visitors’ perception of our room—to the point where they couldn’t fathom a whole system that cost less than $ 10 000. And, despite the clear signpost of “big traffic in cheap

rooms, small traffic in expensive rooms," nobody seemed to be putting two and two together.

Because, after all, if you have an elephant at a fancy shindig, it's gonna eventually break the dancefloor, knock over all the tables, and scare the guests away.

Or, again, I'll say it:

High prices will destroy the high-end audio market.

High prices mean no new audiophiles—they simply can't afford it. High prices mean the entirety of high end has what used to be called the "Buick Disease," where you wonder if each purchase is that customer's last. High prices mean that high-end detractors have no end of stuff to make fun of—and scare away the high-end curious. High prices mean you're catering to a high-maintenance group, which may eat you alive in customer support … or dictate all your future product directions.

Or at least that's how I saw it.

And it got me wondering—am I just overreacting? Are the high end's high prices just an illusion, brought on by the lower-every-year purchasing power of the dollar. Or are the high prices somehow justifiable, as we push forward, closer and

closer to accurately reproducing a live performance?

Good questions. So let's have a look.

History Says: These Prices Are Cray-zee!

Half a decade ago, I had a conversation with a person who wanted to start a company to make "the most expensive loudspeaker cables in the world." Yes. That was the selling point.

To be clear: we're talking $ 20 000 to $ 30 000 *speaker cables*.

And yes, to me, it sounded completely insane. But he brought the data that backed up his business plan. Namely, that there were over 100 different models of loudspeakers on the market that cost over $ 100 000. Given those prices, there was an opportunity to sell uber-priced cables into a certain percentage of installations … and that would create a very nice business.

That company ended up never happening, but that fact stuck with me: *100 different speaker models over $ 100k*.

Guys, this is like having 100 different supercars priced at over $ 1 million. Or, hell, even that

might be more sensible. Or, hell, let's just say it: it's completely ******* insane.

Back when I started down this road, in the dim dark days of 1990, there was exactly one speaker model at the top of the pricing heap. Everyone knew what it was. It was the Infinity IRS V, and it cost $ 50 000.

Or, with inflation, about $ 91 500 today.

Not even 6 figures? Hell, they're not even trying.

And to be clear: *this was the most expensive speaker system out there, by a long shot, period.* So beyond the pale that it wasn't really a consideration; very, very few audiophiles could even think about buying them. Especially in an environment where home equity loans were crazy talk, in the real estate downturn of the time.

So let's talk more-realistic-crazy, 1990 style. How about a pair of full-range electrostatics? Specifically, Martin Logan CLSes. They were $ 2500.

Yes, $ 2500. As in, about $ 4600 today.

Or how about the Carver Amazing Loudspeaker—four 12 inch drivers per side, plus a 6-foot-tall ribbon. (Yes, *six feet*.) Those were $ 3000. Or about $ 5500 today.

And, let's be clear. There was no internet direct in 1990. Those are prices for products sold through dealers, with a traditional 50 % dealer markup. Sold direct, those products would cost a lot less.

"Wait a sec," you're saying. "You mean, you could get super-exotic speaker designs like that, for, like *$ 5k?* Through a *dealer?* What the hell happened?"

Exactly. *What the hell happened?*

So how about electronics? Surely there were uber-expensive electronics at the time, right?

Well, yes and no. Adcom was what was considered to be entry-level high-end. They sold a 2×200 W PC amp for $ 799. Sumo was considered to be a step up, our 2×240 W PC amp was $ 1499. That's $ 1400 to $ 2700 today. And yes, there were more expensive options. Krell (the de facto high end of the high end of that time period) had a 50 W Class A amp for $ 1900, but it's big boy, the KSA250 (250 W Class A), was $ 5700. Or, $ 3500 to $ 10 500 today.

Wait a sec? Only barely into the 5 figures? In a day of $ 120 000 DACs, they're not even trying.

And again, those were all dealer prices. *Not direct.* So it's likely you could chop those prices in half if direct had been feasible.

And … the elephant's offspring, something nobody seems to talk about much in audioland … consider that today, we actually have some newer, more efficient ways to put things together, like surface-mount parts. We also have high density heatsinks, which are much less expensive than their old, chunky counterparts for the same power dissipation. And parts, in general, are less expensive.

Yes, the actual components used to make electronics are less expensive than when I was buying them in 1990. Thank globalization, or more efficient production methods … but the bottom line is, our costs are *lower.* Sometimes lower even in non-constant, inflated dollars.

And don't get me started on DACs. Theta's PCM63-based DACs used four $30 D/A converters and three DSPs that cost about the same, in a chassis built like a tank with multiple transformers, boards, wiring looms, and insanely over-spec'd power supplies … and their top end was $4000 to $5000. Now, with a good stereo delta-sigma DAC with multiple built-in digital fil-

ters costing literally $ 2 in quantity (the AK4490), tell me why DACs based on off-the-shelf audio tech so expensive.

All of ths means … electronics could cost *even less* today.

Think about that. And ask yourself again: what the hell happened?

So, Are These Prices Somehow Justifiable?

Okay, before I get burned at the stake, let's be clear. Some costs are justifiable. Perhaps not to the extent that some companies would like them to be, but there are certain things that just won't ever be cheap.

I'm talking specifically about discrete R2R DAC designs, big power amps, and realistic speakers.

Discrete R2R DAC designs are a special kind of crazy. They're, in a word, heroic. It's a crazy, brute-force approach, using mind-bendingly expensive discrete resistors and precision switches (at least), and perhaps even reaching into oven-controlled temperatures and other efforts to increase linearity. I won't argue whether such an approach is "right" or "wrong," but it won't ever be cheap. Though at least one DIY discrete

R2R DAC is within the reach of almost anyone (the Soekris design), so that just became a bit of a gray area. Especially if you're looking at a high-5-figure discrete R2R design.

Big power amps also are never going to be cheap. At least not unless they're Class D. And sorry, guys, I've never heard a Class D design that outmatches a conventional linear amp. Yes, even the ones that claim to do so. Call it what you will—confirmation bias, whatever—but I will take a linear design over a mathematically optimized switching approximation of it. Same as I'll take multibit over delta-sigma. But I digress. Big power amps mean big components—big power transformers full of steel and copper (these ain't cheap, or if they are, run), big heatsinks (even high-density designs have their limits,) big output devices, big protection systems, big chassis, etc. Bottom line, it adds up. But again, it's all relative. Does a 200 W PC stereo amp need to cost $ 5000? No. But it will probably be more than $ 500. And $ 50 000? Well, that's nuts.

Realistic speakers. As in, speakers that can reproduce concert-level volumes and real low bass. Physics dictates that speakers like this won't be small or cheap. Period. Especially if you want them to be accurate. Especially if you want an

exotic design, like a planar, or a ribbon, or an electrostatic speaker. (And again, some will bring up DSP and room correction and throwing a billion watts of Class D power at some corrected driver arrangement in a small box, but again, call me a dinosaur, but no thanks.) Realistic speakers are going to start at 4 figures, and not at the lower end of the range … they could easily get into 5 figures. And this is the way it should be—most of your money should go into the transducer. Or at least so says this dinosaur.

Okay, so there are some cases where high price can be justified. This shouldn't be surprising—expecting plastic, injection-molded speakers with $0.59 drivers to work as well as ones using $300 drivers in exotic, nonresonant cabinets is obviously notgonnahappen.com, no matter how much DSP and correction magic is applied. Inherently high cost—to a point—is not always bad.

But sometimes when I start ranting about the price of high-end gear, two other justifications/excuses are thrown out:

1. We've made significant progress in audio reproduction, getting us closer to a live performance, and this progress is costly.

2. Now, the products are so much better built, and more beautiful, and that precision and beauty is costly.

Now we're getting to be on much shakier ground. Let's address both of these statements in turn.

First, the "we've gotten closer to a live performance" claim. Hmm. Can you really say that with a straight face? Especially if you've been in this field for a few decades? Can you tell me to my face, with no trace of irony, that this is true?

I expect you can't.

While I believe there has been progress towards better audio on many fronts, I don't think we're really much closer to fooling people that they're listening to a live performance. Not in 2-channel, not in surround, not anywhere.

While I think the basic accuracy level of most systems these days is much better—better speakers thanks to better measurement, better understanding of materials and resonance, better, non-parabolic horns (in the case of compression drivers) especially—I don't think the absolute level has changed much.

In fact, the average audiophile system may be less realistic than the past in terms of being able

to reproduce realistic levels, especially in the bass. This is thanks largely to the current fetish for small, slim tower speakers and "lifestyle above all" systems that are tiny, convenient, do-all, and (to be frank, and IMO) not very good at doing much of anything beyond being tiny, convenient, and do-all.

Bottom line: physics is real. You can't beat it. You can only cheat it.

Second, the claim of better-built beauty. Sure, you bet. I'll readily concede this one. Thanks to the widespread availability of CNC milling, the historical extension of a trend towards massively overbuilt chassis that have no functional advantages, and the current fetish for machined-from-one-piece billet designs, yeah, gear can look better, fit better, and be made in more and more baroquely complex shapes than in the past.

And this crap costs money. Big money. I was recently shown a very elaborate faceplate that was machined from a single block of 1 inch thick aluminum (the whole thing was about 6 inch tall, 17 inch wide, 1 inch deep. The machine time on it was so extreme that it cost about as much as an Yggdrasil to produce (in 100 piece quantities.) So, you could have put the (completely non-

functional) faceplate in a box and sold it for exactly the same as an Yggdrasil—$2399—at direct prices. From a dealer? Double that.

Yes, for *a faceplate.*

So, being a grumpy engineer, I have to ask what functionality would this faceplate add to a product? The answer won't surprise you: *absolutely none.* There's no reason this product couldn't have a 1/8 inch thick painted steel faceplate … and cost over $4000 less at retail.

Except, of course, it wouldn't look as unique (I won't say, "as good," because to me, the faceplate was entirely hideous. But that's a personal judgement.)

Or, if you wanted substantially the same look, you could die-cast the faceplate. Total cost in aluminum would be about 40 *times* less than the CNC nightmare. Again, you just saved over $4000 in retail cost through a dealer.

But I'm getting ahead of myself. Consider that I concede this point. If you want crazy-looking, overbuilt product made in tiny, tiny quantities, sure, the cost is gonna be high.

Just don't expect it to perform any better.

So, What Happened?

What happened in high end is, to me, a story of justification.

In the dim dark days of high-end audio (think, tube and transformer-coupled amp days, color-tv-is-new days), prices were set largely due to technical qualifications—the ability to play louder or lower, to deliver more watts, to do so with lower noise—because we were still pushing the limits of what we could do with the technology of the day.

When Marantz, Sherwood, Harman started making gear, it was not really all that much more expensive than the other products out there. Price wasn't used as a shock factor—technical capability was. Audiophiles would justify their investments (sometimes very much DIY in nature, like bass horn speakers that used entire basements as cabinets) by the technical capability—the ability to flap pants legs with bass, etc.

As our technical capability increased to the point where 200, 300, and higher wattage ratings were common in solid-state amps, and digital audio started making an appearance, offering perfect frequency response, zero wow and flutter, no maintenance, and a vanishingly low noise floor

for just a few hundred bucks, that's when we started to see an increasing emphasis on boutique parts and overbuilding. *Does the amplifier have the right kind of capacitors in it? Does it use the latest fancy binding posts and RCA inputs? Does it use 36 devices per channel and run in Class A up to* 500 W? If it didn't, well, it probably wasn't high end. And, you know, those boutique parts and overbuilt output stages were not cheap … which meant prices went up. And with higher prices came more justification, because the more you spend on something, the less you want to be wrong about it.

Today, we're post-boutique and post-technical-overbuilding. It's been a long while since I've heard someone comment on Tiffany RCA jacks or WBT binding posts or even Rubycon Black Gate capacitors. It's also been a long time since I've heard about how many thousands of amps an amplifier manufacturer's 72-device output stage can source. Instead, now all we talk about (it seems) is how slim the chassis is, how it's milled out of a solid block of aluminum, how the finish is so good (snark: almost as good as Apple gear, for many, many times the price!), how big the touchscreen is, what kind of weird millwork profile (that's been mathematically-generated)

makes the product look unique, whether or not it takes Bluetooth input from a phone (snark: Really? Really? And double snark: which will be obsolete in a couple of years when the standard changes), how tiny everything is, what kind of room correction DSP it has, etc. And, at even higher prices, the justification continues to intensify. *These systems are part of your life! Part of your DNA! You can't be without them! No inexpensive system can even come close! How dare you even suggest it!*

Higher price, higher neurosis. Remember, humanity is not a rational animal. It is a rationalizing animal. And when you've put a significant chunk of your net worth into a machine that plays music (as opposed to, say, shelter or transportation), you better bet you have just bought into a significant neurosis.

And with that, we're now in a self-sustaining loop, where price is the only metric.

Oh, you have a $ 50 000 speaker system?

Mine is $ 100 000.

Ah, well, $ 100k isn't really the end-all be-all, mine is $ 500k."

Guys, this is ludicrous. And it drives a destructive feedback loop. Higher prices mean lower production numbers. And lower production numbers mean higher prices. And higher prices mean lower production numbers.

Aside: lower production numbers mean higher prices. Repeat this 75 times until it really hits you. Producing 100 million iPhones allows for crazy stuff like rounded glass, machined-from-solid-billet chassis, and cutting-edge displays … at semi-reasonable prices. Producing 100 amplifiers is now a huge production run for many high-end audio companies—fully one million times less than Apple's most popular products. Make those big dogs out of billet, and now you see why mid 5-figures is a normal price these days. And when you stop to consider that some manufacturers are now producing products in 10 s (not 100 s, not 1000 s), it's very, very easy to see how prices get out of control. Plus, ah, dealers. You're not going to sell a $ 20 000 preamp or $ 50 000 amp direct, sorry, no way, no how, you really need a psychologist … er, wait, I mean salesperson, to sell you on what such a machine will do for your love life.

And this feedback loop goes on and on, until

you're making, what? One $5 million dollar product for one person per year? A $50 million product? So that one person can lord it over all other audiophiles, simply on the basis of price?

Until … what? Until they die? And then where are you?

So What Do We Do About This?

Well, it's obvious, isn't it? We need to take up arms and storm the castle!

Except, well, there is no castle, and there is really no need for a violent revolution. The end-game of high-end is plain for all to see—ever-increasing prices for an ever-shrinking market, until the elephant crashes all the way through the dance floor.

What we need to do, as manufacturers, is keep our wits. The siren call of higher prices is a huge temptation. As the super-high-end hyperinflates, the sensible high end will want to raise prices, too. Hell, they'll be *called on* to raise prices.

Think I'm kidding? I can't tell you how many times I've been told that we should raise prices because we won't be taken seriously if we don't.

Well, here's my response: **** you.

I give exactly zero ****s about being taken seriously. I care about only one thing: bringing high-end within reach of as many people as possible and building the ranks of audiophiles, so that when the ultra-high-end implosion comes, there's plenty of sensible music-lovers out there.

And that's what manufacturers can do: resist the temptation to inflate up, as the ultra-high-end bubble spirals out of control. Remember how to make an inexpensive, attractive chassis. Remember old production methods like sheet metal and casting. Remember that higher production numbers will reduce your costs. Remember to pass that along to the buyer. And consider—if you're not already there—going direct. That's the biggest benefit to the buyer, really, since it effectively cuts prices in half.

So what can you do, individually? You can, of course, choose less-expensive components. Like, well, duh. You can also be more pointed in your questioning—ask manufacturers where the cost goes. What percentage of that high price tag is the chassis versus the electronics? Hint: if it's fancy, it's a lot. What unique technology are they bringing to the table? If it's off-the-shelf

chips and implementation, how do they justify the cost? Hint: R&D on unique tech is a lot higher than a datasheet implementation, so R&D amortization really doesn't fly there.

But there's more. I want to do more at Schiit. So consider this our manifesto: we will continue to bring the highest-value products, at the fairest-possible margins, in the largest-quantity runs (for even lower cost) to assure that as many people can enjoy high-end audio as possible, and to grow the audience for the future. And we'll continue to do this both in personal and desktop audio, and in the world of speaker amps, preamps, and other gear, as we can. You'll see our 2-channel products this fall, and, if the reaction of various industry people is to be believed, they may have an even bigger impact on that market than we have on the desktop.

Because, you know what? We need to do more than *talk* about the elephant. Talking is only the start. What we really need to do is get the thing out of the room, entirely.

Here's to a sane, affordable, and high-end future!

2016, Chapter 10

Because We Can

Answering the question of "why a Modi Multi-bit?" seemed difficult at first. I mean, what other high-end audio company would introduce a *significantly* downmarket product—one using its flagship technology—only a year after that technology debuted?

It *does* sound a little crazy, doesn't it?

Then I realized … well, for us, it isn't crazy at all. It's *what we do*. Like some other high-end companies might charge more for a product "because that's what the market will bear," or other audio companies refuse to "jeopardize their brand reputation with low-priced products," period. That's what they do.

This is a choice we make: to develop new technology (or at least new, unique platforms) and disseminate it as widely as possible.

We do this because we think that you shouldn't have to spend a ton of money to get great sound.

We do this because we think that the more people who can enjoy great sound, the better. We do this because more people in the market benefits everyone.

And yeah, I gotta admit: we also do it, at least in part, because it makes some organizations out there a little bit uncomfortable, a little bit unsettled. It makes them feel like the audiophile world is moving on, changing, and becoming something they never expected it to be. And, honestly, if you told me, back in 1992, that I'd one day be selling a true Multibit DAC with a closed-form digital filter on a powerful DSP engine for $145—the 1992 equivalent of $249—I would have told you that you're barking mad, it would never happen … so yeah, the world *is* changing.

Or, to sum up, why do we do a Modi Multibit?

Because we *can*.

The Road to Modi Multibit

Now, this isn't to say that the process of multi-bitting the Modi was easy, straightforward, and fun. Unlike all of our other Multibit DACs, this one was a bit of a lark. We'd never designed the

Modi expecting that we'd one day turn it into a Multibit DAC, so the conversation started as a joke.

"Now you just have to do a Modi Multibit, and the line is complete," I told Mike and Dave, shortly after they showed me the first working Bifrost Multibit board.

> And this was a total jest—I hadn't expected the fact that Mike planned for a Bifrost Multibit, so that one was already a bit of a shock. The idea that we could cram all the guts of a Bifrost Multibit into a box the size of a Modi was 100 % blue sky.

I expected our two digital maestros to respond with nervous laughter. Instead, they surprised me by exchanging knowing looks.

"Weeeeeellllll," Dave said.

"We've been thinking about that," Mike finished for him.

"What?" I said, mouth hanging open like someone who'd just been told cigarettes are full of vitamin C and prevent cancer.

"It might be possible," Mike said. "It depends on how good Dave's layout skills are."

“Wait a sec,” I said. “Are you telling me we might be able to have a Modi Multibit?”

Mike nodded, giving me his signature Schiit-eating grin.

“As in, a real Multibit DAC?”

“Right.”

“Not with the closed-form digital filter,” I said. Because that was completely nuts. Totally impossible. Mike must be thinking about using a sample rate converter, like he’d talked about doing with Bifrost Multibit, early in the process.

“Yes, with the burrito filter,” Mike said, using his verbal shorthand for our unique, closed-form digital filter.

“Though just a burrito, not a mega burrito,” Dave added, meaning that the filter would be like the one in Bifrost Multibit, not the one in Gungnir Multibit or Yggdrasil.

“Wow,” I said.

“It’s not done yet,” Mike reminded me. “But I think it’s doable. It’s just that all of the inputs take up a lot of space … ”

"Wait a sec," I said. Until Mike mentioned "inputs," with an "s," I thought he was talking about something more like a super-Modi—a USB-input-only device, powered by USB. "Are you talking about taking a Modi 2 Uber and making it Multibit?"

"Exactly," Mike said, grinning.

For a long while, I couldn't say anything. A Modi Multibit with coax, toslink, and USB inputs and the closed-form filter would essentially be a mini Bifrost Multibit. Functionally, it would be 100 % the same.

What would that do to Bifrost sales? I wondered.

"How much could this sell for?" I asked.

Mike looked up into the sky and made a show of counting on his fingers, muttering about the additional cost for the DSP, the D/A converter, the support ICs, etc. "Maybe two-fifty," he said, finally.

"Two hundred and fifty dollars?"

"Right."

"What about Bifrost?" Dave asked, echoing my own thoughts.

Mike shrugged. “Bifrost isn’t disposable. Plus, power supply. Plus, general layout. Plus, modularity. Plus, this hasn’t happened yet. Don’t count the chickens.”

Dave and I looked uncertainly at each other. I was OK with the basic idea. What Mike was saying made sense. Even if the Modi Multibit had all the same inputs as Bifrost Multibit, even if it had the same digital filter, it wasn’t a future-looking device. It had no upgrade capability at all. When technology changed—a new USB input, a new Multibit D/A converter—well, if you wanted that, you threw away the Modi Multibit, but simply upgraded the Bifrost Multibit. Plus, the power supply in Bifrost Multibit was far and away better than the one in Modi, and always would be. (It’s not possible to beat a custom transformer with multiple secondaries, plus many more stages of general and local regulation, plus the fact that the transformer alone was almost 3× the size.)

But …

But a lot of people would just look at the price. And $ 249 always beat $ 599. Period. End of story.

Was I okay with that idea? At the time, it made me a bit uneasy.

But, as Mike said, it wasn't a done deal. The Modi Multibit could end up stillborn, as many of our projects have been over the years.

So, we'll sit back and see, I told myself.

Modi Multibit Challenges

It actually didn't take long for Dave and Mike to show me a green prototype board for the Modi Multibit—and it was every bit as crazy as I expected it to be. Literally every square inch of surface—on both the top and bottom—was packed with parts. It was, by far, the most dense-looking board we've done.

Aside: dense-*looking?* Yes. Dave and Mike prefer to work with larger, more manageable surface-mount resistors and capacitors—in technical parlance, 0805 sized parts—unless a smaller size is dictated by the layout requirements of a fancy part that needs close-coupled bypassing. I tend to shrink the overall components to fit—Fulla, and another upcoming product, freely use 0402 and 0603 parts. 0402 parts are very, very small—but not as small as they could be. There are 0201 and 01005 parts. No thanks—those parts would easily pass through

a salt shaker.

> **Aside aside:** So why do Mike and Dave like the bigger parts, and I like the smaller parts? Mike and Dave still like to do a lot of their own prototyping work. When you start getting to 0402 part sizes, it's really, really hard to do any prototypes by hand. It's not impossible, but you quickly start wondering why you didn't just hand it to the assembly house and have them deal with it. 0805 is pretty easy to work by hand. But even Mike and Dave are changing … the proliferation of new parts that have unsolderable lands underneath them (for example, there's a great dual voltage regulator chip that has 14 pins—several of which are completely underneath the part—in a 3 mm×3 mm square. Yeah. That kind of thing you don't do by hand. You send it to the assembly house and have them use solder paste and a reflow oven. Our DSPs are also very little fun, with a buried power pad that also can't be done by hand. Times are changing, you have to change with them.)

At first, all I saw was this crazy dense board—no parts on it. Dave had to have the assembly house put the DSP chip on it, since he couldn't do it by hand, so he had to wait for that … and then do

the rest of the assembly.

Eventually, Dave had something ready to show us.

It looked a lot like most of our prototype boards—covered with flux, with some parts stacked on each other, and some parts flying in the air. In fact, one entire regulator (and its support passives) was tacked to one of the clearer areas of the board, with wires trailing across the board to the place where it was needed.

"Regulator got too hot," Dave said. "We needed a second one."

Mike and I nodded. That wasn't surprising. The DSP and the industrial-strength D/A converter use more power than the consumer-grade chips they replaced.

But Dave still didn't look happy.

"It also needs an output mute," he said.

Mike and I groaned. This wasn't surprising, given the fact that we weren't using a made-for-audio D/A chip that had nice little functions like soft mute. But it was a pain in the ass. Because an output mute meant a relay, and relays were big. And the board had precious little real estate.

“Use one of the NEC relays, a UB2, that’s minimum footprint,” I told Dave.

Dave hadn’t used those relays before, so I got him a couple to play with. But even then, looking at the board, it didn’t look like he was going to be able to make it fit.

“I was thinking,” Dave said. “What if we got rid of the microprocessor?”

Dave was referring to the Microchip microcontroller we used to manage the input, output, digital filter, and D/A converter in the Bifrost Multibit. It was a relatively large part—over 1/2 inch on a side.

“And replace it with what?” I asked.

“The DSP,” Mike and Dave said in unison.

“Ah, gotcha,” I said, realizing that yes, we already had a microprocessor in there—it just happened to be a super-fast digital signal processor (DSP) more commonly used for mathematical operations (like our digital filter). It could also be used for some housekeeping functions, without compromising the digital filter in the slightest.

And hey—if Dave thought it could be done, then it was probably doable. He'd have to do the code, in any case.

And so, the Modi Multibit board disappeared for a while. And, once again, I had to wonder, So what happens to Bifrost Multibit if this works?

Modi Multibit for Real

When Mike and Dave came back to me with a Modi Multibit that had only a couple of extra parts hanging off of it, it was very early in 2016. As in, very early. Think February. Sales were still going strong after the holidays—and, inexplicably, even picking up strongly over last year's sales, despite the lack of any new product launches.

And, in February, Bifrost Multibit was only 4 months old.

Four months isn't a lot. It's longer than a quarter, yes, so some corporate CEOs may sniff and ask, "What have you done for me in this reporting period?" but in the overall scheme of Bifrost, it was only about 1/12 of its total lifetime on the planet.

So, the question was: *when did we schedule the intro of the Modi Multibit?*

Or did we do it at all?

From the beginning, I figured we'd do it. The question was, when? With a working product in February, we could launch in May at the earliest (factoring in the lead times for metal, boards, parts, scheduling at the assembly house, maybe one minor metal screwup, maybe one minor production glitch, etc …). But of course there was always the chance that one or more parts would be hard (or impossible) to get, and, of course, the chance of a large metal screwup, or a much larger production glitch.

"June," Mike said finally, after thinking about it for a bit. "Shoot for June. July, maybe, if it slips."

That sounded logical to me. The only thing that was left was to do all of the busywork—the product description, the press release, the product photos, the FAQs, etc, etc … all the usual stuff that's part of a typical product intro, the stuff that nobody thinks about.

That's my job.

And I like that job—hell, sometimes I write product descriptions and FAQs before the prototypes are even laid out. This is my way of testing the logic of the product. Does it sound compelling

enough to buy? Does it make sense? Does it fit with the line? If there are any problems answering these questions, they'll show up when you go to do the copy.

I wondered what kind of problems I'd run into with Modi Multibit.

It turns out, I shouldn't have worried. Modi Multibit makes total sense—if you look at it from the Schiit perspective. We're not here to toe the line, or to sell things for more 'because that's what everyone's doing,' or to fit neatly into a reviewer's box, or (to be blunt) make our competitors feel warm and fuzzy. We started by shaking things up on the desktop, we continued by shaking things up with upgradable DACs, we continued by shaking things up with new ideas like LISST and intelligent amps, continued by shaking things up with Multibit DACs, and so why shouldn't we continue doing exactly the same thing now, by introducing by far the most advanced affordable Multibit DAC?

Aside: As I wrote, it struck me just how far Modi Multibit is in advance of other affordable Multibit DACs. Yes, there may be inexpensive offerings out there, but every single one of them either uses obsolete 16 bit "pull" audio DACs and

audio digital filters (or are non-over sampling, or NOS). Modi Multibit is the only Multibit DAC that's built on a modern platform, using medical/military grade D/A converters and our own closed-form digital filter running on an Analog Devices DSP chip—and it is, by far, the most advanced thing anywhere near its price, from any company, from any country of origin.

And that is why Modi Multibit fits seamlessly into the Modi line, and why it fits seamlessly into our philosophy. Now, you can experience a significant percentage of what makes our Multibit DACs special, for less than the price of some fancy USB decrapifiers. Hell, less than the cost of some USB cables, for that matter.

Not a huge believer in Multibit? That's no problem for the Modi line, either. Save $100 and get a Modi 2 Uber, with the same AK4490 D/A converter that's used in products 10× to 20× it's price. Everybody wins.

Aside: Yes, we made a running change from the AK4396 (a great-sounding D/A converter) to the AK4490 (an even better-sounding DAC, in our opinion) in the Modi 2 Uber.

> **Aside aside:** And yes, it is Modi Multibit, not Modi 2 Multibit. We decided the names were getting a bit long and unwieldy, so we dropped the "2" from Modi Multibit. Yes, we know that technically means that Modi 2 Uber should be Modi Uber. But that would be confusing. Deal with it. We aren't perfect.

Hitches, Glitches, and Other Absences

Astute readers who are familiar with our past may be asking, "Hey, wait a minute, are you leaving something out? It seems that the development of the Modi Multibit went surprisingly smooth … there were no big screw ups, or panics, or last-minute production showstoppers!"

Right. I'm not leaving anything out—there simply weren't any of the usual big hitches or glitches. There were a few parts that took longer than expected to come in, but that's about it. Really boring, when you get right down to it.

But that's how you want a production rollout to be: as boring as possible. Exciting production rollouts are very much not good. Because the excitement is never of the positive variety, I'm afraid.

Other astute readers may have noticed that boring production rollouts have become pretty typical around here. That's 100 % intentional—and that's 100 % a good thing. The less strife we have getting things into production, the better we are at bringing things out in a timely, orderly, and non-disruptive fashion. Which means we can concentrate on more products, or improving current products, rather than trying to patch up a bad launch.

And yes, I know. Boring launches don't exactly make for exciting stories. However, I hope you're excited about the result—the new Modi Multibit—and I hope you're excited to have us keep shaking things up in the future.

Because … well, as game-changing as Modi Multibit is … well, it's just the beginning this year.

2016, Chapter 11

The Road to Jotunheim

So, if you've been following the book to date, you know that Ragnarok had a long and very painful gestation—that is was, in fact, one of the first products I wanted to make back in 2010. Introducing it in 2014 was a huge milestone for us.

But what you don't know is that there's one product that has arguably gone through a even more arduous journey on its way to production … complete with false starts, two complete product redefinitions, seven prototypes, and more "first article" metal than we've ever run for a single product.

And that's Jotunheim.

Since you're reading this chapter shortly after the intro of Jotunheim, you already know the gist of it: a do-all, balanced in/balanced out, single-ended in, single-ended out amp that also happens to be configurable with the use of an add-in card—and

that this configurability is what makes it unique in its field: the only upgradable, updatable amp/DAC or amp/phono pre or just-an-amp on the planet.

Yes, we're seriously proud of Jotunheim, as we think it makes the entire concept of non-upgradable DAC/amps completely obsolete. Why get a combined product—especially one that costs much more than Jotunheim—knowing it will be obsolete in a couple of years, as digital audio continues its inevitable changes?

But I'm editorializing. I think Jotunheim is a big deal. But I'm biased. And the fact is, Jotunheim started for much less sweeping reason than "redefining a market segment."

It started, quite simply, because I was staring at boxes upon boxes of very expensive balanced volume pots.

An Inauspicious Beginning

Now, this was a long time ago—late 2012, in fact. We were still in the noisy, dusty, dirty, held-together-literally-by-tape-and-baling-wire Schiithole on 6th Street in Newhall. Moving into our current palatial (ha) Schiitbox wasn't even a

dream. (Well, except maybe to Alex, because we were running out of space at the Schiithole.)

So, about those pots? Yeah. We'd started making Mjolnirs with a stock of 400 balanced pots, which I'd managed to find ahead of the production run. As the first Mjolnirs flew off the shelves, I placed an order for another 1000 pots. And then, since they were very long lead-time items, I placed a second order.

These two orders were probably the most expensive single-part orders that we'd done to date. They were also, in engineering terms, "NC/NR," or, "noncancellable, nonreturnable."

And, of course, I placed those two pot orders right before Mjolnir sales started to nosedive.

Yeah. There you go. It seems that *not* having a single-ended output on the original Mjolnir *wasn't* one of my better ideas.

So I was standing there, looking at these boxes and boxes of very expensive pots, wondering how we would ever sell them all, and I had a thought: *Hey, if we made a less expensive balanced amp, we could use those pots in there, too.*

Yes, I know. Not rocket science. And not the best reason to develop a new product. But it was

either that, or sit on a pile of very expensive parts, like, forever.

So, in my spare time, I started working on a new amp design. I knew kinda-sorta what I wanted to do with it. Of course, what I wanted to do at the time isn't what Jotunheim turned out to be, but we'll get to that. If there had been a design brief, it would have read something like this:

1. A cheaper balanced amp!
2. Oh yeah, and it has to have single-ended outputs too, because it's already clear that wasn't the hottest idea for Mjolnir
3. If it's gonna be cheaper, it should fit our smaller chassis size, the same as Asgard and Valhalla and Lyr, because I don't want a million different sized chassis
4. Not a circlotron, because those require expensive, complex transformers and don't have an easy way to get single-ended output
5. And while we're at it, let's go for a real, linearized, zero-output-impedance style amplifier, to really drive the heck out of those popular, power-hungry planers (at the time)

Note on #4. This is the key to led us to an entirely new topology … one that Dave didn't even believe would work at all, and one that even I had trouble with accepting, even after

extensive testing. This is our new "Pivot Point" differential current-feedback topology, which (as far as I know) is unique. More on this later.

Note on #5. This is what you call either "ego talking" or "straight out delusional." Of course, I didn't know this at the time. I just figured it would be a nifty cool thing to re-introduce the Hawksford-style linearization that we'd used at Sumo in a headphone amp. Hawksford linearization is a really neat trick that helps manage the transition between N-channel and P-channel MOSFETS, or NPN and PNP transistors, which, of course, are never truly complementary, and always have some transconductance linearity problems near zero-crossing.

Aside: if your eyes are crossing due to the heavy dose of engineeringese, look at it this way: the battle for better amps usually comes down to one thing–how the output stage behaves. This is why some designers choose to make hot, heavy, inconvenient, inefficient Class A amplifiers … because they sidestep the zero-crossing problem by never having the output turn off at all. Of course, this dramatically limits the kind of power output you can have without self-immolation (assuming a true Class A design, but that's a

screed I've already written.) Using Hawksford linearization is a way to make the output stage behave.

And there you go. #4 and #5 are, together, why you should give yourself the time and freedom to explore new designs … and to kill your babies when they are really, really ugly.

At the Pivot Point

Staring at those boxes of pots led to me playing with breadboards, late in 2012. The first order of business was to nail down a new topology for the amp. Like I said, no circlotrons need apply. Given that I wanted this to be a balanced amp, I was rather limited on the choices I had from the list of "usual suspects:"

- **Supersymmetry.** Nelson Pass' patent had just expired for supersymmetry. I could use that idea. It was a neat way to create an inherently differential topology, and it had the potential for using a single output phase as the single-ended output. So why didn't I use it? Well, while I'm less "not invented here" than Mike Moffat, I still like to think I can contribute to

the art. A direct rip-off of Nelson's ideas didn't sit right.[1]

- **I could use totally separate gain stages for each phase, as some balanced amps do.** Of course, this would mean that a single-ended input would not produce balanced output. Not ideal.
- **I could use a balanced-bridged topology,** like Sumo used in the Andromeda II. But these topologies are kludgy and typically don't sound very good. Andromeda II was an exception. But I didn't want to be shooting for an exception.
- **Finally, I could start looking for a true differential topology that sidestepped the problems with the rest—**one that had high-impedance inputs, plus good places to hook the feedback, plus low distortion and high performance. And to do this, I went back to look at something I was never really able to

1 Now, let's be clear: there's lots and lots and lots of "design reuse" going on in audio. More than 95 % of the power amps out there are essentially Lin-based design. Magni and Magni 2 are. There's nothing wrong with that. The Lin topology is perfectly fine for a speaker amp or a headphone amp. It's been tested, busted, iterated, refined, explored, augmented, and examined every which way. But … I like to see what else is out there, when I can.

deploy for Sumo. In the late days of Sumo, I was playing around with a differential topology that used current feedback. It had all the hallmarks of what I needed, but my notes from the era were vague and inconclusive. I didn't know if I'd ever built it. And, even if I had, I didn't know if it would work in a topology that wasn't fully complementary, like we used to do at Sumo.[2]

To make a short story short, I modified the never-used Sumo-era topology, pasting in a noncomplementary input and DC servos, and hooked it all up on a breadboard. I figured that it would probably:

- Have huge DC offset
- Have huge problems with distortion due to noncomplementary nonlinear inputs

2 Yes, I know, more engineeringese. Sorry about that. But let me try to explain. Sumo amps were what they called "fully complementary." The input stages used both NPN and PNP BJT transistors. Now, the problem with BJTs is that they have a tendency to rectify RF noise. Say, like from a cellphone sitting next to the amp. That's why Schiit tends to use JFET inputs. But finding a good complementary pair of JFETs these days is kinda like looking for a lost Rembrandt in your granny's attic–good luck with that. So our input stages are not fully complementary. Would this topology work on a noncomplementary input?

- Blow the input transistors right off the board at high output

But, surprise surprise: *it worked.* DC offset was low, distortion was very low, and the transistors didn't fly off the board.

I blinked a few times, re-ran the tests, verified that yes, it was working, and yes, it was all hooked up right.

And when it all checked out, it was time to go to layout.

But not before Dave's protests.

"It Won't Work."

I was feeling pretty full of myself with this new topology, especially after some internet searching didn't find anything like it. *Hell, this might be unique,* I thought. It might even be new enough to be patentable.

Aside: now, before the Internet lawyers get their panties in a twist, no, we don't patent stuff. Even if it was patentable. Because, before we pursued a patent, we'd have to do a much more thorough search. And even if we went through and obtained a patent, it would buy us nothing

more than the license to sue people who were using it. It would also reveal the entirety of the topology to the world, including parts of the world where US patent protection has no meaning. So relax.

So how full of myself was I feeling? Enough to propose it as the output topology for the then-still-in-early-design-phase Yggdrasil.

"But it won't work," Dave told me, when he first saw the topology.

I laughed. "That's what I told myself when I first came up with it."

"But … you're not telling me it works?" Dave asked.

"Yes. It does."

Dave frowned. "I don't see how it can. What does the front end do when the output's at 20 V and you have a gain of 1? It's seeing all of that voltage!"

"Yes, and it rides up on it. It makes the front end operate much more like a constant-voltage stage."

Dave shook his head and sighed. "But … it can't work."

“Test it,” I told him. “It measures better than anything we make.”

“I don’t see how,” Dave said.

“Test it.”

Dave muttered something and gave up. I’m not sure if he ever tested it, but we did use it in early Yggdrasil prototypes. The reason we didn’t go with this topology for Yggy is simple—it didn’t need it. It turned out that all it needed was a simple two-JFET buffer.

I was also confident enough to start laying out the first Jotunheim board. This board resembles the final product only vaguely. Like the final product, it’s on a board that fits in our 6 inch×9 inch chassis. Like the final product, it uses a C-core “flat pack” transformer, a first in a Schiit product. Like the final product, it has a Pivot Point topology gain stage.

But it doesn’t have any provision for an optional DAC or phono card.

Aaannnd it has one thing the production product lacks: the Hawksford linearized gain stage.

So, why no optional DAC/phono card? Simple. We hadn’t thought of it yet.

And why doesn't the production product have Hawksford linearization? Well, two reasons:

- It uses a ton of parts. It made the gain stage so complex, it barely fit on the board. The layout was pretty compromised, especially since the early boards were 2-layer. (Production boards are 4-layer).
- It can weld single-ended headphone plugs to the Neutrik jacks.

Yeah. Seriously on #2.

"Wait, what the hell are you talking about?" Someone is asking right now. "It welds the headphone plug to the jack? How the hell does it do that?"

Well, to understand why, you need to know two bits of engineering—one on the electrical side, one on the mechanical side. On the electrical side, the Hawksford linearization does exactly what it says: it linearizes the input and output. Aggressively. As in, any deviation between input and output engages a servo loop to correct it. As in, you can literally tune the circuit down to zero output impedance.

Aside: this is what allowed Sumo amps to deliver 200 A (that is, two hundred AMPS) peak

current for short periods of time. Like I said, *aggressively.*

On the mechanical side, you need to know how the common headphone plug works. It actually shorts the outputs as you insert it. This is why some amps say "turn it down before plugging in or unplugging headphones."

Engineers are already wincing at the "zero output impedance + shorted output" statement, but let me spell it out for the nontechnical audience in very simple terms:

With the first Jotunheim board, if you plugged in a headphone to the single-ended jack with the volume turned up, and there was a significant difference between left and right channels, it was possible to arc-weld the plug to the jack.

This, in turn:

- Destroyed the jack.
- Usually made it so you couldn't pull the headphone plug out of the jack—it was welded in that well.
- Typically blew up the output stage.

Now, even in early 2013, when I was working with this first Jotunheim board, I knew that putting a "hey, turn the volume down" disclaimer on the

Jotunheim would do exactly two things for its reliability and for customer satisfaction: jack and schiit.

And so there you go. Design revision #1: ditch the linearized output stage.

But that wasn't the only revision. Oh, no. Not by a looooooonnng shot.

Design Adventures, Cul-De-Sacs, and More Ego Talking

The **second prototype** of Jotunheim still wouldn't be very recognizable to someone looking at the production version today. It had no space for an optional module, for one thing. It also still had a lot of parts on it that reflected my incomplete understanding of what the Pivot Point topology was. In short, I was still building this new topology like a conventional amp, in some ways.

But it was the first prototype that we actually spent a lot of time with, and for that reason, it was important. This was where the tough stuff got worked out: compensation, stability, stress testing, and initial measurements.

Stability was interesting in this case, because of the inherently "fast" nature of current-feedback amps, coupled with the overall complexity of the stage. Add in the fact that many current-feedback topologies don't like to have a pole in the feedback network (if you're an engineer, think about it—it'll become obvious), and I had a number of adventures in major oscillation before everything got quieted down. Not quite as bad as Dave's speaker power amp that took out Channel 28 on UHF (back when analog TV was a thing, yes, we're old, deal with it, but also look up what frequency band UHF is and go a little pale).

But, once the amp was stabilized, we had our first working, playing Jotunheim. This was probably around summer of 2013. We started running it by some of our trusted listeners, and they were enthusiastic.

But they also asked questions like, "Yeah, but if you bring this out, what will it do to Mjolnir?" and "Hmm, it'd be a tough call on this or Lyr."

And that got us wondering: *hey, you know, we have a lot of amps in the line, so where does this one fit?*

Call it a crisis of faith. But for good reason. We had Magni, Asgard, Valhalla, Lyr, and Mjolnir.

Ragnarok was coming. I kept thinking about another amp, a more powerful tube amp, as well. That was a buttload of headphone amps. Especially if you threw Jotunheim in there as well.

Why is this important? Because most people don't like choice. Give them too many choices, and they'll choose "none of the above." The amp line had already departed from the tried-and-true three-choice "good, better, best" scenario, so more complexity wasn't going to help it.

And, at the same time, I knew that if we didn't move forward, someone else would.

Aside: this is probably the core of any successful business that operates in the real world—how do you know when to "not fix schiit that ain't broken," and when to "kill your babies?" I don't think we have any 100 % answer for this.

In any case, those uneasy questions didn't start me moving into a **third prototype**—this one intended to be the final one, the one that we could sell. Maybe by the end of the year. Yes, I know, stop laughing.

The third proto was pretty much a done deal, too—we'd pulled out the extraneous parts, added

the compensation components, and it ran just fine. Except, of course, for the humming transformer and backwards relays.

I used that excuse to put it on the shelf and get back to other, more important things … like Ragnarok.

When it came off the shelf again, though, Mike and I had the first real conversation about where Jotunheim would fit in the line.

"It's a whole new ball game," Mike told me. "A new topology. And it sounds great. We should treat it like Theta—start at the top, work our way down."

I shook my head. "Ragnarok is a done deal, it's not going to go Pivot Point now. Unless you want to see it in 2016 or something."

Mike nodded. "You have a point. How about as a Mjolnir replacement? Or Uber-Mjolnir?"

"This thing makes any Mjolnir replacement pretty much a stupid idea," I told Mike. "The only Mjolnir replacement that makes any sense is one with tubes."

And then I sat back, stunned silent. *Because that was the only thing that made any sense—a tube hybrid Mjolnir. Which would be pretty cool …*

"So put it above Mjolnir, below Ragnarok," Mike said.

I shook my head. "We already have too many amps."

Mike groaned and shook his head. He knew I was right. And then he said, "Put a DAC in it."

"What?" I asked, not sure if I'd heard him right.

"Make it an amp/DAC," Mike said.

"But … amp/DACs suck," I said. "They're obsolescent as soon as digital technology changes, and they have digital noise running around in the same chassis … "

"So make it modular," Mike said, and stopped cold. We both had a moment when we just looked at each other.

Because, holy schiit, that was it. If Jotunheim was modular, the obsolescence issue was moot. And if I ran the card vertically, with parts only on one side, with decoupling from the main power supply and ground … the noise might be a non-issue.

Or, even better, look at the whole thing like a power DAC. No different than a Bifrost discrete output stage. But with more power. Run the D/As with passive filtering, pass it through to the gain stage, and do the ultrasonic filtering throughout. And, if the USB powered the DAC, it wouldn't be running at all unless plugged in. That gave had the potential to be a zero-compromise amp and zero-compromise DAC, all in one.

I explained that to Mike.

He nodded. "That's a game-changer."

And that's how we got to the **fourth prototype.** That version would look familiar to you today. It introduced the idea of the vertical DAC card, with pin-header connector and relay switching. It added a three-position switch so you could select the output of the DAC, the balanced inputs, or the single-ended inputs.

The timeline? Middle of 2014.

"So why didn't you introduce it then, ya buttheads?" someone is likely asking. "It's middle of 2016 now. What, you sit on things for two years after they're done?

Hey, bite me. I have two words for you: Ragnarok and Yggdrasil.

Yeah, the flagships put a dent in the timeline for Jotunheim, too. Both of them were a ton of work. Mike and Dave were full-bore on Yggdrasil, so they didn't have time to lay out a DAC card for Jotunheim. I was full-bore on Ragnarok, so the whole thing went back on the shelf for a while.

Well, except for the time I took to do the layout for the DAC card.

Fun fact: the DAC card is what led to Fulla. In the process of doing a very compact layout, I wondered how small I could go … and if it could be made into a complete DAC and headphone amp. And that's how we got Fulla. I just couldn't say that at the time. Oh yeah, and the first Jotunheim DAC card was based on the AK4396, just like the Fulla.

Ironically, that first DAC card was never built. It used two AK4396 D/A converters, and we'd just started hearing about how those were end of life. So I held off sending that prototype to the board house, while we went completely nuts.

Nuts?

Yes. The **5th prototype** was completely insane.

Why? Because I'd just gotten the Ragnarok working, and it was very cool. And I was happy

and giddy and a bit full of myself. And because of this, I had a very stupid idea. That stupid idea was this:

Hey, wouldn't it be cool if Jotunheim used the same microprocessor management system as Ragnarok?

Yeah. I am an idiot.

But there's no force in nature greater than idiocy, so I convinced myself that this would be a great idea, because hey, we already had a DAC in there, and because hey, eliminating the DC servo would make it sound better, and because hey, we could claim direct lineage with Ragnarok, and because hey, I am a big dummy.

And yes, we actually designed and built this thing. This was in early 2015. It was exactly as nuts as you'd think. We had to go to smaller 0603 parts, we had to add a ton of power supplies and bypassing, we had to add a ton of stuff to interface the microprocessor with the gain stage and the relays, including a bunch of flyspeck diode arrays.

But we did it, and it worked.

Now, it never worked right. The microprocessor put the current draw on the low-voltage supplies over the edge, the switching noise got everywhere,

and making measurements and managing DC offset never got good enough for preamp and headphone use.

Yeah, there's the cul-de-sac. And the ego.

> **Aside:** Know when to walk away from a failed design. You don't get points for persistence.

Which meant the **sixth prototype** was similar to the 5th, except for a few details. We were still playing with a couple of different gain-switching designs, so that prototype got both of them, with a switch on the board so we could compare how they sounded. The intent was to listen to both, then do a final production-ready board with only the winner.

The 6th prototype was also where we first got to plug in a DAC. With the re-layout done for the AK4490, and a different power supply arrangement (we decided to take the power supply from the main rails, which can provide amps of current), we finally had a complete, production-ready product.

By this time, it was well into 2015. Too far in to think of this as a Christmas product. And hell, we already had plenty of plans for 2015. So Jotunheim got back-burnered yet again. We

figured it would be a 2016 product, and that was that.

And that's when I had a very interesting idea. It seems simple in retrospect.

Because you'll notice, up until now, I always referred to Jotunheim as an amp/DAC. Of course, it could be used as "just an amp" too.

But … it had a modular card slot in it. *And that slot could be filled with things other than a DAC.*

Which then made it *an entirely different kind of product.* Now, it was a configurable desktop amp, or modular desktop control center, or protean desktop widget … hell, we didn't really have the words to describe it.

Because, in addition to a DAC, we could do a phono input. Or maybe a wireless card. Or maybe just another input.

And that was *completely new.*

Now I was excited.

So, as I did the small revisions for the **seventh (and final) prototype** board—really just a board to confirm that production was 100 %—I also started designing an alternate module … a phono board.

"Phono?" you may be asking. "Why phono?"

To which I say, *Why not?* There are a ton of people out there with turntables. There are also a ton of people who already have our multibit DACs, and won't want to give them up. And there are a ton of people who will want turntable input, DAC input, and preamp outputs … as well as headphone outputs.

Aside: the design of the phono board is a variant of Mani, minus Mani's extreme flexibility. It has a single gain and single load, appropriate for most MM cartridges. MC carts will need a step-up transformer or pre-pre gain stage. Ah well, you can't have everything. However, in the process of simplification, we have gained some significant advantages as well. The Jotunheim phono input loses a gain stage and coupling caps. Although the RIAA is passive, it uses only two gain stages, and a DC servo, unlike Mani, which uses three gain stages and is capacitor-coupled. We can argue the sonic benefits of either approach …

And, for both the 7th prototype, and for the phono board, there was little drama. The 7th started up and ran, as expected, since it was just

a minor change from the 6th. The phono board worked just fine, and sounded very good.

Electronically, everything seemed set to go for mid-2016 …

A New Chassis and Last-Minute Complications

Sharp-eyed readers will notice that the Jotunheim chassis is a pretty significant departure from what we've done in the past. What's with the L-bend (rather than the U-bend? What's with the lack of fasteners on top? And what's with the press-in feet? (HOLYCRAPWHUT?)

In short, yes, the chassis is a story in itself. But, as this is an epic chapter already, I'll keep it brief.

The chassis started with the most noble and pure of motives, kinda like the original design of the Jotunheim itself.

Or, in other words, the chassis was intended to reduce cost.

Or, put bluntly: to be cheaper.

Yes. From the start, Jotunheim was intended to be an inexpensive product. Unlike many high-end

audio companies, we don't see any problem with this. In fact, we think the chassis should be as inexpensive as possible, so we can deliver higher value on the inside. After all, it's the electronics that make the sound, not the box.

> **Aside:** and yes, I understand, some are gonna argue with me on this one, but I don't think anyone can argue with wanting to put the money in, say, a $150 transformer, rather than a $150 front panel.

Of course, there are limitations. I wouldn't want to do an ugly box, or a "typical" box. You know typical: a front panel on a steel chassis, or front and back caps on a standard extrusion.

But, bottom line: an L-bend aluminum piece is less expensive to do than the U-bend we've been doing. And it still looks really good—it carries the front panel visual seamlessly to the top of the product. So this evolution shouldn't be too shocking.

The invisible fasteners … that was a lark. I wanted to see if we could do it. We actually did prototypes that had both screw-on tops and fastener-free tops. The fastener-free tops won … but we had no idea that they would work at first.

You see, the key to invisible fasteners is a blind PEM … a fastener that fits a mating keyhole slot on the chassis. After our chassis supplier proved they could insert the blind PEMs without marking the chassis, we had these custom-made for us, from a nonstandard alloy (so they could be anodized with the chassis), and for a nonstandard chassis thickness.

Of course, no story is complete without a last-minute complication, and Jotunheim is no exception. When we got the first boards in from the assembly house, they were shockingly hard to assemble—as in, almost impossible to get into these new chassis. A quick measurement confirmed that the volume pot had been moved over by 0.050 inch. Why? I don't know. I don't remember doing it. But what was done was done. Luckily, we had enough time to scrap the boards and re-order correct ones … and the timeline was preserved.

And the feet? That was intentional. Mike doesn't like stick-on feet (not that I blame him, it made sense when you could choose to run an Asgard vertically or horizontally, but it doesn't make much sense for the line now.) So I expect you'll be seeing a lot more press-in feet as time goes on.

And … that's the story of Jotunheim.

Consider this our third generation. Or maybe even a little beyond.

Now, you get to let us know what *you* think.

2016, Chapter 12

Word of Mouth at the Speed of Light

So, two weeks after the introduction of Jotunheim, Google is returning 186 000 results for the very specific phrase, “schiit jotunheim.” Note these results aren’t for “schiit” or “jotunheim.” Just the complete phrase “schiit jotunheim.”

Why am I belaboring this (seemingly trivial) point? It’s simple: *because that’s a hell of a lot of results for an obscure word combination that didn’t exist until two weeks ago.*

Still not impressed? Cool. Consider also:

1. Not one of these results is from the traditional audiophile press (AKA “big” audiophile press, dominant audiophile press, etc—you know the magazines.)
2. Only a few of these results are due to mentions on the large social networks. Again, you know who they are.
3. In large social networks like Reddit, the results are due to mentions on specific subreddits—in other words, communities of like-minded

people. You know, like people who are into headphones.
4. And, by far, the vast majority of results came from focused communities around the world—communities like Head-Fi.

Now, to be clear, yes, we had press at the Schi-itShow, and we had nice mentions in some of the less well-known press (thank you, HiFi Plus, Positive Feedback, Audio-Head!). And yes, we did have a great intro here at Head-fi, including a video with measurements. This isn't a screed about how press is now important in this Grand Cyber Age. This is more a look at how the rules have changed in marketing—and what you might want to do about it, if you have a company of your own.

Yes, I know, sighs and groans. Another marketing chapter. Sorry about that. But I find it fascinating, especially in light of the Jotunheim results. Which could have been verrrrryyyy different … but I'm getting ahead of myself.

And yeah, I know, some of you still aren't impressed. "Only 186 000 results," they sniff. "There are, like, 1.75 million results for "beats headphones."

Okay, fine, agreed. We are quite small in the

grand scheme of things. Hell, I could go so far as to say that Schiit is a microscopic company. But then again, maybe you should be asking what kind of tiny company you can create, without the participation of the major media, and in-between the margins of the giants.

Or, I could note those results are for a *whole company* versus *one product*.

Something to think about, hmm?

Old Rules, New Rules—and Early Prescience?

If you've done product launches back in the dim dark days of the 20th century, you know all the old rules. The old rules, broken down and snarkified, look something like this:

1. **For the love of God, make sure the biggest press possible pays attention to it.** This means one (or all) of three things:
   a) Introduce it at a show like CES, where all the press will be
   b) Engage a very expensive PR firm to get your story in the best press they can find
   c) Plan on leaving your company for a few weeks to do the world's most tedious dog-and-pony show, AKA a "press tour"

2. **For the sake of Pete, beg borrow or steal the biggest advertising budget you can get.** Because, let's face it:
   a) The first press mentions don't last long in the magazines
   b) Advertising works, because you have a captive audience at the magazines
   c) Magazines support those who support them*
3. **Kill yourself if key reviewers don't have the product several months before launch.** Because:
   a) They're writing for issues that are 2 to 3 months out (LOL, yep, remember print?)
   b) You want those reviews out concurrently with the launch, or you're gonna be looking at your own ass on a platter when the Powers That Be find out you botched the intro
   c) You need the tastemakers to weigh in on your product and (very hopefully) pronounce it All Good
4. **Make sure the distributors and dealers have stock and sell-sheets,** and are constantly fed new salesmaking crap as reviews come out. After all:
   a) There ain't no other channel, they have you by the nuts (and they will never let

you forget it)

b) There ain't no direct sale without ad budgets you can't imagine

* Yes, I know, this is a controversial statement, so let me explain a bit. Magazines also support companies that they find interesting, but have no advertising budget at all. It's not 100 % a pay-for-play scenario. Nor do I think that positive reviews have any real correlation to ad spend. Schiit has benefited from attention from the mainstream press, long before we bought a single ad.

In today's everyone's-staring-at-their-friggin-cellphones-at-a-stoplight, Snapchat-obsessed, Pokémon Go-enabled world, you can blow up every single one of those rules, toot-sweet.

That is, *if* you're operating in *certain regimes.*

If you're a company trying to sell car-priced DACs and condo-priced speakers, it's probably best to follow the old rules to a T.

Aside: or, no. Actually, what I'd do, if I was doing some crazy-priced products, would be to blow up the rules anyway. I wouldn't go to shows. I wouldn't get a PR firm. I would spend almost 100 % of my marketing budget at

flying key press and key dealers out to some very desirable and exotic locations, all expenses paid, about once a year. Much better than sharing a busy show schedule with a bunch of other companies. And, every dealer would get a prepaid, premium "hotline smartphone" to contact us directly for any concerns. Hell, depending on the price of the system, the owners might get the phones, too. (But don't worry, we have no interest in doing this—we don't kiss people's ass very well.)

But, if you're selling products that are relatively affordable and reasonably accessible, you can shred the old rules.

Why?

1. The press probably doesn't care too much. What's left of the traditional audiophile press seems to be more focused on very expensive stuff. Why bother chasing what you can't catch?
2. Your advertising budget is better spent on measurable venues. I've talked about this before, so I won't belabor it. Think AdWords.
3. Key reviewers may not hit your target audience. Or your target audience may simply dismiss them out of hand.

4. Distributors and dealers mean you can't compete on value. If they control your sales channels, you're not in control of your company.

So what do you do? What are the equivalent "new rules?"

To figure this out, we have to take a look at why the old rules worked. The old rules were based on a centralized, top-down model where information came from a limited number of trusted sources, and products came from a limited number of vetted retailers. A prospective buyer would first see the product in a PR blurb or show report, then read a review or two, see an ad or four, then decide to go down to the dealer and listen for themselves. If everything went well at the demo, he could go home with the product (or another one—many dealers, like Amazon, aren't real picky about what they sell, as long as they sell *something*).

So, if you didn't have the press attention, the reviews, the ads, and stock at the right dealers, you were done.

Today, the rules have to be based on a decentralized, bottom-up/top-down/sideways-sideways

model where information comes from an unlimited number of forum posters, acquaintances, friends, casual reviewers, serious reviewers, big press, alternative press, wags, shills, and marketers-in-disguise, and products come from Amazon, or an online shop, or an online retailer, or maybe even eBay, Etsy, Rakuten, or a dozen other "click to buy it now" options. And yeah, maybe dealers too.*

* Let's be clear. Dealers are great if there's one nearby and they have the stuff you're interested in. If that's the case, by all means get your ass over there. And BUY THE STUFF THERE, not online. Not even if it saves you $17, you cheap bastard.

And here's the thing: in a decentralized model (in terms of popularization, review authority, and channels), there are a lot less opportunities for rules, period.

But I may have anticipated a couple of them, way back when. One of the first things I wrote in this book (on page 11 of this mega-thread, to be exact) was:

Corollary 5: on the other hand, micro-social almost always works, unless you're a dick.

Finding the small, specific, passionate communities that are interested in your products, whether they are barbecues, espresso machines, audio gear, or high-end bicycle accessories, is almost always worth it. Going out, joining these communities, answering questions that come up, and not selling at all is a wonderful way to get the word out. But don't think you're King Salesman of the Universe out to convert the masses, or start attacking other brands, moderators or forum members. One problem: most agencies are too lazy to do this hard work. And it is hard work. Pay lots of attention to micro-social, and be prepared to post, respond, meet new friends, piss some people off, delight some others, and become part of your specific niche.

And, perhaps even more importantly, I made this decision before the first Asgard shipped:

Ship orders first, reviews can wait.

Why is this decision so important? Because it's so alien to the Old Rules. In the Old Rules, you would never skimp on the review units. You'd always make sure they were out first. Not sending review units was unthinkable.

So, you might be asking, how did this decision come about? Actually, it was fairly deliberate. I

knew that we'd be selling direct, and that there were communities of people who would talk about the product, and the owners would determine whether we lived or died. Plus, I figured (and maybe this is the real breakthrough), *if we can make the owners happy, and they talk about how happy they are, any future reviewers will be less likely to crap on their fun.*

Now, I know this isn't universally true, but I do think it's an important decision, and a correct one.

Aside: so how did Jude get that first Asgard for review, if we didn't give out review units? He bought it.

So, if you tasked me with synthesizing a new set of rules, it might go something like this:

1. **For the love of God, pay attention to micro-social.** This is where the action really is, for any niche product. Each community will have its own key commenters, pro reviewers, and opinion leaders. You don't need to know all of them, but you need to be willing to go in, talk to them, listen to what they have to say, and apply what you learn there.
2. **Make sure you can sell stuff as soon as you introduce it.** Or, at the very least, soon

afterwards. Maintaining the excitement for longer than 4 hours is a sign you should see a doctor—no, wait, that's something else—I mean, keeping everyone interested for months as deadlines slip and slip and slip isn't gonna win you friends. It matters less where you sell it, just have stuff for sale.

3. **When the time comes, carefully choose your reviewers.** There are gonna be a ton of people who want to review your stuff (usually keeping said stuff for free). But you should be choosing based on the reviewers' level of trust (in their community, or in their magazine—look at comments on their reviews), not their willingness to review everything.
4. **When the time comes, start advertising using measurable media first.** Yes, we know, a full-page ad looks sexy, but do you know what it's selling for you? Start with AdWords, track back to sale, learn what works … then go for the magazines when you're very comfortable with the cost.

And, in today's decentralized world, I'd add a fifth rule, which may be the most important rule of all:

If what you do is working, do more of it.

Because I don't think I have all the rules down.

Things keep evolving. We keep learning. And, let's face it, the bottom line is the bottom line.

The Walls Come Down

If you ask a typical tech wag about why things are so different these days, they'll be quick to state the obvious: the old models of one-to-many broadcast advertising and one-to-one sales are gone. The walls are down. It's all many-to-many.

And yeah, they're right. That's the big deal. The key difference between today and All The Ages That Came Before is that everyone can talk to everyone else.

Incessantly.

And they can post pictures.

And videos.

And they can all comment on the posts, pictures, and videos.

And they can do this in their cars, at stoplights, and forget to actually frigging move when the light turns green, making getting around town a blasted pox … no, wait, I'm ranting, reel it back in there a bit.

But the walls coming down and enabling peer-to-peer communications is an easy, facile answer. It's really only the start. I think there are at least three more reasons things are different these days:

1. **Pricing ghettos are blown up.** In the past, magazines tended to compare products of roughly the same price—which meant that nobody would seriously compare, say, an Adcom amp to a Krell amp. In a decentralized, many-to-many world, there's no such pricing ghettos. Now you can compare a Magni to a $10 000 tube amp if you want to.*
2. **Opinions are equal if language is equal.** If you can express yourself well in written prose (or in video), your content has equal footing with any "professional" review—it can be seen by the same audiences, and judged on the same merits. You don't need to make it to a rare staff writer position in one of a handful of publications.
3. **There are exceptional talents and voices who deserve to be heard.** Let's face it. There are a ton of great listeners who either (a) wouldn't pass the sniff test at a magazine, (b) do not have the time or inclination to devote their lives to audio reviewing, or (c) are interested in stuff so far off the mainstream

that the magazines wouldn't publish them anyway. Now we get to hear from a much wider—and more interesting—range of people.

* Why is this important? Because products can overperform their price points. And sometimes comparing a product to one that costs several times as much can result in some interesting inversions. That is, if the reviewer's own confirmation bias allows it.

So, should we all be happily skipping and dancing our way into a brilliant new future where everyone is equal and everyone has a voice and we're all unique and perfect snowflakes forever and ever?

No. Because buttheads and cellphones. And other things.

The Dark Side

All this wonderful many-to-many communications is great, but that doesn't mean it's perfect. Not by a long shot. Maybe a hundred years from now we'll have this all figured out.

But for now, word of mouth at the speed of light isn't perfect. Let's look at how it breaks down.

- **The problem of identity,** or the Invisible ******** Syndrome. Let's face it, when some people can hide in online anonymity, their personality can change—and usually not for the better. It's easy to snipe, bitch, and carp about everyone and everything when nobody knows who you are, and there are no real repercussions. These people don't care about contributing, they care about getting a response. I don't envy the moderators of this site.
- **The problem of agendas,** or You Don't Know Who They Work For. I'm using this category to cover both paid and unpaid shills. Paid shills are pure evil, and they are probably much more prevalent than you think. They're also becoming automated, and much harder to detect as nonhuman. Welcome to the future. Unpaid shills are people who have decided, for good or bad reasons, that One Brand is the Best Among All. And yes, I know there are (unpaid) Schiit shills out there. We didn't make them. We don't encourage them. And yes, they are as misguided as shills for any other single brand.
- **The problem of bias,** or My Watch Is More Expensive Than Yours Folk. Frequently seen with people who can afford very pricey gear.

To them, obviously an inexpensive component can never approach the performance of the Holy Stack, so don't even bother comparing them. Or, if you do compare them, know that the outcome has been decided in advance—and that is that the more expensive component is, of course, better.

- **The problem of last comments,** or What I Read First is Most Important Disease. Once a thread has run its course, it can devolve into new users asking questions that have been answered two dozen times before, or (worse), users posting something they read in one post elsewhere as The Whole and Gospel Truth. Either has the effect of stirring up the thread to no good effect, and making it impossible to get an idea about the product or topic from reading the last page of comments (which is where many people start—to "cut to the chase.")
- **The problem of buried content,** or Low Signal To Noise Ratio. Due to all of the above, our wonderful disintermediated many-to-many culture results in a lot of noise … which then buries the important content (how does this perform, does it work with the THX-1138, how do you install the new firmware, etc) under tons of posts about dogs, computer interfaces,

lawnmowers, “did you have this problem … oh, I forgot to plug it in” and other cruft and nonsense.

So should we go back to the old days? Not a chance. I believe that a start-up company has a much better chance of success (and a much, much higher chance of staying sane) than in the “good old days.” Yes, the magazines (at their best), were a potent source of distilled information, and dealers (at their best) were a great place to try our gear. But note the caveats. And ask yourself: *How often were they at their best?*

No system is perfect. This is what we have now. There’s no going back.

So, how do you make the most of it?

Making Connections

First, a comment: the subhead above isn’t a throwaway phrase. I mean it, in all senses of the term. Your mission is to make connections. Lots of them. On every level.

Because, in a many-to-many communications environment, getting to the Most Important Person In The Room isn’t the end-game. There are lots of

important people. And you probably don't know a hundredth of them.

So, setting up for success starts with planning to reach as many people as possible, on as broad a swath as possible. In the case of the Jotunheim introduction, we used our most potent vehicle for this: the SchiitShow. This once-a-year event invites regular folks from communities like Head-Fi—and press—to come out and see what we have in store. We don't say what they're gonna see beforehand, but you can be assured we save it for important announcements.

Now, we didn't have to do a SchiitShow. We announced the Modi Multibit with no show at all (not even TheShow), just a press release and done. And Modi Multibit is an important product, too. It's selling very well.

But, in the case of the SchiitShow, we had an opportunity to build interest—about a month's worth of run-up where people could wonder what we had to show. This is more powerful than simply announcing.

Why? Because it's mysterious.

And people like mystery. Or they don't. The important thing is, they have an emotional con-

nection to it. And emotional connection is absolutely important. Although I'm gonna hear some howling about this, I think that people don't care very much about equations and specifications. Not really. Not when it comes down to it.

But people do like companies that do things for a reason—as long it's a reason they can connect to on a visceral level. Like our manifesto about blowing up the high prices in high end, a few chapters back. Have you wondered why it's gotten more likes than any other chapter? Because it's written from the gut, emotionally, it has a cause that many people agree with, and can connect to on an emotional level.

Still don't believe me that equations are bunk and emotion is all? Consider that we would sell exactly zero hybrid cars based on economic analysis. Throw in saving the planet (or, better, some cute seals or penguins) and suddenly you can't make them fast enough.

Cynical? Perhaps. But if you're planning on making your company succeed on the back of cold equations and stainless specs, you're probably gonna be in for a big surprise.

So, setting up for success also includes an emotional hook. For the product. For the company.

For the launch. Don't miss that.

And what happened with the launch? We had some great press. We had a wonderful video intro here on Head-fi. And we also had one of the busiest threads here and on the headphone subreddit.

And we had … some not so thrilling comments. Boiled down, they said, "Well, yeah, it's pretty good, but it's not gonna scare any of the Big Expensive Iron." There was also some grousing about only having Jotunheim at the show. And there was grumbling about how I'd oversold the revolutionary nature of Jotunheim.

Not so hot, huh? Well, that's when you have to remember the third part of setting up for success in a disintermediated age: *step back, don't overreact, wait and see, and respond only when it's appropriate.*

Note that doesn't mean "ignore it." It means pay attention, take notes, and learn.

One of the things we learned is, while the press totally got the revolutionary aspect of Jotunheim (its upgradability), they'd been primed. They got an introduction to the product, courtesy me.

So, the next SchiitShow, everyone gets the same intro.

The other thing we learned is that, well, I'm an engineer, and engineers sometimes get excited about things normal humans don't. Jotunheim's new topology was a big deal to me, both due to the gestation and development, and how it performed. Some early commenters didn't really see this as significant. And yes, I totally understand that. Sound is what matters, 100 %, to the end.

Aside: for the record, I think Jotunheim is our most character-free amplifier. I mean this not as an insult, but as perhaps the highest compliment. I happen to like the slightly warmer, richer presentation of the Mjolnir 2, but I don't think it's necessarily more accurate than Jotunheim.

But, as the less-than-thrilled comments continued, I forced myself to sit on my hands. I knew we were right—I knew this was a revolutionary product, and that sonically it could stand up to some very end-game gear.

Eventually, I got a chance to comment. And I didn't say what a lot of people probably thought I was thinking. Instead, I encouraged everyone to

listen for themselves, rather than deciding based on a quick listen or two under show conditions.

What happened after that? Well, once people started receiving their Jotunheims, the comments quickly turned around. Suddenly the amp was a top performer, and even the DAC was pretty good (it's hard to follow the Modi Multibit with a delta-sigma DAC, even though I think Jotunheim's DAC is quite good … ah well.)

And that brings us up to present day, where Jotunheim is selling strongly … and causing me to wonder about how it affects our other products.

Maybe I'll ask you a few questions about that next …

2016, Chapter 13

Into the 2-Channel World: The Saga of Saga (and, um, Freya too)

Saga and Freya? Yes. As in, both were developed at the same time, so I think it's best to talk about them together.

But I'm really getting ahead of myself. We shouldn't be talking details at all yet. Even product types are more detail than we need. In fact, what we need first is a primer on what "2-channel" is, how it differs from "desktop audio" or "personal audio" or "multichannel" and "home theater."

So let's start with that.

Backing that Schiit Up:
A Brief Disambiguation of Types Of Audio, with Bonus Historical Perspective

Okay, let's start at the start. Once people had figured out that audio was, in fact, comprised of pressure variations in air, they soon discovered that they could record it and play it back. The first commercially viable system was, of course, Edison's wax cylinders. They were created by cutting a recording directly onto the cylinder, and played back using a purely mechanical amplification device—a needle, diaphragm, and horn.

If this sounds less than ideal in, well, all terms—from the slow cutting of audio cylinders, to handling and mounting the cylinders themselves, to the just-a-few-minutes play time, to the inevitable squawks and resonances of the mechanical horn amplifier, to the fact that the playback devices didn't have a volume control (it's mechanical, remember? Want less volume, get a smaller horn. Or stuff a sweater in it. Literally.), it sounds less than ideal because, well, it sucked major balls. But when it's the only thing ya got, well, it's pretty miraculous stuff.

Why do I go so far back? Because, at this time,

nobody thought in terms of more than one channel of sound. Everything was mono. Even when advances came in playback technology—starting with flat records that could be pressed in molds, rather than painstakingly cut on a cylinder—we were talking about mono sound.

Fun fact: 78 RPM records had a play time of only 3 minutes per side. Yes, three. As in, three minutes. Sit back and think about that for a while. And, as an added bonus, the records were full of abrasive. Yes, go back and read the previous line again. Instead of smooth, quiet vinyl, they purposely loaded the record with grit that would cut the needle to fit the groove. No, you can't make this Schiit up. But if you ever wondered why there's so much background noise on a 78 RPM record, well, the fact that it's full of abrasive scratching against a crappy dull needle is the reason.

> And, a corollary to this fun fact: some audiophiles still consider 78s to be the peak of recording fidelity. Why? In many cases, there were no electronics in the signal path—it was a purely mechanical recording. Yeah. Perfect sound forever … well, except for the mechanical resonances of the recording chain, the extremely limited bandwidth of the transducers and the medium itself, and the oh-my-gawd, what-the-hell-is-that-stuck-toilet NOISE. Note to the Luddites out there: sometimes there is real progress in audio. Honestly. Electronic recording and the LP were absolutely two of them.

For a while, everything in audio was mono. Home playback was mono. Sound reinforcement was mono. Movies were mono. Performances were recorded in mono. Radio went out in mono (and, if we're talking AM, extremely band limited—like 4 kHz top end there.) You get the picture. Mono mono everywhere, and not a "stereo" or "2-channel" to be seen.

This started changing shortly after the introduction of the LP (yes, there were MANY mono LPs). The recording industry, in an attempt to siphon as much money out of the pockets of their victims

as possible … er, wait, never mind, let's say "in an effort to increase the overall realism of the recorded musical experience, and to provide a new dimension in sound," introduced the idea of "stereo" recordings—that is, recordings with 2 channels.

Stereo was a fundamental shift in the audio universe—and an absolute boon to the audio industry.

Let's start by talking about the fundamental shift first. Stereo was a huge shift in many ways, but the biggest thing it brought to audio was the ability to recreate a virtual "image" of the performance—the illusion of breadth and depth in the playback. Coupled with the very low-noise LP technology (and the insane luxury of having 22 minutes per side before you had to flip the record), stereo was a revelation.

Of course, many early stereo recordings didn't show off this virtual "image" to its maximum potential, instead relying on gimmicky "ping pong" stereo effects with instruments or performers hard-panned to one side or the other. But the potential for a much more realistic recorded music experience was there.

It just required a complete re-thinking of everyone's audio system.

And that's why it was such a boon to the audio industry. Instead of one channel of amplification, audio enthusiasts needed two. Instead of one speaker, two. Instead of a mono cartridge and mono preamplifier, they needed stereo sources.

So what happened? People went out and bought.

They bought big. They bought all-new stereo consoles that provided an all-in-one way to enjoy this new two-channel phenomenon, or they went and bought an extra amp, another speaker, and new sources. Sales swelled, and everyone was happy.

Aside: the notion of high-end cables wasn't a thing back then, so I'm not counting them as part of the overall economic boom, but yeah, people needed more cables, too—assuming they weren't using an all-in-one console.

Of course, everyone wasn't happy forever. Eventually, stereo became the norm. New companies had sprung up to provide affordable, highly integrated solutions for stereo listening—stereo receivers with AM/FM tuners and phono preamps

built-in. Speakers were sold in pairs. Stereo consoles still did a booming biz. Stereo reel-to-reel tape began to creep in, both for recording and playback.

But sales were no longer exploding.

So what to do? The recording industry, remembering the huge success they had with stereo, tried first to increase the number of channels again—this time to 4. This, they dubbed "quadraphonic." The promise was an even deeper immersion in the recorded music. Which, in its way, made sense. Quadraphonic was kinda the first foray into surround, long before Dolby Pro-Logic came on the scene.

The problem was, quadraphonic records were difficult to encode and finicky to play back, and—the death knell—was that there just weren't that many of them. Audiophiles, by and large, decided to leave quadraphonic—and its associated sales of 2× the speakers, 2× the amplifiers, new sources, new recordings, etc—on the sales floor. It tanked.

From there, the recording industry maintained a healthy distance from formats with more than two channels. The digital audio revolution was entirely a stereo revolution, for example.

But that would change with the coming of home theater. With the advent of Dolby Pro-Logic (and, later, Dolby Digital and other digital surround formats), suddenly, 5 channels plus a subwoofer started to come into its own.

Aside: I think Odeon was the first company to sell speakers in 5.1-packs—5 satellites and a sub—but, as they say, that and $4 will get you a fancy coffee. And, added bonus, Mike's Angstrom company was the first to offer an upgradable surround-sound processor. So, yeah, we know a thing or two about this. And no, we won't be doing surround ever again, thank you. I've covered the reasons why before.

Why 5 channels, and not 4? And what's this with the .1?

It starts with the unique requirements of movies. 5-channel surround was designed primarily for movies, not music. There are two front channels, two side (or surround) channels, and one channel—the center channel—to help keep dialogue and center-of-screen sound, well, centered. Unlike traditional stereo, 5-channel surround is designed so that people can more realistically sit off-axis. The center channel is the biggest part of this.

The .1 channel? That's technically the "LFE", or low frequency enhancement channel—band limited below 100 Hz, it's really just for feeding to a subwoofer to create the thumps, bumps, and t-Rex footsteps that are a staple of movies today. So, in total, 5.1 channels.

For a while, home theater was going to be the Savior Of All That Was Right and Good and Profitable in the audio biz. And, on the surface of it, it made total sense. You needed 5 speakers (or more), plus a subwoofer. Plus a multichannel amp. Or more amps. Plus a surround processor, A/V preamp, or AV receiver. And, since we have now passed through the Gates of Neurosis known as "high end cables," it usually meant lots and lots of cables as well—dozens of yards of fancy wire to connect all the speakers, plus 11 or so RCAs, plus source cabling, plus video cables, plus plus. Great news all around! Everybody makes a ton of money.

Except … well, one little thing. Space.

As in, most people don't have a dedicated space for a home theater setup. And the idea of 6, 8, 13, or more speakers in the living room isn't usually the best idea for domestic bliss. (And, today, with more and more people living in close quarters,

especially people just starting out in their careers, well, the downstairs and next-door neighbors may have the ultimate veto power on the idea of 1000+ watts and 15-inch subwoofers pounding away late at night.)

So, home theater was not exactly the home run everyone hoped it would be. Most people seem content to listen to their latest flatscreen, or (at best) add a sound bar and sub and call it a day. Home theater is a dedicated niche. It may be a relatively stable niche, and it might be a niche where aficionados spend lots on their dedicated systems, but it's a relatively small niche.

Of course, now there are schemes to do 7.1 or 7.2 or 9.2 or 11.4 or 16 or 32 channels (not kidding). Together with 5.1, these schemes have come to be known in the industry as "multichannel." Concurrently with this, stereo has come to be known (at least in the US) as "2-channel."

So, multichannel and 2-channel. Does that make more sense now? Maybe not. But let's stick to them, since they've been defined.

And let's add two more terms: desktop audio and personal audio. These are where Schiit debuted, and they're the industry categories we played in exclusively—at least for a short time. We actually

have had a pretty good presence in 2-channel audio since the first Bifrost, and we expanded the reach even further with the Ragnarok and the Multibit DACs.

But we've never really had products dedicated to the 2-channel market.

Until now.

What is a "2-Channel" Component?

"So what separates dedicated 2-channel products from desktop products?" Someone might ask. "A lot of your products already have preamp outputs, Ragnarok can also drive speakers, and you already said that plenty of your DACs are used in 2-channel systems."

Yep. True. But, to address each in order:

Preamp outputs are not full preamp functionality. Yes, many of our headphone amps also have preamp outputs. This makes sense, since a lot of people use them to run powered monitors on their desktop, or to inject some tube flavor into their solid-state speaker-dedicated systems. But these aren't true preamps. They're missing a bunch of things, most notably:

- **Multiple input handling.** Most of our headphone amps are single-input devices. Preamps usually accommodate multiple inputs. And by "multiple," we mean "more than the 2 of Mjolnir 2." Yeah, Ragnarok qualifies, but that's really not a preamp—that's an integrated amp. More on that below.
- **A focus on super low noise and other typical preamp concerns.** Although most of our amps are quiet enough to be preamps, they weren't designed from the start for this duty. Some come with significant ease of use restrictions, such as Valhalla 2—which requires a specific start-up/shut-down sequence when used as a preamp, and the large-value output caps (necessary for headphone amp use) may cause DC-sensing circuitry on the input of some amplifiers to trigger, even though there's no actual DC on the line.
- **Remote control.** And this is the killer. None of our headphone amps have remote control, because, well, they're supposed to be sitting on your desk, or near at hand (headphone cables are, in fact, not infinite in length). And Ragnarok doesn't have a remote, because it was conceptualized and designed too early in our thought process regarding 2-channel gear. 2-channel stuff (where you're sitting back in

your easy chair, looking at speakers across the room) really benefits from a remote control.

Ragnarok is hardly a desktop amplifier. If you're using it on your desk, you have a huge desk—or you have much greater patience than I do about how much crap you'll tolerate on your desk. It's really an old-school integrated, designed to be put in a rack. Yes, it's a headphone amplifier. And yes, it's a speaker amplifier. And yes, it really should have a remote. We'll get to that. Eventually.

DACs are just the beginning—lots of people like one-manufacturer stacks. Sure, we sell lots of DACs into 2-channel systems, especially Multibit DACs. This makes total sense, since the superior imaging of the "megacomboburrito" filter is much more apparent on speakers.

Aside: haven't tried your Multibit DAC on speakers? Have speakers? Do it. Trust me.

But, back to the subhead. DACs really are just the beginning. There are plenty of people who want to create a system from a single manufacturer's products. The rationale is, "Well, they designed them together, they probably go together better than a cobbled-together system." Of course, an

equal number of people want to pick the best from each manufacturer and assemble a unique system—so different strokes, and all that. But the reality is: if we did more than DACs, we would sell a lot more than we do now.

There's a giant missing component: speaker power amplifiers. This is our elephant in the room. We're not going to go very far in 2-channel if we don't have power amps that can run speakers. Yes, we have Ragnarok, but a 60 W PC amp is not exactly going to light the world on fire. We need three-digit-into-8-ohms-per-channel kinda amps.

But that's a subject for another chapter. We're here to talk about preamps.

And the more I looked at the preamp market alone, the more I realized that there was at least as much room for disruption as on the desktop—maybe even more …

Early Thoughts On Saga and Freya

Preamps are pretty open-ended. They can be all-singing, all-dancing—with tons of inputs, built-in phono preamp sections, hell, with digital inputs and outputs even. They could transform all analog inputs to digital. They could handle all

analog and digital separately. They could do just digital. They could do just analog. They could have a line preamp stage with gain. They could have just a buffer with no gain. They could be completely passive. They could be solid-state. They could be tube. They could be balanced. They could be single-ended. They could be remote controlled. Or they could eschew remote (dumb, but yeah, well, they could.)

So, the logical question: where the hell do you start?

I guess a more logical and methodical person would start by looking at the market, doing a survey of what's out there, compiling a list of features and price points, and using that to create a matrix of threats and opportunities.

Of course, we didn't do that.

In fact, what I more than likely did (I honestly don't remember, which lends credence to this guess) is start drinking and thinking about what would make a good preamp.

Now, I wasn't working entirely in a vacuum. In fact, one member here on head-fi made a very powerful case that doing a passive remote preamp would be a great step for us, based on the very

limited options available in the market. Another data point was the very strong sales of the SYS, a product I did mainly on a lark.

So, "passive" was on my mind from the start.

Aside: "passive" refers to a preamp without a gain stage, or (to be blunt), a preamp without an amp. Yes, I know, the terminology makes no sense. It would be more accurate to call it a variable attenuator. And some companies do. And they probably confuse the hell out of many prospective customers. Hence, we'll keep calling it a "passive preamp."

Also, remote control was at the top of the list. I've had many 2-channel systems, and the ones with remotes are by far better than the ones without. Call me lazy, but jumping up and down a half-dozen times to adjust volume—then doing it again on a different recording with a different average level—is not my idea of a good time.

So, passive, remote. What else?

Well, I was able to strike a couple of ideas off the "big list" right away—digital (as in DAC, ADC, or both functions), and phono preamp.

“Whaaaa?” Some of you cry. “But I’d love to have a product that would handle both analog and digital, and man oh man, there’s a lot of turntables out there these days.”

Yeah. And we have DACs for digital reproduction. And adding both ADC and DAC functionality to a preamp means it isn’t really a preamp anymore—it’s a mixed-signal processor. Or at least that’s how I see it. If you transform everything over to the digital domain for processing, it’s fundamentally changing the analog inputs. If you handle both separately, wow, you’re talking a big, complex product that’s entirely full of digital noise, even if you do it right.

So digital? No. Sorry. Not even as a card like in Jotunheim. Jotunheim is designed to be a compact, do-all product. The preamps I saw more as thoroughbred, end-game stuff.

And phono? Yes, there are tons of turntables out there. And we already have the Mani phono preamp. And if we wanted to do an even better phono preamp, it’s better to devote an entire, much larger box to that functionality. One box, one function gives you a better chance that that single function is well-served. Or at least that’s how we see it.

But even with those two big decisions made, we had lots of choices to make:

- **Passive or active?** Passive works for a whole lot of systems, especially ones with short cable runs and high-gain amplifiers. But there are other systems that need gain … especially if you're an aficionado of music that's recorded without excessive compression, like many classical titles. And if your cables are long (like, say, to monoblock amplifiers), you may need an active gain stage to drive them. It's really best to have both, if you can—but that's a really bizarre design decision.
- **Balanced or single-ended.** Argh. Balanced means a lot more parts, more expensive connectors, a larger chassis (the connectors are large.) But balanced works well with our products that have balanced output (which are usually larger products anyway). Single-ended can be simpler, smaller, and a lot less expensive. Again, it's really best to have both, if you can.
- **Tube or solid-state?** Again, tube is great for most systems. But some people like the certainty of solid-state. And both can sound very good. This is another hard call, and one that is much more a judgement call than passive or active.

- **Motorized pot or relay attenuator?** The most important function of a preamp is volume control. Remote volume control can be done in three ways: with a motorized potentiometer—that is, a motor that actually physically spins a volume pot, or with a relay-switched stepped attenuator, like we did in Ragnarok, or with a volume control chip. I won't even discuss the third option, though it's been showing up in many previously "high end" brands. Sorry, that's simply lazy. And when it comes down between a motorized pot—and all of their associated vagaries, like poor tracking at low volumes and eventual wear—and the relay-switched stepped attenuator, which has near-perfect channel tracking at all volume levels, relays rated for 15 million cycles, and nothing more than a couple of resistors in the signal path, well, it was really no contest. At least one decision was easy.

Since you're reading this chapter after the introduction of Saga and Freya, you already know the punchline. Instead of making the three hard decisions above, we decided, mainly, to do it all:

- **We did passive and active, in both products.** Because, as explained above, different approaches work better in different systems. I

expect many, many people will be perfectly happy with the passive approach, but if you need an active stage, you need an active stage.

- **We did balanced and single-ended, in two different products.** The hyper-affordable Saga, and the still-way-cheap Freya. Of course, Freya does both balanced and single-ended—and, with tube gain, it converts everything to balanced output.
- **We did tube and solid-state, in both products.** In Saga, you can switch between passive and a hybrid tube buffer stage. In Freya, you can choose passive, a JFET buffer, or a balanced tube gain stage.

Yeah, I know. Sounds obvious, right?

But it really wasn't so obvious in the beginning. Hell, I had a completely solid-state Saga laid out before scrapping it and deciding to commit to a tube stage. And even that went through several revisions before I found a tube hybrid output stage that I felt good enough about to make into a prototype.

Freya took a bit of agonizing, too—I thought very seriously about making that a tube hybrid as well, so we could use just two tubes (instead of 4). The all-tube stage won in the end, mainly because Freya already has JFET solid-state outputs.

"Why tubes?" Some of you are asking. "I want solid-state!"

Yeah, and that's why there's the JFET buffer on Freya. Or wait a bit, and I'm sure we'll have octal LISST available eventually.

"But seriously, two tube preamps?"

Well, I just explained how you can run them solid-state, and passive, but here's a longer (and more mercenary) answer: because there are a helluva lot of good affordable solid-state amps out there, but not so many affordable tube preamps. And when you start saying, "affordable, remote-control tube preamps," and "affordable, remote tube preamps from a credible US manufacturer," and, dare we say, "an insanely low-priced preamp, period, even if it was just a remote attenuator," well, the playing field gets down to a field of two very, very quickly.

No. Seriously. Think about all the solid-state amps out there from companies like Emotiva and Outlaw. Think about even more that sell for a pricing tier just higher than those two companies. Then ask yourself how many tube preamps these guys make. Like, zero. I think you'll quickly come to the conclusion that I did … that the market to sell a tube preamp to these possibly "tube-curious"

audiences is much, much higher than a solid-state preamp.

Yeah. There you go. Call me a businessman. That's OK.

"Okay, then, why these tubes? Why not use the same ones as in Lyr 2, Mjolnir 2, and Vali 2?"

Aha, now that's a better question, and one that deserves a serious answer. It's because we hate you and want you to buy a whole new series of tubes, because we recently acquired all manufacturers of 6SN7-type tubes.

Kidding, of course. Dang, you're so serious.

Here's the real reason: because they're better tubes. 6SN7-style tubes are by far the most linear tubes I've ever worked with. At typical line-level output, Saga measures almost like an op-amp based design, with triple-zero distortion figures (as in, 0.000*X* %). That is 100 % insane.

Aside: and yeah, we'll have the usual nonsense about how "this ain't really a tube amp, get a tube amp with tons of distortion so you gets the reeeel tube sounds." That's fine. You want a soft, fuzzy, blurry tube preamp, there are plenty of them out there that provide loads of second-

harmonic distortion. The reality is that Saga and Freya, while the distortion is small, are true tube preamps, with typical tube harmonic structure (mostly 2nd harmonic, decreasing rapidly into the noise on higher-order distortion). They're tube preamps—they're just very clean tube preamps. And Saga is even cleaner than Freya.

So yes, we went with these tubes simply because they perform better. Plus the fact that we can treat them right—with 300 V on them in Freya, and 200 V in Saga—these are real tube preamps, not plate-starved designs.

In short, we used these tubes because they were right for no-compromises preamp design. Please accept my apologies if this kicks off a new tube acquisition spree for you. The good news is that the rolling options are rather more limited (6SN7 and 6SL7 types only, including the Russian 6N8C and 6N9C). Best to stick to 6SN7 only for Saga, unless you want more distortion—the hybrid part of its tube output stage is transconductance-matched to the 6SN7. 6SL7s will work fine, but they'll have measurably more distortion.

The Final Decisions, and First Prototypes

Okay, this chapter is getting ridiculously long, but let's talk about the tech decisions a bit more, the first prototypes, remotes, light pipes, and where we went from there.

For both Saga and Freya:

I mentioned that we were set on using relay-switched stepped attenuation from the start. I didn't mention that we decided from the start to do it just like Ragnarok—that is, with a microprocessor sensing a DC value on a real potentiometer, so the volume knob would operate exactly like a "real" one, with stops at the top and bottom of the range. This eliminates the need for LED light bars, light circles, screens, or other devices to indicate the volume level. It also impacts the remote. More on this later. We also decided to give Freya 128 volume steps, double that of Ragnarok and Saga—and 0.5 dB steps, for finer control.

- **The remote …** we played with the idea of using pretty much every remote technology out there, before coming back to one of the oldest and most boring:

- **Ultrasonic.** Yes, just like grandma's old TV. Huge advantage: it's completely nondirectional. Huge disadvantage: it's pretty much a from-scratch design, including the remote handset. Big pain in the rear end, that—and costly. Also unknown reliability. We didn't go far with that.
- **Bluetooth.** As in, "hey, cool, you can use your phone and an app to control your Schiit!" On more research, also as in, "Hell, I don't want to find my damn phone to use the remote, give me a separate handheld instead." Yes. Even amongst our youngest customers—the vast majority hated the idea. Also kinda costly, and it would require us to produce and maintain apps. We went pretty far with this.
- **Proprietary RF.** Like ultrasonic and Bluetooth, this is also a nondirectional technology. However, we've had experience with very flaky proprietary remotes in the past, and, well, remotes most importantly have to WORK.
- **Infrared.** Yep, the old standby, same thing you use for your flatscreen and receiver and such. That's what we did. It's cheap, it's easy, and it's pretty much painless.

· **The remote, part 2.** Yeah, it's a cheap plastic

credit card remote. Consider how much the products cost. Would you like to add $100 to the cost of a Saga for a machined aluminum remote? Seriously, if enough of you want this, I'm sure we can make it happen. But it simply wasn't a priority. I've seen machined remotes that probably cost more to make than Gungnir. These aren't those kind of products.

- **The remote, part 3.** I've mentioned before that the Ragnarok didn't get a remote because we were unsure about how to manage its behavior. When you use the remote, the volume knob won't turn. Which means your system can end up with a different volume level than the knob indicates. This could possibly lead to some problems if you were careless about it. The way we addressed this on Saga and Freya is to have a clear indication when the remote is in control—and software to force you to turn down the volume below the remote level before the knob takes control again.
- **The chassis.** Yes, these are the new-style, Jotunheim-esque chassis. There's no need for a lot of heat dissipation (well, aside from the tubes, which hang outside the chassis), so there's no need to use the chassis as a heat sink. This saves our butt in assembly. See next

point.

- **The LEDs.** Even though Saga and Freya are minimalistic devices, they have a lot of LEDs on the front of them. So many that we thought about doing custom light pipes (molded plastic tubes to channel the light from surface-mount LEDs to the front panel). In the end, we didn't, mainly because light pipes have their own vagaries:
  - They can't be in groups of 5 or more—they have to be singles. This means they are a pain in the rear end to put into the board, and result in very little or no labor savings.
  - This means that they're never going to align perfectly with the front panel, even if we heat-stake them to the board, even if the front panel is machined to have "guiding cones" to get them in the right place.
  - The per-piece price was still relatively high, even at 100 000 pcs.
- The Jotunheim-style chassis allows us to align the individual LEDs from the underside of the product. This is the main reason. If it was a U-bend, like the Asgard 2/Valhalla 2/Lyr 2, we'd be toast—there would be no good way to align the LEDs that are placed near the center of the chassis.

Cool. So how about Saga alone?

I mentioned Saga's hybrid tube buffer stage, but I didn't go into any real technical particulars. Suffice to say, it's a weird choice. It stacks the tube on top of a PNP transistor. The main reason I did it was to reduce output impedance even further … resulting in Saga's about 170 Ω output impedance. It also reduces distortion. And, in listening tests, I think it sounds better than a simple cathode follower. But a lot of people have told me I'm crazy, too.

And Freya?

I mentioned Freya's JFET buffer, but I also didn't give any real details on that. Bottom line: it's pretty much the same as the buffer we use in Yggdrasil. It's interesting to switch it in and out of the circuit in real time, while watching the display on the Stanford. Interesting in that it's essentially impossible to tell when it's switched in. Except for a 1 dB to 2 dB increase in noise floor, the measurements are exactly the same at line level. The JFET buffer far, far outperforms the Stanford's (very) low-distortion signal source, in terms of THD. The only peaks in the measurement are from the Stanford.

I mentioned Freya's differential tube stage, but again, I didn't give any real details. Let me

correct that. First, by "differential," we mean "differential." This is a real, fully balanced, differential tube stage with tube buffered output. No solid-state in sight. No balanced-to-single-ended conversion. The relay attenuator is fully balanced as well. With a rail voltage of 300 V, this differential stage has insane headroom for a line-level (2 V to 4 V RMS) device.

And in case I missed it, fully balanced and differential all the way through. Run a balanced input into Freya, and it's never converted from balanced and back again. Run a single-ended input into Freya, and it's automatically converted to balanced by the inherent properties of the differential tube gain stage. Use the JFET buffer or passive options, though, and it remains single-ended.

All of these decisions were in place when we went to do the first prototypes, late in 2015. Unlike many of our products, Saga and Freya stayed pretty much exactly the same throughout development.

As usual, when the prototypes showed up, I put them together, fired them up and made sure they worked DC-wise … and that's where the problems started.

The problems were in the transformers—all three of them. Saga's transformer was monumentally screwed up—producing only 18 V on the high-voltage rail, when it should be more than 200 V before regulation. I looked back at the specs, thinking I'd sent the transformer manufacturer the specs for the solid-state version of Saga that never was. Nope. It seemed right. There was no reason for it. I contacted the manufacturer and they agreed to do another prototype.

Freya? One transformer was screwy, producing only 1/2 the output voltage it should (this one was for the JFET buffer stage—it was running like ±10 V rather than ±20 V.) Again, no explanation on this. Maybe the transformer guys were just having an off week.

The other Freya transformer was my fault. I'd under-spec'd the tube heater rail current (by about 1/2), which was pulling down both the heater and the high-voltage rail. It's kinda hard to regulate to 300 V when you're starting with 260 V, you know.

I also screwed up the transformer pin outs—a classic Jason screwup. Sometimes it's because I do the board before I get the transformers, and sometimes it's because they end up changing the

transformers, and sometimes I'm just an idiot. Just like I get power switches backwards from time to time (that's just lazy, not checking my work.)

In any case, the Freya was limping, at least, (and it had a 5 V rail that worked fine), so I handed it off to Dave to do firmware. While he was still working on it, I got a replacement transformer for the Saga that appeared to work fine, dropped it in, and gave him that one to work on, too.

A couple weeks later, Dave came back with some hilariously copper-taped and big-flying-component versions of the preamps. I took one look and laughed.

"I screwed up the power dissipation, huh?" I asked him.

"Just a little bit," Dave agreed.

The problem turned out to be mainly in the heater supply. I'd planned on having regulated DC heaters. That's great, it's what we do in Lyr 2 and Mjolnir 2. But the plan starts falling apart when you're talking 2.4 A of current in the Freya. Even the 600 mA in Saga was tough on the transformer.

“Go to AC heaters?” I asked, doubtfully. AC heaters are fine, but, if not laid out well, could hum.

“I never had a problem with hum on AC heaters with 6SN7s,” Mike said.

“But these are preamps,” I protested.

“Line preamps,” Mike sniffed. “Not phono.”

“Build it, see if it works,” Dave said, ever the pragmatist.

“Okay,” I said. “Then I’ll get new transformers, and you can get the software done.”

Dave blinked, surprised. “The firmware is done.”

“What?” I asked. The boards looked so rough it didn’t seem they could possibly be done.

“They work,” Dave said.

“They work? As in work?"

“Right. The transformers just get reallllllly hot.”

Dave was right. The high-voltage transformer in Freya would literally double as a hot plate after a few minutes. But it worked. The inputs switched. The volume volumed.

“All I need is the remote,” Dave said. I hadn’t received the sample remotes for approval yet.

Aside: we handed off the production of the remote to a turn-key remote manufacturer—we just specified what type of remote, what codes, and what buttons. If we’d done it ourselves, we’d have risked getting tied up in the design and production of it. This way, it was a really, really easy handoff.

> And, like I said, if you want some fancy milled machined aluminum remote, tell us. The more who want it, the higher probability we’ll do it. But we weren’t going to do it to start.

“Do you think it’ll just work?” I asked Dave, miming the remote. I was still kinda floored by having two working (well, limping) preamps, just a couple weeks after handing them off for firmware.

Dave shrugged. “IR is easy. It pretty much always works.”

“If you say so.”

Maybe I’ve been through too many painful product development cycles, but a little voice kept whispering: It can’t be this easy.

Well, surprise surprise: it pretty much was that easy.

I tried Saga and Freya with AC heaters, and they worked just fine—the hum measured no differently than the DC versions, and it was well below what it needed to be for a line preamp. Just as Mike had predicted.

Board layout is critical with AC heaters, though, so I did another set of prototypes—now with new transformers designed for AC heaters. They measured the same, just fine. Though the voltages were still a bit low, argh. So we went back for another round of transformers.

Aside: why so many problems on these transformers? No clue. The manufacturer seems to have internalized (finally) our directive to "make sure the damn things don't hum," but now it seems like they may be sacrificing performance to make sure there's no hum. Ah well, as long as it works in the end.

And—here's the capper—when the remote control samples came in, I handed one to Dave … and he handed me back two preamps a week later, both happily working with the remote.

Like. Friggin. Magic.

Yeah. Almost like we're getting good at this (but as soon as we think this, well, that way lies ruin. Best to expect more problems than you have. There is no power of positive thinking in this one.)

A Seismic Shake-Up in Value?

So what are we left with here? Two preamps, both remote-controlled, both with relay-switched stepped attenuators, both of which can be run passive, both of which ship with tube or tubes, both of which can be run completely solid-state.

And (drum roll) … one that costs $ 349, and one that costs $ 699.

If this doesn't sound like anything else out there—if this sounds insanely, crazily, stoopidly cheep, well, that's because it is.

And it's that way because that's what we do. We design something, we figure out how much it costs is to make, and then we apply our standard margin to it. We don't price things "at what the market will bear," or "higher so people will take it seriously," or "higher because there's literally nothing like it, anywhere near the price."

Call us old.
Call us crazy.
Call us reactionary.
But that's how we think it should be.

We think these are absolute, barking mad deals. Even more now that we've emerged from the design phase and spent some time looking around at the market. Hell, there are passive stepped attenuators (only, not with tube stages or remotes) that cost a lot more! When you start talking 6SN7-based preamps from credible companies, you start adding zeros really, really fast. And when you start talking about relay-switched stepped attenuators, you are talking really, really rarefied air. That's completely beyond-the-pale cost-no-object high-end design. Hell, it's beyond a lot of "high end" companies that have given up and simply use volume control chips now (barfarolaomaticdeluxe, is how I feel about that, but I'm biased.)

So, do I sound proud of our new preamps? It's because I am. I'm thrilled to be bringing some real value back into high-end 2-channel audio. I'm pleased that the price is so shockingly low (I had to run the numbers twice myself to believe it.)

But in the end, it’s up to you.

2016, Chapter 14

Why You Can't Always Get What You Want

Whenever we introduce new products, there's the inevitable "ah hell, why didn't you do XYZ?" comments.

Why didn't you add a third balanced input, why didn't you include a 12 V trigger, why didn't you make the default color black, why didn't you put the power switch on the front?

The implication being, of course, that (a) the addition of this feature (or features) would make the product absolutely perfect in every way (at least for the person who wrote the comment), and (b) it would mean another immediate sale (or sales), and *sometimes* (c) Schiit ain't the brightest light in the firmament for overlooking this obvious and easy addition.

And, literally every day, we get emails from people who would like to have just one little change, like two additional Toslink inputs, or a digital

pass-through output, or a Jotunheim that takes both phono and DAC cards at the same time, or a Ragnarok with no speaker output at all, or a Bifrost with balanced output, or a red chassis, or a knob with a more prominent indicator (that lights up) or variable LED intensity, or maybe even a full custom design to the person's specifications.

Now, I'll be the first to admit that (c) we may not be the brightest people on the planet, but it *also* sometimes seems like people have this crazy idea that the Schiit production floor is something like a cross between Willie Wonka's Chocolate Factory and Doc Smith's Krazee Kreations Lab, where we can just whip up an infinite variety of fun and exciting stuff on a whim, anytime we'd like.

Or, alternately, maybe they think that our factory simply has infinite capacity—that there are hundreds of skilled technicians standing by to build something exactly like what they'd like, from raw chassis on up.

Unfortunately, neither of these scenarios is anywhere close to the truth. So, it is with a heavy heart that I embark on a chapter that might be somewhat akin to telling a kid that there really ain't no Santa Claus, and give you a look at how

modern mass production really works, or Why You Can't Always Get What You Want.

Mass Production in Two Steps

Okay. Let's cut to the chase. This is how literally all mass production works, in every company on the planet, in two steps:

1. We guess what might be a good product for you.
2. We make lots of them, as efficiently as possible.

"Buh, buh, buh, but … big companies do tons of market research and focus groups and use case testing and other big words as well, that's not guessing!" somebody is saying right now.

Wrong.

Market research is guesswork. Market research cannot tell you what the next breakthrough product will be. Market research is not even all that well correlated to success. Something like 1 in 120 new food products make the grade after introduction. Every single one of them tested well.

Focus groups are guesswork. Or worse. The number of ways a focus group can go offtrack is uncountable. All it takes is one loudmouth, and

the whole shebang is invalid. The original iPhone concept would have never survived a focus group, especially one comprised of the "correct" "target audience"—that is, smartphone users.

Use case testing is guesswork. Not only can you not account for crazy use cases you haven't foreseen, but there's no way to predict how other factors will influence users to find workarounds.

So yes, other companies have more guesswork. But in the end, step 1 is still step 1: *We guess what might be a good product for you.*

And note the most important word of Step 2: *lots.*

"Lots" does not mean, "ten thousand variations of 10 pieces each." "Lots" does not mean "each one individually customizable for cosmetics and feature set." "Lots" does not mean "a total run of 75 pieces."

Also note the second most important word of Step 2: *efficiently.*

"Efficiently" does not mean, "each one hand-carved out of a single sustainably-sourced log of Egyptian Amberwood," or "each one CNCd from a solid block of virgin aluminum" or even "each one made ready to paint and finish to the exact specifications of the customer."

Yes, I know. Not very sexy. But there are good reasons for this focus on functionality, which I'll get into later. Because I know there's at least one guy out there saying, "But that ain't right! I heard 'mass customization' is the wave of the future."

Yes, that may be. But it certainly ain't the wave of today.

The Fallacy of "Mass Customization"

There have been any number of articles about what a wonderful world it would be when companies start having the ability to make highly custom things for each individual customer (usually enabled by various half-baked or pie-in-the-sky ideas like 3D printing or full-bonkers nanotechnology.*)

* Before I am eviscerated for this, let me remind you that (a) I was involved in marketing several 3D printing technologies, starting in the dim dark days of stereolithography, and (b) I am a science fiction writer who is familiar with nanotech-as-it-exists-today. 3D printing is half-baked, period. The finish is (largely) crap, and the physical properties are not congruent

with modern consumer expectations, to put it politely. Also, it's expensive. Nanotech, as in the fantasy 2058 Popular Science article "Nanoassemblers Make Everything You Want for Free" simply ain't real enough to discuss. Nor may it ever get much beyond, say, plant life. Look at the energy densities involved, say, in "growing" a gun out of scrap car parts.

However, despite this half-baked nature, the idea that "mass customization" might be a good thing has taken hold in some minds. And yes, I understand the allure. It's really tempting to imagine a day when you can simply tweak a product to your needs, and still expect it to show up in two days via Amazon Prime.

But we're a long way from there.

How long? All it takes is a quick look at the iPhone. Here's a consumer product that's made by one of the largest and most valuable companies on the planet, using some of the highest technology on the planet.

And, last year, Apple sold (roughly) 230 million of them. That's nearly *one quarter billion devices.*

And they made them in how many variations? Twelve.

Yes, 12.

Two sizes, three colors, two memory sizes. (IIRC) That means, roughly, about 20 million devices per variation. Or, in other terms, Apple makes about the same number of iPhones per variation per year as there are cars sold worldwide, in total, from all manufacturers.

And let me know what happens when you ask them for a headphone jack.

Yeah. Thought so.

Bottom line: if one of the most advanced devices, from one of the largest companies, made in numbers that are absolutely mind-boggling, cannot offer customization beyond a couple sizes, colors, and memory capacities, well, we are *one hell of a long way* from true mass customization.

> And yeah, I know, oversimplifying, there are carriers and such to force the need for other variants, and yeah, now they have five colors, but yeah, still, let me know what they say about that headphone jack. Or a button that actually clicks.

Annndd … the trend in most things is actually towards *less customization*, rather than more.

Skeptical? Look at cars. Today, a popular car may have 3 to 4 interior colors and 8 to 10 exterior colors, maximum. And by "interior colors," I mean. "Change the seats and door panel inserts, leave the rest alone."

Contrast this with a vintage Mustang, which had 16 to 20 exterior colors and 10 to 12 interior colors. And by "interior colors," that meant, "Every frigging piece hosed down in turquoise, if that's what you want."

So, yes, we've actually *lost* customization options.

Why? Well, when your industry has gone from vertically integrated to one that works with sub-suppliers, and when any significant change might trigger the need to re-certify the car with all the various government testing agencies … well, it's not as easy to make as many variations.

On the other hand, cars today usually have many more features, are much more pleasant to drive, are beyond-the-pale more reliable, and don't cost that much more (a 1965 base Mustang, in constant dollars, was about $18 600—and that's for a 3-speed, inline-6, no-power-steering, no-power-brakes, no-power-nothing, no-air-conditioning, no-seatbelts coupe.)

"But why?" you might ask. "I want my mass customization! Someone must be able to do this!"

Well, yes, of course. Hell, they were doing it in the middle ages. When someone went to get a set of armor made for themselves, they didn't go to the Ye Olde Mart of Wal and pick a size 40 long. The suits were made exactly to the person's individual measurements and body vagaries.

Of course, they cost as much as a Ferrari.

But this isn't the only downside to mass customization. Consider:

- **What are the workflow ramifications?** If you intend to do mass customization, how is it done? Do you stock a whole lot of different panels, each with different holes, or in different colors, or both? Do you have PC boards assembled to a point, then fill them in on the order? Or do you have a complete machine shop and plating facility in house and do stuff to order? Do you do the same with boards?
- **What is the time cost?** How long does it take to do the customization? Certainly longer than pulling a closed box off a shelf. If you're starting from scratch with each order, how long does it take to produce a final product?

How do you cost out each variation? What do you do if the raw materials aren't available? What happens when you have 350 purple knobs, but you only need 10 silver?

- **What do you do about testing?** Variation implies changes in specifications and capabilities. Are all variants fully qualified and certified by all applicable agencies? Are test procedures established for all variations? Are you expected to make it up as you go along?
- **What do you do about returns?** While you may love the idea of a green Mjolnir 2 with pink knobs and a limited-edition 1973 Rolling Stones litho in white across the entire chassis, what happens if you don't love the sound, and it comes back to the shop? Are we supposed to find a buyer for this unicorn? Or are we expected to repurpose it to meet the next customization request? What effect does this have on price?
- **How do you provide service?** Multiple variations also imply changes in Bill of Materials (that's what goes into the product), schematics, specifications, capabilities, ATE scripts, etc. Is all of this documented for every variation? How are the documents stored and tracked? Will the technicians have the docs easily at hand when it comes time for service? What

impact will that have on warranty length and out-of-warranty cost?

- **Who deals with the future?** So Bob of Bob's Amp Shack (Proudly Making Whatever You Want For Money Since 1986) decides to hang up his shingle. Does that mean you're hung out to dry on your custom product? On the other hand, a large company with documented designs will (at worst) be acquired if the founders lose interest, which means there'll still be support.*

* This is not a hint. Mike and I aren't bored. Not by a looooooonnggg shot.

The Nuts and Bolts of Mass Production

The reality is that modern mass production is the end-result of a lot of research, design, engineering, purchasing, subcontracted assembly, and internal procedures. By the time we get to making a single product, we've burnt thousands of hours and invested tens to hundreds of thousands of dollars.

And, when we make that product, it is well-specified and understood. It has gone through several rounds of prototyping. Production testing

has been defined. Pass/fail guidelines have been established.

And … that's also why we need to make them all the same.

If we had to prototype each variation—and adjust testing for each variant—and do it on the fly, according to customer request … we simply wouldn't be able to do it. At least not in today's volumes and at today's prices.

Consider that when we're making something, we have:

1. Hundreds to thousands of chassis from an external supplier in-house—all the same. Per product.
2. Hundreds to thousands of PC boards assembled by an external supplier in-house—also all the same. Some programmed and tested too. Also per product.
3. Innumerable small parts—tens of thousands of knobs, setscrews, hundreds of thousands of binned screws.
4. Tools to do the programming, testing, and assembly.
5. A standardized programming procedure (if the product needs it.)
6. A standardized assembly procedure.

7. A standardized test procedure.
8. A standardized burn-in time.
9. A standardized re-test procedure.
10. Hundreds to thousands of boxes and inserts to pack the product into.
11. Space allocated for packaged goods that allow for efficient packing and shipping of orders.

All of this is what allows us to make a Magni 2 at $ 99, or a Freya at $ 349, or a Modi Multibit at $ 249.

Aaannnd … repeating for effect … this is for *each product,* and *each variation.* And yes, that includes variations like Bifrost 4490 vs Multibit, or Jonada vs Jono bs Jodac (Jotunheim alone, Jotunheim Phono, or Jotunheim DAC). And we also have some black chassis from time to time.

Or, in other words, this is already a *lot* of variation.

Can you imagine what it would be if we allowed even more customization? It would be, to be frank, complete and utter insanity. We'd need double the staff. Prices would have to go up.

Hell, I can already make a very strong case for culling the line.

Before you panic, let me say this: no plans are in place for a cull. We're watching sales in the

aftermath of the Jotunheim intro, but that's all we're doing. Waiting and seeing. So far, every product has its fans, and there are no real stinkers. We'll see what happens next year.

"But I'm not interested in variations, I'm interested in you adding (this particular feature here) to all of your products," you might say.

Cool. Got it.

Unfortunately, the same rules apply. Note the statement above about "once we start production, the blueprint is already in place, parts are already made, etc." In other words, the ship has left the dock. We're not going to be able to turn it around.

And, once we're running, we don't really want to turn it around. Making changes on a production device—no matter how seemingly small—can have devastating ramifications on production. Oops, that change accidentally moved a hole, and all the new chassis are unusable. Now we're 12 weeks out from shipping. Oops, that change actually had a negative impact on the product performance and it doesn't meet spec. Ah crap, that may mean scrapping the boards—with a ton of expensive parts on them. Oops, the new and old metal got mixed, some of it kinda-sorta fits, and some of it doesn't. Production pandemonium!

That's why we usually eschew running changes for measured change—in the case of the amps, add a "2" and include a bunch of features people have been asking for, or which we think are now a good idea. In the case of the DACs, announce an upgrade. But not every month. And not every decade. And make sure there's a smooth process in place for the upgrade.

The reality of modern mass production is that there are a lot of moving parts … for every individual product. So that makes change difficult, painful, or impossible.

And the other reality is that we don't guess right all the time. Which makes for feature sets you may not consider ideal. I understand. We are not perfect.

However, we do listen, and we do get better. Look at an Asgard, an Asgard 2, and a Jotunheim. Personally, I think that's real process.

And, we'll also continue to improve processes. It's possible that we should look at some more vertical integration, at least in some very limited areas. That might happen. That might increase our flexibility. But that flexibility doesn't mean more variations. More likely, it means "less time

spent out of stock" and "more rapid product development."

> Of course, that last one also means "more products," which, well, is "more variations." Ah well.

How To Improve Your Chances of Getting What You Want

"Okay, so if we can't have mass customization, and the ship already sailed on your current product features, how do I get what I want?" you may ask.

Well, I wish I had a better answer for you, but it comes down to: adapt or wait.

"Adapt!" you yell. "I'm not here to accommodate you! I'm here to get the thing best suited to my needs!"

Yes. I understand. And if some other product is more suited to your needs, you should absolutely buy it. That simply makes sense. Blind brand loyalty—to anyone—makes no sense.

However, if you love how a product sounds, how it works, how it looks, you may just find a way to make it work for you. Even if it is not 100 %

perfect. And especially if it's a small fraction of the cost of the competition.

Alternately, you can let us know what you're looking for, and you can wait. If you've noticed, our updated products usually address the things that people most wanted in the design, within the limits of the chassis, of course. (Changing an upgradable DAC to have different inputs/outputs, for example, is more challenging, because it necessitates a chassis change.) We do listen, and we do respond, especially to good arguments.

Unfortunately, there are some things that are difficult to make happen. These are usually due to sound engineering reasons. For example, the home theater bypass mode on the preamps. Yes, we could actually add an "invisible" home theater bypass mode that is enabled by holding down one of the buttons for 2 seconds, like the Ragnarok output selection switch. But this has several problems:

1. How do we indicate it's in HT bypass mode? There's no light on the front panel for this. I guess I could have all the LEDs on the front panel blink on and off, but that might be slightly irritating.
2. Will you accept you can't trigger it from the remote? I'm not sure we can determine what

is a long button press on the remote from a single press.

3. What happens when your kid triggers it and passes full volume Rammstein through to your 2500 W Class D amps, and you're dodging voice coils?

Yeah. Kludgy. I only added the same kind of kludgy mode switch on Ragnarok because it seemed like a good idea at the time. I can guarantee you that the next generation of Ragnarok* will show what's connected (speakers, headphones, both, neither) on the front panel, and that it will have a front-panel button for it, and that it will have remote control. All of these are kinda "like duh" stuff that we missed the first time around …

… because we guessed wrong.

* No guesses on when this will happen. My plate is really full. Let's get Vidar perfected, produced, and shipped. And a couple of other things. Then we'll see.

So let us know what you need. I can't guarantee that we can accommodate every request (and I can't guarantee that my obstinate attitude won't kick in when I read another hyperbolic screed about what a massive tragedy it is to have the

switches on the back, or that I'm blinding your unborn child with the LEDs on the front), but we do listen … and we do change.

And, who knows? You may absolutely Get What You Want.

2016, Chapter 15

The Vidar Chronicles, Part I

Why am I writing about Vidar *now* if it isn't finished, you ask?

Well, in part *because* it isn't finished. I'm sitting here, waiting for some information on heatsink availability and pricing. Which I need before I know how much space I'm working with in the chassis. Which determines the size of the board (or boards). So, while I wait for more data, I'm in limbo.

Again.

I already did this once, when we threw out the fan-forced horizontal heat tunnel approach. Now, I'm doing this one more time, as we throw out the passive vertical heat tunnel approach.

Wait, what? Why are we starting over again? Are we incompetent? Crazy? Stupid? Obsessive? Scared?

Well, maybe a bit of all of the above. To find out why, let's pull back and discuss …

Why Power Amps Are Different

Power amps (as in, the kind that drive speakers, not headphones) are different than virtually every other audio component out there.

How so?

Let's start with the obvious: power. Speaker power amps are way, way, way, way beefier than headphone amps. Delivering a watt or so into 32 Ω is in no way, shape or form like delivering a few hundred watts into 4 Ω—especially when that 4 Ω is really 2.1 Ω at 80 Hz and has significant reactivity and back EMF on it from two giant woofers gyrating like grandma's jello casserole. Speaker power amps are expected to do this without breaking a sweat. This means that speaker amp output stages are way, way, way, way overbuilt compared to a headphone amp.

But there's more: protection. Speaker amps are also expected to *not* self-immolate when you accidentally clank the two heavy-duty palladium-plated spade connectors together as you are hooking up the speaker while playing Megadeth

at full volume. This means they need to have some kind of over-current protection in them, which is (a) unusual for a headphone amp, and (b) unobtrusive enough that you can still deliver enough current to throw serious sparks.

And more: heat. If we're not talking Class D (and we are not), speaker amps have to dissipate quite a lot of heat. Unlike headphone amps, which probably cruise in Class A most of the time, a Class AB speaker amp will spend a lot more time on the Class B side. This means that it will have variable heat output, which will ramp way, way up when you're running, say, Magneplanars (4 Ω and inefficient) rather than Zus (12 Ω and efficient.) This is why speaker amps usually have significant heatsinking. More on this later.

And even more: regulation. So many manufacturers engaged in so much asshattery in the 1970s that the Federal Trade Commission stepped in and actually issued rules regarding power output claims for audio power amplifiers (at least ones over a handful of watts). These rules included rating with both channels driven, continuous RMS power, and the ability to run the amp for 1 hour at 1/3 rated power. That "1/3 power" test is brutal, as it is run at the highest dissipation point for a Class AB amplifier. The 1/3 power test

has been supplanted by a 1/8 power for 1 hour + 5 minutes at full power test by the FTC, but Stereophile still uses the 1/3 power test. More on this later.

And finally: it's a make or break component. Good power amps make companies—see the glory days of the past, and look at the history of companies like GAS, Hafler, and Adcom. Those companies built their reputation on power amps. Modern companies have built their reputation on power amps as well, most notably (on the affordable, Class AB side) Emotiva. But even as good power amps make companies, bad power amps kill companies. This is what happened with GAS (Great American Sound). They made powerful amps. Unfortunately, many of them also blew themselves up. The service load from these amp failures eventually sank the company.

So yes, when it comes to speaker power amps, a degree of paranoia is warranted.

And that's probably why you see so many amps based on standardized Class D modules these days. The hard engineering work is done. The protection system is built in. Throw a couple in a box with a power supply and you have an amp, with minimal chance of grandiose failures.

And, it's also probably why you usually see the same Lin-type voltage-feedback topologies on the Class AB side of things. Probably 95% of the Class AB power amps on the market today are this classic topology. Again, using this topology minimizes risk. It's well-known. It's been tested and busted literally thousands of times. The strengths and weaknesses are well-understood. Biasing, protection, compensation, output stages, etc … they've been honed and refined over several decades.

At this point, you're probably thinking, *Any sane company aiming to make a low-cost power amp would probably do Class D with standard modules or Lin, right?*

Right. But maybe we aren't sane.

Balancing Sanity and Insanity

Some of you are out there, groaning into your coffees, thinking, *Ah hell, Schiit's gonna do it again. They're gonna overcomplicate something like an inexpensive power amp, and make it a crazy, late, nontestable device I want nowhere near my stack of gear.*

No. Not quite.

If we were completely crazy, we would have made Vidar a circlotron design with the Ragnarok control system and called it done. Of course, that wouldn't have hit our cost targets, nor allowed the flexibility we wanted, nor, well, a whole bunch of things.

So if we weren't completely insane, what exactly were we shooting for with Vidar? Again, let's pull back and take a broader look, starting with the market.

If you look at the power amp options out there, you're confronted with a dizzying array of choices. Slick boxes with great cosmetics from respected names. Small outfits doing interesting stuff with lots of customer accolades. Components with eye-watering price tags. And gear that's a whole lot cheaper.

However, when you start saying, "Yeah, but I want a Class AB amp, and I want it under a grand," the options list suddenly gets a whole lot shorter. Add the requirement for a linear power supply, and the list gets chopped again.

And if you add, "Oh yeah, and it'd be nice if it was made in the USA," well, you get crickets.

That realization drove our first must-haves for Vidar:

1. **Powerful and inexpensive.** As in, at least a 100 W stereo amp for under $1000. Way under $1000. This is a significant decision, because power amps typically have some very expensive components in them, including the power transformer and heatsinks. So cost becomes a primary driver of what you do.
2. **Old school.** As in, Class AB with linear power supply. This is also significant, because this means you have BIG heatsinks and a BIG transformer. Not cheap.
3. **Made here.** Which means it has to be simple to build. It can't be a wiring nightmare or many-board monstrosity.

At the same time, I knew that the Ragnarok approach (circlotron and first-generation intelligent control system) wouldn't be the best solution here. Circlotron transformers are expensive due to the number of taps and quadfilar winding. The first-gen intelligent control system might be a bit much for an inexpensive amp.

And yet, I still wanted it to be something special. So, I added the following must-haves:

1. **No-compromise design.** As in, no coupling caps, no servos, no IC gain stages, no shortcuts,

no corners cut—a true, high-end design that anyone would be proud to use in any system, regardless of cost.

2. **Versatile architecture.** A 100 W stereo amp is a good starting—and ending—point for many systems, especially if it delivers a solid 200 W into 4 Ω as well. Add the ability to use balanced input, and it also becomes a 400 W monoblock. Vidar could be two amps in one.
3. **Interesting topology.** And here's where I went off the rails a bit. But I wasn't completely insane—this is what would have gotten thrown out if it didn't work. I had a complete, voltage-feedback Lin-topology schematic for Vidar in case what I had in mind didn't work out. But I'd gotten intrigued with current-feedback topologies during the Jotunheim development, and I wanted to see how it would perform in single-ended form. More on this later.

With these must-haves in mind, I set out on the first sanity checks with Vidar: getting costs for the transformer and heatsinks.

The transformer had me sweating. If it cost as much as the transformer in Ragnarok, we were sunk already. If we had to go to China to get it, we'd be breaking one of our own internal rules. Luckily, with some negotiation of size, mounting

style, and connector types, the transformer came in at a price that made the amp feasible … while still being made in the USA.

But then we came to the heatsinks.

Heatsinks can also be very, very expensive. Especially if they are cosmetically finished and proudly positioned where everyone can see them. So I thought that going with a conventional heatsink design was out of the question, and didn't even bother exploring that route.

Note to self: never dismiss something out of hand.

Vidar The First: Fan Follies

Early in 2016, I had identified what I thought would be a good heatsink strategy for Vidar: a long, horizontal heat tunnel formed from two high-density heatsinks, backed by a variable-speed fan under microprocessor control.

On first glance, it made sense. Horizontal heat tunnels with fan-forced cooling have been done in tons of audio gear. It was a proven, well-known strategy. Furthermore, the small, lightweight, high-density heatsinks were inexpensive. Even better, you could hide the whole mess inside the chassis, so it didn't have to be cosmetic.

And, heck, I already had a microprocessor in there, so doing real-time temperature sensing and running the fan would be no big deal. Hell, most of the time it could probably be off, or just ticking over slowly. You'd never even hear it.

All in all, it seemed like a great way to go. Even though I knew I'd hear some groans and moans about having a fan in the amp, I knew that the complaints would fall away when people found out that it would be dead-quiet. Or nearly, anyway.

Furthermore, I had a slick gain stage drawn up—a single-ended, fully-complementary interpretation of Jotunheim's current-feedback topology. I'd done some smaller prototypes of it, so I knew it was fast, precise, and very low-distortion. I figured I'd graft it on to a linearized MOSFET output stage (so I could run very low bias, keeping the fan off even more of the time) and call it a day.

The chassis was a different matter. It took quite a bit of going back and forth to arrive at something that seemed logical: a half-width product that lined up with Saga in terms of front aspect, but could be used side-by-side as monoblocks on a typical shelf as well. I used that general idea to

draw up concepts and get dimensions. However, I didn't bother getting prototypes done, because I figured we should get the board working before we went any farther.

Hey, at least I made one good call.

The heatsinks put us behind, though, since they were done from an unusual die. We paid for finished parts … and waited … and waited … and waited.

But eventually, they came in. I did the usual building-a-prototype kind of things, stuffing the board, doing an initial power-up, doing some static measurements, etc.

And during this time, I started to get a sinking feeling about the heat tunnel approach.

Why? Lots of little things:

1. How do the heatsinks attach to the board? I hadn't really thought of that.
2. How does the fan attach to the heatsinks? Again, oversight. Eventually I designed a complex metal girdle that solved both of the above problems, but it would add cost I didn't expect.

3. How would the fan interface with the chassis in order to minimize noise and ensure the air went where it should? Again, another custom part would ne necessary.
4. Would the whole assembly survive shipping, or would it need additional structural support under the board?
5. Would the fan be as flipping loud as it seemed like it was going to be, or did we need to embalm it in sorbothane?

But it was Dave's problem for a while, as I sent it off to have firmware done for the oversight, management, and protection of the amp. This version of the amp used something very akin to the Ragnarok control system, with full active bias oversight. That proved to be both a good and bad thing.

When Dave brought the amp back, he wasn't smiling.

"I blew it up," he said.

"And?" With Dave, this is to be expected.

"And the input devices are really small for their dissipation … " he said.

"I'll fix that."

" … and they blow up if you short the amp."

“Huh.” That wasn’t good. That wasn’t expected. Not at all.

“And … ” Dave turned on the amp. It sounded like an AM radio tuned in-between stations with the volume up full. “The fan is loud,” he shouted.

“Can’t you turn it down?”

“It needs to run pretty fast to pull air through the tunnel.”

Crap. I sat back and took a look at the amp again. Dave was right. The high-density heatsinks needed a veritable Hoover to suck air through the long tunnel.

“And it has noise from the fan PWM on the ground,” Dave said.

I shook my head. Yeah, we could go to a different extrusion. Yeah, we could get a quieter fan. Yeah, we could make all the bizarre little parts we needed to make a heat tunnel really work. All the crap about why fans suck came rushing back to me.

And I decided, right there, it wouldn’t go any further.

And that's why Vidar the First never even got fully operational. We killed it before it played a single note.

Second note to self: you should have listened to it. Or tried to.

Vidar The Second: Heat Tunnel Hell

One of the reasons I was so eager to kill Vidar the First was simple: I'd already been thinking about an alternate approach to the heatsinking. That thought process went something like this:

Instead of using a fan, maybe there was a high-density heatsink that could be used to create a vertical heat tunnel. A vertical heat tunnel should pull air through it via convection, doing the work of a fan … silently.

Yep. Neat idea.

And, not just a neat idea—there was an affordable, compact, high-density extrusion that offered exactly the same surface area as the old Sumo 120 W amp heatsinks. Coupled with an aluminum chassis, we'd be golden. Not only that, the heatsink was in stock. I had some pieces cut to size and shipped to me in a few days. Suddenly, it looked

like we were on the fast track to getting an amp that worked.

I laid out a new version of the amp—this one designed to have a big hole in the middle, like a donut. This time, I designed in larger input devices, increased some resistor sizes to eliminate the short-it-and-it-blows-up problem, and made one very big decision:

This amp wouldn't have 100 % *bias oversight, like Ragnarok.*

This was a big decision, because it meant that the amp would need to have its initial bias set manually, and would allow for some bias variation during operation. However, the microprocessor would still allow for the elimination of the DC servo, and would provide oversight of temperature and protection. *Aaanndd … the amp would be testable under the standard* 1/3 *power regime used at Stereophile.*

Soon we were looking at another assembled prototype. This one went much more smoothly than before. Even the heatsinks fit! I ran it through the usual DC tests and sent it off to Dave for firmware.

But again, Dave came back not looking happy. “It oscillates,” he said.

“How bad?”

“Bad. Like many amps of current bad.”

Ah. That was very bad. “When?”

“As soon as you turn up the bias.”

Hmm. I asked Dave to try a few oscillation-killing tricks (bypassing, compensation, etc), but he brought it back and plopped it on my desk, saying it was still a no-go.

Fine, I figured, I’d dig into it.

So, that weekend, I stripped the gain stage down to its bare, uncompensated, open-loop form, and measured it. (It’s important to see how the stage operates before the feedback loop is closed, in order to get the compensation right.) There were no huge surprises, so I closed everything up and slowly turned up the power.

Again, everything was just fine. So, I slowly began increasing the bias. And that’s where everything went wacky. As soon as the outputs had any bias on them, they went into violent, device-melting oscillation. I couldn’t get the amp anywhere near the target bias.

So, I spent some time re-doing what I'd told Dave to experiment with—compensation, bypassing, etc. None of that worked, so I went to more radical solutions—measuring the inductance of the output resistors, swapping them over to film versions (completely non-inductive) for temporary testing, increasing the gate stopper size on the MOSFETs, adding additional compensation poles, moving the compensation around, compensating the output stage itself, bypassing the linearization, eliminating the linearization entirely.

Nothing worked.

So I went even crazier. I pulled the gain stage back to what I knew worked fine—a basic stage that ended at the drivers. It worked fine. I pulled off all the MOSFETs (3 pairs) and went to a single pair of MOSFETs. That worked fine, too.

But when I added a second pair of MOSFET outputs, BOOM … back to the output-frying oscillation.

So, cue another montage of going back and forth with bypassing, compensation, hacking up the board to eliminate routing variables, etc, until the whole gain stage looked like I'd dumped a pile of random parts on it as the solder was cooling.

And none of that worked. Not with more than one pair of MOSFETs, anyway.

Now, I was over a week into hours-every-day work on this thing. And it was incredibly frustrating, because I'd used the same kind of MOSFET output devices in every Sumo amp I ever designed. The Sumo amps never had these kinds of problems. So why was this different? I tried a different batch of MOSFETs, I tried tighter matching, I tried even crazier ways of damping oscillation … and in all cases, the results were the same: two pairs of MOSFETs (or more) and boom.

Why did it act like this? No idea. Maybe the MOSFET design itself had changed. It was a similar part number, but it wasn't the same … and it was from a different manufacturer … and manufacturers remix their parts from time to time. Or, maybe the combination of an ultra-fast current-feedback stage (it has a bandwidth of several megahertz before an input filter is applied) and MOSFET outputs really wasn't fated to work.

So I decided to do something even more radical: replace the MOSFETs with bipolar transistors (BJTs.)

This is something I'd contemplated when first designing the amp, because bipolars could poten-

tially reduce complexity. Bipolars wouldn't need the transconductance linearization, for example. But they would also require some new, beefy drivers—which would have to be on the heatsink, which I had limited space for.

So I drilled a few more holes on the heatsink, added big drivers, threw together a bias network, and installed three pairs of Toshiba BJTs.

And the amp just fired up and worked.

Like, completely stable, no problems, like it was standing there with crossed arms, looking at me like I was an idiot and asking, *why didn't you do this before?*

Yeah.

Sometimes life is weird. I wasn't going to question this development—except for the heatsinking problem, going to BJT simplified the amp—and probably improved its measured performance at the same time. Fine. Win-win. We had a show coming up anyway (RMAF), and I really wanted to show the amp there.

So, I re-laid-out the board for bipolars, for the big drivers, designed up a heat spreader to attach the drivers and VAS stage to the heatsinks, and got a couple of chassis put together by our supplier.*

The idea was, that if these worked well enough, they'd go to the show.

Aside: our chassis guys deserve kudos for this. The need for the heatspreader was last-minute, and they delivered. They also delivered an alternate design in a couple of days when I found out the part needed to be bigger to dissipate more heat. They did a quick-turn anodize on the heatsinks that improved their heat dissipation. They did finished-looking chassis for the show, on time. This was a wonderful showing by our metal house.

To make a long story short, the amps worked. Yeah, I had to tweak the compensation a bit, and Dave had to do some new firmware, but a week before RMAF, everything looked good. Except …

"Aren't we going to torture test these?" Dave asked.

"Absolutely not," I said.

"What if they blow up?"

"They won't. We'll torture-test them after the show." The spectacle of blowing up the amps right before the show—and having nothing to use—wasn't appealing. And, in reality, the amps would

be loafing at the show. Typically, running amps into speaker loads is much easier than the "run it to full power and short it" test, the 1/3 power test, or anything else we do to stress an amp.

Foolhardy? Not really. If you've been following my bleatings here, you know the amps made it through the show just fine.

They've also survived an impressive array of torture-tests after the show, including:

1. Running to thermal shutdown to test the thermal measurement.
2. Running to clipping within and outside the audio bandwidth to check for nasty stuff like simultaneous conduction.
3. Shorting at clipping and below to test the protection system.
4. Running for hours into Magneplanars at high volumes.
5. Being handed off to employees and friends to see if they survive normal handling (accidental shorts, etc).
6. Validation of performance into reactive loads at high levels.
7. 1/8 power and 1/3 power long-term testing.

And that's where things go off the rails a bit—with the 1/8 power and 1/3 power testing. As of

this writing, the amps will make it through the 1/8 power for 1 hour plus 5 minutes at full output that the FTC mandates. But they make it only about 15 minutes into the 1/3 power testing before the thermal protection shuts them down.

If you're sitting there saying, "Well, I don't see the problem there, you pass FTC, right?"

Well, yeah. Barely. In a 22 °C room.

And Stereophile still tests at 1/3 power.

So, for a while we joked about adding a "panic fan" to the design that would only come on with exceptional thermal loads, such as 1/3 power testing. And, you know, that might work. But the fact is, the Vidar prototype is running much warmer than expected, especially when compared against a Polaris 2 with the same amount of heatsinking.

In retrospect, this may seem obvious. A passive heat tunnel needs to have an extreme thermal gradient to be effective. Plus, the two heatsinks are facing each other—they are radiantly heating each other. Heatsinks on the sides of chassis radiate into an effectively infinite space. Big advantage. Plus, the high-density design of the heatsink actually means there is limited thermal

gradient to work with. So, the design isn't taking advantage of heatsinking as well as it could.

So what do we do? We look at other options.

Vidar The Third: The Charm?

These "other options" mean "new heatsink design that will probably end up being structural and external." Which means we may even end up saving some money, since the chassis becomes simpler. But that's all yet to be seen. I'm waiting for some quotes right now, both on price and availability, that will determine our future direction.

"Oh no, so we're gonna see Vidar in 2018, that's what you're saying, right?"

Not at all. In fact, I don't see any reason to revise our timeline of shipping in Q1 2017. Even with having to go to a new heatsink design. Here's why:

1. **The hard parts of determining topology, etc are done**. We have something, we know what it is, we know how it performs, and it's both a logical extension of our Pivot Point topology and the Ragnarok control system.

2. **We have an amp that's electrically fantastic**. It overperforms on power output, delivers very low distortion, is dead-stable, and has a great "Generation 2" control system that eliminates the stuff you don't want in an amp (coupling caps and DC servos) but *doesn't* de-bias on 1/3 power testing.
3. **It's been torture-tested.** Severely. The fact that it gets VERY hot and still doesn't fail is a big plus, because it won't get that hot in final form.

All we need to do is get rid of the heat.

And that, I've done before. In going back to external heatsinks, this is something that I'm completely familiar with. It's also an amazing boon, because it allows us more area to spread out the parts, reduce heat concentration, and eliminate heat spreaders. The amp gets simpler … again.

So now, I'll close this part one—and look forward to seeing where the heatsink options come in …

To be continued.

2016, Chapter 16

Even More Fulla It

Ah, the poor, forgotten Fulla. It's amazing to me that over 2 years have gone by since we introduced (what we thought was) a giant-killer: a powerful, good-sounding dongle-DAC with an analog volume pot, available at the then-crazy price of $79.

Doesn't sound so giant-killerish? Well, considering that dongles were still selling for $199 and up at the time, it was a fairly big deal.

But the world moves on. Now, companies like Audioquest have great-sounding dongles for only $99—dongles with their own in-house, low-power USB interface (which makes connecting to smartphones a lot simpler). At the same time, other, less … ahem, upstanding companies are busy ripping off their earlier work and selling it for even less than the original Fulla. And if you move up the chain, there's no shortage of dongle options that offer even more power and higher performance than Fulla could ever imagine.

In short, it's a different world, only two short years later. The market has become dongle-saturated. For those who like to add batteries for a portable solution, there are plenty of options there as well—many with screens and storage. In fact, that whole market is moving so fast, it can be difficult to keep up with all the new options.

Call me old-fashioned, or call me non-competitive, but all those options told me one thing: *Fulla had to die.*

Or, it had to be *reborn* … as something entirely new.

A Fateful Conversation

About 18 months ago, I had an interesting (and insightful) conversation with a Schiit fan at TheShow Newport. He really liked the Fulla … but he wanted it to be something very different.

"It should have way more functionality," he told me.

I nodded. Of course, I'd seen dongles with more functionality—preamp outs, gain switches, things like that. But Fulla was already bumping up against costing too much to produce. If we

added functionality, it would have to justify a higher cost.

"Like a preamp out," I offered.

"Or a fixed out," he said. "So you can use it as just a DAC—a fixed 2 V out for your main system, so you don't have to worry about where the volume is set."

I frowned. If someone plugged a set of earbuds into a fixed 2 V out, that might not end up so well, not even with resistors in line to raise the output impedance.

He saw my uncertainty and added, "Of course, you could have a preamp out too. And maybe an analog in to go with it."

I laughed. The mental image of a tiny box covered in RCAs wasn't pretty. "That's a totally different product."

"So?"

"So it'll cost more," I said.

"If it's within reason, I'll pay it," he told me.

"Yeah, and that's the problem: if it's within reason," I said. "I don't know if it'll be within reason.

I'd have to do some design to figure out what it would cost."

"I'm sure it wouldn't be much more," the guy told me.

"No clue," I said, knowing that seemingly small changes on a chassis could cause significant price differences … not to mention the fact that I wasn't sure that any part of it was a good idea. Even 6 months after launch, we were already seeing price erosion in the dongle-DAC space.

And that's pretty much where I left it for another year. During that time, dongles got cheaper and better. And Fulla went nowhere, of course. It couldn't. We had a fixed design, and a fixed number of parts and chassis—a *very large* number of parts and chassis.

Worse, we began seeing some of the uglier aspects of dongles as Fulla got older. They are tiny products. It's easy to crush one. Or to break off the small potentiometer. Or have issues with the USB connector. The failure rate wasn't higher than industry standard, but it was higher than *our* standard.

But eventually, Alex came to me and said, "It's time to make a decision on the next Fulla." The

look on his face seemed to say, *And I hope it disappears.*

Death or Rebirth?

And Alex had a point. A perfectly rational answer to the question of the next Fulla would be, "There's no next Fulla," and call it done. Fulla wasn't lighting up the sales charts, there were lots of options priced near it, and it would take some pretty *serious design work to make something that stood out in a crowded market.*

And that's almost what we did. Like I said, it made sense. We didn't really need Fulla.

But …

Spending time on the forums, reading about people's experiences with Schiit products (and others), and making sure to pay close attention to people's first audiophile purchases, I knew there was something we were missing. Specifically, a one-box, flexible product at a price significantly less than a Magni and Modi combo.

> **Aside:** for those of you chuckling about how something should cost less than a Magni and Modi, give thanks that life is good for you. The

reality is that $198 is a significant chunk of money for a lot of people. It's even more if you're going to school and don't have a full-time job, or if you are just out of school and staring down the barrel of some serious student loans. In addition, the various connections, cables, etc needed to get a system working can be confusing when you're first getting into it as well—and it doesn't take much confusion for a sale to turn into an abandoned cart.

So what price point would make sense for such a product? Significantly less needed to be *viscerally* less. Really *real.* As in, not $169 or $149 or even $129.

But $99? Now you're talking.

And that's when that conversation came back to me. Not just a $99 DAC/amp, but a $99 *do-all* DAC/amp. *That* would be interesting. Especially if it could be quite a bit more powerful than the original Fulla … and use the great new top-line AK4490 DAC.

Aside: people who have followed my blatherings to date know that everything after the "but" in the paragraph above is called "ego talking," and it's usually a great way for designers to get

> themselves in big trouble.

However, there was still one big problem: the fact that one of the reasons dongles had changed was that sources had changed—in that many of them wouldn't supply the 500 mA (or more) that such a powerful device would require.

But I put the misgivings about power away for the moment. Because sometimes the best thing to do is simply to build it, and see how well it worked. Even if it ended up being just a desktop device, even if it had problems with some low-powered USB ports, it was worth exploring. I decided to go ahead, do the design and make a prototype—and sort out the whole power situation later.

And that's where I was in Spring 2016: with a basic idea (for a new Fulla), a decision on direction (it's a full-featured desktop product now, damnit, not a dongle) and a looming problem: how to interface it easily with phones and such.

Good enough to start. I figured there would be other problems to work out—and there were …

Decisions, Decisions

Okay, so the second-generation Fulla would be designed for a desktop. But how big would it be?

The original Fulla was a tiny product. It had to be bigger than that. At the same time, it should be smaller than the Magni and Modi stack. In the end, I chose a footprint exactly half the size of a Magni/Modi stack—and shaved off a quarter inch of height. So, 2.5 inch wide, 3.5 inch deep, and 1 inch high.

A mockup looked pretty cool. It even could use the same potentiometer and the same volume knob as a Magni, which was a big plus. The potentiometer on the Magni was a lot better than the one in the original Fulla in terms of tracking and durability. And the knob would save us having to do another custom part.

With the size decided, I needed to make some decisions on features. First up was how to get more power out of the product. Fulla's 200 mW output into 32 Ω wasn't really worth getting excited about. It was plenty for most headphones, yes, but I wanted to do better.

Doing better, though, meant both a better power supply and a more beefy amplifier stage. This wasn't as simple as it might sound, though, because unlike all of our other amps, Fulla uses a switchmode power supply, and op-amp output stage.

Horrors! Some of you are clutching your pearls and looking for a safe space. *How can Schiit, bastion of linear power supplies and discrete design, sully ourselves with such inferior technology?*

Well, here's the super-complex rationalization for such a drastic decision: price.

Yes. It is that simple. Horses for courses. When you're asked to create a power supply that can deliver a reliable ±5 V from USB's single 5 V rail, and to make the most of the 500 mA to 1 A that USB ports can supply, you're looking at a switchmode supply. Period. End of story. No magic fairy dust or wishful thinking will get you around that fact.

Now, of course, we could do much better than the previous Fulla's power supply, which used a charge pump to create the −5 V rail, but passed the +5 V along from USB (after filtering). This asymmetrical arrangement meant that the −5 V rail couldn't deliver the same amount of current as the +5 V rail, and it was the main reason why the original Fulla's output power was limited. With more space in Fulla 2, we were able to use a symmetrical rail generator with a better, inductor-based switchmode supply for each rail, delivering more than 3× the current of the original Fulla's

supplies. With some filtering, it's a good, solid, quiet power supply.

Similarly, when you're asked to create an output stage half the size of a postage stamp that can swing the output from rail to rail, you're not going to do it with discrete parts. Enter op-amps. Specifically, the same exotic, high-current op-amps we used with the first Fulla, the LMH6643. This fast, high-current, rail-to-rail device has only one downfall—its current output is still less than we'd like. So, in Fulla 2, we simply used two of them per channel to double its output current capability.

Now, we have an amp stage that can deliver 550 mW into 16 Ω—much better. Still not a Jotunheim, or even a Magni, but there you go. Sometimes you have to know when to say "When."

Beyond power output, though, what kind of I/O should Fulla 2 have? I waffled a bit, wondering about having just a preamp out, or not having an analog input, but in the end, I threw them all in: fixed and preamp outs, as well as an analog in. Fixed so you could use it as just a DAC, variable so you could use it as a preamp (it amplifies the analog input, too), and an analog input in case you have a phone or something (that still has

a 1/8 inch output) that you'd like to connect to Fulla.

The catch to all of this I/O? It all has to go in and out on 1/8 inch jacks. Yeah, I know. But 1/8 inch to dual RCA cables are easily available.

Bluetooth? Nope. Yes, it's convenient, and yes, it sounds like ass. Yes, I know, it doesn't sound completely like ass, but it doesn't sound as good as a cabled connection. Buy a cable. Don't be lazy.

Aside: I have been told by friends at very large audio companies that Bluetooth is also still (a) not entirely reliable, (b) quite variable in terms of quality as the data rate changes over distance and direction. I also know how much it costs to implement and to qualify. Fulla 2 wouldn't be a $99 product with Bluetooth. So, given the choice of adding an ass-sounding input that would cause additional customer consternation and add a lot of cost, we decided to pass.

> **Aside to the aside:** when Bluetooth can reliably pass an uncompressed 16/44.1 stream, we'll revisit this decision … we're stubborn, but not crazy.

For headphone output, I went with a full-size 1/4 inch TRS jack, because, let's face it, most serious full-size headphones still use this connector. And Fulla's intended to be a serious product. And 1/8 inch to 1/4 inch adapters are common. And yeah, I know, I heard from tons of people who said, "But you can do balanced output on TRRS for even more power!" But that's also a recipe for frying outputs or frying headphones when someone ignores all the warnings in the manual and the big warning sticker and giant plug in the TRRS jack and plugs a TRS headphone into it anyway. Or vice-versa. No thanks. Shoot us for playing it safe in an affordable product.

And that's where we ended up with the first prototype. It felt pretty good. It looked pretty cool. I threw together a board and sent it out. After it came in, I sent the board and a bag of tiny parts out to our PCB assemblers. And when it came back, I was eager to hear how it sounded.

Unfortunately, it was a bag of crap.

Lessons in Tiny Parts

Designers working in high-end audio have it easy. We get to work with parts that most other electronics engineers would call "rustic." They're

big. Easy to handle. Lots of them still have *leads* on them. You don't have to worry about them blowing away when you sneeze.

Hell, even when Schiit uses surface-mount parts, we're using big parts—like 0805 resistors and old-school SOICs. Yeah, I know, you can stop laughing. Because even in our realm, tiny parts are creeping in, together with parts that are simply impossible to solder by hand (parts with pads underneath them, for example.)

And, in the case of Fulla 2, we were still sufficiently size-constrained that it has a lot of 0402 resistors, as well as some exotic stuff like 4 mm×4 mm parts with power pads and 24 pins, as well as 100-pins QFNs and … well, let's just say that it isn't what anyone wants to put together by hand. That's why the prototypes go to our PCB assemblers, rather than to me or Dave.

And when you screw up a bunch of traces going to the new QFN DAC, it's kinda hard to fix it.

In fact, it's so hard, we gave up. I tried it, Dave tried it, our assembly house tried it … and none of us could de-screw what I had screwed up. So that was the ignoble end of the first prototype: it never worked, and ended up dead in the trash.

While I was fixing my (entirely avoidable) board screw-ups, though, I got to thinking about the USB power problem. Because Fulla 2 would end up drawing even more current than the original. It would be a 100 % no-go with iPhones, iPads, many Android phones, and some power-limited USB ports. And that wasn't ideal.

But now, I thought I saw a way around it. Since we were using a rail-generator chip to do both the +5 V and −5 V rails, and since it operated down to 3 V or so, we didn't need to have the full USB 5 V into it. Which meant we could, theoretically, simply add another micro USB input that provided power only … which meant that the main USB port wouldn't need to draw any power at all … which meant that our power problems were solved.

Except for one problem: USB power reporting.

You see, USB ports don't (usually) measure the amount of power a device draws. They rely on the device to tell it how much current it needs. USB devices can request 0 mA (no power), 100 mA (low power) or 500 mA (full power).

Aside: yes, I know, this is monumentally more complex now thanks to USB 3.1, but let's leave

that alone for now. The example I'm using is relevant to audio devices, so roll with it.

Fulla 2, like Fulla, reported as a 500 mA device. Which meant trouble with a whole lot of products. Apple iOS devices refuse any device that wants more than 100 mA. Some USB ports claim to be able to provide full power, but they really don't. Some Android phones will provide 500 mA for a while, but not when the battery falls below a certain point.

But the problem was, even with a separate power input jack for Fulla, it wouldn't matter—because it would still report as needing 500 mA.

Unless there was a way to customize it to detect where the power was coming from, and report as 0 mA *if the second USB input was used.*

And if we could do that, it would be perfect! You could just use something like a phone charger for power if your USB port couldn't run the device. Or, even without a compatible phone charger, 5 V AC adapters with micro USB out were a dime a dozen.

I knew it could be done. The CM6631A USB receiver we use has many general purpose I/O ports (GPIO). If we could pull one high when the

auxiliary power was connected, and if we could alter the firmware to change the reporting, we'd really have something.

But I'm getting ahead of myself. Because I also have to talk about the knob.

The knob. Yes, the most recognizable aspect of Fulla 2. The big vertical knob. At the same time I was fixing the board and figuring out how to do power switching, I was also becoming unhappy with the cosmetics. Because Fulla 2 looked OK, but it wasn't really all that interesting. And the small knob and shorter chassis made it a little awkward. I wondered if we could do something about that.

A little research, and I found out that yes, we absolutely could. Alps also made a vertical version of the pot we used in Magni, the same exact part with the same exact quality. And the guys who make our knobs also make aluminum shell knobs of the same exact style in much larger diameters. Put the two together, and suddenly Fulla 2 looked pretty cool … and it was exactly in keeping with the original, which also had a knob on top.

With these changes made—but no definitive answer on whether or not we could have the USB-

reporting-altering firmware in place—it was time to go for a second prototype.

And We're Golden

Why go for a second prototype when we didn't have the power-switching in place yet? It's simple:

1. Moving the cost of entry for good sound from $198 to $99 would be a significant achievement, allowing even more people to hear "what the shouting was about."
2. Even if the power switching didn't work, it wasn't like we had anything like that on Modi, and Modi did very well.
3. I wanted to get the product out this year, and sometimes, you have to fish or cut bait. This was fishing time.

And so, we got a second prototype back … and this one worked great. It even worked when powered by the auxiliary USB jack. Unfortunately, the CM6631A still reported as a 500 mA device, even when the power USB was connected.

Could we make the power-switching work? In my mind, this was the difference between a good and great product. Fulla 2 would be plenty good if it moved the cost of entry down, but had

problems with low-powered USB ports and we had to tell people to use a powered hub or other workarounds to make it work. But it would be a great product if we were able to say, "Low power USB port? No problem. Just plug in your phone charger and we're all good."

Long story short, with C-Media working on the firmware, and a couple of hardware tricks, we were able to do exactly that. Plug Fulla 2 into your computer, and it reports as a 500 mA device. Plug a separate power USB into Fulla 2, then plug it into your computer, and it reports as a 0 mA device.

Of course, that took a couple of hardware hacks … so with one more prototype under our belt, it was time to move on to production.

As (seems to be) usual these days, the first articles were uneventful. They worked as expected, tested well, and the new firmware performed like a champ. Alex pronounced them good. Tyler (our head of finance and admin … think of him as Hermes from Futurama, a *good* bureaucrat—hough he has yet to show up in a flying desk) said, "Wow, the bass is really good on this, I'm going to buy one as soon as they're out." That

puzzled me a bit. “But you already have, like, a Lyr 2 and a Bifrost and other stuff, right?”

Tyler nodded. “Yeah. But this is tiny. It’s great for a small desk. And it sounds good.”

I grinned. That was exactly the response I wanted. A small, do-all product that anyone could be happy with. Everything seemed good for the launch.

Except …

Except sometimes there are idiots. And sometimes those idiots is me.

I’m an Idiot

About a week before we were scheduled to announce Fulla 2, Alex told me, “Wow, this is a really great device. But I thought the preamp out was supposed to be variable.”

Deep in Vidar guts, I didn’t even look up. “Of course it is. There’s a fixed output too.”

There was a long pause. Finally, Alez said, “Yeah. The outside one. But it’s variable.”

*Ah, ****.*

“Wait. Are you saying the screen is wrong?”

“Um, well … ”

“Check it out, please. And let me know.” Argh. If it was wrong, we had a whole bunch of metal that was screened wrong. Which meant we couldn’t ship …

I realized I could check myself by looking at the PCB layout. A few seconds later, I was staring at a layout that said, *Yep, you’re a bonehead, you reversed the screen.*

Crap.

What to do? Push the launch, of course. Call the metal guys and see what we could do. And scrap the manuals, because they were done for the wrong screen.

Another long story short, our metal guys proved themselves again. They had a batch of parts in process that they could screen and get to us in a few days. This would hold us until they could refinish the previous run. I got ahold of our printer and got them new artwork for the manual, which we’d have in a few days as well.

Crisis averted, but yeah, sometimes when you get too cocky, you get taken down a peg or two. This was a reminder: stay humble.

And stay silent—as you've already seen, my confident prediction of shipping Freyas in October turned out to be a crock as well. We'll see about Vidar. As far as the products after that, well, I'm keeping my mouth shut.

In the end, everything worked. Now, the cost for a do-all desktop system is $99, not $248 (remember, you need Magni 2 Uber for preamp outputs). Hopefully this will bring even more people into the market.

Will the market change significantly in the next couple of years, obviating the Fulla 2? I don't know. I don't have a crystal ball.

And, in the end, only one thing is important: what *you* think.

2016, Chapter 17

A Directional Assessment

Wow. Sounds pretty serious, right? "A Directional Assessment." Eeek.

Don't fear, though, we'll keep this light and fun. One of the reasons I decided to name this chapter so oddly is that the year in itself was odd. We didn't know where we were going when it started, but we came out of it with a whole new direction.

Yes, I know, I'm getting ahead of myself. So let's go back to that first statement: *we didn't know where we were going when 2016 began.*

100 % truth.

On January 1, 2016, we had only one firm product introduction in mind, and a bunch of nebulous blowhardosity about getting into 2-channel products. And yeah, I mentioned a new topology (which became Jotunheim), and I mentioned that I thought Manhattan would be a 2016 product (LOL) and that we had another non-digital shocker coming (still coming, ha) and that we

were going to do more ads and improve distribution and stuff like that.

But the reality is that timelines slip (on the Manhattan Project and otherwise), and reality intervenes (direct distribution worldwide is vastly complex) and we didn't get a bunch of the crap I yammered about done in 2016.

So, back to January 1, 2016. The one product we knew we'd be introducing? Modi Multibit. Yes, it was already running. Or limping, more like it. We knew that 2016 would be the year of Multibit for the Masses.

But Jotunheim? Well, we didn't really have a solid timeline for that.

Saga, Freya, Vidar? They didn't even have names, much less designs. All I had was a quote for a remote control.

Fulla 2? Not even a glimmer about that. I figured we'd just sell out Fulla and be done with it, at the beginning of the year.

So, let's call 2016 *The Year Without a Plan*.

Life Without a Plan

Life without a plan really isn't all that bad, for all the gloom and doom that many corporate naysayers will spout. Schiit long ago passed the point of being "big enough," so we're not in the crazy grow-grow-grow at all costs mode.

Aside: in my opinion, that path ends in only one of three ways: acquisition, insanity, or death, none of which are particularly interesting to Mike and I.

And, since we're not out to *Meet Our Aggressive Revenue Targets* or even to *Maintain Our Momentum,* we were able to take a well-deserved breather for the first 6 months of 2016. During that time, nothing new came out of Schiit. Absolutely nothing.

And, you know what? The world spun on.

What's more, sales *increased.*

Yeah. Go figure. Sometimes less is more. And in the first half of 2016, that's where we were. At least from the perspective of the outside world. In reality, we were busy. I was finalizing Jotunheim and working on the first prototypes of Saga, Freya,

and Vidar. Mike was working on a whole slew of new stuff, some of which you'll see very soon.

And somewhere in there, Alex asked me about Fulla, and I remembered the conversation about a next-generation Fulla, and we got a Fulla 2.

And also somewhere in there, I got a couple of other ideas which I'm still playing with. Big ideas. Ideas that I haven't talked to anyone about (in one case, not even anyone at Schiit besides Mike and Dave.)

And also somewhere in there, Mike and Dave discovered they'd need some heavy-duty talent on the Manhattan Project … and ended up dramatically expanding the scope of what they're taking on.

In short, even though it looked like nothing was happening, the first half of 2016 was really important. It was a time to sit back and think in broader terms, and even start wondering, a bit "what we want to be when we grow up."

> **Aside:** we actually decided not to grow up, but then again, you knew that already.

And, if I am allowed to be a bit grandiose, the first half of 2016, in many ways, ended up being

the foundation for the future—maybe even many years to come.

The Concrete Bits

"Okay, okay, I know you're gonna tease us about what you can't tell us," you may be saying. "So tell us something you can.

Okay. You want concrete, let's talk concrete. Let's talk about what we did in 2016, and why it was important. (We'll get to the future a bit later.)

So, a summary of 2016 in products:

1. **Modi Multibit.** AKA "the least expensive Schiit Multibit DAC." AKA "the most advanced multibit DAC on the planet for the price—from any country of origin." AKA "Multibit for the masses." Why is this important? It should be obvious—this modern multibit DAC is built on an entirely new, proprietary platform that doesn't rely on obsolete or salvage parts, and includes our unique digital filtering technology. And that means that more people can get a taste of what multibit is about, and decide if it's something they are interested in. We're not going to get into technology screeds about what's better or worse in terms of measurement and perception—if you want

to argue about that, I'll remind you that the Modi 2 and Modi 2 Uber are both still available at $150 and $100 less.

2. **Jotunheim.** AKA "the first sane DAC/amp," AKA, "the first affordable, configurable, do-all balanced desktop amp/DAC/preamp, AKA, "the first pivot point amp." This one's important because it eliminates the key drawback in combined products—that is, the problem of obsolescence—by using a modular, upgradable architecture. It also introduces our new, no-compromise Pivot Point gain stage, which allows for seamless integration of balanced and single-ended inputs and outputs. This is a key product that can easily be end-game for many systems.
3. **Saga.** Our first dedicated, remote-controlled preamp. Also featuring a cost-no-object relay-switched stepped attenuator and passive or 6SN7 tube buffer modes. In a market filled with 4- and 5-figure preamps that use compromises fit for A/V receivers (like integrated volume control chips), the $349 Saga is a hard reset. This one is moving very well, even in the absence of a matching amplifier.
4. **Freya.** Twice the phases and 4× the tubes of Saga get you a balanced preamp with tube gain, JFET buffer, and passive modes

for $699. Even if you use only one of the modes, this slaughters the price/performance ratio in preamps. And yes, not yet shipping as of this writing. I'm hoping we can start shipping this year.

5. **Vidar.** In a world full of $2k, $5k, and $20k amps, Vidar is a weird outlier at $699. Especially when you consider that it is a cost-no-object style design with linear power supply, huge transformer, massive heatsinks, and overbuilt Class AB output stage. Run one for 100/200 W into 8/4 Ω, or run a pair with balanced input for 400 W into 8 Ω. Although this isn't shipping yet, it's a vitally important product for us, especially considering that the list of great sub-$1000 amps is very small, especially if you add the caveat of "made in the USA."
6. **Fulla 2.** Surprise, surprise. Developed from a conversation and a desire to see what we could do to move the barrier to entry down by half, Fulla 2 is blowing up our sales. For $99, it's a taste of what good audio can be—and it's versatile enough to span all the needs of your desktop. Taken together with Jotunheim, we now have two formidable DAC/amp offerings.

Six products. Not a bad product introduction schedule, if we'd spaced them evenly throughout

the year. But hey, we still haven't shipped two of them, so, well, there you go.

But consider what these product intros have done for us:

1. **They've given us the most formidable multibit DAC line around.** Multibit now goes from $ 249 to $ 2299 here at Schiit—something for everyone.
2. **We've gone from zero to hero in combo products.** Last year, Fulla didn't really count. It was just another dongle-DAC. Now we have an uber-affordable step-in to great audio on the desktop with Fulla 2, and a versatile, upgradable platform with Jotunheim.
3. **We've reset the bar in the entry-level desktop market.** Specifically, we've cut the bar in half—from $ 198 to $ 99—with Fulla 2.
4. **We've shown a full 2-channel system.** Yeah, I know, we haven't yet delivered it, and the Vidar is still planned for Q1, but the reality is, we had a great-sounding, all-Schiit 2-channel system running at RMAF.
5. **We're now in the 2-channel market.** Yes, still only with a single product at the time of this writing, but it will soon be two. And it gets us into remote controlled products as well. All of this can be applied to the future …

and the 2-channel line can continue to grow. So, you know what? For a year without a plan, I think we did just fine.

Looking Into the Crystal Ball

Crystal balls are notoriously hard to use, fuzzy, and unreliable, so perhaps they're an apt metaphor as we look forward to the year. Because a lot of this stuff might happen … or it might not. We'll see.

The near future, though, is pretty clear. We've been talking about the Vidar amp, we've been showing you sketches, hell, I wrote a chapter on how it wasn't done yet, so it should be no surprise that Vidar is coming.

Still Q1? As of this writing, it's still looking good. I have the new heatsinks in-house, I have prototype boards in-house, and I have unfinished first article chassis in-house. This means I'm not far off of final integration, including thermal testing. Then ordering of long-lead parts like heatsinks and chassis can be turned on in earnest. I expect we'll be ordering at least some of these parts this month, so the clock will start ticking very soon.

Beyond Vidar, don't be surprised if you see a couple of other small things in the first quarter, including a couple of products that are new to us. Details are murky, though, in the crystal ball, so I'm sorry I can't give you precise product types, specs, and renderings. I will say that one of them is more than a simple product release … and it addresses a pain point you've started asking us to work on.

Beyond the first quarter, things get much more interesting … or not.

"Or not?" you ask. "What the hell does that mean?"

It means that beyond Q1, we're working on a lot of big ideas. These big ideas are inherently more, ah, unstable than regular products. They could work great, and show up early. They could need a lot more development than we expect. Or they might not work out at all.

What can I tell you about these new ideas? Well, perhaps at least the number. We're working on four big ideas, including the Manhattan Project. Beyond that, I have to get very cagey.

Let's see if I can give you some hints:

- The **Manhattan Project** is the most speculative of the big ideas. It remains something that has never been done before, a product that does not fit into any category that has existed in audio. It has grown into something much more ambitious than originally expected, though, requiring the hire of a full-time Ph.D mathematician to work on it. We are getting results, but if anything slips far out the timeline, or fails entirely, this will be what fails. But if it works … it changes audio. (Without a ransom.)
- The mentioned **Other Not Digital Thing** from last year is still a thing, and I believe we will be shipping that this year. Buckle up. It won't be for everyone, but it should be a seismic shift in one market.
- **Another new thing** is actually very old. Or at least moderately old. When people start inherently rejecting ideas for no good reason, sometimes it's time to look at why they are doing so, to question the wisdom behind the rejection, and to see if it's time for some new thinking. And that's all I will say about that. I'd expect a family of products to come of this … but I'm still playing with it. That's all I can say for now.
- And, **one last thing** that's so obvious, I'm

surprised nobody has done it before. But it's a very different market for us, and I still need to build the prototypes, though the general idea has been validated … so I'll shut up now.

Nebulous? Hazy? Infuriating?

Yes, perhaps all of the above. I know, I know, I wish I could say more, but hey, even the simplest intros sometimes slip (cough … Freya … cough). And none of these things are simple. None of these things are "business as usual." All of these things are pretty darn big ideas. Some of them can literally change the audio business.

"Change," you say suspiciously. "I don't like change. Other people talk about change, but then they want a bunch of money."

Yes. We understand. We know all about ransom-based schemes to remake audio. We hope they don't succeed. Because if they do, Mike might just piss off to make deep fryers or something like that. We don't need ransoms. We don't need the Infinite Wisdom of the Anointed Few to light the One True Way for great sound. We don't need the studios, artists, listeners, and manufacturers bound into proprietary standards that can change arbitrarily to "drive" demand for new stuff. We

don't need to be like home theater or medical devices.

What we need, at least in our mind, is more people doing interesting things, unique things that actually move the industry forward. And that's what we'll continue to do.

And nothing we do will have ransoms, or standards.

They will all be about choice. Your choice. To embrace and enjoy as you see fit … or to listen happily without them.

That's where we're going in 2017. Our new direction. Towards bigger ideas. Towards bigger goals. Goals like … *changing audio.*

For the better.
We hope.
And … without a ransom.

2017

2017, Chapter 1

Why Do We Do What We Do?

How best to start the year?

In 2016, I started with a bunch of blather about marketing. And I do have some thoughts on that subject for this year. But I think it's best to hold those for a while, until after we make some announcements. Then we can talk a whole lot about the futility of (much) marketing, where it's gone wrong, what matters, what doesn't, and how to make it right(er).

So you get to hear about the other thing that I had on my mind when I sat down to write this, and that's about why we do what we do. Or, to put it in more grandiose terms, about the philosophical underpinnings of Schiit Audio.

Specifically, I'll be discussing the *whys* that affect our products—why we choose to design what we do, why they are the way they are, why they have one feature, but not another.

This chapter is largely prompted by one kind of questions we get, both on these forums and via support, that boil down to:

1. Why didn't you include XYZ, I really wanted XYZ?
2. Why does it have to work that way, I think it should work the other way.
3. Why don't you support (insert FOTM technology here).

Because, from the inside looking out, there's a lot of method to our madness. But it may not look that way from your side of the screen. Especially if you're looking only at a single product. Especially if you don't have the time to grok us in fullness. Especially if it simply comes down what's best for your needs.

> Or, in other words, if you are a normal human being who has not absorbed so much of my blather that you can read my mind … you know, like 99.999 99 % of the planet.

Now, before I begin, a cautionary note:
Just because there are good reasons our products are the way they are, it doesn't mean you'll agree with those reasons.

Go back and read that again.

In fact, you may disagree with one or more of our *whys*. You may reject the philosophical basis of our products. And that's fine. This isn't about us telling you the One True Way. This is about explaining the reasons we do what we do. You can decide for yourself if it makes sense … or if it's BS.

Cool? Okay, let's get down to the *whys*.

Nuts, Volts, Switches, Metal, and Code: Design Whys

There are a million ways I can organize these various whys (technology, user interface, business philosophy, etc), but there will probably end up being some unclassified stuff, as well as some supercategories. So let's just dive right in to the biggest and most diverse category I can think of: design.

Design is really broad. It covers every part of digital and analog topology and implementation, every bit of why our products look the way they do, everything about how they operate. So, I'm going to see if I can break this down into top-level categories. Hopefully they'll make some sense.

Why #1: because sound is what it's all about! Sounds obvious, doesn't it? And yet there's a lot of gear out there that gets introduced first by the extravagance of their chassis, the size of their touchscreen, and the laundry list of "smart integration" features that they sport. There's no mention that the audio path includes things like volume control chips proper to smartphones and $199 receivers. There's no mention the designer used electronic switches, rather than relays. There's no mention that the analog inputs are converted to digital with a $7 ADC chip, or that the amplifier outputs are sonically approximated in Class D.

You know, because audio should integrate with your lifestyle. Products should be beautiful and unobtrusive. And because that's really the future—seamless, perfect, high-integration stuff with all the right standards and logos discreetly etched on a polished panel.

And, well, yes, the mindblowing cost, that's because this high integration demands new processes, new thinking, huge R$D costs, blah blah, BS BS …

No. Sorry. We aren't having any of that. Sound is *what it's all about.* The audio path shouldn't be

compromised. There's no reason to compromise it, especially at nosebleed prices.

Consider just one example: Routing audio through a black-box IC for volume control or input switching is a travesty when compared to the use of relay switching and relay attenuators.

- In the first example, you have literally no idea what the signal is running through—there could be hundreds of active elements in the path, or there could be entire amplification stages.
- In the second, you know exactly what it's running through—a set of mechanical switches and individual resistors.

So how is this "sound is what it's all about" philosophy reflect on our designs? In myriad ways:

1. **We use discrete analog stages whenever possible.** Discrete design allows us to tailor the amplifier stage exactly to its purpose. Discrete design also allows us to break out of the largely-similar forms of amplifier available on a chip. Discrete allows us to easily integrate tubes in meaningful ways. Discrete allows us to create gain stages with low or zero feedback. Discrete enables higher power output and

greater voltage swing. In short, we believe that discrete design, done right, provides superior sonic results. So we do discrete whenever we can … but, of course, there are times we can't. These times usually are constrained in cost or size … or both.

2. **We employ digital architectures with no missing codes and closed-form filters whenever we can.** Mike didn't spend a good chunk of his lifetime creating new platforms for multibit DACs because we wanted a buzzword to market with. We *literally* believe multibit sounds better than the delta-sigma alternatives, and that it is the *correct* way to approach digital audio for maximum fidelity. Of course, there are those who will argue with us, and that's fine, and of course, there are times when the more costly multibit approach doesn't fit into the budget … but as technology improves, who knows?
3. **We use Class A and AB topologies only.** Yes, we know that Class D is better than it has been, and that you can get like three billion watts out of a matchbox-sized device (just exaggerating a little.) But we also think that it, like delta-sigma, is a mathematically compromised approach that approximates the input signal, and is not yet (or perhaps ever) a candidate

for ultimate fidelity. Sure, if size and cost are constrained, it has a place. But if you don't have to use it, why seek it out? Understood some will disagree, but like I said, this is how we feel, not The Ultimate Truth for Everyone.

4. **We use linear power supplies in almost everything.** Except for Fulla 2, everything we make has linear power supplies. And yes, we're aware that switchmode supplies are much better than they have been, but again, why seek noise and complexity when none is needed? We understand some will disagree, but again, this is how we feel.
5. **We use good old mechanical switches and relays.** If you're using a switch on our products, it's connected to a switch using mechanical contacts, not a mystery-meat IC for switching or control—unless it's connected to a relay, also using mechanical connections. Call us paranoid, but we like to know what our sound is passing through. And mechanical connections and discrete resistors are the most innocuous things in the path.
6. **We use real analog potentiometers.** No volume control chips were harmed in the making of a Schiit product. If our products don't use a relay stepped attenuator, they use real analog pots—quality stuff from Alps.

Why #2: because the simplest interface wins. Or at least it should, in our opinion. That's why our stuff uses simple knobs, switches, and pushbuttons for control. Turn a knob for volume. Press a button to step through inputs. Flick a switch to change gain. Our goal is something that's easy enough to hook up, plug in, and use immediately, without a lot of education or fuss.

This is why you won't see things from us like:

- Redundant controls for the sake of controls (like digital filter switching)
- 156-page owner's manuals
- Cranky, slow touchscreens with 5 levels of menus
- Displays that don't do anything some lights couldn't do
- Meters that may never move, or provide no useful information

Is this too simple? Again, this isn't about being right for everyone. This is about why we do things. If you like it, cool. If you don't like it, that's cool too.

Why #3: because it should look good, too … Yes, yes, sound is what it's all about, but you know what? It doesn't really cost anything to create something that's reasonably attractive.

In fact, if the chassis design is smart, it might cost even less than the dominant paradigm of "sled+top+front panel" model. It might even cost less than something that uses standard extrusions and front and back panels.

What this means is that we spend a reasonable amount of time making our stuff look good. At least to us. No, it won't ever be the most complex, bejeweled product on the planet, but we consider it worth the effort to make it look good.

So how do we keep looks from leading us down a path to ultra-expensive gear? A few tricks:

1. We have a small number of basic chassis designs, which we modify slightly to create new products. This saves us tons of time and effort—otherwise we might want to create something unique for each product, which would be a disaster.
2. We don't go for known-expensive production techniques—things like milling the enclosure from a solid block of aluminum, or using 1 inch thick front panels, stuff like that. Everything we make is bent sheetmetal or turned aluminum.
3. We keep everything as simple as possible. Most of our chassis are 2 pieces, as compared to 3

for most other products. Believe it or not, this is really significant when it comes to cost.

Why #4: because half-baked technology sucks. One of the things we're asked about most often is why we're not supporting things like Bluetooth, or WiFi, or I^2S, or Ethernet, or why we don't make things with batteries. Well, all of those fall into the category of "half-baked technology," at least to us.

"But wait, batteries?" you ask. "You gotta be kidding, everyone's using lithium polymer, it's a well-known technology, it's literally everywhere."

Yes. Tell that to Samsung's Galaxy Note 7.

Sure, lithium polymer is pretty well worked-out today, but schtuff happens. And sometimes it's not pretty. As a much smaller manufacturer (who would want to do something totally bonkers like using two 11.1 V LiPo cells), we'd be taking on a much greater risk of failure. And even if we stuck to a tried-and-true single 3.7 V cell and switching rail generator, you're probably still looking at capping the warranty at 1 year. So, okay, maybe LiPo tech is 3/4 baked, but excuse us if we decide we have other, more interesting stuff to design first.

- Bluetooth? Still has problems with stable connections and doesn't support lossless audio. Two strikes against that one. Again, for those saying "it's worked out," talk to Apple and the iPhone 7 about their Bluetooth issues. Excuse us for sidestepping that customer-service nightmare.
- I^2S, Ethernet? Fringe use cases, complex and painful to use. Anything that might blow up if connected to the wrong thing (I^2S) or needs a lot of care and feeding to get working right (Ethernet) gives us hives. I think if I asked Mike, he'd say, "better to fix USB first." But that's up to Mike.
- WiFi? It means you're selling a computer. Have fun entering your WiFi domain and password with one knob and a screen. Or no screen. Yes, there are ways around this, but nothing that makes us want to wade in.

And … let's go a bit deeper. Because USB itself may be called half-baked technology, because its standards keep changing. Or even all of digital audio … after all, it keeps evolving as well. And this is why, if you have to use a half-baked technology, or sell products based on half-baked technology, we feel that you also have to address obsolescence. Because razing your wallet when tech changes really sucks.

And that's why our more expensive DACs are upgradable.

Why #5: because bad ideas should be actively resisted. Whether it's HDCD in the 1990s or MQA today, we believe that we shouldn't lie down and add questionable technology in fear of losing a few customers. After all, look what happened to HDCD. All that panic … for nothing.

While we have no problem with MQA as a format that can be decoded by a software player, we have extreme objections to them knowing, and, worse, dictating, aspects of our DACs' code. We have even further extreme objections to any attempts to make audio more like home theater, with its dizzying arrays of requirements (I am told that one standard requires 1200 separate test tracks to be run through a multichannel Audio Precision), customer confusion, and lack of independent progress (other than proscribed by the All-Knowing Pundits of the Standards Bodies).

So yes, we're actively resisting what we consider to be bad ideas. Call us crazy, as we look forward into an uncertain future (we have no idea if Tidal will survive, what "high res" format Pandora and

Napster will support, much less if Apple or Spotify will join the fray.)

Bottom line: we'll see how it shakes out. But I'd expect Mike will be designing deep fryers before he designs a DAC for MQA.

Bureaucrat Level 36 Report: Biz Whys

Tyler (our head of financial stuff) really should be writing this, rather than me. He's actually proud to be compared to Hermes on Futurama. Hell, sometimes he wears ties and fancy shoes to work. Despite this, he's actually a cool guy.

But, since Tyler doesn't write these things, you get me. I'll do the best I can, but in actuality, this is the part of the business that I like the least … the actual business.

> And I know, in reality, I really can't ignore it.

Why #1: Because thinking big beats thinking small. Huh? This might sound a little weird, but let me explain. I've worked for companies that thought small. How small can this run be? How few can we make? How can we avoid inventory? How can we do it CNC and avoid tooling? How can we shrink the time it takes from conceiving a

product to getting paid. Companies like this are usually looking to conserve cash, because doing big runs and keeping inventory is costly. Ask Alex and Tyler how they feel when a whole bunch of big product runs line up in a month.

Despite this, we'll think big any day. Thinking big—as in, large runs of products, investing in tooling, not worrying about doing things just-in-time, keeping inventory, buying ahead, etc—allows us to provide much higher overall value. Even though this means we have to be extremely fiscally conservative, Mike and I would prefer to re-invest in the company, increase runs, invest in tooling, and find new ways to make things better. We work directly with distributors and manufacturers to increase quality, combine parts purchasing, and to maximize what we can put into a product.

Does this work out all the time? Of course not.

As I write this, I know we're out of stock on some critical products. Blame the holidays, and the sustained high sales rate that historically follows the holidays. Or blame our inability to predict how popular some new products will be.

The thing is, we will get back in stock—and sooner than ever, since we're now working with

our suppliers more effectively. They know they can count on us for reasonable requests (not arm-twisting for lower and lower prices every time), large orders and fast payment, so they work harder for us. And they ramp up, to meet our larger production needs. It works well.

And, at the same time, we'll get back into stock quickly, thanks to a highly talented and motivated team. While other companies think small (as in, "how little can I pay someone,") again, we think big. Schiit operates beyond "first world" wages. This type of investment pays off in increased productivity, greater flexibility, and increased ability to scale.

Sounds idealistic? Maybe a little old-school? Yes. Perhaps. So be it. This isn't about maximizing weekly profit—this is about building something that will stand a long, long time. Or at least we hope so.

Why #2: Because cost should be as low as possible, but no lower. This one might sound a bit weird, too. Because isn't it a business' goal to drive down cost? The short (-sighted) answer is, "Yes." The longer answer includes knowing when to back down on your rabid goal of shareholder value.

We do a ton of things to ensure that the cost of our products are as low as possible, and that everyone gets a deal they're happy with. Some are internal. Some are external.

Let's look at a few things that we do to keep costs down:

1. **We standardize as much as possible and buy in large quantities.** I mentioned this before—simple and similar chassis, ganging parts buys, scheduling for additional discounts. All this stuff makes sense.
2. **We keep staffing lean.** We're a very efficient company. If we staffed at "industry norms," we'd have 5× the employees. Because we have well-paid, motivated people, we can do more with less—and keep costs down. And it doesn't hurt that we don't have salespeople or anyone with a marketing title.
3. **We don't waste time and staff on sales.** No sales, no special deals, no discounts, no points systems, no loyalty rewards—believe it or not, stuff like that takes a ton of time to create and administer. As an added bonus, everyone who didn't get the product at the sales price feels like they got screwed. As an added added bonus, sales stop until the discounts happen, leading to a never-ending cycle of sales.

And here are some things we don't do:

1. **Push suppliers to the breaking point.** When it's time to reduce costs, it's fashionable to go and yell at your suppliers. We don't do this. If we need to reduce costs, we have a discussion with them. If it's not feasible, we increase the price of the product, or elect to take a lower margin.
2. **Decontent the product.** On the other hand, we can go the Taco Bell route, and continue taking the meat out of the meat until it's not meat anymore. Hell, in the 1980s we were calling it "puppy chow in a tortilla." I don't want to know what it is today. Look at the progression of, say, receivers in the 1990s. Prices held steady, as faceplates got thinner, then eventually turned to grained plastic, transformers got smaller, chassis got weaker, etc. Nope. Sorry. That's not us.

Why #3: Because you need to resist the creep. You know what I'm talking about. It's that moderately-priced product that, in its second generation, put on a few pounds (of billet) and ballooned 50 % in price. And, in its third, put on a whole lot more weight and added a zero to the tag. That's the creep.

Yes, we know. It's tempting. Wow, it's a hot

product. You can move the price up, right? And if you add a few more features, you can move the price up again, right? And if you add some fancier cosmetics, the price can go up again, right? Because you're moving upmarket. You're getting in front of buyers with more disposable income. You're playing in the big leagues.

Yes. And you also may be killing your company.

Good luck when the winds shift, or when the current demographic dies out. Good luck with ever-shortening runs in a world where mass production is key to reasonable cost. Good luck explaining why this makes sense with digital products, in a world where digital costs keep dropping.

We'll stick to what we've always done: building a product to a reasonable price, applying a standard multiplier, and setting the most reasonable price we can. If it's a hot product … the price stays the same. When we redesign it … we don't add bling. And, when the cost of production drops … we pass it along. Like with Bifrost.

Resisting the creep—that's just what we do. We may be crazy. Who knows?

Are You Guys Crazy? Personal Whys

Now, all of the above has probably been touched on at one point or another in this book. It was probably past time, though, for an omnibus post about why we do things from a design and business point of view.

What I haven't talked about much is some of the more intangible whys—the personal whys. I can speak to my own, and I can infer some for Mike, but if you want to know his whys, it's probably best to ask him directly. Mike and I both love music. That's kind of a given. Mike knows a lot more about (more kinds of) music than I do, though, and he's much more deeply involved with acoustic, live, "real," music.

But beyond that, I think we also both love tinkering with things. We are engineers, after all. I know that one of the best things about Schiit, for me, is the freedom to experiment. Sometimes it doesn't go anywhere—hell, most of the time it doesn't go anywhere—but when it does, it's the best feeling in the world.

And … I really, really like the fact we're making things that people can enjoy. I spent a long time in marketing, where high praise is "well, it's not that offensive," or "that's kinda clever." And that's

if you got it past the client (more on this in a future chapter.) And, in the same way, I love the fact that most of what we do is affordable. Changing the economics of high-end is important to me. (And yeah, I like to tweak the gold-plated Bentley crowd a bit, too.)

I don't know about Mike, but I'm very happy to have found a place where I think I really fit—where I can use marketing, and writing, and engineering, and pull it all together, and create products that make audio better.

Hopefully.

If I don't mess it up …

2017, Chapter 2
“Obsolete.”

On Friday, January 20th, I was miserable.

Well, maybe not completely miserable, like you are when you have a crappy-ass flu, or pneumonia, or something like that, but definitely spacey, low-energy, and under the weather. Of course, I had nobody to blame for this except myself: I was 36 hours into a multi-day fast.

Yes, fast. As in, no eating. As in, at all.

Yes, I am an idiot. Or at least a little crazy. I like to try different biohacks from time to time, and see if they help or not. I’ve been tuning my diet, exercise, and supplements for many years now. Hell, I order my own lab tests to check my progress. And yes, I understand this sounds like a bucket of nuts … until you factor some notable results, including eliminating my high blood pressure and bringing my cholesterol numbers down into a happy range, without medication.

So, fasting was something I wanted to try for a while. Lisa was away with Jen on a writing retreat, so I had some days when I wouldn't have to cook, so I jumped at the chance.

And so, when AndreYew posted:

> Now the Stereophile Facebook dude is calling Yggy's DACs "obsolete." :/

I was not, perhaps, in the best place to read about it.

Say Some Nice Things About Stereophile

Okay. Let's be clear. This is not a "dress down and insult Stereophile" chapter. In fact, even today, my feelings towards Stereophile are, in general, warm.

Why? Because they sought us out and reviewed our products, before we ever spent a dime in advertising. They reviewed Bifrost, SYS, and Mani before any Schiit ad hit the magazine. I know Stereophile gets accused of "pay for play" crap all the time, because I read the comments on their own website. And the fact is, they don't operate that way. That speaks to honest curiosity and journalistic integrity. And that's significant.

> Contrast this with another industry mag that kept asking us to advertise with them … think, like, every month … without ever asking to review a single component of ours. After a couple of years I snapped and sent a testy email telling them that if they'd showed just a tiny bit of interest in our products, I would show a tiny bit of interest in advertising.

So, you know what? The fact that Stereophile went out of their way to find our products and review them—especially Bifrost, when we were a

tiny, tiny company—is deserving of recognition and respect.

But, on Yggdrasil, I think something very, very strange happened.

Back to the Dark Side

Okay. But what about that "obsolete" slur on Stereophile's Facebook page?

Yeah. Rage.

I sat there at home, all thoughts of a nice, leisurely stroll through social media before a nice, leisurely drive into the office forgotten. My low-blood-glucose-addled, adrenaline-pumped brain was on fire. The Yggy review had been very positive. Herb Reichert had loved the DAC. So why had they taken a cheap shot at Yggy on Facebook? What'd we done to piss in their Cheerios? What had we done to them?

It hurt even more, because I'd just confirmed a new ad placement with Stereophile's ad rep yesterday. All sorts of stupid ideas ricocheted around in my skull, but first and foremost was something like this: *Cancel the ad! Cancel all the ads! Make them feel the pain!*

And so, the first thing I did was to send an email to our Stereophile ad rep, copying John Atkinson, Stereophile's editor, saying that I found it hard to support them when they posted such blatantly inaccurate stuff, and educating them on what "obsolete" means in terms of engineering.

Aside: one of the things that lit me up the most was that "obsolete" has a clear definition in engineering. A part is obsolete when it is "not recommended for new designs, or 'NRND.'" The PCM1704 is obsolete. It is NRND. The AD5791 is not obsolete. It is in current production.

After that email, I posted this on HeadFi:

> Yes, this gets my blood boiling. Just commented on Facebook and sent a letter to the ad rep and editor.
>
> The reality is, for those just tuning in: AD5791BRUZ DACs used in Yggdrasil are 100 % current and used in the most mission-critical applications on the planet, including medical imaging and defense. They are much higher spec in terms of INL and DNL than any audio

DAC, ever. The Yggdrasil uses 4 of these blindingly expensive DACs.

Then, it was time to jump over to the Stereophile Facebook page and post much this same thing on their feed (which was already starting to take flak from some readers.)

By the time I'd gotten done with that, I'd gotten a response from John Atkinson, saying (paraphrasing) that he still thought "obsolete" was an appropriate adjective, given that we "live in a world of 24 bit audio." Still, he said he'd add a clarifying line on the Stereophile Facebook post.

Wait a sec, I thought, my addled brain kicking into high gear. *Does this mean that John Atkinson does the Facebook posts for Stereophile himself?*

That was a really weird thought. I always figured that their Facebook presence was overseen by some minimum-wage "social media specialist" in-house (or by the same person at an agency.) I was very familiar with that model, having done it for many companies at my marketing agency. And, to be perfectly frank, social media mavens weren't always the sharpest tools in the shed. It would be understandable if a low-level guy took a shot at us. But the editor of Stereophile himself? It made no sense.

So, yeah, the "clarifying line" got posted—"obsolete in the context of 24 bit audio."

Stereophile Magazine Like Page
5 hrs ·

Schiit Audio goes their own way, using obsolete 20-bit DAC chips - obsolete in the context of 24-bit audio data - in this converter. Yet Herb Reichert was very impressed by the sound.

Schiit Audio Yggdrasil D/A processor

Right now, I swear, Schiit Audio's Mike Moffat and Jason Stoddard are sitting there in California, smugly smirking at me and John Atkinson. While JA was struggling to properly measure Schiit's Ragnarok (Fate of the Gods) integrated amplifier for my…

STEREOPHILE.COM

Lee Scoggins and 103 others 7 Comments 3 Shares

Like Comment Share Top Comments

So I had to add to my response:

Jason Stoddard An additional comment: "in the context of 24-bit audio data" is somewhat disingenuous, since delta-sigma DACs process all audio data (16 or 24-bit) at 2-5 bits natively and then apply noise-shaping to improve static performance—and delta-sigma is used in the vast majority of DACs today. Of course, there are many paths to great sound, and the proof of the pudding is in the eating.

Like · Reply · Just now

Jason Stoddard Obsolete? Not at all. The 20-bit DACs used in Yggdrasil (specifically, four AD5791BRUZ) are 100% current, much more tightly spec'd than any audio DAC in terms of DNL and INL, and are used in the most mission-critical applications that include medical imaging and defense. This is a disappointing misrepresentation of the product.

Like · Reply · 6 · 59 mins

By this time, reaction to Stereophile's post was starting to percolate through the internet. At

least 90 % of it was negative—as in, critical of Stereophile. In fact, going back to the original Stereophile Facebook post today, there is not one single positive comment on it. There are some neutral comments (people looking for more information on associated components, etc,) but there are no supportive comments.

So, if John's intention was to generate controversy and get eyeballs, he definitely did that. But, in my professional opinion, zero supportive posts ain't a great outcome, even in the contentious and trollsome world of social media.

I sent an email to Alex saying I'd be in a little late, and spent some time looking at the reactions to Stereophile's post. And thinking about our history with Stereophile and with the results of their reviews in general. (One thing great about fasting is you don't spend time eating, so you can spend time thinking. Also, fasting apparently also floods your system with adrenaline to keep your blood glucose in range.

Had we done something to piss them off, really? No, I sincerely don't think so. Was "obsolete" a fair characterization of the DAC, something I should just leave alone? No, not in my opinion, anyway. Mike wasn't awake yet (he's a late

worker and late riser, I get up with the dawn), but I thought that would be his opinion as well.

Or, bigger, was this something to worry about? Was this something that, if we kept grinding on, would affect us negatively?

That was a tougher question. Especially coupled with the fact that magazines, despite their attentions, had never really moved the needle for us.

And especially, especially coupled with the fact that people were sticking up for us—supporting Schiit's viewpoint—across an increasingly large swath of the web. People didn't like the fact that Stereophile was attacking us. And not just people in this forum.

And that really sealed it. I realized that the things that mattered were what we were already doing: participating in forums, directly talking to customers, telling our story on Schiit Happened. The negative reaction—and the speed at which it spread—and the real-time changes it caused in our metrics—really underscored that the real action was no longer on the page, but on the screen.

Aside: and that's what I really think is the future. The replacements for the current industry magazines aren't going to come from a reorganization of the magazine, or a change in format of the magazine, or a move by the magazines online, or a concerted effort by the magazine to be present on social media … the replacements will come from places where people go to engage and interact organically. Places like Head-fi. This is the future. Just as commerce changed from the browsing at the local shop to point-and-click-and-get-it-in-2-days on Amazon, the voices of the experts are now online, in forums—and there are more experts, with more opinions, reviewing more gear than any magazine can hope to compete with. Yes, it's a bit like the wild west at the moment, but it's becoming clear where this is going … and it's going to be big.

So, after making a post or two about how "you guys are the future," I finally got up off my hunger-addled butt and went into the office. Just another crazy day at Schiit, best to put it down for now. I knew I could fight about this forever, bringing up points like the fact that delta-sigma DACs are, at best, 5 bit, so were they obsolete for missing 19 bit? And "living in a world of 24 bit music" is

pretty hilarious, when it's less than 0.1% of all recorded music at best, and the vast majority was still 16/44. Or even Stereophile's own guest editorial/manifesto by Bob Stuart, which clearly stated:

> Principally because of the combination of environmental noise and microphone self-noise (plus tape noise with analogue masters), very few recordings achieve let alone exceed 16 bit dynamic range. Add to this the fact that we can hear signals within noise only to about 10 dB below the noise level (see olive curve in Figure 21) and it follows that bits 19 to 24 carry no useful information.

So with more than 18 bit, sounds like we're OK, right? (Thank You, Bigro, for posting this.)

But fighting forever is no fun. Designing new products is fun. So I got ready and headed into the office, figuring that was the end.

Making Lemonade

Except it wasn't.

Because, somewhere between my shower and sitting down at my desk at Schiit, as the debate

raged on at Head-Fi, and on Stereophile's Facebook post, and elsewhere on the internet, I had a horrible, evil thought. At that moment, I knew exactly what the Grinch felt like as he hatched his plan to steal Christmas.

This was the thought: *"Obsolete" would be a great ad headline.*

Hell, I could see the ad, complete in my mind: a single word, "Obsolete," floating over a beautiful picture of the Yggdrasil. And copy saying something like "This is what Stereophile called Yggdrasil," and reasons why it wasn't obsolete.

And so, by the time I came into the office, I was chuckling to myself.

"You know what they say about people that laugh by themselves?" Alex asked.

"Yep!" I told him. "But that's OK."

"What happened?"

"Stereophile called Yggy 'obsolete.'"

Alex frowned, looking suddenly concerned.

"And I need to see if I can change the Stereophile ad," I said.

Alex's frown deepened momentarily, then he nodded and laughed. He knows all about my history in marketing, and he knows that I like to stir things up. He got it right away. "Oh, schiit," he said.

"Exactly," I told him, running upstairs.

Now, here's the thing. Like John Atkinson at Stereophile, I wear a lot of hats. One of the things I do is all of Schiit's marketing and advertising. And you could say that most of our marketing is simply a reaction against the stupid, boring, homogenized, oh-gawd-don't-offend-anyone crap I've been forced to make for people at Centric for over 20 years. I'm hands-on, and very good with Photoshop, InDesign, and Illustrator … and I have every photo from every shoot of Schiit products at my disposal … so doing a new ad took, literally, minutes.

Now, the question was, would they accept the ad?

After all, I'd confirmed our "Reset" ad for the preamps yesterday. Stereophile could easily say "done deal, sorry, we have your artwork and confirmation of it, run along now." And that would have been just fine.

But … you gotta try. I sent an email to the ad rep, simply asking if we could change the artwork, with the new "Obsolete" ad attached.

After that, I figured it would be time to sit and wait. Because they might want to talk amongst themselves, debate it, even ask for changes, or even reject it altogether. I'm not stupid. I know it's a (a) a bit of a dick move, and (b) deliberately pissing in Stereophile's own breakfast muesli, which might not be the brightest thing to do in terms of our future relationship.

Aside: but hey, Stereophile may have thought they were pissing on our Schiit's campfire with the "obsolete" comment … until they found out they were mutants that piss kerosene.

So, I posted the ad I wanted to run right here, and spent some more time on the forums. To say Stereophile was taking it on the chin was a bit of an understatement. Some of the top comments included:

> Wow. Just wow. 'Obsolete' is a rough term. Just how many mega buck components that Sphile reviews are packed with "obsolete" parts and designs … ? Probably 80 %.—DigitalFrontEnd

just looked at the page, the poster is really snide. I've never seen Stereophile post an inflammatory headline like that on a product positively reviewed.—mithrandir38, before "the poster" was known to be John Atkinson

The more I think about this, the angrier it makes me. The first time anyone puts any new effort into DAC design in 20+ years and Stereophile labels it obsolete.—dmckean44

Utterly bizarre. Another reason I'm not renewing Stereophile when my subscription is up. The misinformation doesn't end.—belgiangenius

I am rolling in the aisles over the obsolete comment. I wonder what Stereophile thinks of tubes and vinyl?—AudioBear—this was when I started going, "wait, well duh," and wanting to change the copy in the ad. Yes, I am an idiot.

And then, finally, Mike Moffat joined the fray and posted, "Perhaps tubes are obsolete as well … " And the comments from DigitalFrontEnd and AudioBear clicked, and my brain stripped a gear.

I should have thought of this from the start! Hell, everything we do is obsolete!

So, I posted this on Head-fi:

> Ah hell, I totally missed this.
>
> Discrete design: obsolete.
> Tubes: obsolete.
> Class AB amps: obsolete.
> Real switches, no software UI: obsolete.
> Multibit DACs: obsolete.
>
> Hell, everything we do is obsolete!
>
> I love it.

And—wonder of wonders—at that point, Stereophile's ad rep had gotten back to me, accepting the ad, and asking if it was final text. I told him I wanted to make a quick change to it, and sent him new artwork.

It was done.

And—kudos to them—Stereophile hadn't even blinked. They'd accepted our dick-mode ad without a second thought. In fact, the ad rep told me that he really liked what we were doing.

Or, in more wise posts:

> First they ignore you, then they laugh at you, then they fight you, then you win.—Mohandas K. Gandhi (KoshNaranek)
>
> Never mess with the guys who have no sacred cows. Especially when you consider yourselves above reproach.—FrivolsListener

Things Change Fast

In the old days, something like this could have been disastrous for a company like Schiit. To be branded with a disparaging term, "obsolete," with no method other than the slow, painful back-and-forth of letters to the editor or manufacturer's comments, could have hurt us.

Why Stereophile felt the need to do this—as a banner over a very positive review by Herb Reichert—is something I may never know. I get the feeling that some of the staff over there don't like us. Whether it's our name, our demeanor, our products, our undermining of the sacred order of Holy Magazine, Anointed Dealer, and Reverential Customer, or something else entirely, I'll never know.

But it's not 1987 anymore. Hell, it's not 1997 anymore (1997 is pre-Google, if that makes you

feel a little twinge of future shock.) And, really, in 2007, Facebook was only starting to take the mantle from MySpace and Second Life seemed like a smart thing to get into.

No. It's 2017. And in 2017, things move rapidly.

By the middle of the day, Stereophile's comments had spread far and wide. Posters brought up other equipment that didn't measure so well, or failed in Stereophile's tests, and had gotten a (paraphrasing) "Great job!" in testing. Several people who really respected Stereophile's testing started openly expressing doubts. Now, Centric no longer subscribes to sentiment analysis tools (software used to keep tabs on what people feel in social media), but I'm sure the net result wouldn't have been pretty.

But, by the middle of the day, my fast's grogginess and general malaise had gone away.

Aside: apparently this is normal after fasting for a couple of days. For the curious, I ended up taking it till Sunday afternoon, then eating again. The trippiest thing was going to the store before breaking the fast. Jury's still out if it was a good idea or not.

I wasn't hungry at all, nor tired, nor spacey. I

was happy I'd turned around something ugly into (what I thought was) a neat new campaign. So when I responded to posts online, it was usually just to thank Stereophile for the ad idea.

And this wasn't snark—this was sincere thanks. Sometimes you have to have someone kick you in the ass to get you out of a rut. And our advertising, while OK, wasn't exactly something I thought would set the world on fire. This new direction I like a lot more. Including "Obsolete" shirts at Axpona.

For Stereophile, though, things continued to be served up with additional snideness. When someone said they liked the sound of Yggy, he was told that he might simply like the sound of truncation (truculently delivered, after I had disambiguated rounding from truncation both via email and via Facebook post), and reiterating their opinion that Yggy was "sub-optimal engineering." Later, they would dismiss some of the most interesting discussion of Yggy—which included independent measurements (from 2 different sources) that didn't show the same behavior Stereophile noted, as well as some interesting proposed experiments, such as checking the musical content of the 4 LSBs—with one word: *amusing*.

But, back to "Obsolete."

I think Stereophile handed us a great theme that can serve us for several years to come. Because, you know what? It's time to skewer some sacred cows. And this campaign gives us a perfect platform for that. Imagine follow-up ads with headlines like:

Cheap.
Some people say we should charge a lot more. Here's why we don't.

Idiotic.
Yes, but our name is a big part of our success. Find out why.

Silly.
That's what we're called when we insist on Class A and linear power.

And so on. There are tons of other hurtful adjectives that we can throw at ourselves, unpack, and defuse. The words are simple, strong, memorable, and create ads with stopping power. And, once deployed, they can never be used in a negative way against us, ever again.

So yes, we're obsolete. Proudly obsolete. Even our prices are from another era. Our business

model, based on actual costs and what the market will bear, is completely out of whack with what they consider current today. Our products are almost unbelievably throwback, using power supplies and topologies considered wasteful and frivolous today. We offer not a single product with a screen, not a single product with a soothing artificially intelligent voice, not a single product with a wireless interface, hell, not a single product with a battery or soft-touch buttons.

And maybe that's what Stereophile sees. Maybe, to them, we are a veritable dinosaur wallowing in the primordial ooze, roaring in futility against the fire of the world-shattering meteor that ended their reign.

Or maybe it's the other way around.

As with everything, we'll see how it goes.

2017, Chapter 3

Leaving Marketing

It's funny how perspectives change.

When I was in college doing DIY speakers, I'd have laughed if you'd told me I'd ever do anything other than design speakers at my own company, and build a giant company on our idiosyncratic designs. Of course, reality intervened, and the speaker dream went nowhere.

When I was working at Sumo, I would have guffawed if you'd said I'd turn my back on engineering and have my own marketing company—to the exclusion of technical design. But opportunities changed, and that's exactly what I ended up doing when I founded Centric.

And when I started Schiit, I told everyone that it would be a fun & "hobby company," something that would be comfortable in a garage, or maybe a small structure in the backyard. But again, it took on a life of its own. It grew much faster, and

became much larger, than I expected. And my perspective has to change yet again.

And that's why, this year, I'm leaving Centric, the marketing company I founded and ran for 23 years.

Because now it's time to focus on Schiit.

Wait A Sec …

"Uh, I thought you were focused on Schiit," someone's probably saying right now. "I mean, you design all the analog stuff, you have this book, you do the ads, you go to (some of) the shows, you must be focused on Schiit, right?"

Well, yeah. For the three days a week I've been spending at Schiit, anyway.

You see, I've been sharing time between Centric and Schiit. At first, I spent all my time at Centric and did Schiit on the side, in the evenings and weekends. Then I started taking two days out of the week for Schiit, then three days. It worked very well, at least until recently.

"Wait a sec!" someone howls. "You've been doing this part-time?"

Exactly right. And now it's time to do it full-time. Because, let's face it, we're getting bigger, our line is expanding, the products need to be better than ever, we need to deploy new ways of making things, and we're going to need to pay more attention to customer-focused stuff (like "alert me when something is in stock" automation, reorganizing the site to help people find what they're looking for, being more strategic about outreach, actually taking the time to look at metrics, and, in general, growing up.

A bit. Don't panic. We're not going to go all big-corporate evil and Monsanto-y on you.

"Oh hell, if you've done this much part-time, how crazy is it going to be full-time?" someone asks. "You're going to take over the world!"

Actually, no. Maybe not even our little corner of it. Because you never know what the future holds. We could fall on our face, no matter how much time I put in. But the reality is, I'm having the time of my life at Schiit, which is something I couldn't always say at Centric.

And if I'm having the time of my life, why not do more of that?

So What Happens to Centric?

That's a fair question. Centric has had a long, long run as a marketing company—over 2 decades. In dog years, that's about 140. Centuries, that is. Most marketing companies either implode (founders' egos can be big problems), fail, or get acquired long before that.

Aside: for disambiguation of what "a marketing company" is, refer to previous chapters about marketing. Centric was, to be precise, a full-service marketing agency with a focus on creative and technology. Which meant we had in-house designers, artists, etc and in-house programming and technology staff. Which translated to a whole lot of online work from about 1998 onwards. But which also encompassed items like complete rebranding, trade show strategy and implementation, and even print, outdoor, and broadcast ads.

And Centric had a lot of firsts: We were one of the earliest agencies online, one of the first to do websites, one of the first to do online marketing, one of the first to do SEO, one of the first to do Flash sites, one of the first to do social media marketing, one of the first to do virtual reality and augmented reality stuff. In fact, I think we

still did the largest ever VR event—the 12Avatars contest—and a whole lot of first-ever stuff online, including a kids' virtual world for Bandai.

Looking back on it, I could summarize Centric as growing and thriving by constantly pushing the limits of change, in a constantly changing environment. When we started, hell, people were still doing color separations (look it up) and scoffing at computer typesetting. We did one of the first huge catalogs using digital photography (done with a monochrome sensor and 3 rotating filters) in an era when people were still using drum-scanning. We embraced the web very early, which led us to working with some huge companies like HP and Compaq. We were first in line when we learned about SEO, creating a department to manage it. We experimented avidly with social media from the MySpace and Friendster days onward. We played with Second Life and augmented reality. We deployed novel insurance-backed contests that got a lot of attention. We built our own consoles, took our own metrics, made our own reports and recommendations to clients …

… but today, the pace of change has slowed. There are, or seem to be, less new opportunities. Marketing has hunkered down into a highly

metrics-driven, test-and-iterate based system, or "big creative" that can only be afforded by incomprehensibly giant companies.

Or, again, so it seems. Perhaps the pace finally got to me. Perhaps marketing is only the province of the very young, who can devote 14-hour days in 7-day weeks without blinking. Perhaps we're missing the new frontier.

And that's a good lead-in to what's going to happen to Centric.

Because Centric won't die. It'll just move on. Our current staff is taking the company and running with it. It looks like Centric will end up being more metrics-driven, more testing-based … but there will still be a creative core. I've signed up to spend time helping them through the transition this year (for free), but from what I've seen, they aren't going to have any problem without me.

And I want them to be successful. I sincerely hope they can navigate the problems with today's marketing, and the problems with today's agencies, and create something truly stunning. Because the potential is there. Right now, there are many problems to address. Which means, for the companies that get it right, there will be huge opportunities.

Problems? Yeah, let's talk about problems. And how Schiit helped me see some of them a whole lot more clearly.

What's Wrong With Marketing

What's wrong with marketing? One word: *Influence*.

That's a book title. By Robert Cialdini. Look it up. Read it. And suddenly understand why all marketing has a familiar (and sickeningly sweet) flavor, why you're treated a certain way, spoken to in a certain way, and offered to in a certain way. Sometime between the publishing of that book and today, marketers inhaled it, absorbed it, and it became part of their very being, and that's why marketing is the way it is today.

Because, for all the stats and numbers Cialdini cites, they are all from an era pre-Influence. Today, after several decades of Influence's tricks, people are beginning to see through it. They're beginning to look askance at too-tricky, too-obvious techniques, because they've seen them too many times before. They're beginning to tune out the soothing patter of the corporate script.

And, in this type of environment, where everything has been calculated to influence, where everything is so plainly fake, *anything* authentic is an amazing departure—and people flock to it.

I didn't realize how terribly fake, so sickeningly upbeat, so easy it was to drop into that mindset and write copy that sounds exactly like everything else these days, that is designed to soothe and cajole and persuade … until I started Schiit.

And Schiit is authentic. Schiit stands out.

And people notice. I think this is a big part of our success—because we broke out of the sickeningly fake mold of "we're so happy to have you as a customer, you are immensely valuable to us, current wait time on hold is 96 minutes, I'm sorry we can't cancel your account without an early termination charge of $185, did you know you may elect to choose our wonderful new offering of Product Insurance which guards against most likelihood of loss for only another $58 while you wait." The right words, the perfect soothing voice … and you're screwed in the end, or sold something. Schiit is not as nice, smooth, right, soothing, or perfect … but we get you taken care of. Which shows *real* caring.

I've said it before, and I've said it again, the logical end-game of this is the McNugget phrase: *made with white meat.*

Use this phrase on a non-marketer, and they usually either:

a) Smile and say, "Yeah, it's made with all white meat chicken, it's good for you."
b) Look at you blankly and ask why you're so terrified.

Use this phrase on a marketer, and you'll get one of two very different responses:

a) They'll look increasingly more disturbed as the implications sink in, finally saying "Wow, I'd never eat that," or something to that effect.
b) They'll shake their head and laugh at how great a marketing phrase it is.

Why this difference? Because the phrase is engineered to seem maximally positive to the widest audience, without actually saying anything. To a non-sensitized, non-marketer, this sounds like a positive phrase, something to be proud of.

But let's break it down.

1. It does not say how much white meat is used.
2. It does not say what kind of white meat is used.

3. It does not say that there is any kind of white meat, at all, in the product.

So, this phrase could equally apply to a 100 % soy-based chicken substitute nugget being made entirely in a factory, with the human pushing the button while eating a pork chop as the only "with white meat" statement. Or maybe it's a 100 % automated factory that has a pork chop sitting on the machine. Still with white meat.

Made with white meat.

Sound a little different to you now? There you go. The pinnacle of "reduces the appearance of fine lines and wrinkles," (appearance?) "save up to 90 %," (on one sock, everything else not so much), "maximizes the potential for a positive outcome," (lolwhut), "virtually everything," (virtually, ouch) and so on.

But the Cialdini book isn't the only signpost of bad marketing.

Let's add in benchmarking against the competition, paralysis by analysis, focus-grouping to death, and the nail-biting, gut-clenching possibility of actually being noticed. All of these are subsets of fear, and fear is the other big problem in, well, pretty much everything to do with big corporate environments. Nobody wants to take

a chance, because if you fail, you're done. And that's why we see things like:

- **Obsessing over what the competition is doing,** leading to me-too products, differentiation by feature proliferation, same-old-same-old ineffective advertising, races to the bottom (because, hey, if customer service sucks at the competition, why bother doing anything other than incrementally better?), and, basically, all the stuck-in-a-rut **** that ruins companies. Ask yourself: do you really think your competition is so much smarter than you? If so, then maybe best to not get into business in the first place.
- **Paralysis by analysis,** leading to late or compromised product introductions. I've seen this a thousand times—when companies are not 100 % sure of a product, they'll stall, they'll wait for more research reports, they'll tweak the description, they'll tweak the advertising, they'll go back to engineering to ask for more features, they'll iterate and iterate and iterate until the product is a complete non-starter. Bring something out. Don't do more if it doesn't do well. Fix it if it has a problem. And if it does well, make it better and do it again.
- **Focus-grouping to death,** leading to the end of some very good and deserving ideas. Guys,

> I can't say this any more forcefully: most focus groups are 100% complete BS, done to massage the egos and wallets of the marketing group. One loudmouth invalidates them. Lack of vision within the group kills great product. I have never seen a focus group provide information that wasn't completely obvious to someone with a modicum of functioning brain cells. I have seen them kill great concepts in products, service, and advertising.

So, in short, what's wrong with marketing is simple: it all sounds the same, and nobody wants to take a chance. Which is why it all sounds the same. And the merry-go-round of benchmarking, research, and focus-groups means it trends towards a mean. A mean that I think, on my more cynical days, looks kind of like Comic Sans—yes, it's a font, and yes, you can read it, but no, it's not something I'd ever want to see on my documents.

What's Wrong With Agencies

"Well, I gotcha, corporate life sux, I ain't gonna be a drone, I'm gonna work at an agency so I can do great things," somebody is probably saying right now.

And, you know what? Maybe that'll work for you. Because working at an agency, coming up with amazing creative ideas that really do stand out, that are really different, that really move the needle is a wonderful rush. You're expected to be "out there." You're expected to challenge the client. You're expected to shake things up.

So, yeah, if you want to do great marketing work, I won't talk you out of working for an agency. Or starting your own, like I did. It can be a ton of fun. It can also be quite lucrative.

But …

> You knew this was coming, right?

But let's start with the lucrative. Centric did very well for many years. Unfortunately, those "many years" were not consecutive, nor did I get to pick the ups and downs, nor the magnitude of the ups and downs. '97 was great, '98 we almost died, '99 thru '01 were great, '02 to '03 not so much, '04 to '05 our biggest time, '06 to '08 a slow slump, '09 almost fatal, '10 to '14 steady, '15 to '16 a slow decline.

Why such variations? Chalk it up to both us and to our clients. When times are good, it's tempting to put all of your effort into maintaining good

relationships with your current accounts, and forgetting about adding new ones. The problem with this is that your current clients:

a) May be heading into an industry downturn, and therefore turning off the advertising spigot (and there's NOTHING you can do about this)
b) May be about to have their lunch money stolen by a newer, bigger, better, badder bully competitor than they ever expected
c) May be on the verge of getting acquired, and the new management has its own pet agencies
d) May be on the verge of a marketing shakeup that shakes you out.

Here's the brutal reality—you can be doing great work, selling the crap out of your client's products, everyone there can love you … and it can end with essentially no warning.

And, when that happens, and your revenue shrinks by $1.5 million instantly, and you're staffed for that level, well, you know where this is going.

Or maybe you don't. So let's spell it out. It's going to be layoffs. Because you're not going to suddenly go out, sell your heart out, and spin up new clients immediately. Getting new clients is a

6 to 18 month cycle. So if you've been ignoring new client generation for a year or two, there's nothing in that cycle. And then you're majorly hozed.

And, when you have layoffs in a small or mid-sized organization, they're deadly. Because, up to the layoff, everyone knows everyone else. They're more like a family than a big corporation. They go out to lunch with each other. So when you have to say, "Sorry, but you have to go," and "Congratulations, you get to stay," you're creating a bitter group of ex-employees who may stay in touch with your current employees for years … feeding them poison. Plus, anyone left knows there may be additional cuts, so you should plan on losing them … the most talented first. The fallout from layoffs can take many years to work itself out. It's never clean, it's never easy … and it may never be over.

"So easy, we just keep selling all the time, right?" you ask. "Getting new clients all the time."

Yeah, and maybe some organizations can do that. At Centric, we always lost sight of the ball. Whether it was simply trying to do the best for our current clients and not selling new ones, or

being so insanely busy there was no time to do the work, we always ended up on a rollercoaster.

And, now that I mentioned "busy," let's talk about that. Agencies are high-pressure environments. There's a reason most of the people working in them are young. They have a higher tolerance for the 12 to 14 hour days and 6 to 7 days per week. They're more capable of staying overnight to re-do a presentation for a critical new client. They're more willing to go all-out, over and over, even when great stuff is shot down.

Now, is every agency like this? No. But many are. We eliminated this from Centric in the 2010 re-org—and, at the same time, eliminated the need for anyone to come into the office other than one day a week. Not coincidentally, these were our most stable years.

And—here's the real reason that many people leave agencies, forever, in a decade or less: no matter how amazing your work is, it might not matter. Your client might not like the color orange, so it's done, no discussion of changes. Your client may not understand they're selling to a 16 to 24 demographic, even if they are 56. Your client may be an alcoholic, certifiable, or simply abusive. And your options for dealing with these clients

are limited. Because, after all, they're clients, you need to serve them well so the agency will be stable

> until it isn't

and that's the way it has to be. Yes, you can defend, you can cajole, you can bring numbers and charts … and maybe that will work. And maybe it won't.

Be ready for the latter.

Had enough? Wait, there's more. As an agency, you also have a limited ability to choose who you work for. There are lots of good creatives out there, and not so many companies that are willing to spend good money on outside creative work.

> **Aside:** And this should be the first warning sign, shouldn't it? If a company needs an outside agency to tell them how best to market their own product, shouldn't you run screaming? Yes, I know this is always rationalized with phrases like "needing a fresh perspective," and "getting an outside look at things," but the reality is—if you don't know how to best portray your own products, you got problems.

So, even if you are a vegan, you may be working on the National Cattlemen's Beef Association, or if you're a low-carber, you may be working on Pepsi, or if you're an organic aficionado, you may be working with Monsanto. And so on. You'll be expected to muster great enthusiasm for products you might never consider using, or even find abhorrent. And you'll be expected to do great work for them.

With all of the above, do you wonder why many agencies do "pro bono" (free) work for clients they like and believe in? First, they get to pick good causes. Second, since it's free, the causes will frequently accept the work, without even picking it apart first. Pro bono is sometimes an agency's best work, and it can feel wonderful.

But it doesn't keep the agency fed. And that—literally—is the bottom line.

How to Fix Marketing, and Keep Your Sanity

So how do you magically fix all of this stuff I've just spent 3300 words blathering about? Hell, I don't know. I don't have all the answers. If I did, I might have stuck with the agency side, ha!

Seriously, though, I may be able to put up some signposts. These aren't 100% of the answer. They're probably not even half. But it might help …

For companies wanting to do better marketing:

- **Trust your gut.** Your first reaction to an ad is most important. Don't overthink it. Don't overanalyze it. Don't show it to your spouse/cousin/co-worker/dog/fish. If it stops you and has the right message, it's the right marketing.
- **Do something fast.** Don't spend a billion years agonizing over the "perfect" ad, website, copy, etc. Get it out there. Iterate. Keep it fresh and different … so by the time you've been benchmarked, you've moved on.
- **Measure and optimize.** Find out if what you're doing is moving the needle. Do more of what works. Drop stuff that doesn't. Test variations to see what does better. Don't let metrics take over the creative, but know what it's doing for you.

For people working in agencies, or starting their own:

- **Be bold.** We never hired timid or mainstream people at Centric—and our clients never picked us because we were safe. Better to be remembered than lost in the crowd.

We've had clients come back two or three years later, after making a mistable with another agency (usually someone larger).

- **Mind the tenure.** Agencies are high-pressure environments. How long can you keep running at full speed? Do you have a plan for the future? Or are you comfortable doing this ad infinitum? Be honest with yourself.
- **Forge new paths.** The pace of change may seem to have slowed, but I believe there are multiple ways to blow up the old agency model and come out ahead. Some of the ones we discussed (but I didn't take forward, because they would be an entirely new company, and I need another company like I need a hole in the head) include going to a 100 % price-list model (clients hate hate hate hate hate the uncertainty around invoicing) and focusing on specific measurable platforms, such as specializing in helping companies sell, and market, on Amazon.

To all that choose to go down the marketing path, good luck! It can be a whole lot of fun. Just remember to keep your sanity.

Am I Sad?

A little, yes.

I started Centric over 23 years ago. Anything that takes up 2 decades of your life is more than just a passing whim.

I'm very happy to have created an agency that employed great people and did some wonderful work. I'm even more thrilled that our current staff is taking it and running with it … and hopefully turning it into something even better. But it is sad to part ways.

But it's a happy time for Schiit. Because now it's my sole focus. If I keep learning, designing, and making things better, that should be a good thing.

Should. If I don't mess it up.

2017, Chapter 4

Deprogramming

Welcome, comrades!

Welcome to this wonderful new age of audio glasnost!

This is a beautiful new era, sure to be heralded as the beginning of a new chapter in the great adventure that is the audiophile hobby!

To help you, dear comrade, prepare for this new age, we have prepared this scientifically-designed and thoroughly-tested deprogramming text. Because, most devoted brother (or sister), we understand that you have endured many years of programming by entities bent on twisting the very way you think. Bent on nefarious aims such as convincing you that yes, everyone should pay for satellite radio to get pretty much the same radio we used to have before they crapped it up with 500% more ads and "personalities" who won't shut the hell up. Bent on changing black to white, up to down. Bent on …

But we ramble. Please study and absorb the following text, in order to prepare yourself for the new age of audio transparency and plain-speaking. Re-read as many times as necessary, and please share with comrades near and dear.

WARNING: This deprogramming text may challenge what you consider to be incontrovertible audio truths. You may experience strong emotions. This is normal, comrade.

Welcome to the new age.

Deprogramming Part One: Power Poppycock

In the design and application of power supplies, we have found a vein of propaganda that runs deep and wide. It is this vein we will mine first. Please, dear comrade, note this is not a complete compilation of all power-related terms. Please feel free to submit your own to the Ministry of Audio Transparency for inclusion in a future revised deprogramming text.

Following are some terms and their deprogrammed equivalents:

LPS. AKA “Linear Power Supply.” Usually portrayed as a dramatic upgrade from a “switching”

supply. However, if the product already uses a linear power supply, and is properly engineered, it's unlikely that a LPS will make much, if any difference. Also note: another way of saying, "linear power supply," is "same old power supply used in pretty much everything since the beginning of electronics, and before our tolerance for power-supply noise was exceeded by our desire for tiny stuff with good battery life." Not very sexy when described like that, right? Also, note that putting "precision" or "audiophile" in front of "LPS" adds absolutely nothing in terms of quantitative description.

Linear Power Supply. See LPS.

Audiophile Power Supply. Like LPS, but says even less. "Linear" is at least a quantitiative descriptor referring to the fact that it is not a switching supply.

Precision Power Supply. Like LPS, but says even less. "Linear" is at least a quantitiative descriptor referring to the fact that it is not a switching supply.

External Power Supply. Often portrayed either (1) as a critical design decision that differentiates the product from proletariat versions which house the power supply and audio circuitry in the same

chassis, or (2) as an optional component that increases the performance of the product. Both of these assertions will be addressed separately.

1. **As a differentiating design decision.** Sometimes, designers move the entire power supply outside of the component. While this may have some benefits in terms of radiated noise, it is frequently obviated by stacking the component on top of its power supply. Many transformers have much higher radiated field from the top and bottom than from the sides, so stacking the components eliminates any benefit that might come from separating the power supply section from the audio circuitry. It also more than doubles the cost of the chassis, due to the need for two chassis and interconnection cabling and connectors. It can also be hazardous if the interconnection cabling carries high voltage. It can also result in lower performance if the regulation is done in the power supply chassis, rather than in the audio chassis—regulators are ideally placed as close as possible to active circuitry for best performance.
2. **As an optional component.** Sometimes, you have the option of purchasing a component with a "basic" internal power supply, or adding

an optional external power supply, typically with the promise of higher performance. Leaving aside the wasteful nature of having redundant power supplies, ask yourself how poor the design must be if it changes drastically with the addition of an external power supply—especially if it is a line-level component running in Class A, and therefore having no additional peak power demands.

Toroid Transformer. Often portrayed as a "better kind of transformer," toroids do look cool. They're round and kind of retro-futuristic, like a cyber-donut. Some comrades are so programmed that they frantically search, lemming-like, to ensure a component uses toroids, and consider them a "must-have." In reality, toroids are expensive, may require adjustment to "null" out their radiated noise (meaning more production expense), and are frequently misapplied. For example, the high-frequency response of toroids are better than conventional (EI-core) transformers, which is exactly what you don't want when specifying a line transformer. You'd much rather have the transformer reject the high-frequency noise from your computer's power supply, than pass it through. Toroids have their uses, but many times they are spec'd for no other reason

than cosmetics. Sometimes potted into cylindrical cans with big stickers on the top calling out various imaginative qualities of the transformer, such as "custom" or "precision". Again, note lack of quantitative meaning.

R-Core Transformer. See Toroid. Similar idea.

Low noise. Many power supplies are described as "low noise," with the propagandist's hopes that the subject (you) will feel a net positive reaction to these words, perhaps conflating it with "low noise from the headphone or speaker output." In the absence of more information, however, such assertions are meaningless. Low noise in terms of radiated EM field? Low noise in terms of RF radiation? Low noise in terms of no audible hum from the transformer? Low noise in terms of DC output from the supply? Let's break this down:

- **Low noise (EM):** all products that use magnetic components (such as transformers or chokes) emit electromagnetic field. Especially if they have aluminum chassis. There are no hard-and-fast standards for this. However, the propagandist was probably not talking about this anyway.
- **Low noise (RF):** well, we certainly hope so—both the FCC and CE have standards for this.

If you don't meet them, bad stuff is coming your way.)

- **Low noise (no audible hum):** well, all transformers hum a bit, but it shouldn't sound like the 50-year-old chest freezer growling away in the corner of your garage.
- **Low noise (DC voltage output):** this is probably what the propagandist is talking about. However, without numbers, "low noise" means nothing. Good low-noise regulators can get down into the handfuls of microvolts these days. Claims of nanovolts should be met with guffaws. Because physics. And, as an added bonus, even with numbers it may be meaningless—if the power supply rejection ratio (PSRR) of the amplifier stage is very high, even a noisy supply might not matter. Bottom line: the best way to evaluate low noise is in use—do you hear noise from the headphone, or from the speaker at the listening position? If so, then you may need to use products with a lower noise floor. However, this may have nothing to do with the power supply—it may do with the inherent noise floor of the device.

High current. This is a term usually used very similarly to "low noise," in that it is intended to convey a vague positivity in the absence of numbers. "High current," used without numbers,

is meaningless. Even when used with numbers, the meaning of "high" varies with the product and application. A 1 A (one amp) power supply would be considered high current for a preamp, but very weak for a speaker amp. So, even with numbers, you need context to make this phrase meaningful.

- **Related term: VA.** This refers to volt-amps, a common measure of a transformer's power output. Large numbers, such as 500 VA or 1000 VA, are usually used for speaker power amps. This is a reasonable way to get an idea of the overall size of the transformer used in a conventional linear power supply, but it is not a be-all or end-all in terms of system performance. Look at the rated power into different loads—this gives the overall system performance, and is more meaningful.

Fast. Many power supplies are described as "fast." Again, this is a meaningless term used to convey a vague positive feeling. In reality, power supplies are only as fast as their charging time constant—a number which is never supplied, because it's probably not very impressive. Nor is speed related to the size or number of components. Two gigantic capacitors will charge just as fast as

the same value achieved by paralleling dozens of small capacitors. Again, physics rules.

- **Related term: ESR.** Sometimes, one encounters claims of "low ESR," where "ESR" stands for Equivalent Series Resistance, a quality measure for capacitors. If provided with ESR numbers, this can be a helpful metric—however, without numbers, you guessed it … it's meaningless.

Deprogramming Part Two: Balanced Blather

"Balanced" is another term that many propagandists love. In reading some programming texts on the benefits of balanced components, one may quickly come to the conclusion that balanced is the One True Way, that All Components Should be Balanced, and that Balanced Is the Path of Nirvana.

Wrong. Look at recording studios. Those are all 100 % balanced, and they have produced some mightily crap-sounding recordings.

The reality is, there is little magic in balanced for balanced's sake alone. It's in the implementation. And that's something the propagandists don't talk about so much. There are also many other

things they don't like to talk about, like how many balanced components aren't really balanced, how all inherently balanced components need to sum the balanced output to single-ended for best performance in SE, and how many components convert the signal for every input.

The reality of not-really balanced. "Balanced" is a lovely word, since it implies all gear that isn't balanced is a bit off, a bit, well, unbalanced. The siren call of "balanced" has led many propagandists to apply it to all manner of gear, even if the gear is not really balanced. Unfortunately, doing balanced properly is expensive. It requires more components, exotic components, and/or exotic circuit topologies to do it right. So, it's much easier to do one of two things:

1. Hang balanced connectors on the box and ignore the inverted balanced phase entirely. Yes, this happens. And it is a convenient way to get balanced connections when you need them. But it is in no way truly balanced.
2. Convert balanced to SE for internal processing, then SE to balanced at the output. This is a common way to do it, since it eliminates the need for exotic topologies, lots of parts, and exotic components like 4-gang potentiometers. However, in this case, the signal is going

through two conversion stages (many times with IC op-amps), which is something the propagandist is usually loathe to discuss. If the product has balanced connectors and a 2-gang volume pot, it's either ignoring or converting.

The reality of summing. Summing is the process of converting a balanced signal to single-ended, such as on the front end of a product using #2 above. Summing preserves the common-mode rejection ratio of balanced operation, and causes common-mode signals such as correlated noise to fall out of the signal. Summing, therefore, is the right way to derive a single-ended signal from a balanced signal. Some inherently balanced products use summing to derive single-ended outputs as well. Ironically, the propagandists are silent on summing when it is used to convert balanced to SE on the input of a non-balanced "balanced" product, but may try to use it as a negative when it is used on the output of an inherently balanced product. The reality is that the non-balanced "balanced" product has two conversion stages—twice as many as the second example.

The reality of conversion. Even if the product is inherently balanced, how does it interface with single-ended signals? Again, this is something

the propagandists don't want to discuss. Some products use an inexpensive "phase splitter" IC op-amp to derive the second balanced phase. Other products use inherently balanced, differential stages that can only output a balanced signal when presented with single-ended input. Bottom line: It's good to know how your balanced gear is implementing balanced.

So, how best to choose balanced gear, in this morass of confusion? Choose it if it sounds good to you, and works with your system. Don't worry about interfacing it with single-ended gear—remember, great sound is the goal, not matching interfaces.

Deprogramming Part Three: Classless Claptrap

We have gone through extensive deprogramming with respect to amplifier classes in previous texts. A fast refresher: propagandists try to apply Class A to everything, even if it is clearly not Class A. When faced with Class A claims, narrow your eyes and ask how they are defining Class A. And if you're faced with clearly propaganda-driven constructs like "Class A-D," remember, it's D in the end.

Deprogramming Part Four: Damping Drivel

Progandists like to blather on about “damping,” even when it is criminally misapplied—and, of course, there are never any numbers to back up the claims. There are many classes of products where damping is absolutely critical—products such as speakers, headphones, and turntables. In other words, products that move.

When something moves, damping becomes an important metric. You wouldn’t want to have a speaker made out of 22-gauge sheet steel, because they’d literally ring like a bell and sound awful. Similarly, you wouldn’t want a turntable platter made from carbon fiber, because you’d rather have a lot of mass to smooth out any variations in speed, as well as to help provide an inert platform for the record. That’s why you tend to see thick walls on speakers, internal damping applied to headphones, and massive turntables. Damping is important for these products.

For a product that just, well, sits there? Not so much, comrade.

Putting a DAC in a CNC lead billet enclosure will do approximately zero for its performance. Wrapping that same DAC entirely in sorbothane (a damping compound) again, will do pretty

much nothing. Same goes for an amp, a preamp, or pretty much any kind of electronic component. Despite this, propagandists frequently go on about features such as "isolation feet," or "CAD-designed nonresonant chassis."

So, here's the quick summary:

- **For products that move: damping is important.** Manufacturers know this, and may be able to provide quantitative data on how they have damped their products. If not, the informal "knuckle rap" test on speakers to hear how "dead" they are provides some information about damping. In terms of turntables, physical and material construction provides clues—heavy metal platters and plinths usually are better than thin plastic, for example.
- **For products that don't move: damping is not important.** This includes products that are moved, such as portable devices. If they don't have a spinning or vibrating component, damping doesn't really matter.*

* Really. Even though some ceramic capacitors can have piezoelectric properties (that is, they generate voltage when vibrated or squeezed), the magnitude of this effect is tiny in proportion to the voltages involved, and will be obviated

through the impedance of the overall power supply, as well as the circuit's inherent PSRR.

Deprogramming Part Five: Magnitude Madness

Finally, dear comrade, let us discuss the propagandist's greatest vice: hyperbole. To listen to the propagandists, you may have been led to believe that changing from one component to another is a life-changing experience. To read their texts, tears will stream from your eyes at the sublime beauty of your recordings, finally revealed in their ultimate glory. You will bask in the warmth of all-encompassing love as you are transported in ecstasy.

In reality, the magnitude of changes is relatively small, unless you're talking about two things:

1. **A transducer.** Speakers and headphones sound markedly different. Nobody ever compared Audeze LCD-2s to Grado GS1000es, and said "Hmm, I really can't tell any difference.
2. **A broken component.** If your amp or DAC are literally broken, then yes, they may sound so different that the propagandists' terms are correct.

Now, comrade, "relatively small," does not mean, "inconsequential." Some small differences can make your listening experience subjectively much better. You may be highly attuned to these small differences, as well, and perceive them as being much bigger than the average listener.

However, if you don't hear marked differences, don't panic. We are all different human beings, with different perception.

A Final Warning

Be aware, comrade, that this deprogramming guide cannot cover every permutation of the propaganda you have been exposed to. It also cannot counter decades of indoctrination. We apologize if you feel distressed or upset, but we assure you: deprogramming is the first step into a shining new age of audio.

Also, please be aware, dear sister (or brother), that forces still conspire to skew your senses, to re-program you into accepting that Only the Finest CNC-billet Chassis with Balanced Teflon-Oxygen Interconnects, Selectively Damped by Hand-Selected Holographic Materials, with Optional External Ultra LPS with Low-ESR Capacitor Upgrades is the Only Way to audio nirvana. These

forces, known as “marketing,” are pervasive and, unfortunately, outside the scope of this deprogramming document.

Again: read, re-read, and share.

Author’s note: Yeah, this was brought about by having to hear the acronym “LPS,” said in a non-ironic way, for about the one trillionth time. It was also prompted by seeing a transformer in a box (like Cthulu, but 6× the price) described as something truly life-changing.

Consider that Cthulu could be described such:

“Integrated, High-Current Precision LPS. An audiophile power supply that dramatically redefines your audio experience. Offering fully 562 % more electromagnetic-based energy reserves than a standard Schiit Wall Transformer, Cthulu provides the power needed for amazing bass extension, effortless musical crescendos, and unsurpassed impact—while delivering smooth, linear AC for all your Schiit wall-powered devices. Thrill to the finest details of your favorite musical pieces with Cthulu, your indispensable Schiit Linear Power Supply. Order now and be one of the first to experience this marvel for only $ 124!”

And yes, sure, I can write promo copy. I certainly did when I was in marketing. I also got very good at making things sound really, really good, while promising nothing. Go back and re-read that paragraph. It does not claim a single audio benefit from using Cthulu … but it certainly implies it, right? Because many times, if you expect a difference … well, there it is.

> And, conversely, if you expect *no* difference … well, there it *isn't*.

2017, Chapter 5

Back to the Past

I've said before that you'll think we're crazy. And, given that it's April 1 as we announce this, you may think this is a joke … even though it isn't.

But with this announcement, you'll really think we've lost it. Because this is about the wildest thing we've done … I mean, we're talking a serious turn-around, a trip right back to the 19th century, a complete inversion of how we do business, a 100 % super-surprise retro thing that even I never expected we would do.

Aside: this is the fourth Manhattan-project level event that I mentioned in passing in the end-of-2016 wrap-up chapter. I'd just had an epiphany about how to make it work … and I'm thrilled and a little shocked that this is the first of our truly wacky ideas to become reality.

And yet, even though I never expected to do what we're doing, I'm oddly excited about it. In a year of exciting product launches, this thing has

me thrilled. Maybe it's just that it's so different. Maybe it's just the potential. Maybe it's just–

(loudly, from the back) "Ah, shaddup already and tell us what it is! You've bloviated for, like 12 tweets now, and we still have no idea WTF you're talking about!"

Okay, here it is: we're opening our first retail store.

Yes, you read that right. Retail store. As in, a place where you can go in, sit down, and listen to a whole bunch of Schiit. As in, a place where you can hear both desktop systems and speaker systems. As in, a place where pretty much the whole line is set up, and you can compare to your heart's content.

> Yes, I know, bring the smelling salts.

"Way way way wait a second!" The voice at the back bellows. "Are you talking about bricks and mortar?"

Yep! Or, well, more accurately, 2×4s and stucco, with a fake kind of southwestern glaze on the outside. Because earthquakes. We don't really do bricks that much in SoCal. But if by "bricks and mortar," you meant, "physical store where people

can buy stuff," yes, that's what we're talking about.

> And no, none of us have been institutionalized lately. At least as far as I know.

"But, wait!" That voice at the back comes again. "A retail store? As in, that thing you've done so well at avoiding? Why start now? Why do you think you'll do any better at this than any other dealer on the planet? What can you possibly hope to accomplish with this?

Okay. In order, yes, yes, because it's something people have been asking for and we're a little crazy, we don't, and to learn lots about this whole new (to us) side of the business.

And, an aside: it's *not* a dealer.

But I'm getting ahead of myself. First, let's talk about the whys.

Why Go Back to the Past?

I know I've been hard on dealers at times. I've railed against the distribution model and its now-unnecessary-in-the-21st-century inefficiency and costs. I've attributed a big part of our success on our avoidance of dealers, on selling direct.

But, in all that time, I always gave credit to dealers for one thing: providing a place where people can come and compare a whole bunch of different gear, before they buy. There's definitely a value in that.

> **Aside:** assuming the deck isn't stacked to favor the products with the highest margins. Yes, I know, honest dealers don't do that. That's why, if you buy from a dealer, you choose carefully.
>
> > **Aside to the aside:** and what we're talking about *still* isn't a dealer. But let's get through the whys first.

So, first why: Because it would be nice to allow people to compare our products before buying. Especially with the line getting big. We now have a wide selection of desktop products, DACs, preamps, accessories … and we'll be selling affordable stereo amps before long.

But there's more to it than that. We also get an increasing number of people who ask if they can pick up an order at the factory, or if we have a demo room at the factory. This is a problem, because there's always a lot of next-generation stuff running around the main SchiitBox (and

it might even be sitting on Alex's desk, right next to the entrance), plus people talking about what's coming, or meeting with vendors who are involved in next-generation products, or … you get the point. Picking up products is usually fine, albeit disruptive, but having people in the shop for any length of time is not the greatest idea if we want to leave the leaking of new products to myself and Mike alone.

And there's even more, but this one is more about attitude and curiosity: I've always wondered what it takes to create a successful audio retail experience. I've talked about it with some other audio manufacturers, and the talks have always ended in a kind of stunned exhaustion. Because once you start adding up the costs to open one retail store, including fixtures and staffing, and then replicate it a dozen or twenty times, you get into some seriously intimidating figures. It always gets too big, too fast. And we shut down the discussion, walk away, and have a drink.

But it always comes up again … what if you could do retail? What if you could really make a go of it, without making the costs so high it makes the company non-competitive? What if it worked?

And that, of course, was on the back of my mind,

as we started to wind down Centric last year, and I sat staring across the desk at Tyler. And that's where this thing really took off.

"We Should Do Retail": A Convo With Tyler

Tyler, if I haven't introduced him before, is our head of finance and admin. Finance, because we need someone who's responsible to head that up, and admin, because there are lots of loose ends in our business, loose ends that are best not left to Mike and I. In reality, this dual title means that Tyler does our bookkeeping, interfaces with the bank and the accountants, does wires, some POs, HR stuff (famous Tyler quote: "I got dental and vision coverage for the employees for free, I'm telling you so you know, but I know you don't care, because the overall coverage costs are less."), and a whole bunch of miscellaneous stuff that I'm sure he'll have to shovel off on to a lieutenant some time in the future. Tyler has always been of the opinion that we need to do retail.

"I've been looking at lease rates in Old Town (Pasadena) for a SchiitShop," Tyler told me one day. "It's not that bad."

"Plus staffing, plus fixtures."

"Set it up so it sells coffee during the day, hookah at night, cover the costs with that," Tyler shot back. We'd discussed stuff like this before: having some kind of hangout attractor that could cover the basic costs, so this wasn't so surprising.

But still, inwardly I groaned. Because having a location in Pasadena—32 miles away—would be cool and all, but it would also involve staffing, fixtures, security, a long-ass lease, and who knows if it would work.

"You know you want to do it," said Tyler. "Make it the polar opposite of the Apple Store, all black and aluminum."

"Apple has, like, a trillion dollars to spend on stores," I told him.

"They had to start somewhere."

"Yeah, they started with only ten billion or something like that." I shot back. But Tyler was right. There were lots of people asking if they could listen to our stuff, even locally.

"Plus, what happens when we open one store, and everyone wants one in New York. Or Chicago? Or Atlanta? Or San Francisco? Or, hell, worldwide? It's a can of worms," I added.

Tyler laughed. “That’s when you tell them tough, it’s in beta.”

“Yeah, that’ll go over well,” I said.

But still, in the back of my mind, something stuck. Something about that idea of “beta.” Just like Gmail a decade ago: Gmail (beta). It was good enough that people used it, but it got some slack by not claiming to be the be-all-end-all-perfect-product.

What if we could do something like that? I wondered. Maybe not in Pasadena, but right here? Centric was winding down. We had a cool space in Old Town Newhall, which was starting to really get turned around, with a gastropub and a winery and a wine/cheese/coffee shop and a smokehouse and a brewery … and we’d been leasing that space so long it was month-to-month and cheap. And that space was really kitted out for high-end retail …

But then I started thinking about what it would cost to staff it, and how much we’d have to sell to cover the costs … and it all fell apart again.

That is, until December 2016, when I had my epiphany.

In December, my decision to leave Centric was made. I'd seen the numbers, and it was time to move on. Which meant clearing out the old office, since the guys who would take the company from here worked remotely. Which meant, I was sitting there one day, alone, writing the year-end wrap up for Schiit 2016, staring at the reality of that very cool space, and wondering, *How can I make retail work?*

This time, though, the answers started falling into place.

An approximation of my thought process would have gone something like this:

The big problem with retail is that retail is costly. The three greatest costs are in staffing, rent, and fixtures. We had fixtures largely covered, because the person who'd had the office before Centric had kitted it out as a high-end wine shop, and spent, literally a hundred thousand dollars in solid custom hardwood shelving and cubbies, polished concrete floors, hand-finished plaster, and on and on. It was a really nice place to spend time. A few thousand dollars worth of table, couch, chairs, and signage, and we'd be done with fixtures.

We also had a head start on rent, because, as I mentioned, we were already month-to-month and the overall rent was pretty inexpensive. Still, it wasn't an insignificant cost. If there was a way to bring it down by half, that would make it a no-brainer. We'd have fixtures and a retail space for a year for the cost of three or four ad placements in the hifi mags.

Or, I realized, we could have the SchiitShow at this location, and that would cover fully half the rent! SchiitShows are expensive. Really expensive. But if we moved them up here, and expanded out into the patio space we shared with the restaurant next door, and had them cater the event … then the remaining rent really looked like a pittance.

Epiphany Number One: you can offset rent with the SchiitShow.

Wow. Now, that was exciting!

But it still came back to the ever-painful sticking point: how did we afford to staff it? Even one full-time staffer would be much more costly than the rent. That was a real problem, and that was something we'd have to deal with, for real, if we wanted to make this work. I mean, we couldn't just pull someone off the line and say, "Hey, this

is your job now, have fun selling stuff." Even though lots of our staff would probably be happy to spend a day just sitting around and listening to our gear.

Cue me sitting back thunderstruck. Because, really, we had enough staff that we could cover the retail space with people working a half day a week. And it would serve as a way for everyone to get to know the products better, which a lot of people wanted to do. We could see if we had enough volunteers to cover the time.

Epiphany Number Two: we might not have to staff!

And if that worked, we now had a way to open a retail store—at minimal cost—and staff it—at minimal cost—and experiment in this brave new world that we knew nothing about.

I sent an email to Alex and Tyler, outlining the idea and asking Alex if he thought we'd have enough volunteers. Alex's reply: "Hell, *I'll* spend one day a week there."

And, as we thought the idea through, more good ideas came out: we could also use the space for seminars (Mike or I talking about various aspects of digital or analog stuff), we could have product

intros there, we could bring the press there, we could have "tech nights," where we'd have a technician there to hook up bizarre combos of products or try new permutations with customer gear.

Done right, the retail space could be much more than a retail space.

But one thing was certain: it definitely wouldn't be a dealer.

Introducing the Schiitr: (Not) A Dealer

"Schiitr? Is that what you're calling this? What a dumb name," someone mutters at the back.

Yeah, well, lots of people thought our name was dumb when we started. Including lots of "big" audio companies that are now looking uncomfortably over their shoulders at us (or staring at our backs.)

Aside: now, don't take that as a guarantee this crazy idea will work. It may crash and burn. That's the way with betas. It may not be around a year from now. It may break even and become an important way for us to interface with local clients. It could go big and we could roll it out

in multiple locations. Any of these things could happen. Or something completely unexpected. We'll see.

So why isn't the Schiitr a dealer? It has to do with our decisions. And those decisions include:

1. **We'll only be selling Schiit.** No other brands. We won't be carrying headphones or speakers. We won't be selling music servers or disk spinners. We won't be pimping $3000 cables and turntable mats and magic stones and hologram stickers. Of course, we'll have a selection of headphones and a couple of speakers to demo, but nobody is going to be taking those home. Sticking to selling only our stuff saves us a massive amount of headache and admin work for those times when people decide to return stuff that isn't ours, or when they get something that isn't ours that doesn't work. That kind of admin work can eat us alive. We have well-developed procedures for dealing with our own returns and exchanges, but not so much for anyone else. Plus, selling only our stuff saves us from the inevitable fisticuffs that will result from not stocking (insert your favorite headphone brand here) or not stocking (that one special-edition model of your favorite headphone brand here.)

2. **We will not be selling, period.** The people who staff the Schiitr will vary—from hardcore audiophiles to guys who test the product, to people who screw them together. They're not salespeople, they are not expected to be salespeople, they are expected to babysit the products, answer the questions they can answer, and ask you to leave if you spend all day there drinking the free coffee and never buying anything.
3. **We will not guarantee an answer to every esoteric, tech-nerd, or audiophila nervosa question.** As mentioned above, some staff will know the line backwards and forwards. Some will not. Some will have strong opinions about what sounds good. Some will not. Some will be highly techy. Some will not. Hell, I'll be there from time to time, and I can't guarantee I can answer all your questions. So there.
4. **We will provide free drinks.** No, not alcohol. Holy crap, I don't even want to think about the liability there. But if you want (good) coffee or tea, feel free. But remember, there's only one bathroom, and the walls are thin.
5. **We will leave you alone to listen as long as you want.** Unless, of course, there are a bunch of people waiting. Or we're closing. You get the picture.

"Wow, that sounds fantastic!" Some people are saying. "What's the catch?"

Sure. You bet. Lots of catches. Switches in the back and all that. So here you are:

1. **There's only one Schiitr.** It's in Newhall, CA. That's about 30 minutes away from downtown LA on a good traffic day. However, it is literally right across from the Metrorail station, so if you want to ditch the car, you can take the train. You can also look forward to hearing the train's horn blaring in the background as you listen. Yes, we know, this is unfair to everyone who doesn't live in Southern California. But hey, look at it this way: it's also unfair to people who DO live in Southern California, given our traffic and our crappy public transportation. But until we invent a teleporter, well, that's the way it's gonna be. And if we invent a teleporter, expect me to be a little too busy being, oh, the first multi-trillionaire for a while.
2. **It's not going to be open all the time.** Our initial hours are Tuesday to Saturday, 10 AM to 6 PM. Yes, I know, this is entirely unfair towards Mondays and towards people who work normal hours. But that's part of the compromise necessary to keep staffing costs

down. Ah well, we'll see how it goes.

3. **We won't be perfect from Day 1.** It's an experiment. Beta. We won't have every single permutation of our products available at every moment. We won't have every single detail worked out. This is a test. Give us feedback. We can improve. But everything hinges on this location working out first, before we go any farther. Remember, we still don't have a firm plan on how to expand beyond this one location (though we're talking about a couple of bizarre ideas that, if they work, will be very, very interesting). And remember, we will not impale ourselves on this idea—if we can't make it work, so be it.

So if you'd like to come out and see us, we'd like to see you. Come in, sit down, relax, have a beverage, and enjoy some stacks of Schiit. Join our mailing list and we'll keep you up to date on what's going on there, from Tech Nights to seminars to the next SchiitShow.

Talk to us, and let us know what's working, and what isn't. We'll tweak and improve … and hopefully start an entirely old revolution, in the revival of the company store.

2017, Chapter 6

Conversations With A Stubborn Engineer

Back in my first job for Magnavox Advanced Products and Systems, there was an engineer that, well, all engineers know. He was, in short, a caricature, a downward-looking, muttering, rumpled-shirt, wrinkled-old-black-tie, pocket-protector kind of guy who literally (a) lived with his mom (in his 40s) and (b) had to be reminded to periodically cash the paychecks he'd leave in his desk for months at a time.

I don't know if I ever really knew what he did for us. I think he was a mechanical engineer, because he never seemed to spend any time in the electronics lab, but he was always in the review meetings. Sadly, I don't even remember his name, because some of his bewildered, innocent questions echo in the back of my mind to this day.

I can rationalize this, of course; Magnavox was a

serious government job working on secure communications for alphabet agencies; we weren't encouraged to fraternize; discussing the wrong thing with the wrong person could be very, very bad; the lab itself fostered paranoia, as you had to go through 2 scrambling keypads and an armed guard to simply get into the cube farm—our unnamed hero of this story, the Stubborn Engineer, pointed out a failing in that system that I, in my naivety, demonstrated to the staff, and instantly earned the suspicion of, well, pretty much everyone in the company, and probably delayed me getting my Secret clearance.

Aside: the security flaw the Stubborn Engineer pointed out was the fact that, for all the unmarked doors and scrambling keypads and anterooms with guards watching you punch in your passcode, you could go straight from the corridor outside to the secure lab by lifting up the ceiling tiles. He was smart enough to simply mutter about it; I was young and a smart-ass and I decided to demonstrate it. Not the brightest move. But then again, we live and learn. Or certainly hope so.

So, let's stick with that moniker: *The Stubborn Engineer.*

The “Do You Mean” Question

The Stubborn Engineer’s specialty was his “Do you mean … ” question. This question was usually delivered after an hour’s worth of presentation by the team lead, or 45 minutes of soporific blathering from a vendor rep.

And it was deadly. Because the Stubborn Engineer had a way of distilling the most opaque buzzword blather and corp-speak ******** into simplest terms, and then turning it around a hundred and eighty degrees to make something seem like the silliest thing in the world.

Examples?

Let’s start with the one that really frosted my ass, because it affected my pet project.

I was doing the software programming interface for the communication system we were working on. I’d suggested using a standard Sharp pocket computer of the time that ran a version of BASIC. Easy to program, inexpensive to buy, simple to document, lots of programmers (at the time) who could make changes as the hardware evolved.

Stubborn Engineer only half-watched the demo, leafed through the software user manual I’d written, sighed, and said, “So, do you mean

on the project where we're spending multiple millions of dollars and doing custom hundred-thousand-dollar ASICs to create a truly secure communications interface, we're counting on a consumer product you can buy at Radio Shack to set them up with? Will Sharp be making these over the next decade of service life? Or are you planning on buying a truckload of them and convincing DC to store them in a spare room of the 5-sided building there?"

I went red and snapped, "Better than using your foil-wrapped UNIX terminals that cost $35k and can't be moved out of the room."

They were wrapped internally in copper foil because we were paranoid about people remotely sensing the EM from an operating computer, so even cheap PCs were blindingly expensive.

Stubborn engineer pursed his lips and shrugged. He didn't care. He'd asked his question.

Now, everyone in the room was muttering about our problem with the Sharp portable computers. What had seemed a great idea for the past four months was now fragged like a model town under a nuclear blast. Suddenly we were drawing up plans for a purpose-built handheld and building

new timelines. I never saw how it turned out, because the Sumo job came up about a month after that, and, well, the rest is history.

But he did that to pretty much everyone.

"So do you mean you're designing in a single-source connector that costs $1685 each, when you could use a standard SMA that costs a couple of bucks?"

"So you want us to engineer up a complete hand-held enclosure for something that's still on eighteen 10 inch×12 inch boards—would you like to be first up in review when we have to scrap it?"

"So do you mean you want to run microwave tests in the hall and take the chance of frying someone's future kids because it's too much trouble to haul the stuff out to the remote test area?"

It was funny when it wasn't you. And he made you stop, look around, and consider the big picture. In my time at Magnavox, he probably saved more money and time than any other complete team put together. (When they listened to him—they didn't always.)

And he was right about my damn Sharps. They would be a problem in the future. I wasn't looking far enough ahead. And I wanted to use something

that I was familiar with. And I wanted to show off a little. But that didn't mean the answer was right.

So why dredge up this ancient history? Because, I realized a few days ago at CanJam, we need another Stubborn Engineer. Or a few of them.

Because, when you turn things around, a lot of what we're doing looks really, really silly.

An Audio Epiphany (or Two)

I remembered the Stubborn Engineer first when I was sitting on a panel about "Headphone Audio in a Changing World." This is a panel they select people from the audio biz and ask them to combobulate various trends with high-end personal audio (and, hopefully, come up with some interesting things for the audience to consider.)

Now, when I was asked to participate, I almost passed. Because the first thought that came to mind was *Ah, great, Apple EarPods, Lightning connectors, Bluetooth crap, and the Internet of Things … not exactly high-end, hmm.*

But, in the end, I accepted, thinking, *Well, I gotta keep an open mind, let's see what's coming down the pike. Plus, maybe they'll need a voice of reason.*

And, sigh, yeah, they needed a voice of reason.

The panel started out with a question about wireless headphones. One of the headphone manufacturers on the panel was super-excited about this, and praised the new Apple tech as a game-changer. The fact that it was still a compressed standard, and that there were multiple different and confusing standards for Bluetooth connectivity, was never even hinted at. The other headphone manufacturer said they were also working on wireless tech, but mentioned some of the downsides, including battery life, and the recent episode of exploding wireless headphones.

I decided to not mince words.

"It's funny," I said. "We spend all this time trying to get the highest quality sources—the best masters, high-res material, stuff like that—and then send it down a quality-destroying, MP3-quality connection? Let's be clear, Bluetooth wireless headphones are quality-destroying and potentially dangerous technology for the sake of convenience only."

That took some of the audience aback. Some muttered to themselves, but I also saw some knowing nods. Nobody had ever couched it in such stark terms before.

(What would the Stubborn Engineer have said? Wait. I'm getting to that.)

Then, of course, they talked about other typical stuff—DSP and the challenges of small battery-powered amps, etc … all the stuff that implied the future was just another endless romp down the assumption that lives would get ever-more-frenetic, even more on-the-go, with no fixed place to relax and enjoy … as if the entire audiophile market would be turned over to overcaffeinated transients in the next 10 years.

I said a few things about how tuning to personal preferences is fine, but that we, as a scrappy manufacturer of cheap stuff, were focused on making sure that people could get a taste of what the high-end could do, for not very many dollars.

That theme came back later when someone in the audience asked, "But how do we get outside of the audiophile ghetto? How do we expand the audience into the mass market?"

I sighed. Because this is one of the questions that always comes up at audiophile events. The whinging, hand-wringing, oh-gawd-we're-all-gonna-die-if-we-don't-go-mass-market positioning that is, well, 100 % crap.

"First," I said. "We have to make sure that if people want to take an audiophile path, that it's affordable. And I'm not talking "affordable" in terms of $5000 amplifiers. I'm talking 2 figures. I'm talking, at or below the fashion-accessory headphones. And that's what we do."

"But," I added. "We also need to let it go. There are forums online where people can discuss $1500 steak knives and $15 000 bicycles and $150 000 cars, and there's nobody whining about how they need to be able to move into the mass market. Why do we think we're special that we have to proselytize and convert people who aren't interested? Bottom line, some people will love high-quality sound, and some won't. It's OK. Let it go."

Now, this was something they'd never heard.

I was challenged by someone who said something about advertising, so I reminded them that I had an agency for 20 years, that we had done work for the high-end, that the question of reaching the mass market almost always came up, that (in the early days), we took the money and tried to do that, and that it always failed.

I brought up the economics of mass advertising (they were thinking that "a million dollar cam-

paign" was a lot of money, I told them that was a 3-month buy in the San Fernando Valley). Big brands are built on 9-figure-per-year campaigns that span decades. Bottom line.

Nobody in high end can do that.

And—to be clear—nobody in high-end SHOULD do that. There's plenty of room in high-end for all—from entry-level companies like us to ultra-high-end audio jewelry. It's just that maybe, right now, the spotlight is blazing too brightly on the jewelry.

More From The Show

The Stubborn Engineer came up again later when I was talking to a designer at another company. Said company was thinking about doing something like the Audeze Lightning cable for their headphones—essentially a phone-powered DAC/amp dongle with a Lightning connector on one side.

I said, "So when Apple goes to USB-C next year, do you throw all the Lightning cables in the landfill?"

The guy stood there looking dumbstruck. "Ah. Oh yeah. That's right! Apple probably will go USB-C. And then … "

"And then you're boned."

"We'd have to do a Lightning version and a USB-C version at the same time," he said, brows furrowed. "Oooh, that's expensive. Not good."

"Maybe better to wait," I offered.

He nodded and muttered to himself a bit, while I stood there and wondered how anyone couldn't see that one coming. Apple is a mile-a-minute on USB-C, as evidenced by the dongle-bonage of the new MacBook Pros (yeah, I have one, but I like my Surface Pro better … yes, I know, heresy, but Apple is not really Apple anymore.)

Apple is absolutely gonna bone everyone with USB-C on phones. Just like the 30-pin to Lightning transition. Perhaps they can put all the iPhone 7 Lightning earbuds in a landfill next to the trillions of AOL disks from decades past.

And, I realized, this is where the Stubborn Engineer could help us. He could get us to turn around and look at some of these crazy audio tropes from an entirely different angle.

And that might be, well … illuminating.

What Would The Stubborn Engineer Say?

Okay. I'm doing that thing where I stick my finger in the stat amp running 500 V rails … while dripping wet from the swimming pool. Some of you won't like where this is going. Some will accuse me of being fundamentally anti-customer. Some will say I'm the Stubborn Engineer.

So, let's take a look at some of the latest audio fetishes, when run through the lens of the Stubborn Engineer.

On Bluetooth: "So, you want us to spend thousands of hours in engineering to create an interface that is lower-quality, less reliable, and requires a de facto agreement to tens of thousands of hours for support, patches, and upgrades in the future, because you're too lazy to plug in a wire?"

On WiFi: "So, you want us to pick one of several half-baked standards—any of which might go away at any time—and spend thousands of hours in hardware engineering, plus become a de facto software company producing apps for setup, or add a screen and alphanumeric IO to the product,

plus the tens of thousands of support hours in the future, plus updates on both the hardware and software until the end of time, because using one of the well-known and reliable wired analog or digital interfaces is just too much trouble? Oh, and remind me why we have multiple digital formats and connection standards?

On wireless headphones: "So, you want us to spend tens of thousands of engineering hours to produce something with lower sound quality, higher noise floor, customer frustration due to lost connections, the inconvenience of another device to charge, additional weight, and potentially dangerous batteries, based on a fast-moving standard we have no input on, because customers can't just plug the headphones into their phone?

On streamers: "So you want us to become an actual computer company, producing custom hardware and software, with the tens of thousands of hours associated with it, to release a product that will require the largest amount of customer service ever seen in the history of the company, to replicate something any $300 computer already does?"

On DAPs: "So, you want us to build all the infrastructure for advanced handheld manufacturing,

starting with micron-tolerance machined chassis and tough glass technology for large screens, plus advanced battery management, phone-level microprocessors, and custom operational software—which will be buggy as all heck, no matter what we do—in order to create a device that is bigger and does a small subset of what your phone already does? Oh, and have you ever seen the customer support for smartphones? You're crazy, I quit."

On MQA: "So, you want us to submit to a proctological examination of all our DACs by a third party and rely on their timely approval to certify them, with the associated hundreds of hours of engineering necessary to include their code in our devices, plus the thousands of hours in the future to update it, plus the thousands of hours of customer support to explain the various options on the one shaky streaming-music service that supports it, for the limited amount of music that is MQA-encoded, plus how the MQA-encoded CDs will play through the system, and take the chance that this is (a) what will become the de-facto standard for big streaming services like Spotify and Apple, and (b) will not simply be perceived as an irrelevant consumer standard like SRS?, and (c) has not painted itself into a lose-lose

corner?"

Aside: MQA's lose-lose corner: in my opinion, MQA has two ways to lose and zero ways to win. Lose 1: Get picked up by Apple and Spotify. At that point, we can just say, "Well, it's another compressed consumer standard, not relevant to true HiFi," and ignore it. Lose 2: Not get picked up by Apple and Spotify. In which case the studios will stop seeing dollar signs and it'll go away. But that's my opinion. I've been wrong before.

On DSD: "So, the vaults never opened, and your sales are great without it … why would you expend any effort on it?"

Aside: I think he might be able to be convinced on DSD, since, unlike MQA, it is another open standard. He'd probably grimace and groan about the lack of content, though. And, no, this isn't a veiled hint that we'll be doing anything with DSD.

I can keep going—home theater tech (surround sound) is another place where you have little control over where the industry is going, and a limited number of companies that support the actual implementation of important standards

like HDMI. High integration products that don't allow for upgrading are also vulnerable as digital tech changes (so you have a cool Bluetooth/Wifi/speaker/headphone/all-singing combo now … but what happens when better tech comes along?) but I think you get where I'm going.

The point being: charging in to new tech has a lot of gotchas … a lot of gotchas that aren't easy to see from the outside. This is why we're going to be slow and cautious when it comes to new standards, especially when they're fundamentally a convenience, and especially, especially when that convenience comes at the price of quality.

We also need to understand when The Stubborn Engineer is just plain wrong. Because nobody is perfect, no matter how sharp their tongue.

We understand that means our products won't always be for everyone, and we're OK with that. Again, we have been wrong many times before. The market—meaning you—will let us know when we're lagging too much. And then we just have to be agile and open-minded enough to adapt.

And I think that's something we'll be able to do.