



TRANSCRIPT


Former Product Manager at Amazon Web Services

 Amazon.com, Inc. [AMZN](#) | Posted July 12, 2022

Former Product Manager at Amazon Web Services

Expert Details

Compliance

 Former Exec

Former Product Manager at Amazon Web Services, leaving in November 2019. The expert can speak in-depth to AWS and their work across the world of Cloud Computing, with a focus on Machine Learning, Database, and Analytics services.

Current Director Of Technology at Imply. In this role, the expert is responsible for a wide range of technological decisions and strategies.

Prior to Imply, the expert worked in Professional Services at Snowflake Inc., leaving in February of 2022.

Prior, the expert was the Product Manager at Amazon Web Services, leaving in November 2019. In this role, the expert worked with global teams to improve and evangelize services for data and analytics, while helping customers get the most from their use of data in the Cloud. Some of this was product management work, defining offerings, managing development teams, roadmap, budget, support, and execution. Some were evangelism.

The expert was one of less than 1% of Amazonians to be granted the top tier of speaker certification, which they use to present keynote and other roles at conferences worldwide. Some of this was strategy, collecting information from customers and researchers to try to understand what the market would need in 3-5 years to be sure we would be able to deliver it.

All of it was across the world of Cloud Computing, with a focus on Machine Learning, Database, and Analytics services.

The expert was involved in the creation, launch, and development of a number of services in that space. The biggest ones are elastic cash, aurora, redshift, elastic search, and the creation of open search.

The expert can speak to data analytics and data systems. In the software space that would include vendors such as Oracle, Microsoft, IBM, and Snowflake.

Interview Call on June 10, 2022

Tegus Client ▶ 00:00:00

Hello, thanks for speaking with me today. So, we are digging into cloud computing and some of the hyperscalers so the usual suspects of Azure, AWS and GCP. But I think a good place to start is probably just get a quick overview of your background, your career, how you got to AWS, why you left, what you did at Snowflakes, that would be great.

Former Product Manager at Amazon Web Services ▶ 00:00:19

Okay. So, I began as a physicist at NASA in the 1980s. And I ended up running IT operations at Johnson Space Center. And ever since then, I've been doing things in IT and with data. So that includes a wide number of roles and jobs. And without giving every individual one, I ended up being recruited to a company called Sequence.

Sequence was acquired by IBM in 1999, and I stayed there until 2014, where I left as Acting Vice President. Technically that was an Executive Director. In that role, I did a number of interesting things with cloud. And one relevant to this conversation as I led a team that was hired by Amazon.com Inc. in 2006 to help create AWS.

And it was an IBM team that was hired by Amazon to advise on crude and driving systems and other things that Amazon needed to know. And the side anecdotes, they also asked me to write up. I think it was 12 or 13 white papers about how Amazon does things. And they never actually used any of the processes I gave them.

So, in 2014, seeing where IBM was going and not wanting to spend the next year, essentially laying off people I had spent a decade developing, I retired from IBM. Founded my own startup, I ended up selling just over six months later to Couchbase. So, I spent a year at Couchbase, their leadership changed. We mutually agree, believe, and I went over a friend of mine at Amazon asked me to join them there.

So, from 2016 to the end of 2019, so three and a half years, I did go to market for database analytics and machine learning services. And towards the end of that period, my stock invested, I received it all. The company was changing the way that was not profitable with, and I was ready to retire. So, I retired. And that lasted about two months.

And I ended up joining Snowflake, where I did strategy and professional services. And at the end of last year, I had finished the job that I had hired in the Snowflakes to do. I was looking around Snowflake for what I wanted to do next. So, I left Snowflake and I joined Imply, which is a really exciting company in real-time data analytics, where I was employee #220.

And I think we're at about 260 now. We just got our Series D. So, we're just in time before obviously, the gates for that funding seem to close down now. So that makes us technically a unicorn with a 1.1 valuation.

Tegus Client ▶ 00:02:48

Great. So, when doing diligence on some of these large opaquer companies, I'd like to just get a sense of what is the actual ore chart and where you sat and that's why I have a sense of how it all flows. So, let's spend five or 10 minutes just on ore nothing too substantial. Just give me a sense of how is AWS actually organized in terms of Selipsky all the way down or I guess, during your time was Jassy and that kind of where you fit within that?

Former Product Manager at Amazon Web Services ▶ 00:03:11

I completely missed Selipsky. He went off just before I started, and he came back after I left.

Tegus Client ▶ 00:03:17

Right. Where did you go and he went to Tableau, right?

Former Product Manager at Amazon Web Services ▶ 00:03:20

Well, Tableau, which they got bought by Salesforce. So, he was a Salesforce executive for a while. So, the topical, you have Amazon and what's called the S committee, which stands for senior leadership with the STM. And so that's now Jassy when I was there and his directs. Plus a few people who are second level but are running very important units. So there, you'll find the heads of the major Amazon units, AWS, commerce and marketplace, fulfillment and logistics, Amazon Prime as well as key functions like finance, HR and legal.

And then each of those S team members essentially has the same thing with a similar Steam for their unit. So, you have the CEO of AWS. And then within there, you have a group of EVPs who are each running a key area. So, in AWS, the EVPs are compute, storage, database and it's interesting that became database analytics, machine learning and blockchain, but it's still referred to as DBS for database services.

Because otherwise, the acronym was unbuilding. Operations, which essentially are the people who run the data centers. Sales and marketing, engineering and then again, a finance function and HR function. And actually, I should have said, engineering at that level, storage, compute and database report to engineering, but they're important enough units that they also sit on the S team for AWS.

Tegus Client ▶ 00:04:43

So, it's almost like compute, storage and database, database analytics, machine learning watching, they are a subcomponent of engineering, but they have a presence on the S team?

Former Product Manager at Amazon Web Services ▶ 00:04:53

Yes. Well, and they're also the source of all revenues. Fairly important there. And then within each of those, so I reported to the EVP for database services. So, I was only an Executive Director. So, it was a relatively low-level person. Most of my peers on the org chart were actually much higher than I was in the organization because they were VPs who covered large sections of database or these other services.

So, there was one VP who was responsible for relational databases, another one for I know SQL databases, another one for analytics, another one for machine learning. And then there was another one who I don't know exactly what this title was. What they ended up doing mostly was working with the classified customers because they had a security clearance.

So, it's a fairly large chunk of Amazon's business, deals with defense and intelligence. Well, not just pubs, but specifically defense and intelligence, and you have to kind of isolate those people out. So, I also have a security clearance, but I don't have TSCI like you have got it. And then there was me, doing go-to-market and interfacing with marketing and other teams and my teams.

And then if you go in the other direction, the people I had reporting to me, which were either doing marketing functions or sales enablement functions or other weird miscellaneous things like data labs, which was a program I set up for customers to come to cities and do kind of intensive one-week codevelopment program. So, I ended up at the end with about 100 or so people reporting up to me.

So, in Amazon, Empire building is very much round upon. It's all about objectives. So that's where I was in the org chart. So, I was three layers away from Andy Jassy and four layers away from Jeff Bezos. I do not have a lot of interaction there. I met with Andy Jassy at least twice a month in one meeting or another.

Tegus Client ▶ 00:06:41

Okay. I feel like I hear this term a lot, but it's actually a little bit nebulous at least in my head. So, when you say you were doing go-to-market for database, analytics, machine learning, what exactly did that entail? Like what was your actual responsibility, upselling customers on sort of the higher margin software and platform sort of product solutions, or what did you actually do?

Former Product Manager at Amazon Web Services ▶ 00:07:03

Yes. So, part of it was working with AWS marketing around what marketing things we wanted to do that were related to the data services. A good part of my team was responsible for events and the data services things at events. So, Amazon puts together this huge reinvent show at the end of every November, which is a wonderful example of Amazon frugality. It's always the week after Thanksgiving because that's when Las Vegas is cheapest.

But they also do more than 100 one to two-day mini versions of this called AWS Summit in cities all over the world. So, I was helping both organize those events, staff those events, what will we do, how do we make sure that the leads are getting set into the sales organization. There's some fairly strict rules about who's allowed to speak publicly at Amazon.

So how do I make sure we have enough speakers, making sure that our product managers or technical people, whatever had gone through and had passed a fairly difficult test, you have to pass to be able to speak publicly. I spent a lot of time with alliances and working particularly with global SIs about what their data practices would look like and how did we organize that and accelerate each other's businesses for customer success.

I spent a lot of my time dealing with launching new services and how do we bring in new service into market, I do that. I spend a lot of time on growing existing services and how can we do this. And a lot of it, how do we communicate to the Amazon sales force to make sure that they are aware of and able to position and promote the database services.

I also ended up running the data analytics team, which were the ones who are pulling in all the data and giving each service their daily and weekly updates on all the statistics, which are mostly financial statistics, but also included statistics on, say, incoming and closed service calls. So, I did something different every day, basically, which is very much the Amazon way.

You create small service teams, they're responsible and accountable. And when you see something needs to be fixed, you just go fix it. So that results in a lot of duplication of effort and some redundancies, but that's considered acceptable cost for the benefit of rapid innovation and quick response.

Tegus Client ▶ 00:09:10

Right. So, let's start from the bottom. Let's have a discussion on the unit economics. I mean, my understanding is it's a super elegant business because it has a ton of sort of capital moats around it. You can't just go out and build these massive zones and stuff all around the world, and it's fairly intense initial capital outlay. And then once that's built out outside of some basic maintenance, it's a usage-based pricing structure that just kind of has some sort of payback and then it's just all returned post that payback.

Former Product Manager at Amazon Web Services ▶ 00:09:38

I wouldn't go quite that far because first of all, the machinery needs to be replaced on a regular basis. So, the vast majority of the CapEx is not the building. So those are just cheap warehouses. It's the hardware you stick in the building. So, when I was there, the hardware usually had a 42-month cycle.

Because hardware is more reliable now and the business is more mature, that's been extended out to a 60-month cycle. But that still means I'm replacing 20% of my fleet every year, even if I don't grow it at all. So, the CapEx never goes away, although the intensity reduces from the initial thing.

Tegus Client ▶ 00:10:13

This 42-month sort of replacement cycle being pushed out to 60. Like is that honest, you actually think that the hardware efficacy has improved, the durability has improved and sort of that's a real economic advantage that the company has?

Former Product Manager at Amazon Web Services ▶ 00:10:26

Yes. And a bigger reason is the rate of change has decreased. So, one of the things about how AWS operates, which was revolutionary at the time, but it's now kind of standard, is you don't fix anything that breaks. So, you install a rack of computer equipment. And when a hard drive fails or when a processor fails, you don't have somebody come in and replace it, you just abandoned it in place and keep using the rest of the rack.

And in the early days of AWS, they would replace the rack when more than 40% of it had failed. Well, by the time we roll the clock forward to about 2012, 2013, the systems have become durable and reliable enough that instead of replacing racks because of failure, usually, they were replacing racks because of obsolescence.

That Moore's Law and its equivalent in storage and memory had advanced to the point that buying new equipment was cheaper than continuing to operate the existing equipment in terms of price performance. And that was happening after about three to four years, that's the three and a half-year cycle that was there so it's a 36-month cycle from 2007 to about 2012. You went to a 42-month cycle from 2012 to last year.

And then they extended it to between 42 and 48 depending on the type of equipment for the next decade. And then last year, we're looking and saying, when we look at the math of when it makes sense to replace things now, it's not 42 or 48 months anymore, it's 60.

Tegus Client ▶ 00:11:49

And so that was an interesting comment, the rate of change comment. You're effectively saying that the amount of advancement you had on sort of rack infrastructure and the sort of server farms, it's no longer exponential, but not the same degree it was. So, you effectively don't have to replace for obsolescence as frequently.

Former Product Manager at Amazon Web Services ▶ 00:12:06

Exactly. So, I'd like to pick an easy example to put numbers on, take a look at hard drives. And about half of them is on storage purchases are hard drives and about half are solid state drives. So hard drives 10 years ago, the hard drives that were being purchased were usually holding about one to two terabytes. And by five years ago, they jumped up to 10 terabytes.

And today, they're at about 18 and going to 30. So, in terms of raw number of bites for drive, it's going up pretty fast. But percentagewise, it went from tripling every five years to maybe 50% every five years. So where, say, five years ago, I replaced the rack, I would be having three times as much capacity. Today, when I replace the rack, I only have about 50% more capacity. So, it doesn't make sense to replace the racks as quickly.

And I mean that's hard drives. But if I look at solid state drives, if I look at memory, processor is a special case because that one, the main processor provider, basically stopped innovating which required a change in processor technology. So that's why we're seeing both a lot more AMD and Graviton now.

Tegus Client ▶ 00:13:14

What happened to the processors?

Former Product Manager at Amazon Web Services ▶ 00:13:16

Well, are you familiar with Intel and the Tick-Tock world from them?

Tegus Client ▶ 00:13:19

I know Intel, the company, but I don't know if the use case you're referring to.

Former Product Manager at Amazon Web Services ▶ 00:13:23

So, Intel, for about 20 years, like a clock, every eight or nine months, would have a tick or a tock. So, a tick was moving the chip to a new smaller feature size and then a tock was putting advanced firmware and software capabilities on the chip. And every time you had a tick or tock, you had a price performance improvement of between 10% and 20%.

So, this was driven by a combination of Moore's Law and Intel's engineering. And that meant that if you were a provider, if you're Amazon or Microsoft or Google or Alibaba or whatever, all you had to do is buy the new chips every nine months or so and you got a 10% to 20% price performance increase. So, this was great, until roughly 2018, when Intel proved incapable of continuing that kind of advancement.

And partly that was due to a bad engineering choice, but mostly that was due to an Intel leadership that was more focused on stock value than engineering excellence. So, they started around 2015, 2016, doing stock buybacks instead of investing in the next-generation engineering. And I presume you're familiar with the IBG, YBG theory of executive management. That stands for I'll be gone, you'll be gone.

So that's when leaders basically make decisions that pump up the stock price look great for two or three years, they know it's not a good long-term strategy, but they're going to leave in two years and cash out. And that's exactly what happened at Intel. So, Amazon, among others, could see that this was going to happen. And so, we started increasing their investment in their own ASICs and in this case, in a CPU from an ARM design, which they called Graviton, which initially launched in 2020.

And the second generation of which is now getting out there and the third generation of which has been announced. And essentially, it let Amazon get back on the track of the sort of price performance increases. Although it's not perfect because it's not every workload that will run on Intel will run on Graviton. But roughly half will.

And Graviton right now the chip is about 50% cheaper than Intel, which makes the system about 25% cheaper since the processors are roughly half the cost of the server. So that resulted in a short-term big bump as the Intel stopped getting cheaper and then all of a sudden, there is this 25% drop. But with Graviton, they're able to get back to a 10% to 15% every year kind of a cadence.

So, ARM design. And it's ARM diverse with Amazon's engineering during modifications to make it more Amazonian. And then meanwhile, the other player here is AMD, which right now is at least for server chips, has an intel chip that's much better than Intel's chip.

So, while the vast majority of the current Amazon fleet is Intel, if you look at what's being purchased this quarter, Intel is now well under half. with a big chunk AMD and a big chunk, Graviton. And to be honest, the only reason that Amazon is buying any intel at all right now, is because supply chain constraints mean they can't get as much Graviton or AMD as they would like.

And we're seeing similar things at other hyperscalers. So, Google has done way in with AMD, Microsoft is doing both AMD and then their ARM chip, they're reselling a company called Ampere. They're using ARM chip, right now they are AMD, ARM. We see companies at more differentiated names.

In case the bottom line is when you're at Amazon and you're doing your service and you're doing your annual OP-1 review, operational plan one year, which is when every service at Amazon has to justify their budget, justify their head count. And say what would you do with more headcount, what would happen if we took your headcount away. One of the questions on that for end that everyone has to always answer.

And you also get interrogate on this at least once a quarter, is how are you going to lower your cost. Because a big part of the Amazon approach to success, not just AWS, but Amazon is that flywheel that Jeff Bezos wrote on the United Airlines cocktail napkin in the '90s, which essentially says lower prices mean happier customers, means more vendors, mean lower prices. And the only way to lower your price sustainably is lower your cost.

So, Amazon is relentless about lowering costs, and that means lightweight teams, hire less people, gives them a lot of autonomy, innovate quickly and drive the cost down and offer the lowest sustainable price you can. The original name that Jeff Bezos wanted by the way, was relentless. And right now, if you go to relentless.com, it will resolve back to Amazon. His mother convinced him that was too aggressive a name for bookstore.

But that relentless focus on cost reduction is a big part throughout Amazon. They squeeze their suppliers somewhat coolly. But another big part of the Amazon culture is customer obsession. But it's kind of the Amazon attitude that, look, here's Amazon, we're going to buy a lot from you, but we're going to always be pressuring you for lower price. And if you can figure out how to stay in business that way, we'll find another supplier.

Tegus Client ▶ 00:18:11

Right. So, it seems like so long as Moore's Law holds and you continue even at the same pace Moore's Law, but just the secular improvement in production per whatever cost equivalent unit, you should have some sort of price performance improvement. You should have actual margin tailwind in less Amazon invest that back into price, which I think they've been doing as the other hyperscaler.

Former Product Manager at Amazon Web Services ▶ 00:18:31

So back in like 2018 they did their own study. And said that AWS could raise prices 10% to 15% across the board and losing the market share. And of course, for everything we teach in business school, you should do that. That's more value for the shareholders. And the responsive time from Andy Jassy is we promised our customers the lowest sustainable price, and that's what we're going to offer.

We are more interested in the next 100 years than the next 100 days. Now immediately, when you're super profitable, you can afford to have attitudes like that, and they do. But part of what this means in the overall market dynamics is effectively, Amazon sets the pricing, and everyone else struggles to keep up. If Amazon didn't exist, we would still have cloud computing today, but it would be much smaller, and it would cost a lot more.

Tegus Client ▶ 00:19:20

Right. I mean, simplistically, I know there's infrastructure platform software, like, one, what's the composition of AWS revenues across those three? And then two, what are the unit economics? How should I think through how a dollar of top line flows through to profit, EBIT, whatever metric you're familiar with?

Former Product Manager at Amazon Web Services ▶ 00:19:38

Sure. So, compute and storage are high-volume, low-margin services. Database and other platform services are lower volume but much higher margin. So, compute and storage are the big hitters for volume.

So, if I look at a \$1 of AWS revenue, I would speculate that somewhere on the order of \$0.50 of that dollar is coming from compute services, which is mostly in Amazon terms the EC2, which is renting a machine. Some of it is some other compute services like Lambda, which is serverless processing or Kinesis, which is stream processing. That's a service for managing streams.

So, some of it is video streams, like your live streaming a football game. And some of it is data streams. When I'm moving a click stream analysis and counting every click from every customer from my website. So, Kinesis helps move that onto the databases. The most common stream service out in the world is something called Kafka, and Amazon also has a managed service for Kafka, which is part of compute.

But most of the compute is that EC2. So, \$0.50 are compute, about \$0.25 are storage and about \$0.20 are database analytics, machine learning. And then \$0.05 is everything else AWS does. So professional services, security services, other kind of smaller niche things. I'm actually familiar with the unit AWS elemental for example, that would be in that 5%.

Now when I look at earnings and margin, compute and storage has margins that are in the low 20s. The database and other services have margins that are in the 60s. So, when I look at actual earnings, somewhere between 60% and 80% of the earnings in any given quarter are coming from things like database and analytics and other platform services.

Tegus Client ▶ 00:21:23

Got it. So, when I think about the low 20s versus the 60s, where is the delta come in? Is it the contribution margin? Like COGS or are you thinking about a retailer, is the markup just higher for platform services? Or is compute and storage just a lot more OpEx intensive in terms of you need more people to manage the machines and the servers and all that stuff? Or why is there an actual margin inferential between the two?

Former Product Manager at Amazon Web Services ▶ 00:21:47

I think there's two or three reasons. So, one of them is just the whole thing of volume. There's a lot more volume in compute and storage and also a lot less services. So, each service is a much higher volume. Secondly, competition. Compute and storage services are essentially undifferentiated. There's no real difference between Amazon EC2, Google virtual instance or Azure virtual machine, or if I drive into object storage.

There's no huge difference or a significant difference between Amazon S3, Azure Blob and Google Cloud storage. There are some small differences, but they're the difference between Chevy and Toyota. One might have more cup holders than the other, but they're fundamentally similar. So, this is essentially undifferentiated volume business. So, it's very price sensitive one.

And then the other reason is the higher margin services right on top of these. So, if you look at the day-to-day services, they are built using EC2 and S3 and other compute and storage services. So, when you look at the platform services, those are much more differentiated. Each of the providers has quite different services from each other. They are also higher value.

And frankly, even at those relatively high margins, they're saving customers a tremendous amount of money versus licenses that historically, they were paying to companies like Oracle or IBM. So, it's a world where they can differentiate and have a higher margin yet still be a huge cost advantage to the customers. So, if I do the TLVR version of everything you just said, they have higher margin share because they can.

Tegus Client ▶ 00:23:21

Yes. So more granularly, kind of three factors. One is because computing storage is higher volume, there's probably going to be more fixed cost per transaction. There's more cost per unit volume transaction. So that's a drag. The second is there's just more competition in compute in storage so that's pricing compression. And the third, I didn't fully appreciate. You're seeing the higher margin set of built on top of computing storage. So, you're saying they can just extract the compute and storage margin and layer more on top of that because it's basically two products in one, or what was that third point?

Former Product Manager at Amazon Web Services ▶ 00:23:55

Yes. That it's important for them to build up a very large compute and storage base. Because there's definitely economy of scale there. Your one million server is much cheaper than your 1,000 servers, you're doing the same thing, effectively. And when I have millions of servers, that lowers my cost of the higher end services.

Tegus Client ▶ 00:24:14

Got it. So, is across the different hyperscalers, is there a meaningful difference in this rough numbers, like you basically said 75%, 25%, high level? Are Azure and GCP materially different? And if so, does that afford them a competitive advantage versus AWS or in your mind, no?

Former Product Manager at Amazon Web Services ▶ 00:24:33

So Azure, the earnings they get from their high-level services are lower, it's more like 50% of the earnings. And partly as a less of them and partly because Microsoft has a very different go-to-market model than Amazon. Where Amazon and Google are both a come, use the services you want, pay for the services as you use them, pay more pay less as you use more used less. Microsoft is all about signing three to five-year giant enterprise contracts that lock customers into levels of their spending.

Tegus Client ▶ 00:25:02

Got it. And then probably concede on some price per unit as a result?

Former Product Manager at Amazon Web Services ▶ 00:25:06

Somewhat. Plus, at Microsoft, the highest-margin things are places where Amazon doesn't compete at all, things like Office 365 or Microsoft Dynamics. Amazon doesn't really offer anything to compete with Word and Excel and PowerPoint. And for Microsoft, they're making margins above 90% there because they don't face serious competition, except possibly from Google, G Docs, who are also getting 90% margins from those sorts of services there.

Tegus Client ▶ 00:25:34

Interesting. So, your point is basically the mix for Azure, let's call it, infrastructure versus platform and up. It's similar, but Microsoft just has lower margin on the platform and up, is that right?

Former Product Manager at Amazon Web Services ▶ 00:25:46

They both have a lower margin and a lower percentage of their overall business as platform. Amazon platform is about 20%. Microsoft, it's more like 15%. Also, Microsoft is not as good at reducing costs as Amazon is. They have a higher cost basis. They have Microsoft has more salespeople than AWS has employees, and that's part of their design. Microsoft has a swarm of people sleeping over customers, talking to people, taking people to lunch, whatever.

The Amazon approach is, here's a better service and cheaper And both of those models work. As a general rule of thumb, the more regulated in industry is, the better the Microsoft model works. The Microsoft tends to be much better in industries like health care or government. While Amazon tends to dominate in lightly regulated industries, particularly SaaS and PaaS companies.

But the single largest users Amazon.com, but Salesforce is only slightly behind them. But if you look at other really big Amazon customers, a lot of them are something as a service company, might be Consumer-as-a-Service like Netflix or Airbnb or it might be Business-as-a-Service like Snowflake, like Salesforce, like Workday. Microsoft has almost no presence in that part of the market.

Tegus Client ▶ 00:27:01

Right. Based off what you're saying, my impression is that in a recession-type scenario, where people start tightening the belt. You have a lot of VC funding and some of the irrationality you've seen in the last decade that sort of drives up as people get more focused on discrete unit economics and maybe cutting some of the fat. From what you're saying that does that imply that AWS is more exposed to a recession than Microsoft? Just given that Microsoft has these long-term contracts with enterprise customers fixed in nature?

Former Product Manager at Amazon Web Services ▶ 00:27:28

No. I would argue exactly the opposite. Because what's driving cloud growth and AWS growth is right now, Amazon, Microsoft and Google are not really competing with each other. They're competing with data centers. They're competing with all that infrastructure that's sitting out there that organizations built and put into their own data center.

And the current economics are even if you just built a new data center yesterday, you will depending on what you're doing in it, you will save somewhere between 20% to 60% in per year for the next three to five years, if you burn that data center to the ground and move everything in it onto the cloud.

Plus, if you live in like New York or Chicago, you might get an insurance payout when you bring the data center to the ground. So, in a recessionary world, if anything, that increases the pressure of these global enterprises to move on to cloud more quickly and recognize those savings. That well will dry up around the end of 2026.

Tegus Client ▶ 00:28:24

In terms of lift and shift, that low-hanging fruit has gone?

Former Product Manager at Amazon Web Services ▶ 00:28:27

So, the short answer is yes, and to throw a few numbers at it. Right now, right around 30% of global IT is on the cloud. And that number is slightly inflated by a small number of digital native companies like [Amazon](#) and [Google](#) and [Netflix](#) that are 100% on the cloud. But at current rate and speed of adoption at the end of 2026, the global cloud is going to be about 85% of IT. After that, the last 15% for various budgetary regulatory, financial workload reasons will move very slowly or maybe never.

I mean we've seen this pattern before, look at mainframe. Nobody's got a new project on mainframes for 20 years plus, but there's still a \$12 billion business. It takes a long time for that to completely go away. But when the cloud hit 85% or so, all of that low-hanging fruit is gone. And for se, [Microsoft](#) to grow, they now have to peel customers away from [Amazon](#) or [google](#) or maybe [Alibaba](#). Wow, that's a lot harder and much more expensive.

So, this is another thing that's driving the economics is whatever your market share is at the end of '26, probably what your market share is going to be in 2030 and 2035 and 2040, give or take a few points. So, the number one objective for Adam Selipsky right now and the number one objective for Tom Kurian at [Google](#) and the number one objective for Satya Nadella, who technically runs all of [Microsoft](#), but Azure is his focus, is grow, grow, grow.

[Google](#), in particular, has made a decision to say we're not even going to try to be profitable. We're getting so much money from advertising, we don't care. We just want to grow our market share. So, they're spending money like a drunken sailor, trying to buy business, hiring people from other companies trying to grow out their market share before we hit that deadline in four and a half years now.

Tegus Client ▶ 00:30:17

Got it. So, with that said, how do you think about the revenue concentration of [AWS](#) across large caps versus mid-caps versus start-ups? That's a tough question to answer. And maybe another way you can answer and just x number of customers have x percentage of [AWS](#) revenue. I'm just trying to get a sense of how exposed they to individual clients.

Former Product Manager at Amazon Web Services ▶ 00:30:37

So, the short answer is that the clients they're most exposed to as individuals are these large as a service company, with one big exception. So, I'll take that in a moment. But as a service company are how they're addressing the small market and much of the medium market. There's not a lot of dry cleaners who are directly using [AWS](#).

But there's a lot of them that are using companies like, say, Intuit to do their accounting. And Intuit is running Accounting-as-a-Service on AWS, or I can point to a lot of other service companies like that. Avalara is another good example. So, they are the leader in sales tax as a service and corporate tax as a service.

So, they will tie into both cash register, POS on site and e-commerce and take care of collecting and reporting all of your taxes, which is really hard. So, you want to do that as a service. And a company like Amazon doesn't use the service like Avalara, they do it themselves, but people who have little shops on Shopify or Etsy or whatever, they are absolutely using Avalara, one of their competitors.

And they're not huge. They're about a \$2 billion company, but they've got about 20,000 customers, and they're all SMB. And there's several hundred of these kind of service companies. Others are providing as a service to a niche industry, it's like veterinary office as a service, that's never going to grow into a multibillion-dollar business, but it's a nice little niche. And they're running it on usually AWS.

Not away some of it's on Google, a very small amount on Microsoft. So, the direct sales are going either to these as-a-service companies, some of which are SMBs themselves, but many of which are enterprise or two enterprises. But with one exception, none of the non as a service company are hitting above 1% of Amazon capacity. There's a few that are quite large that run a lot of their business on Amazon like Capital One, JPMorgan, Chase or Nike.

Tegus Client ▶ 00:32:31

Sorry, you're saying nobody is more than 1% of revenue.

Former Product Manager at Amazon Web Services ▶ 00:32:34

Well, the biggest single customer is Amazon.com., which today is only about 9% of revenue or 9% of workload. But of 100 Amazon servers, probably five of them are running commerce. That number, of course, will go up during Prime Day and Christmas and down the rest of the year. And another four or so are running other Amazon services, Alexa, Twitch, Amazon Prime, Amazon Advertising, other stuff, fulfillment by Amazon.

So that is probably the customer of AWS is that's most exposed to recessionary pressures. Thus, people have been buying a lot on Amazon in the last few years, and they're mostly discretionary purchases. So that number could go down a lot. So, if I'm Amazon Commerce, which once upon a time was half of AWS, today is 5%, and really doesn't contribute much to revenue because of the internal transfer effect.

Tegus Client ▶ 00:33:26

Right. So non-AWS workloads, nobody is more than sort of 1%?

Former Product Manager at Amazon Web Services ▶ 00:33:31

No one in corporate is more than 1%. Salesforce is about 8%. Snowflake is about 2%. Some of these other SaaS companies are 1% to 2%. Adobe is about 2%. Netflix is about 1.5%, something like that. But Netflix is a good example of a company that is a discretionary purchase that is highly exposed to recessionary pressures.

But for the ones that are doing B2B services like Salesforce, again, they're much cheaper than the legacy solutions. So, recessions are actually good for them. They accelerate people and they're doing transitions that otherwise would have taken longer.

Tegus Client ▶ 00:34:09

So, like net-net, where do you shake out on this? Like on the one hand, I kind of get the argument that look in a recession, new business starts consumer spending, the traditional things that get crushed in a recession. On the other, your point is, maybe you don't think about it on a same-client basis, just think about the runway. A recession might prompt CTOs to be like, hey, I need to get on this cloud thing today. So maybe revenues accelerate. So, with those two sorts of competing dynamics, where do you shake out? What's does the growth in your mind, look like in a recession?

Former Product Manager at Amazon Web Services ▶ 00:34:42

I think a recession accelerates cloud adoption. It also puts even more pressure on pricing. But to be honest, you can't really put more pressure on pricing than Amazon does every day anyway. So given Amazon's focus on market share, I think macroeconomics and a recessionary phase will be good for all of the hyperscalers, will be very challenging for the second and third tier.

They're really going to be crushing for companies like Digital Ocean or Oracle Cloud. The ones that can't quite get up to these hyperscale levels. And they're going to be really hard for discretionary cloud companies like Facebook. But for the ones that are mostly doing business for business services, which are the three that we've mostly been focused on, to be honest, recession is good for them.

Tegus Client ▶ 00:35:25

I wanted to touch base on the point I was making earlier, right now, a lot of people have been spinning up EC2 and whatever compute storage, all these different instances on like projects for 2030. These growth areas, moonshot things that probably are the first to get cut when people get religion on profit.

So, I don't even know if you have any sense of the competition, given that most of Amazon's revenue is just variable nature. You can spin it up, spend it down pay as you go. Do you think that is a relevant mix of Amazon's revenue composition? Or do you feel that compute spending is like the last thing people will cut and it's mostly mandatory, not discretionary spend?

Former Product Manager at Amazon Web Services ▶ 00:36:00

Those R&D loan shots tend to be a trivial part of the actual total spending of an enterprise customer. So yes, they will be cut, but they're like 2%, 3%, 4% of their total compute spend. The big dollars in compute spend are the transactions to keep right on and the analytics to keep the regulators a big and also provide insight for the business. And those are things you can't cut. You can run them a little more efficiently.

And there's a roughly \$2 billion business right now in third-party cloud economics companies whose whole pitch is let us show you how to minimize your cloud spend. And almost always, they're able to pay for themselves pretty easily by finding inefficiencies and ways to do the cloud spend more effectively. But what almost always happens is when they show efficiencies in the cloud spend, IT department does want to give that budget up.

Instead, they pour that budget into cybersecurity or analytics or more cloud spend. So, I see those efficiencies happen, but I don't see them actually lowering the customers' total spending. It does increase the bang they get for the buck, but they reduce the number of bucks they spend.

Tegus Client ▶ 00:37:06

Sure. So, the last thing then is, if I told you, AWS is running at like \$70 billion of revenue today. \$72 billion is the annualized quarterly run rate. I know they do it a little bit differently, but that's kind of what I'm looking at on reported numbers. At \$72 billion that the premise there is 30% cloud option, you're saying that can effectively triple. So, is the right top line to think about in '26 \$220 billion? Or how do you think about what that business can be?

Former Product Manager at Amazon Web Services ▶ 00:37:36

Right. So, pulling numbers out of a hat. If I do my math and looking at all of this, the global infrastructure and platform cloud. So now not counting Software as a Service. Right now, globally is right around \$220 billion. And when we hit that mature cloud at the end of 2016, it will be about 900. And I expect Amazon to maintain or increase its market share, which means Amazon's revenue in there should be \$310 billion, give or take, about \$30 billion.

And I'm kind of leaning towards the high side of that because as we get towards the end of the low-hanging fruit of the data centers, the almost as low-hanging fruit is the second and third-tier cloud provider, which currently among them have about 20% of the market. If you look at everything that's not the big four, it leaves about 20% of that market. And I think they're either going to get a quiet or crushed out between 2024 and 2027.

Tegus Client ▶ 00:38:31

So, you are bullish hyperscalers, is the short answer? I hope that it's a \$310 billion top line business in 2016, that would be a good outcome.

Former Product Manager at Amazon Web Services ▶ 00:38:40

I am absolutely bullish on the ability to grow the business and do that much revenue.

Tegus Client ▶ 00:38:44

You are less constructive on profitability.

Former Product Manager at Amazon Web Services ▶ 00:38:47

Exactly. Because even increasing level of competition. And so, there's three major black swans that could throw all the stuff off. So, the most likely of the three is government regulation, especially in Europe, that forced some breakup, some divestments. And what we've seen with government regulation and technology industries, they don't really end up changing the overall dynamics, but they slow momentum completely. So, we saw that happen to IBM. We saw that happen to Microsoft.

We may see that happen as both in Europe and the U.S. The antitrust pendulum is swinging towards a stricter level of regulation. Or to be a little cynical, if there's one thing that Democrats and Republicans agree on right now is they both hate big tech companies for different reasons. But there's the old joke that in the U.S., we have an evil part to get a stupid party and sometimes they find something that's evil and stupid is real. And tech regulation might be there.

It's less likely Black Swan, but we're hugely impactful is the microprocessors that are powering almost all of this. Every AMD chip, or a high-end AMD ship, data center AMD chip, every data center AMD chip, every data center Graviton chip, every Apple MT chip are all coming out of the same place, which is TSMC, Plant 14 and Kaohsiung in Southwestern Taiwan.

And two to three years from now, TSMC will have successfully built new plants in U.S. and Europe that can provide that sort of technology. But right now, it's all single sourced. So partly, that's a supply chain constraint. If I go to TSMC and say, I want \$10 million of this new chip, they'll say, "Great, you can have it 1.5 years because it's going to take that long to get in mine." But of course, the bigger problem is Taiwan is in an area that has both geopolitical dangerous and natural dangerous.

We hear a lot about the geopolitical stuff, but people forget that, that scenario that has both earthquakes and volcanos. And we have placed in the last 20 years, seen huge supply chain disruptions, once from a volcanic activity near Kobe, Japan, which essentially shut down the global supply of DRAM for about 1.5 years. And the other one was massive flooding in Thailand that shutdown on supply of hard drives, which are mostly made in the Bangkok area and still are.

So, something happening in Taiwan could be a major supply chain disruption that would stop cloud growth immediately because they wouldn't have components. And then the third black swan, which is a more distant one, if quantum computing is able to successfully break cryptography, which I would say is on the odds of less than 2% likely in the next five years, but it's not zero, then that undermines everything on the Internet.

Because it completely eliminates the ability to send a secure message from one machine to another. And if that happens, then since we have to rethink the whole Internet and the cloud is completely unsecured. It no longer becomes a place that you can conduct business, but that's a whole point of Black Swan. They're low odds, high-impact event. I suppose I can also talk about black swans like Yellowstone erupting or media or hitting the earth, but those are bigger than IT.

Tegus Client ▶ 00:41:52

Yes. I think if Quantum breaking cryptography less than 2%, I'd say extension levels probably less than 0.1%.

Former Product Manager at Amazon Web Services ▶ 00:42:00

Probably less than that. Yes. But I mean, at a more practical level. But Amazon's second biggest region is Eastern Oregon, which has significant seismic risk. So why did they put it there? Because it has really cheap hydroelectric power. But Amazon has situated their data centers a couple of hundred miles apart and have enough capacity that if City Mount Hood erupts, it will take out part of that capacity, but there's enough there to keep running.

Tegus Client ▶ 00:42:25

When the stuff happened in Kobe, Japan, and the supply of DRAM was down for 1.5 years, I don't actually know what does DRAM feed into, is that relevant for cloud computing? And if so what happened with cloud computing?

Former Product Manager at Amazon Web Services ▶ 00:42:38

Well, this happened in 2003. So, cloud computing didn't exist. And by the way, DRAM is memory. It's computer memory, stands for dynamic random-access memory. And what got destroyed in Kobe was the one factory in the world that may be special nonconductive glue you needed to glue the memory chips on to the memory sticks. So, for about 1.5 years, nobody could manufacture computer memory.

And it's amazing how many little things like that, and nobody ever thinks about our single points of failure and the global supply chain. Because it's been demonstrated very well in the last year. We built a global supply chain to obsessively optimize for cost reduction with no thought to resiliency. And it makes things really cheap when it works well, but any little thing goes wrong and the whole process grinds to all.

Tegus Client ▶ 00:43:26

So how far are we away from TSMC sort of dual sourcing and try sourcing the chip production, micro-plastic production, like in the U.S., Europe, et cetera.

Former Product Manager at Amazon Web Services ▶ 00:43:35

Late '23, early '24. So, they're building a plant in Ireland, and they're building a plant in, I want to say, Ohio. I might be wrong about that, but somewhere in the U.S., might have been Arizona. But they're building a plant in the U.S. and a plant in Ireland that will also be able to do these most advanced seven-nanometer and sub seven-nanometer chip technologies. But it's so much that TSMC this year is actually spending more CapEx than Amazon. This will be the first year in the last five that Amazon is not the biggest CapEx spender in the world, outside of government.

Tegus Client ▶ 00:44:10

So, you're obviously very bullish on top line and maybe a little bit north of what I'm estimating, but I'll take either outcome as a good one, is you're saying that you think profit pressure is something that needs to be appreciated. And what is driving that? It's purely competition and within competition, it's Azure, GCP, right?

Like, to your point, the residual 20% that are not the hyperscalers, they're either going to be acquired or crushed out, right? So, what's the concern with profitability? They don't care about near-term economics because, to your point, 26 shares, 20, 45 shares, is it just going to be uneconomic probably get as much as they can.

Former Product Manager at Amazon Web Services ▶ 00:44:47

Right? So, part of it is possibility. And part of it is ability to execute at scale. It's hard to run a \$100 billion business. It's even harder to run a \$300 billion business. Is Amazon able to recruit and retain talent with the correct mindset and operations to keep what has worked for them working at this scale? I'm already seeing that kind of start to break down.

The rate of innovation in Amazon is much lower now than it was three or four years ago. I think part of that is just the sheer scale that Amazon has grown into. I think part of that, frankly, is just the increased maturity. I mean airplanes got faster every year until they hit a certain level and they stop getting faster. You only have a certain amount of that effect.

But a bigger issue, and it's not the main reason how left Amazon was one of them is you're a company in Seattle, you're growing really fast, you're hiring people, who do you end up hiring? You're hiring all the ones who got fired for Microsoft. And they're bringing some of the toxic fee bummer attitudes into Amazon, and I'm already seeing it slowing down and hurting the company. So, for example, I don't know if you ever heard a stack ranking at Microsoft.

This was a policy of the steamer era, which every manager had to rank everyone in their team. Top 10% got a raise and the bottom 10% got fired. Everybody hated this. And among other things, and when you were a manager, you would deliberately go hire some losers, so you had some of the fire, so you didn't have to fire anyone you feared about them. So perverse incentives.

Well, when Nadella came in, he said, this is stupid. I want everyone we hired to be a rock star. And if they're not rock stars, let's get them to leave. But I'm not going to say you arbitrarily kill the bottom 10%. And all those people got fired seem to end up at Amazon, which now says we don't do stack ranking, instead they have something called unregretted attrition, which basically says go hire the bottom 10%.

So, the things that hurt Microsoft are now hurting Amazon. So, it's a small piece. I don't think it's enough to kill them. But it's enough to make me when the mindset, the customer obsession, the focus on innovation, the obsession with lowering your prices that basis built. How many more generations of people are going to be able to execute on that? But it kind of goes against everything you get taught at other companies. It's 180 degrees what I was taught as a manager at IBM.

And I think Amazon is good with that for three years and maybe five years. 10 years from now, I think they're going to look like any other large company to a large degree. And so, you need this kind of mindset to be able to operate in this kind of hyper aggressive hyper-innovative way out of profit. We'll see.

Tegus Client ▶ 00:47:30

Great. Well, thanks again for your time and help today. Take care.

The information contained herein is for informational purposes only, is not to be construed as investment advice, and should not be used as the sole basis for making any investment decision. Tegus, Inc.

("Tegus") makes no representations and accepts no liability for the content herein or for any errors, omissions, or inaccuracies, and will in no way be held liable for any potential or actual violations of United States laws, including without limitation any securities laws, based on the information herein. The views expressed in this document are those of the expert or the interviewer, and they are not endorsed by, nor do they represent the opinion of, Tegus. Tegus reserves all copyright, intellectual, and other property rights in this document and other Site Content, which are protected by the Copyright Laws of the United States and may not be copied, reproduced, sold, published, modified, or exploited in any way without the express written consent of Tegus.