

DR: [00:00] I'm DR, and this is "Telco In 20."

[00:14] Being a contender in the public cloud game requires super deep pockets and a crap ton of money. For example, in 2023, the top three hyperscalers, AWS, Microsoft Azure and Google Cloud, spent nearly \$130 billion on CapEx with \$80 billion of that going to cloud infrastructure. And that's just last year.

[00:34] Since the year 2000, the big three have collectively spent over \$820 billion to build more than 144 regions around the world. By 2025, their CapEx investments are expected to top more than a trillion dollars with \$250 billion of that spend occurring in the last two years, and they're not stopping. With the rise of generative AI, that number is just going to keep going up. As a Telco exec and a customer of the public cloud, you need to keep tabs on how your hyperscaler is investing in its future. Is it keeping up with the leaders, falling back, slowing down? Is it investing in chips, software, or LLMs?

[01:10] This is because as you move workloads to your provider's data centers, you're betting not only on what it has to offer today but also its trajectory for future products and services. Its roadmap is your roadmap too.

[01:22] Thankfully, we keep up with Charles Fitzgerald from Platformonomics, who stays on top of hyperscaler announcements and tracks all this information for us, and has followed the CapEx report. Today, I'm talking to Charles about how much the hyperscalers spent in 2023, how generative AI is influencing hyperscaler investments, and whether repatriation from the public cloud to on-prem is really a thing. I'll give you a hint, it's not.

[01:45] So let's take 20. Charles Fitzgerald is managing director at Platformonomics. Hi Charles. Welcome to Telco in 20.

Charles: [01:57] Hi, it's great to be back.

DR: [01:59] I know. I'm so excited to have you back on the podcast. And so your business, Platformonomics, follows hyperscaler CapEx spending. You recently published your 2023 retrospective, and so what changes have you seen in hyperscaler CapEx spending over the last 12 months?

Charles:

[02:15] Well, the overall combined spending amongst Amazon, Google, and Microsoft was actually flat from the year previous. But we saw Amazon, Google, and Microsoft collectively at the corporate level, so not just cloud infrastructure, those three companies spent over \$127 billion on CapEx, and my estimate is about \$80 billion of that went to cloud infrastructure.

[02:38] But as we look at each of those three companies, they're kind of going in different directions. Microsoft's CapEx spending was way up as they ride the generative AI wave. ChatGPT runs on Microsoft Azure, so they've been in just a mad scramble to find enough GPUs to keep up with demand. So their CapEx spending was actually up 45% year to year to over \$41 billion. And that's not just using GPUs to train models. They're also serving them at scale. So every time you talk to ChatGPT, there are some GPUs that have to go to work. So they have a lot of GPUs that are also doing inference.

[03:18] It really looks like Microsoft has been NVIDIA's largest customer for two or three years at least. My guess is their incremental spend on CapEx for AI last year was on the order of \$12 billion and it's not just their own set of GPUs. Microsoft is also a very large customer of some of the boutique GPU hosters that are popping up out there. So Microsoft is definitely up and to the right.

[03:47] Google was basically flat and maybe there's some signs of an AI-driven CapEx inflection starting in the fourth quarter. They've been kind of disappointing from a CapEx perspective over the last couple of years because they used to do a big refresh cycle every four years and we saw huge spikes in the amount of spending. They've kind of stopped doing that. So their spend in 2023 was up 2%. They spent about \$32 billion, but if you look at their CapEx on a quarterly basis, it did increase quarter to quarter throughout the year.

[04:22] So the trend is kind of headed in the right direction. They are kind of belatedly mobilizing for generative AI and it's definitely going to need some infrastructure spending. So if Gemini takes off, I would expect a lot more spending from them.

[04:36] But the really interesting one is Amazon. So Amazon overall at the corporate level was down 20%, and the spending on AWS infrastructure actually declined by 10%. And that's the first time AWS CapEx has ever declined. Amazon still at the

corporate level spends more than any other non-Chinese company. They spend over \$53 billion, so it's a mix of the AWS cloud infrastructure and also all of the retail fulfillment warehouse spending. And they've been reining that retail spending in after just going nuts during the pandemic. They were well over \$70 billion for annual spend and they're having to rationalize that and that makes sense.

[05:23] But it's much harder to explain the first-ever decline and spending for AWS infrastructure to see them cut the investment as the business is growing is kind of nuts. But cutting the investment during this AI mania is doubly weird, given how much they've been telling us that they're a leader in AI, at least in their own minds. So their actions really don't align with their words.

DR: [05:47] So we know that AI is revolutionizing pretty much every industry, including the telco industry. For example, telco's biggest event of the year, MWC in Barcelona, was pretty much all about GenAI. How are you seeing GenAI impacting the cloud infrastructure business? You mentioned a little bit about Microsoft really spending a lot on that, but just in general, what are you seeing with GenAI impacting the hyperscalers?

Charles: [06:09] The whole technology industry today is all about generative AI. That is the big race. That's where all the investment in IQ is going to. I mean, you definitely see it in Microsoft's numbers, but Microsoft said 6% of Azure growth was due to generative AI, so it's definitely having a material impact on their revenues.

[06:29] Google's kind of slowly waking up. They have a lot of assets they can bring to bear. They mostly just kind of need to get out of their own way, and you could argue they're making it harder than it needs to be.

[06:39] Amazon is the interesting one. Because they've, so far, mostly just run their mouths about how great they are, but they're behind on most dimensions of AI prowess and the spending going negative, deploying all of these very expensive GPUs. But data centers are kind of getting rebuilt because you need different networks inside the data center. You need denser power distribution. So AI really is the focal point across the tech industry. And I assume the telcos will be customers of these big

AI engines as opposed to trying to compete head-to-head with them.

DR: [07:15] Yeah, I don't really see them competing per se, head-to-head with the clouds and building data centers that are specialized for generative AI. But I do see them going out there and building an LLM for the industry. I don't know that every telco needs to build their own LLM. That's a multi-billion dollar investment, but again, the telcos always like to go out and try to build it themselves rather than leveraging what other people are doing and building on top of that.

Charles: [07:42] I think they're going to struggle to keep up. My guess is that the cutting-edge foundation models are probably better at telco tasks than any LLM that the telcos can build themselves. It's too much investment. They don't have the talent they need, and at this point, they're behind.

[07:58] Even looking at a company like Amazon that is two years behind, they're not catching up. If anything, they're probably falling further and further behind. Pick your random industry. The idea that they're going to get anywhere close to where the frontier companies is, I think, is a reach.

DR: [08:14] Yeah, I think it's a better bet. Even if you bet on, like you said AWS being behind, it's still leaps and bounds ahead of where you would be and the investment is massive. So just pick the best player, the one that you think is suited for you, and ride their coattails.

Charles: [08:28] Yeah, I mean I've lived through listening to telcos talk about how they were going to win the internet because it involved networks. And mobile was going to win because it involved networks. And cloud, they were going to win because it involved networks. None of those really panned out.

DR: [08:42] Well, let's talk about a new thing that they're trying. So the telcos are super interested in getting developers to build application for their network APIs. It's a new way to monetize this investment. They're now technically leveraging that asset. For instance, the whiz kids over at McKinsey, those consultants that charge a ton of money but are super smart, recently put out a report that estimates this market of network APIs and developers using network APIs to be \$100-\$300 billion over the

next five to seven years. And so again, how do you handicap their prospects and being able to convert on this opportunity?

Charles:

[09:14] Well one, I would always do the opposite of what McKinsey counsels. But two, I've been listening to the telcos talk about building a developer community around their APIs for decades, longer than I'd care to admit. I can remember Bell Atlantic pitching me on their AIN platform in the 20th century, and I may have coined the phrase, you can't spell brain dead without AIN after that conversation.

[09:39] So the telco API dream is evergreen. It's also much easier said than done, as we've seen by the lack of material success in this space. It's just really hard to build out a developer ecosystem.

DR:

[09:54] Totally.

Charles:

[09:55] You need functionality that you're exposing the API that's compelling, and there's a business model around it where developers can actually make money. And you've got to go and engage with the developer community and they're a tough audience to engage with. They're notoriously skeptical of anything that smacks of marketing. You need great development tools. You need documentation. You need samples. It's very, very hard to build a developer community as a hobby, which seems to be the way most telcos have approached it.

[10:30] And if you look at the successful API-based businesses, the API really tends to be at the core of the business. So I think it's going to be very hard for people who are doing it as one of many initiatives, but who knows? Maybe someone will surprise us.

DR:

[10:47] Yeah, I mean, I keep talking to telcos about this. You guys famously love to go chase streaming, so you buy a media company. Or you're going to go chase cloud so you go buy an infrastructure company, but here you can't go buy developers.

[10:59] And the other thing I see them doing is, "Great. Well we'll just put our APIs and make them available on Azure, on AWS, and those guys have a developer community and we'll just be able to piggyback." I think this is also a mistake because I think the strategic asset here is, as you said, the developer community.

[11:14] And so I don't know how they're going to do it. They need to come up with new ideas to monetize their network. They want to piggyback on other players and be in the driver's seat and I don't know how you do that when the strategic piece is the developers and you're not building it.

Charles: [11:27] You got to be all in if you're going to go after the developer community.

DR: [11:31] So speaking about being all in, I'm certainly all in on the public cloud. And one thing that we saw at the end of 2022 and certainly a lot of discussion in 2023 was this whole idea of cloud repatriation, that a public cloud is too expensive or it's just fancy server rental. And so I think you started a new index where you're tracking, is cloud repatriation a thing? And so, are people really moving their workloads from the public cloud back to their own data centers or other sort of on-the-ground, non-public cloud data centers?

Charles: [12:03] The short answer is no. I mean, this was a 2021 peak zero interest rates idea that came out of the VC world. And there's a lot of angst in VC land about the difficulty startups face in competing with the huge economies of both scale and scope of the big cloud vendors. So the idea was that "Hey, companies should pull their workloads out of the cloud and move them back to their own data centers, and they could save some money, and making the capital investment in that kind of stuff was palatable in a world of zero interest rates."

[12:38] And it really was kind of motivated by the idea that wouldn't it be great for VCs if customers would assemble their own infrastructure from random piece parts from a bunch of different companies, i.e., those that were supported by those venture capitalists.

[12:55] And in terms of what's happening in the world, there are a few anecdotes out there, but we really haven't seen much repatriation. And if you start to look at the numbers, you certainly don't see it happening. And that's kind of where this index that I put together tries to look at the industry broadly because we're seeing the clouds continue to grow. The data center hosters are not growing as fast and those big data center companies should be the prime beneficiaries if repatriation was happening, but they keep cutting their guidance on their revenue growth.

[13:28] So I have this super simple metric where I just look at the growth. Digital Realty and Equinix are two of the biggest data center operators and use them as kind of the data center repatriation standard bearers and compare them to AWS's growth. And it's super simple.

[13:46] AWS is the slowest growing of the big clouds, but it still shows that the cloud is dominating in this space. So it's an interesting topic. I mean, I hear a lot less about cloud repatriation after interest rates went up because the idea of going out and buying a bunch of your own infrastructure so you can operate it has become less exciting, but it's more than just the taking the CapEx back, right?

[14:12] If you're a company who's in a steady state and really all you can do is cost optimization, it might make sense to try to take some of the cost out of the business by owning and operating and optimizing that server infrastructure. But they're very real costs and opportunity costs to spending your time doing low-level infrastructure, right? The time you spend there is not time you're doing higher-value software up the stack. You need a very specific type of talent that most enterprises don't have. You not only need the infrastructure experts, but you're also on the hook for a bunch of other things like, "Congratulations, you're now responsible for your own security." And you know, it's tough for the average enterprise to find talent.

DR:

[14:57] I think that's true for small companies. Let's say my startup to Totogi, not focusing my team on managing the hardware and just focusing on what we're trying to do best, which is the software piece. But it's also true for super big companies like telcos. The opportunity cost of having people and expense on this infrastructure layer that isn't a differentiator in your business. You should absolutely focus on making your public cloud spend efficient and not be dumb about that.

[15:26] And I think we're going to enter a phase where the telcos are going to be like, wow, this is really expensive. And I'm like, well, don't move back to on-prem, and don't build your stuff in a very cloud-agnostic way so that you can move back. But instead, how do you make it more efficient where you go more cloud native and you cost optimize that? And the pricing's

always changing and I think that's where you should spend your energy and effort.

Charles: [15:49] And the values what you do at the top of the stack, not how you cost optimize the bottom of the stack. So pick your priorities. And I think this becomes even more extreme as we move into this AI era. I think the advantages of scale and scope that the big hyperclouds get even greater. It's the economies of scale, the investments, the skill sets, all those things are even harder to do than traditional application hosting.

[16:17] And when you start to look at things like the shortage of GPUs, guess what? The big hyperclouds are going to be much higher on the customer-listed NVIDIA than the average enterprise. And they do things like they own trans-oceanic cables. They're building their own silicon. Enterprise IT is not likely to go down that path.

DR: [16:38] Absolutely.

Charles: [16:38] It's going to be tough to match that. So I think we'll continue to see IT workloads net flow into the cloud for years to come.

DR: [16:48] Yeah, I think the case for the public cloud just keeps getting stronger and stronger, and I do believe it is these three big hyperscalers that just continue to grow and expand and spend on CapEx.

[16:58] And so, rounding out one of the favorite things that you add to your report is your update on what you call, the clown providers. And more specifically IBM and Oracle who like to pretend and like to put out commercials and marketing messages that they're in these big leagues. And so what are some quick stats and thoughts that you have that prove that they really should not be mentioned in the same breath as the big three?

Charles: [17:21] Well, CapEx has been a great way to separate the clouds from those clowns, over the last decade. Real clouds spent real money on CapEx. The clowns talked a good game but didn't make the accompanying investment. So most of the clowns have disappeared over the years.

[17:37] I remember Verizon Cloud...

DR: Terremark...

Charles: ... there were a bunch of others. We're kinda down to a clown car race between just IBM and Oracle, but both IBM and Oracle just fall further and further behind with every passing year.

[17:53] Microsoft spent more on CapEx last year than Oracle has in its entire history. Google spent more last year than IBM has in the last decade. And there was a time, beginning of the 21st century, when IBM was by far the biggest CapEx spender in tech, and now they just trend down and to the right with every passing quarter. The number just somehow gets smaller and smaller, but they're still out there talking up cloud and using the marketing buzzwords.

DR: [18:20] Yeah, the thing with hyperscalers is it really is about putting your money where your mouth is. The CapEx is really important. You got to invest. Like you mentioned at the beginning, AWS is talking a big game on GenAI and you're going to be able to see it in the spend.

Charles: [18:33] Yeah. Going forward, AI is, if anything, even more capital-intensive than cloud. So CapEx spending is going to be a great way to separate the AI clowns from the people who are actually doing it. So we'll keep tracking about who's putting money where their mouths are.

[18:50] But if you look at the combination—Amazon, Google, and Microsoft—not 2024 but in 2025, their collective CapEx investment in the 21st century will be over a trillion dollars. So this is not a game for the faint of wallet.

DR: [19:08] That's for sure. Oh my God. Well, Charles, it was so great catching up with you. You laid down tons of knowledge. I love the facts and it really gives something meaty for telco execs that are listening to the podcast to really compare and contrast who's really doing it. And this was a great conversation.

Charles: [19:26] I enjoyed it. Thanks for having me back.

DR: [19:27] Awesome. Thank you. Stick around because we're ending each podcast with a Telco in 20 takeaway. I have 20 seconds to tell you something you need to know.

[19:40] When it comes to choosing a platform for your applications, one of your first decisions will be whether to go with a private cloud or a public cloud. A lot of telcos make this choice based on the regions available to them today. But you need to realize that moving to the public cloud is a decade-long decision. You need to project out to the future and imagine where the hyperscalers will be in 10 years.

[20:00] The hyperscalers have made it clear that opening regions is the best way to grow their revenue. So you can bet they'll be opening a region in your country or in a place that's available for you to use. This applies to faraway continents, small countries, and islands alike. Africa, Belize, Fiji, I'm talking to you. So don't choose your platform based on the regions available today. Instead, decide what's best for your telco's future, which if you ask me is the public cloud, and start building towards that vision.

[20:28] You need to invest in tools, people, and training. It will take a while, and by the time you get there, I predict the hyperscalers will be ready and waiting for you. So go all in and start building now for the public cloud.

[20:40] Want a handy guide to all the hyperscaler regions across the globe? You'll find a link to our hyperscaler region tracker in the show notes. It includes the latest on the big three's plans and progress from announced regions to those that are fully operational, and we make this information available to anyone who wants to see it for free. In a faraway continent like Africa, I'm visiting you this month, April 23rd to the 25th, I'll be at the NOVACOM Summit in Diani Beach, Kenya. I won't have time to go on safari to see the big five, but I will be hunting telco execs while I'm there. Hopefully, they'll be easy to track so I can pounce on them just like Tigger. If you'll be there, let me know.

[21:18] Until then, tune into more Telco in 20 episodes, like and follow, and leave us a fantastic five-star review. Connect with me on LinkedIn and X at TelcoDR. Don't forget to sign up for my incredibly awesome, totally free email newsletter on TelcoDR.com. And be sure to check out our awesome sauce YouTube channel and smash that subscribe button. Later, nerds.