

Episode 137 | Hyperscalers' \$2T AI Bet (Charles Fitzgerald)
Released March 31, 2026

DR: [00:00] The three big US telcos combined spent about \$50 billion in CAPEX last year. In previous year, you could see telcos in the rear-view mirror a little bit. Now they're not even in the conversation. So what happened?

Charles Fitzgerald: [00:13] At this point, the hyper CAPEX companies have completely surpassed all of those other capital intensive industries. They are the four biggest spenders by a long shot. And if you even look at Meta, who are the smallest spender of those four companies, they spent over \$70 billion last year. So Meta alone outspent all three US telcos combined with \$20 billion to spare.

DR: [00:40] Wow.

Announcer: [00:46] This is "Telco in 20," a podcast that helps telco execs achieve a competitive advantage with AI and the public cloud. It is hosted by DR, also known as DR. Today, we're talking to Charles Fitzgerald, managing director of Platformonomics.

DR: [01:05] Hi guys, I'm DR. The AI race is on and telcos want in. This year, the weigh in is sovereign cloud and AI-RAN. For sovereign cloud, Telefónica is leading EURO-3C, a consortium of 87 members putting in 75 million euros, building out edge and cloud nodes across more than a dozen European countries. NVIDIA is pushing AI-RAN, teaming up with Nokia and T-Mobile on live field trials and pitching millions of cell sites as a distributed AI factory. Bold bets. But can telcos pull it off? Here's what they're up against. Amazon, Google, Microsoft, and Meta spent over \$400 billion on infrastructure last year, up 66%. And they're guiding to over \$600 billion this year, on pace to blow past \$2 trillion cumulative before 2026 is over. US telco CAPEX, \$50 billion, and falling. That's not a race. That's a different sport.

[02:09] My buddy, Charles Fitzgerald, managing director at Platformonomics has been tracking this for over a decade in his annual Follow the CAPEX report. He's our first ever four-time guest and the only person I know who answers every hype cycle with one question, "Show me the CAPEX." Today, we dig into what his numbers say about sovereign cloud, whether AI-RAN is a real business case or a better narrative than a strategy, and whether Oracle's playing the most expensive game of catch-up in tech history. So let's take 20. Charles Fitzgerald is managing

Episode 137 | Hyperscalers' \$2T AI Bet (Charles Fitzgerald)
Released March 31, 2026

director at Platformonomics. Hi, Charles. Welcome back to "Telco in 20."

Charles Fitzgerald: [02:50] Thanks for having me.

DR: [02:52] Well, you're our first ever four-time guest, so that is amazing and awesome. And we're going to talk today about your Follow the CAPEX report that we're always following and reading. And so this is your 10th annual Follow the CAPEX report. You started your tracking before AI became the central force driving these numbers. Now four companies, Amazon, Google, Microsoft, and Meta are spending over \$400 billion in a single year, up 66% from the previous year. So help me understand the shape of this curve with AI. For every dollar they were spending before AI became the dominant thesis, what are they spending now?

Charles Fitzgerald: [03:33] Well, if you go back 10 years ago, those four companies, who I like to call the hyperCAPEX companies, collectively spent \$28 billion on CAPEX, and that seemed like a really big number at that time.

DR: [03:45] Yeah.

Charles Fitzgerald: [03:46] So if you fast-forward to today, their spend has grown over 15X over the last decade. And the AI inflection point was the introduction of ChatGPT in November 2022. So their growth rates of CAPEX spend doubled after that, and their absolute spend has almost tripled. So over the last two years, those four companies have grown their CAPEX over 60%, and they're slated to do 60% again this year. So expected CAPEX is around \$650 billion. So cloud was big, but AI is even bigger.

DR: [04:24] And are they just spending to win? Is it double till you lose kind of thing going on with these guys?

Charles Fitzgerald: [04:30] Well, nobody really knows, but you got to spend to stay in the game.

DR: [04:34] Yeah.

Charles Fitzgerald: [04:34] We'll find out what winning looks like and what losing looks like.

DR: [04:38] Yeah. I think the same thing is going on with the frontier model companies like OpenAI and Anthropic and Gemini, which is also Google. They seem to be in that same sort of hyper spend, just go crazy. And so let's bring this back to telco. So the three big US telcos combined spent about \$50 billion in CAPEX last year, and that number was down 5%. In previous year, you could see telcos in the rear-view mirror a little bit. Now they're not even in the conversation. So what happened?

Charles Fitzgerald: [05:08] Well, over the last decade, I've been watching my four favorite hyperCAPEX companies work their way up the list of the biggest CAPEX spenders in the world. And the telcos were on that list, the energy producers, semiconductor manufacturers, utilities. And I just thought it was kind of cool that companies I quaintly think of as software companies were on that list at all. But at this point, they have completely surpassed all of those other capital intensive industries. They are the four biggest spenders by a long shot. And if you even look at Meta, who are the smallest spender of those four companies, they spent over \$70 billion last year. So Meta alone outspent all three US telcos combined with \$20 billion to spare.

DR: [05:55] Wow. Well, every few years, the telcos try to invent a new way to monetize their infrastructure CAPEX in the form of their networks and their edge spend. The 2026 version is something they're calling sovereign cloud, using their data centers and edge sites to keep AI workloads in country and under local control. So they're really leaning into this data privacy and sovereign data idea. So from your Follow the CAPEX vantage point, can the telcos have an impact against these hyperCAPEX spenders? Or do you think the winner's already been determined?

Charles Fitzgerald: [06:31] Well, every year, we talk about the long history of big telco dreams that didn't quite pan out. And it literally has been decades that the telcos have always had some story for why they're going to own some layer of value above the network and be more than a dumb pipe. So they were going to own the internet because, hey, it was a network. They were going to own mobile computing because they thought they could control the handset. Then they were going to own the cloud because, hey, it was also network-based. And for the last couple years, they've been talking vaguely about doing something on the edge. And we won't even bring up the consumer misadventures in content with things like Warner Brothers and Yahoo. So whatever the

new telco dream is, I think the first question is, why is this time different than all the others?

DR: [07:17] Yeah.

Charles Fitzgerald: [07:18] Sovereign cloud is one new dream, and it's based on the hope that local government regulations will advantage local players, both telcos, but not just telcos, and advantage them against the hyperscalers. But the word cloud in the name is a clue to the degree of difficulty here. You're playing catch up with a multi-hundred billion dollar industry and you're going up against some of the biggest and most sophisticated technology companies in the world. The telcos just aren't tech companies. They don't create technology. They buy technology and then they operate it. You have to make it. So you've got to deploy lots of capital for infrastructure, and then you have to create an incredibly sophisticated layer of software on top of it.

[08:05] And when you look at the AI side of it, if you want to build sovereign AI models, you need vast numbers of GPUs to do the training. And if you want to serve those models, so doing the inference and essentially running the models, you need a set of cloud services like compute and storage and database that sit alongside the GPUs that do the inference, because you can't run a complete application without those other services. So now, you're in the cloud business. And when I do a Google search for sovereign cloud, the results are all hyperscalers. There's no sign of the telcos. So I think they've probably already missed that boat. And it's one thing for governments to mandate things, but they can't mandate things that'll actually work and much less mandate things that are competitive with the hyperscalers. So color me skeptical.

DR: [08:57] Yeah. I think you're totally right, especially on your software point. I think people tend to forget that it's not just the infrastructure that they're managing, but the hyperscalers create really valuable software that helps you create, manage these workloads. And I don't know how the telcos are going to do that. They don't have software capability. And then the other thing is, the hardware refresh cycles on AI, it's so different than what we were doing on cloud round one, where you were thinking like three to five years longevity and timeframe on these machines. Now, we're doing an 18 to 24 month cycle. It's a treadmill at a much higher speed. And you couldn't keep up on the first one. So I think it's going to be tough.

Episode 137 | Hyperscalers' \$2T AI Bet (Charles Fitzgerald)
Released March 31, 2026

Charles Fitzgerald: [09:42] The software piece is really key. Those cloud companies have tens of thousands of software engineers, and there's no place to go where you could just buy those capabilities off the shelf.

DR: [09:52] Yep, absolutely. So the other idea coming out of MWC in Barcelona was AI-RAN. This notion that you run radio and AI workloads on the same GPU accelerated fabric at the edge, and then you're selling this leftover capacity that you have. So from a first principles, follow the CAPEX point of view, is there actually a solid argument for AI-RAN, or is this just telcos and vendors reverse engineering a narrative around hardware they've already bought and they don't have any other ideas on how to monetize it?

Charles Fitzgerald: [10:22] Well, this may surprise you, but you can color me skeptical on this one too.

DR: [10:26] Yeah.

Charles Fitzgerald: [10:26] AI-RAN is really an NVIDIA initiative to sell a bunch of stuff to the telcos. And kudos to them. NVIDIA wants more customers, and they have a secondary goal of keeping AWS, Azure, and Google from dominating the market. And they've spent a lot of time and energy standing up and supporting these specialty GPU neoclouds over the last couple of years. And it sounds like they'd like to get the telcos into that same business and of course get paid along the way.

[10:56] So it's another case of, okay, look at who you're competing with. You're now competing with the hyperclouds for AI services. And the premise for the telcos is they're going to do this as a hobby. So in addition to their day job, they're going to sell some leftovers and be competitive against the biggest CAPEX spenders in the world and the most sophisticated technology companies. And I think if you're a telco, you have to ask yourself, is helping NVIDIA reduce hyperscaler market power really the biggest and best ideas telcos have for their future? And it comes down to you either own your own strategy and destiny or you get to become an actor in somebody else's strategy. And that's what AI-RAN looks like to me

DR: [11:41] I'm with you. I like your comment. In addition to their day job, they're going to do this side project. And I also like the point you made there on NVIDIA is finding new customers to

buy their machines, and then these guys are going to try to operate them. I had Chris Penrose on the podcast from NVIDIA. And I was like, "Well, who's responsible for making this work?" This is complicated stuff, right? Back to software and operationalizing it. And he didn't have a really good answer on that. He was like, "The telcos." And I'm like, "We'll see about that."

Charles Fitzgerald:

[12:10] Yeah. I mean, NVIDIA's very good at low level software, the firmware, the drivers, CUDA, but they really don't play at the higher level stuff. They do a lot of press releases about doing their own software. But fundamentally, they're a chip and increasingly a server company.

DR:

[12:28] Yeah. So your 2025 Follow the CAPEX work shows that this is still a brutally concentrated game besides a handful of platforms plus NVIDIA. So when you look at that, is there any realistic way for second tier clouds, like our favorite IBM, or maybe even Oracle, or telcos to matter? Or are we done? It's a four company market plus maybe one chip vendor.

Charles Fitzgerald:

[12:51] At scale, it's extraordinarily hard to catch up. In order to play this game, you need to build an immense globally distributed computer. The hyperscalers have hundreds of data centers around the world. They own transoceanic cables. They design their own chips. And now they are building gigawatts of power. That's tough to compete with. Maybe there's some niches. But the idea that somebody's going to catch up and go head-to-head with the hyperscalers in terms of both breadth and depth and global footprint, that seems hard.

[13:23] And so we've talked about IBM and Oracle in the past. IBM is completely out of the game. They managed to put up negative CAPEX in the last year, which is a new record, and I'm not even sure how you do that. But Oracle's Hail Mary to catch up in cloud, I think, is super instructive here. They poo-pooed cloud for years. And after they lost half their database market share, they realized they really had to be a cloud provider. So now they're trying to catch up with the hyperscalers from almost a standing start.

[13:54] Amazon, Google, Meta, Microsoft passed a trillion dollars in cumulative CAPEX spend in 2024, and they'll pass \$2 trillion this year. So if we look at Oracle, Oracle's lifetime CAPEX spend over about five decades of existence is less than 90

billion, with a little over a third of that in the last year. They've borrowed mountains of money. They've lit a bonfire of cash. Their free cash flow has gone deeply negative. They're having to offer huge discounts to snag customers, which kills their margins. And they're super inefficient in spending the CAPEX because they're 15 years behind the hyperscalers on the learning curve. So they're in a huge state of financial distress right now. And the number I really look at, they're spending over \$3 in CAPEX for a dollar of revenue, and their cloud business is still teeny.

[14:50] And if you contrast that with AWS who are in the midst of a huge AI build out themselves, AWS only spends about 75 cents per dollar of revenue right now. So Oracle is throwing this Hail Mary to try and catch up and it's incredibly high risk. And Oracle, at least, is actually a tech company that can write some software.

[15:12] So you do have some of these neoclouds around CoreWeave and others that have managed to deploy a few billion dollars worth of infrastructure. They've gotten preferential supply from NVIDIA in a market where GPUs are scarce. I'm skeptical about the neoclouds, but I'm really skeptical about the telcos getting into this game. And it's the same set of issues we've already talked about. They're very late in a world where the infrastructure investment compounds over time. And that's a dynamic they should understand from their network history. They don't have the cash to play in the big leagues of CAPEX here, and they don't have the technology chops to play this game. But in the absence of other ideas, they'll probably do some press releases until the next big thing comes along, and then we'll probably talk about that next year.

DR: [16:00] Yeah. Or even just the culture of speed.

Charles Fitzgerald: [16:03] I mean, they get confused and think that they're technology companies and they're not. They are integrators, they are operators, they market and sell services, but they don't actually create technology. And it's very hard to compete on a technology basis with the people who actually create the technology. You're just at a huge disadvantage.

DR: [16:25] Yeah. Well, that is awesome. I hope every telco exec that is evaluating these options and these new strategies listens to that wise information from you. Like I mentioned at the top,

Episode 137 | Hyperscalers' \$2T AI Bet (Charles Fitzgerald)
Released March 31, 2026

you've been on the podcast four times. You're the first person to do that. And to celebrate that amazing milestone, we're sending you some personalized Apple AirPods Max headphones.

Charles Fitzgerald: [16:49] Wow. All right.

DR: [16:51] Yeah.

Charles Fitzgerald: [16:51] I'm honored, and I'm also amazed you're still on the air after having that many appearances by me.

DR: [16:57] No, people love your podcasts. They get really great listens. And so I think you probably will be on the podcast again.

Charles Fitzgerald: [17:04] Alright.

DR: [17:05] So Charles, thank you so much for sharing your awesome wisdom on CAPEX spend. I really appreciate it. It was all great.

Charles Fitzgerald: [17:11] Great to be here.

DR: [17:12] Awesome. Stick around. We end each podcast with a "Telco in 20" takeaway. I've got two minutes to tell you something you need to know.

[17:28] Charles gave it to us straight. Telcos don't create technology. They buy and operate it. On the other hand, hyperscalers build products that compound in value. Every dollar they invest makes their platform smarter, their services stickier, their moat deeper. So how can operators build an asset that compounds? You build an ontology. Here's where to start. Your rating rules, your provisioning workflows, your product specs, that knowledge lives in vendor code and people's heads. And right now, it's invisible to every AI system you deploy.

[18:04] The Totogi Ontology captures that operational knowledge and makes it executable. That means AI can use it to act across all your systems. And here's why that's valuable. Every time the system acts, resolves a dispute, applies an offer, provisions a service, the context gets richer. The next decision is smarter than the last. And the business impact is real. You retain subscribers you would've lost. You build products that were stuck in the backlog because no one could figure out how to deliver them. You stop paying vendors to do work that your

Episode 137 | Hyperscalers' \$2T AI Bet (Charles Fitzgerald)
Released March 31, 2026

systems could now do themselves. Revenue goes up, cost goes down, and the gap widens every day in your favor this time. Sovereign cloud and AI-RAN are examples of telco trying to win the hyperscalers' game with a fraction of the budget. An ontology is a different game where deep telco knowledge is the advantage, not raw spend. And unlike hardware, that advantage grows every single day. So start playing a game you can win.

[19:04] Want to see what competing on your own terms looks like? DM me on LinkedIn or X @TelcoDR and let's talk Totogi. Until then, tune in to more "Telco in 20" episodes, like and follow, and leave us a five-star review. Don't forget to sign up for my no BS email newsletter on TelcoDR.com. And check out our awesome YouTube channel and hit that subscribe button. Later nerds.