

NETEVENTS GLOBAL IT SUMMIT

SAN JOSE, USA

OCTOBER 3 & 4, 2019

FINAL

Conference Debate Session VII - Hyperscale for Enterprises: You Don't Need to Be Google or Facebook to Achieve Data Center Greatness

Introduced & Chaired by Brad Casemore, Research Vice President, Datacenter Networks, IDC

Panellists:

Mansour Karam	CEO & Founder, Apstra
Kevin Deierling	Senior VP of Marketing, Mellanox Technologies
Kyle Forster	Founder, Big Switch Networks
Mike Capuana	Chief Marketing Officer, Pluribus Networks

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Thank you. My name is Brad Casemore and I'm Research VP for Data Centre Networks at IDC, and I'm pleased to be able to host this panel today. We're going to talk about hyperscale for enterprises. Many enterprises today, as they look at their data centre infrastructure and, particularly, their network infrastructure have a degree of hyperscale envy, and the panel participants today will address how they can satiate that and move on and actually conquer the problem.

As the title says, you don't need to be Google or Facebook to achieve data centre greatness, certainly in a networking context. I'm just going to cue this up briefly, and I'm going to have our panel dive in. Just to cue it up, we've already established that

enterprises are on a journey and you've heard about it before, and digital transformation, that's really the driver. It's an imperative that all enterprises are pursuing.

It's not an option these days. It's something that they need to do. It's something that they must do. It has significant applications for their infrastructure in their data centre and their networks, because typically, it must be modernised. They're dealing with very traditional, outdated architectures. They're dealing with operational models that have become outdated and they need to modernise. They need to modernise architecturally and they need to modernise the way they run their networks.

So it's a design issue, it's a deployment issue and it's an ongoing operational issue. As they look to what's happening, we see cloud as not only a destination for workloads, but cloud has its own operating principles. When you look at the hyperscalers, they've been able to do things at unprecedented scale and with unprecedented agility. They've used comprehensive automation within their data centres for their data centre networks.

They have definitely - they were one of the earliest progenitors of, of course, SDN principles. They have used scale-out architectures in a way that has been tremendously efficient and effective within their data centres. They use real-time analytics increasingly, and streaming telemetry. In order to not only inform their policies, their declarative policies in their data centre, but to run things more efficiently and to be able to remediate problems fast and more effectively.

This has allowed them to move from a reactive approach to a proactive approach within their networks. It's very significant. It's not just about remediating problems, it's about being able to plan for eventualities and to move to a very proactive approach, and they've done part of this through having the right abstractions and through using disaggregation so that they can innovate at speed on an overlay, on the underlay, in software, in hardware. They can decouple those so that they can move very quickly.

Enterprises are hyperscalers when you look at their business, and we'll talk about some of those reasons why, but they want to move in this direction. I want to cue it up. I had initially two slides, but we've got four panel participants and they all have something fascinating to say on this topic. I'm going to have them briefly introduce themselves and then we'll dive right in. Kevin?

Kevin Deierling, Mellanox Technologies

Yes, I'm Kevin Deierling. I was in the last session so you probably know me. I'll probably talk a little bit more about switches than adaptors on this session, because that's one of the areas where we really see the hyperscale guys have embraced disaggregation, where you decouple the hardware from the software. The other thing that I'd like to address is really the difference between an enterprise and a cloud and how the enterprises are going to have to move towards the cloud model. We call it hyperscale for the masses.

If I characterise a cloud guy, they have a protocol stack that - sorry, an enterprise guy, they have a protocol stack that looks like that, that's been given to them by their

vendors, whether it's Cisco or Arista, and they do just a tiny little bit of actually monitoring automation. The cloud guys are the exact opposite. They have a really narrow protocol stack and they do a ton of automation and monitoring. So I think that's something that's important for us to dig into.

Mike Capuano, Pluribus Networks

My name's Mike Capuano, Chief Marketing Officer at Pluribus Networks. We are also focused on networking, as it turns out whenever you're trying to automate a data centre, there's been a lot of progress in virtualisation and automation of the storage and compute layer. The network always seems to be the Achilles heel.

Pluribus is focused on bringing network automation, SDN control, the physical underlay - that's the actual switches that connect everything together - network virtualisation, which is creating an abstraction that's often referred to as an overlay and network analytics. What I'm going to try to talk about is an approach to delivering that network automation that works well in smaller data centres, whether that's a midmarket business who aspires to be like a hyperscaler or even aspires to be like a large enterprise, but it's too costly and too complex to actually get that network automation for that smaller site.

Then, also, we're seeing trends in general around distribution of data centres, so we're going from a model where there's a large multi-football field sized data centre in Iowa to where there's a much higher distribution of data centres in your neighbourhood, and how do you build automation for that type of environment that is cost-effective and scalable across that more distributed environment, so those are the things that I'll focus on.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Mansour?

Mansour Karam, CEO and Founder, Apstra

My name is Mansour Karam, I'm the CEO and founder of Apstra. Apstra pioneered intent-based networking. We build software that sits on top of networking infrastructure and provides powerful automation and abstraction across the entire life cycle of those network services, Day 0, Day 1, and Day 2-plus, helping organisations consume infrastructure a lot more simply and to help them deliver on their business initiatives a lot more effectively.

We've seen customers that have deployed Apstra software see a 99% improvement in their infrastructure agility, a 70% improvement in their mean time to resolution of issues - which is a measure of reliability - and an 83% reduction in operational efficiency, measured in terms of man hours required to deliver on a specific infrastructure task.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Kyle?

Kyle Forster, Founder, Big Switch Networks

Thanks, Brad. My name is Kyle Forster. I'm the founder of Big Switch Networks, and if you'll indulge me, a little story that I think is appropriate for this to set up the conversation. We started Big Switch in 2010 and two years ago, imagine that we're seven years into trying to build the best SDN company that we can possibly build. One of our customers, a Chief Architect, one of our customers, he said this quote and he didn't even say it with us in the room. He said it and it went viral around our partner who was there and it came back to us. He said, I chose Big Switch because it was the closest I could get to Amazon networking on-prem.

It was a seminal moment for us. Seven years in and suddenly it was like the light went off. This became almost a mission statement for the company, and suddenly we had so many things that we had to build. It really opened up at least our eyes and certainly my eyes to say, look, in so many different ways, the cloud has won, and now what can we do about it? Can we build systems on-prem that aren't inspired by public clouds, that aren't kind of like public clouds, but in fact that are exactly like public clouds?

If you start with, hey, why don't we build the exact same experience on-prem that people get in clouds, if you use that as a starting point, then there are interesting deltas off of that, she really, really interesting hybrid cloud components that start showing up. So this one quote two years ago sent us on this fascinating journey, and I hope to have the chance to be able to talk about that a little bit more.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Absolutely. One of the things - and you've touched on it, all of you have touched on it from different perspectives, and that's exactly what we wanted, so I think you all bring a different perspective to this. There's no question that there is a demand for the sort of agility, for the sort of flexibility, for the sort of reliability, for the adaptability that hyperscalers have in their data centre networks, but enterprises are enterprises, right? Many of them have bricks and mortar manifestations, as well as digital.

They are becoming more digital. There's no question about it. They have certain scale sets, they have certain comfort levels, they have certain cultures. They're very different from the hyperscalers in their application mix. So how do you make this consumable and, frankly, simpler for them to be able to implement it in their environment?

So I wonder - I think all of you can - there are different architectures. I know we have controller lists, we have controller base, we have a mix of technologies and approaches here. I wonder if we could maybe get a perspective then, so start with you in terms of simple is very important in this context, right, because they don't have teams of PhDs in most enterprises and they're looking to do this in a way that's consumable for them.

Mansour Karam, CEO and Founder, Apstra

Yes, what we have seen is that enterprise environments are in some ways a lot more complex than your public cloud. The public cloud has a certain set of applications that's quite constrained. They're only building a product for themselves. When we're a vendor and we're going to the enterprise, every enterprise environment is quite different.

There are various mixes of modern, cloud-native applications and legacy environments, and so the challenge is really to build a product that ticks - distils, in a sense, the benefits that automation provided to public cloud providers. Distil those benefits, but then deliver them in a package that applies or takes into account the complexities of enterprise environments. So it means different things. It means that you have to have the ability to support modern applications but also legacy applications.

You need to have an architecture that's quite extensible because, ultimately, you are stitching within an environment that already exists, where there a different set of tools, and so having that ability to extend the product, having the ability to integrate the product within the environment becomes critical. Then you need to have an ability to apply to different domains, so you can't be rigid and say, it's only going to apply to this exact, specific topology, which is a lot of what simplifies the problem for these hyperscalers.

They only have one network and that network is designed in a certain way. In enterprise environments, you have a variety of such networks. You have a variety of different domains, and you need to have a product that, at the core, has an architecture that allows you to adapt to these different domains.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Kyle, how have you dealt with this with your customers from Big Switch's perspective?

Kyle Forster, Founder ,Big Switch Networks

Sure. I think - there's this area that I'm really passionate about. I'm actually so passionate about it, I talk about this so often in the office and with our customers, a couple of colleagues actually asked me to write up a LinkedIn post about it. The LinkedIn post flopped, nobody looked at the thing at all, but the really important part, I grabbed a snapshot of the Google search trends over the last 10 years for AWS VPCs, which is kind of the primary networking primitive on Amazon.

Like, hey, you want to spin up a VM on Amazon? Well, the first thing that you do is you actually spin up a VPC which has a network set up around it, and you see a trend line that just looks like this. Then I did a search for Cisco VLANs and you see a trend line that looks like that. The two of them crossed nine months ago. There were more Google searches for AWS VPCs than Cisco VLANs. It becomes this thing of, how do we present an Amazon experience or an Azure experience or a Google Cloud experience on-prem?

Well, it was kind of obvious for us the first thing we had to do was start by presenting VPCs, not kind of like a VPC, not sort of, but an actual - the exact same object model. I'm incredibly proud now that, when somebody uses a terraform template on AWS, they can use the exact same terraform template on-prem as far as the networking stanzas go with our stuff. We're working on when somebody uses a cloud formation template on AWS, you can use the exact same networking stanzas in that same cloud formation template on-prem. This suddenly made a whole bunch of the - all of the cloud native integrations really, really, really, really easy.

Some of the trickier parts were things like Nutanix integration, vCenter integration. Those were surprisingly slick to make a great user experience. They actually translated into the VPC model much better. There was the VMware vDS or the Nutanix networks, the OpenStack Neutron networks. That translation was much, much easier than I thought. So it's kind of way of saying, can you present - there's the networking team, but then what's the networking team doing to present an option for all of their colleagues?

If we're an extension of our customers' networking teams, this is our job. Like, how does our team present this experience? Our teams are no longer 80 people, 100 people, 200 people. There's an awful lot of the market that's now four to seven, maybe 10-person networking teams, and how do you get there by presenting - when the team has the aspiration to present a cloud-like experience, how do we get them all the way there? I think we found the nubbin of that with VPCs.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Absolutely. Mike, I know talking about customers and their evolution and enterprises, I know you have [unclear] your installed base, some customers - we won't name them, but some customers who've been around for over 100 years and they've obviously gone through many IT cycles in that period. Could you talk about some of the challenges that they've faced and how you've helped them address going from very traditional environments, legacy environments and moving on to a cloud there?

Mike Capuano, Pluribus Networks

Sure. Your original question was how do you make this more consumable, how do you make it simpler? I think the problem is the status in our automation has been realised by, like you said, hyperscalers who have huge IT stats, lots of PhDs, writing code. Even in the larger enterprises who have achieved, say, a software-defined data centre, they have really large IT teams.

So our focus is, how do we literally do what you just said? How do we make software-defined networking, open networking, network virtualisation, analytics, consumable for the rest of the world? So I'll point to something Kevin said in the last panel. How many panels are you on today, by the way? Your highest cost item in your data centre is your CPU, your RAM and your flash, right?

So imagine you go from a large enterprise that has 50 racks down to a medium-sized data centre with four racks or a medium-sized data centre that's moved into four different colos and they've got two racks or four racks in each colo, how do you automate that? Because for the large data centre, you've got all these external servers that you're deploying for management, for controllers, controllers for the SDN layer, controllers for the network virtualisation layer, stacks of servers for network analytics.

In a big data centre, that all just disappears and doesn't matter, but in a small data centre, all of a sudden, that's a big deal. So what we said is, hey, switches being developed by companies like Edgecore, Dell EMC, Mellanox actually have pretty powerful processors in them and RAM in them and flash drives in them. They're being deployed anyway, because you have to physically connect the leaf and spine architecture, so why let that latent server-like processing power go to waste?

Let's leverage that, write some clever distributed code, so literally on the switch you have to deploy anyway, you're running your SDN function, you're running your network virtualisation function, you're running your network analytics function. Then, the other benefit is, now I can pre-integrate all that, pre-test all that. It just works out of the box. So that's why, as a company, we have some large enterprise, but generally we're not pursuing large enterprise.

We are pursuing folks who have small and medium data centres, folks who are moving from a centralised model to a more distributed model, EdgeCompute where you've got these constrained environments, and that's the way we're trying to make it simple. To your earlier question, what did we have to go through with one of these larger customers? The good news is our design, like I said, kind of works out of the box, but for every deployment - and Kyle and I talked about this in the hallway - you've got to have a tack, you've got to give those customers a lot of love and sometimes handhold them through it.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Absolutely. Kevin, obviously, from your perspective, you're a bit different, right? Because you've got network interface cards, you've got SmartNICs, you've got network silicon, you've got switches, so I wonder what you see from this perspective in your engagements?

Kevin Deierling, Mellanox Technologies

Yes, I think the key point here is that the cloud and the hyperscalers have these massive teams that have developed everything from scratch, and that journey that Kyle was talking about is exciting for us because we don't have those massive teams ourselves, and all the software folks that are here reflected at this table, we've been working with them for years. So Apstra's a great partner. How do you get automation if you're not Google or Amazon or Microsoft? Well, you talk to Apstra. That's how you get it. That's the thing.

So for us, having an open platform that we can run different operating systems on - SONiC was something that Microsoft developed in-house for their data centres. They run it at scale, it manages hundreds of thousands of servers in their data centres. That's now available. It runs on an open platform like ours.

So companies like ourselves and Apstra and others are supporting a lot of that infrastructure that, for the first time, smaller enterprises and mid-size enterprises can take advantage of some of the open source code that's out there, things like SONiC, also other open platforms, Big Switch is here, we've got Pluribus, we've got Cumulus, who's not here. All of that stuff can run on these open platforms. So really, you can steal a playing card from the hyperscalers today and deploy things.

For us, it's a double-edged sword. That means that you can - the folks that Pluribus was talking about here are our competitors, okay? When you disaggregate hardware and software, it means there is no vendor lock-in. It's not like Cisco where it's a black box and you get the hardware and the software together. If somebody offers better hardware, you can go buy it from them. If you want to use whichever software you want, you can do it. It's finally all integrated and gives that choice, so that whole notion of open network is really a powerful thing.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Do you want to say something?

Mansour Karam, CEO and Founder, Apstra

Yes, so I completely agree with the comments around disaggregation, and the way I talk about it is this decoupling that one has to do between the hardware and the software. In fact, if there is one piece of advice that I'd give to enterprise organisations that are in the process of transforming their infrastructures, it's please don't start by choosing one strategic partner that is your established hardware vendor that you've been working with for the last 20 years.

If you start there, then you're going to restrict your choices dramatically. You're going to lock yourself into hardware and you're going to lock yourself into a fully-integrated solution that is not best of breed that will slow you down and will not deliver the needs of your business, and the reason why this is critical is because Gartner has seen, has demonstrated, has measured that digital transformations or initiatives are three times more likely to fail if you fail in your process of infrastructure transformation.

So this is why we really ask organisations to look at software first, look at that layer of software. What is that abstraction that you're going to start with as the foundation? Software is key. Software takes the critical role here. Once you've defined that, once you know what these abstractions are, once you've built that layer of automation, that service layer that will interact with your business applications, once you have that right, then choose your hardware. Hardware is critical. It's important, but it needs to be decoupled from that software decision.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

It's a very interesting point. In fact, we've found for traditional enterprises and their environments, they're still on - even in their data centre networks - a five to seven year cycle, but enterprises who are more software-driven in terms of their approach to their networks and software-led, there tend to be hardware benefits, as well, because they tend to be on a faster cycle.

Mansour Karam, CEO and Founder, Apstra

Absolutely. In fact, there are so many vendors out there and the right formula is going to involve a heterogenous environment. There's going to be heterogenous hardware, heterogenous software. So don't lock yourself into just one integrated stack. Allow yourself the ability to leverage or take advantage of all those options that are out there. It's going to be a requirement for your business.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

One thing that I mentioned and definitely is present in hyperscale - and it's definitely moving to the enterprise data centre, as well - is the need for pervasive visibility of a real-time nature. Now, it's not just Day 2, right? There are many environments, and I'm sure you've all found this, with SDN where they don't even - they really don't have visibility into what's out there before they deploy, right?

Panel Speaker – Male

That's a point.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

How important is that insight into informing their policy? I wonder, Kyle, if you could talk to that?

Kyle Forster, Founder, Big Switch Networks

Sure. For a number of years, we built up a lot of actually analytics tooling and troubleshooting tooling by trial and error, which actually got us pretty far. One thing that started, I think, especially as we started doing this journey more and more towards replicating public cloud, we started looking at a lot of the SIGCOMM papers that were coming out from public clouds themselves on their own internal troubleshooting tools.

So I separate out two things. How do actually non-networking professionals troubleshoot the network, the way that us as users of Amazon, Google or Azure do? There are some really interesting tools there, something as simple as CloudWatch, VPC Flow Log. We should be generating those. Then, there's also gone through the

SIGCOMM papers, through recruiting, you can figure out some of the really, really successful internal tools that are being used to troubleshoot these clouds, and then being closely inspired by those, it's very, very far very fast.

We see plenty of teams - the Azure team is a great example where they have a series of incidents that they published, that in 2010 took them about on average 20 hours for remediation, and now they're down to four seconds, because the average is four seconds because so many of them are machine-remediated and a very small number involve any human at all. They're published. This is all published work.

So I think we've actually seen a huge boost here, both by going off to the public offerings that a networking team offers to other people to help them troubleshoot, as well as going after, actually, some of the in-house tools, and couple that with a lot of the trial and error that just comes from not only listening to users, but much more interesting is actually watching users use a networking system and just jotting down where they start to look frustrated or jotting down where they're starting to spend a lot of time. Those user studies, to me, are fascinating. We don't do that very much in the networking industry, but gosh, do they ever tell you a lot.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Mike, I wonder your perspective on this. I know the key feature in your portfolio...

Mike Capuano, Pluribus Networks

Right, instrumentation analytics is critical. Going back to the statement around disaggregation, the value you get in a disaggregated world where you've got hardware innovations happening separately from software innovation, and they're not tied together and one is not slowing the other down is they can both move faster. So we're focused on software. We run on open networking hardware. We put a lot of energy into analytics and, ultimately, let me give you a use case.

So we have, for example, a large school district deployed in the Midwest, 130,000 students and employees and they're distributed, so they've got multiple data centres. They use our fabric and its ability to capture every TCP flow to shut down DDoS attacks. So what they can do is, because we present every flow that goes across a fabric, if they see some unusual spike and they look at this through our graphical user interface, they'll know it's a DDoS attack and they will then use SDN.

They can go to any one switch in their entire fabric, shut down a port and then, automatically, within a second, it shuts down all the ports across the fabric and shuts that DDoS attack down. What that lets them do is not have a \$400,000 a year DDoS cloud subscription. So it gives them that automation, it gives them that visibility and it actually saves them a lot of money, so that's a financial result you could see from having good analytics.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Absolutely. Your perspective on that?

Kevin Deierling, Mellanox Technologies

Yes, so we like to talk about mean time to innocence. When something happens in the data centre and applications start to perform poorly, the first thing you do is go find your network admin and start screaming at him that it's got to be the network. We call that the mean time to innocence. When can I show that there's nothing wrong with the network? It's something else. It's this application that's running on this server which is spewing data and things like that.

So we actually also focus a tremendous amount on the low-level infrastructure to provide analytics, so we have a tool that we introduced called What Just Happened, WJH, so that when things start happening, we keep a record of all of the activities that are happening and we feed it to these guys so that they can show you, so that very quickly, instead of having to wait days and days and days, you can say, hey, it's not the network.

The network is functioning just fine. It's that node over there is spewing massive amounts of data or the database, it's pushing back for some reason that node there, this server or this SAP or Oracle. Those are our customers, so I shouldn't use them as examples, but it's really a critical thing, is to have the analytics.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

It's a great point, and I think we all know, the networking team often gets the blame for anything that goes wrong, and they love to be able to say, no, it's not us.

Panel Speaker – Male

Exactly.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

I think we have some time for some questions. Does anybody have any questions? We have one upfront here.

Audience Q&A

Audience – Male

One of the key points that I think they see from the public cloud is a matter of abstraction. A lot of the low-level stuff you don't need to touch, you don't need to think about, it's automated for you. Yet you've all been talking a lot about producing analytics

and delivering them to the end-user. That's not something that is going to be wanted in that sort of environment. So what do you think about the arrival of new layers of abstraction through things like Anthos coming into the private clouds, into the build-your-own public clouds on your own hardware?

Panel Speaker – Male

It is all about abstractions, you're exactly right. It's interesting that while abstractions are critical to the public cloud, they're even more critical to your traditional organisations and enterprises who have a slew of different vendors and very heterogeneous environments. Getting those abstractions right is not only critical but it's also necessary, it's also harder to achieve.

How it relates to analytics is abstractions is where everything starts. You don't see what's going on underneath but abstractions, describing what services you want from your infrastructure, is the starting point. Then in order to achieve that reliably not only does this starting point drive how you configure all those elements underneath that you shouldn't be thinking about but that the software should, but then also the software needs to measure the infrastructure, keep state of that infrastructure and have the ability to verify that in fact the infrastructure is delivering on those services that you've laid out in your abstraction layer.

You need to have a very integrated approach where these abstractions, this service layer is driving both the configurations and the telemetry that is being gathered and the tests that you're running and your infrastructure to verify that your infrastructure is delivering on what you need.

It's a bit like a self-driving car. Your layer of abstraction is, I'd like to get to this destination. That's what you're thinking about. Then the software is both controlling all the components in the car, but it can't do that if it doesn't know what's going on around. It gathers all this telemetry, measures this telemetry, analyses it continuously to ensure that the car is indeed delivering on your stated intent.

Audience – Male

So you can drive the automation.

Panel Speaker – Male

Exactly. That's the only way. Without it, it's extremely dangerous. If you look at the first generation of automation it was defining abstractions and then pushing configurations. That's how there are so many horror stories you can find online about people losing their jobs and bringing down infrastructures in a way that affects thousands of devices and you have no idea where to start in terms of debugging or troubleshooting this.

Kyle Forster, Founder, Big Switch Networks

One thing we found is, we took an approach - let's use Anthos. I think that's a great case study. Anthos, before you roll out a pod, you roll out a VPC. Google-sensitive VPCs are quite unique and really easy to use, such a beautiful model. But then if something is not working in there or if there is a series of connections or performance issues that you didn't expect the Anthos pods tracks to some outside pod.

Then you have flow log. Users troll through flow logs, try to figure out what happened after the fact. One thing we found is, we started doing a lot of the base-level infrastructure to expose VPCs and flow log on-prem. There's a whole bunch of analytics tools that we started putting on top of flow log on-prem. Think of it as one-to-one NetFlow.

Those tools were so useful and so valuable for in-cloud deployments that for us we just launched this line of products that is actually in Amazon. For a number of years we thought of ourselves as always building switch-side software but actually all the software that we started building on top is so useful that we launched a bunch of troubleshooting tools for AWS. We're just about to launch troubleshooting tools for Azure and for GCP, and they're the exact same troubleshooting tools that we use to troubleshoot stuff on-prem. It's the exact same health check and monitoring check and it's using the exact same APIs.

To me it's one of these fascinating case studies. To your point about abstraction, when you get the abstractions just right the tool chain you build on top, whether something is on-prem or in-cloud. It's the exact same tool chain. To the extent that as humans we're like human middleware all of us in this big tool chain, it kind of becomes the same team. That to me is a really neat thing as we think about the careers of a lot of networking professionals out there who for a long time I think may have thought of themselves siloed on-prem. Can their careers grow to something where they're not siloed on-prem? As a network engineer, they take care of the networks, they take care of connectivity. Whether that connectivity is in-cloud or on-prem it doesn't matter. It's the same team making sure things are healthy.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Final question over here, please.

Kishore Jethanandani, Futurist Lens

My name is Kishore Jethanandani. I'm with FuturistLens. My question is for Mansour. You drew a relationship between digital transformation and the infrastructure being ready for it. Could you illustrate that? Can you hear me?

Brad Casemore, Research Vice President, Datacenter Networks, IDC

I couldn't. There's a lot of echo.

Kishore Jethanandani, Futurist Lens

It was the relationship between digital transformation and infrastructure.

Audience – Male

Yeah. You have to get the infrastructure ready for digital transformation. Can you illustrate that with a few scenarios where digital transformation worked because the infrastructure was ready for it, and another scenario where it didn't work because the infrastructure was out of sync? In what kind of digital transformation...

Mansour Karam, CEO and Founder, Apstra

There are lots of examples where it didn't work. Let me give you an example where it actually worked and it's one of our abstract customers. Essentially it's an insurance brokerage, a Fortune 500 company. The business initiatives that were part of this digital transformation were Tesla's initiative that drivers in Tesla can choose their insurance right from the car by typing - just by clicking on the screen, right from the car so they don't have to call an insurance company or insurance brokerage in terms of figuring out what insurance they want. So right here, one click you choose your insurance and you buy it from the car.

In order for them to deliver on this they had to transform their infrastructure. Their comment to us was, we had two choices: either we increase the size of our team tenfold or we have to embrace automation, abstractions, go through this painful process of changing our operational model. So they went down that path and then changed their operational model in order to deliver in time for this business initiative that was in front of them. In my mind that's a prime example of digital transformation requiring infrastructure transformation.

Brad Casemore, Research Vice President, Datacenter Networks, IDC

Thank you very much, gentlemen, and thank you audience for hearing what these folks have to say. I think we can all agree after this session that there's definitely a need for datacentre modernisation and certainly network transformation in the datacentre in order to support DX. Thank you very much.