



Data Center Technologies and Trends

May 26, 2021

Draft Transcript

Featured Speakers:

Analyst Chair: Baron Fung, Research Director, Dell'Oro Group

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Mark Fox, NetEvents

Welcome, everyone, I'm Mark Fox CEO of NetEvents and delighted to welcome the international press two this session today on data center technologies and trends, chairing this session we have Baron Fung Research Director of Dell'Oro group and Baron is going to introduce his panel in a few seconds, Baron, over to you.

Baron Fung, Research Director, Dell'Oro Group

Thank you, Mark, thanks for the introduction. I am Baron Fung, I'm an analyst at Dell'Oro group responsible for our data center research pertaining to cloud cap x servers and server connectivity. It is my great pleasure today to be moderating this event with our distinguished panel of experts in data center technologies. Joining us today are Doug Gourlay Vice President and GM of Arista, Cloud Network Software Division. Jonathan Seckler, Senior Director at Infrastructure Solutions Group for Servers, and Partners at Dell Technologies. Tony Bishop, Vice President of Platform and Growth Marketing, Digital Realty. And finally, Jacob Rapp, Director and Lead Technologist for Networking Security office of CTO at VMware. Next, I am going to go through some slides, which will help set the stage for today's discussion. So just to start off servers generally comprise more than half of the data center, capital expenditure and drives the forecasts of assumptions for forecasts, other technology areas such as compute, storage, networking, power, and cooling. In this forecast here, we project server shipments to grow at a 6%



- 1 -

CAGR through 2025 growth is projected to be an even across customer segments. Last year the top four US cloud providers, Amazon, Google, Microsoft, Facebook, accounting for more than a third of the global shipments of servers worldwide and will increase their share over time as they continue to consolidate infrastructure. Over time, we are expecting faster growth coming from the top Chinese cloud providers, Cloud adoption in China is still behind the rest of the world. So, there are a lot of upside potential, plus the government in China is making a strong concerted effort to boost infrastructure spending on the long term. The rest of our segment, mostly related to tier two, tier three cloud service providers, growth there could be mixed. We see this segment adopting a hybrid cloud strategy, and some of the larger cloud providers could potentially repatriate workload back from the public cloud to invest in their own infrastructure. While the telco currently occupies the smallest slice of this market, it is expected to be the fastest growing segment driven mostly by new use cases that requires low latencies, applications, and would increase the demand for micro data centers with servers and networks at the telco edge, the enterprises are likely to adopt multi cloud deployment model with regulated on mission critical workloads remaining on premise. The recent pandemic was really disruptive to the industry and drove some enterprises to adopt a more nimble, flexible IT approach. We expect the other, what we call the rest of enterprises, mostly SMB, SMB to mostly migrate to public cloud, as their preferred IT strategy. Next, I'll touch on some information about deployment of it, locations. In this slide, is the same forecast data as the previous chart, but shows where infrastructure will be deployed. Today most of the servers and infrastructure are deployed in centralized data centers, including Point of Presence in networks by the cloud hyperscale companies and large enterprises. Servers deployed at the infrastructure edge as multi access edge compute nodes or MEC is still a small part of the data center footprint today. This is potentially a fast-growing market with triple digit CAGR over the next five years, IT and networking vendors, and ecosystem providers have been active in developing and evaluating solutions today. The Edge could be a next driver in infrastructure spending in the long term. Then the question is, which set of customers will be investing in deploying these solutions will the telecom operators lead the way? Or will the cloud service providers expand their reach to the network edge? We are also anticipating the higher enterprises to deploy edge solutions over their own private networks. As we move to the next slide, I'll show how data center technologies are allocated in terms of CapEx. As I mentioned earlier, servers will continue to be a dominant spend in data centers, and the allocation of servers to rise relative to other areas, given several drivers: one is the successive kind of server CPU refresh cycles, which will increase server content, increasing the adoption of accelerated computing servers for AI and machine learning, and the continued use of servers to replace proprietary hardware and data center appliances through virtualization. Other components of the data center from network infrastructure to physical infrastructure will also need to evolve as well. All these factors would increase the size of this data center capex pipe to nearly \$300 billion by 2025. In this next slide. I am going to talk about how the Top 4 US is projected to invest in their capex this year. So, we are expecting the top for us cloud to grow capex by 20% in aggregate, in 2021, the spending here can fall into three main buckets: servers will be the primary capex driver with server deployments to new data centers and servers, used to replace expired servers. A combination of higher server content with more memory, more cores, and with Intel Ice lake ramp, could potentially drive commodity costs higher as well. For networks, the top four cloud will also undergo an upgrade to the next generation of networks based on 200 and 400 gigabit connectivity. And lastly the cloud providers are continuing the expansion

- 2 -



of data centers globally and will also drive the cost of physical infrastructure costs, higher. As we can see in the next slide. This chart shows the total number of availability zones for top three public cloud providers. The yellow bars show the cumulative total number of availability zones in 2021, and the green bars indicate the cumulative total number of zones. Over the next two years, these cloud providers are set to expand the footprint of their own data centers, or through colocation providers, the total number of availability zones for each company could surpass over 100 in the next couple of years. It is really a race to see which provider has the most extensive network of data centers, with the purpose of improving network performance, increasing redundancy, and increasing the number of data centers in different regions to meet data sovereignty requirements. In the next slide, I'll be talking about some of the accelerated computing trends. So, we have seen a proliferation of specialized coprocessors such as GPUs FPGAs and custom accelerators, built by the hyper scalars, with growing demands in AI and ML related applications. I believe we are still in the really early stages of transitions, and there's a lot more room for growth. But in order for the industry to adopt greater employment opportunity computing in both centralized and as data centers. The industry really needs to come together and simultaneously innovate in areas of CPU design system architecture related to memory storage interconnects heat dissipation, networking, and application development. In the next slide, I will talk about another type of acceleration with Smart NICs also referred to as DPU or data processing units. Smart NICs are Ethernet adapters with an onboard programmable processor, such as the ARM chip or FPGA. Some of the major cloud service providers are fully deployed Smart NICs today, we can see in a few more years, the cloud, and other cloud companies enterprises and edge data centers, potentially deploying Smart NICs and we project as much as a quarter of the servers in 2025 to be equipped with this kind of device. Smart NICs, designed to offer the wide range of services, such as networking, storage, and security protocols from the host processor, which could potentially help improve utilization, and reduce network latencies. Smart NICs also give control to end users and empower them to optimize the servers and networks to own software optimizations. This my slides last slide. And next we'll move to some of our discussion topics. So, the first topic is pretty wide open, some of the current trends and advances of technologies in a data center. So, before we begin the session, I'd like to take the opportunity for each of our presenters to give a quick introduction of themselves. We can start with Doug.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Awesome, thanks. Doug Gourlay VP GM for software at Arista, I get to look after a lot of our cool product lines from network observability to how we're building wetted edge routers, using some of the chipsets and systems we have available today to cloud networking, putting normalizing connectivity abstraction and multi cloud environments to a lot of other fun ones. So it's been pretty awesome, I joined Arista in 2009, came back couple years, almost three years ago, and it also did a quick stint, building the data center product division out when I was at Cisco for 12 years. It's a pleasure to be here.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Hey yeah this is Jonathan Seckler, I am from Dell Technologies. I work in the infrastructure solution group marketing group, I lead the solution and server marketing organization for Dell Technologies and been with Dell for about 10 years, and prior to that was at an Advanced Micro Devices.



Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

Hi folks, it's an honour to be here. My name is Tony Bishop, my background has been in distributed infrastructure running data centers and it for large institutions, and multiple startups throughout my career. I'm part of a firm called Digital Realty which is one of the best kept secrets in the world is the largest multi-tenant data center provider in the world, running in six continents 24 countries, almost 300 data centers and 47 metros. We are where the cloud and the enterprise come together and happy to share those insights today.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Hi my name is Jacob Rapp, and I work for VMware. I've been here for about six and a half years. I started trying to scale out the NSX business to where it's at today, and currently I'm working in the networking security office, the CTO on various new technologies and projects for the company.

Baron Fung, Research Director, Dell'Oro Group

Alright, thanks a lot. So, we'll begin our discussion. Like I said this is a pretty wide-open topic, current trends and advances in technologies for data center. And for today's discussion let's begin with some of the recent trends on servers, given the wide range of platforms that we have really seen introduced recently. So let me start with Jonathan. So, Intel ramping production of ice lake servers this year and obviously the cloud providers are looking to increase their spending significantly, mostly on servers as well as this year. This new platform is a type of more cores and deeper memory supporting the latest interconnect standards. From your perspective, what are some of the immediate benefits some of these cloud providers can derive by adopting this new platform in terms of delivering better ROI and enabling new applications, etc?

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Yeah, I think it's a great question and it's an interesting phenomenon that we have, I have seen this. Now, over the years quite a few generations of server technology, where the very famous quote of we're going to double CPU performance every 18 months is generally true we have seen a real proliferation of additional cores, additional memory bandwidth and other technologies related to that. I think though that the real benefit to the cloud providers and the real benefit to the enterprises who are also investing in data centers, is the increase in automation capabilities that their platforms are going to be giving them as well. Here at Dell we recently did a couple of studies, one was with Forrester. And one of the things they talked about was that one of the biggest benefits that enterprises who invest in their own infrastructure have found, surprisingly enough, is the is the benefit in increased automation. And that's been really important these past few years because automation and the need for a light's out type data center, being able to work remotely is critical. And I don't think that's going to go away. The other thing that we found out with the separate research study with IDC was the criticality of investing in new infrastructure over time we actually discovered that the total cost of ownership of a server, that is a server infrastructure that's three years old, is actually less than keeping that same infrastructure around for six years. Meaning that the benefits and performance that you get from these new architectures, the benefits and automation that you get from these architectures, etc. more than pay for themselves,



over the course of a typical life cycle and this becomes an important factor as new technologies like ice lake and the AMD epic platform are becoming the market.

Baron Fung, Research Director, Dell'Oro Group

Okay, so the point you make about benefits in reducing the lifecycle. Do you see that kind of translating to the enterprise customers as well, because generally they keep their servers for a little longer while?

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

They do, and I think I would probably argue it's probably a false economy right. It is certainly easier to do that and the amortization schedule on computer hardware is probably longer than the three years that I just mentioned, but yeah absolutely enterprises, will reap the benefit of a faster refresh cycle, mostly due to the improvements both in infrastructure costs, the reduced management cost and the improved performance that they're getting over time.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Baron, one other thought on that the ice lake transition is actually looking at some of the core technologies that Intel's embedded into a set of the ice lake processors that actually creates conflict with one of your predictions. So the advances in smart NIC, which is predicated on the ability to offload some of the protocol handling to a smart NIC or a GPU as a co processing element, if I'm Intel I hate that. Right, because it's like hey wait I only get five or \$600 per CPU and you're telling me something's gonna drop \$1,000 on a smart NIC and 15 \$100 on a badass FPGA or GPU or something, I'm frustrated when I see that. And if you look at some of the options within the ice lake family they've introduced some of the switching logic they got from I believe a fulcrum acquisition, to having a multiport switch and there was significant accelerators and de-cryptor library handling and QA T for formats faster DPDK plane and so on, as well as a pretty wide set of PCI interfaces which does allow it to very quickly interconnect with some of the adjacent co processing elements. But if I'm Intel I'm gonna do everything possible to try to shove into my ice lakes, and the next generation of processor families targeted at the server, reasonable enough acceleration to try to forestall the adoption of smart next, because I want the money if I'm Intel right. And that's a smart business decision on their part to create products that diminish the value of off boarded GPUs and DPS and again other accelerated co processors. So, it's kind of a co-op petition environment,

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

It is but you know it's interesting, I'm gonna pile on and say that it really then that it becomes an issue for the software community ISVs, because they've got to develop to these different architectures would be interesting, I don't know Jacob if you have a point of view on this? But the whole benefit of smart NICs seems would go hand in hand with hosting the virtualization layer, farther closer than network.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMware

Yeah, a while back we crossed the threshold of the CPUs outpacing some of the network interfaces, I mean the network interfaces are kind of going on this stair step right of like one gig to 10 gig 10 To 4040 to 100 and so on. But the processors themselves, even what's on there today without the smart NICs on



hand, are more than enough to handle all the networking and security virtualization techniques that we're doing. I think when we look at smart NICs, it gives us an opportunity to extend that into different environments. So, say you have these really high-performance databases that you want to not run any type of agent on because you're very much focused on the performance of that platform. Smart Nix can kind of air gap out a virtualization tier that we can kind of develop on and distribute the networking security functions directly there as well, but think there's other opportunities beyond just the standard compute to go with as well.

Baron Fung, Research Director, Dell'Oro Group

It's a good point, we actually have a slide to talk about smart NICs in the next section on (inaudible) computing but now that we're talking about it, maybe we can spend a couple more minutes on it. Yeah, Doug made a really good point about how processors and smart NICs could be in conflict with each other. In fact, a lot the vendors that are not producing CPUs, network vendors or chip vendors have some kind of smart private portfolio. Now of course there's also the counterpoint in which processors AI CPU is higher, give me more cores, you know, so maybe the core professors they're getting more numerous. So the benefits offered become potentially less beneficial. With this mining so this is something we're paying close attention to see how that evolves. And regarding smart NICs these see potentially, it'd be able to offload some of the network functions because there are some vendors solutions with smart NICs, that could take some of the automation from network to the NIC itself. So, what do you think about that Maybe from Doug or Jacob, any comment?

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

I think the race is effectively between a smart NIC vendor or a GPU vendor. Whatever the dedicated co processing vendor is to build something that is more differentiated and significantly faster or more feature capability or whatever, let's just say has more business and technical value than that which is delivered in my stock processor, as long as they continue to keep a reasonably wide gap there, and they're delivering things that customers perceive of high value, there's a market. Now the market as the CPU continues to add more capability on its own, they're gonna always try to take the 80th percentile, right and doubt and maybe more if they if they do a great job, but there's a gap and it's like hey I'm a hyperscale guy I might need this I mean even back in the late 90s We saw web front ends being built with SSL coprocessors, because the crypto library processing was significantly more required in a big web scale company, and then they, you know, fast forward three years and it's built into a load balancer Fast forward five years you have decent crypto lib support in the CPUs and the SSL co processing market just completely disappeared over a five year horizon. I think there's always an opportunity and the it's incumbent upon the CO processing vendor to try to stay ahead of the curve of what can be done in a sort of system on chip design, coming out of Intel or coming out of one of your ampere one of the ARM library of companies as well which when you look at it delivers probably an order of magnitude more actual processors to the market than Intel does and we've certainly seen the beginning of inroads both in Amazon and now just announced this week in OCI and other cloud providers, of the ARM architecture is being offered as well it's another thing to just be aware of and pay attention to because the per thread, perform per thread execution performance at lower power draw is significantly advantaged on those risk based architectures and kind of back to Jonathan's earliest

- 6 -



point on, hey, if I want to have a highly efficient data center it's all about power, and it's how much power I can shove into an environment, how much I can pull it, and therefore the most compute I can get out of it as possible efficiency is the game and a lot of these large providers and I think you know Arm has an advantage for certain workload types there. Given the broad ecosystem of development driven out of the mobile and IoT worlds, and now with an obvious move into the compute area it's something to also pay attention to, it's happening.

Baron Fung, Research Director, Dell'Oro Group

(inaudible) just wondering from an eco system perspective, what needs to be done to come together to see more Arm adoption from the server vendors networking software and so on? We've seen some nice traction where it's coming from, AWS ampere has made some significant announcements recently. So, like to see what play.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Yeah, one thought for you Baron, is the thing that Apple developed that allowed the real time emulation and translation of Endian dependent x86 architecture code to be executed on top of the M one processor. If there was a really reliable open source project of that and, hint to Simon and Drew and all my buddies at Arm, that would be pretty valuable, because people like me who might have an NDA independent operating system with a lot of user space code would probably get a kick out of it, and might go to find some value there. Right compliant, okay, I'm sure there's others.

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

I've been really enjoyed this geek out and I'm sure the audience has, we've been talking about the what MABLEY kind of pull it to the why. Also, what I would pose for us all is to is you hear the driving of it. Driving intelligence down convergent density, etc. in the awesomeness that you just heard from the esteemed panel. We kind of pull it back to the why. I think taking into current trends, what's driving this, and what I would pose to all of us think about is, let's understand the impact of digital technology, social, mobile, analytics, cloud things to business processes across all companies, that's what's driving it. We see that as a place where Arista, Dell, all the wonderful folks that you see on IBM where technologies etc happening with enterprises and service providers. It's really what I would pose the data creation lifecycle, and the data creation lifecycle is now changing where it originates creation is happening at the edge right, it's happening in many different places by users sensors and things. The processing just to give you a kind of stats we did a study called data gravity index because we start to see a lot of customers changing how they build their infrastructure and more data centric way. Here's a couple of interesting stats that makes all that geek out very relevant. Number one is on the data creation rate currently in enterprises G 2000 Just the Forbes Global 2000 enterprises alone. They're creating data at over 600,000 gigabits per second. Today in 2021, going to more than 1.51 point 5 million gigabits per second in 2024, that's just data creation. Then let's get to where Doug and the team are heading on the processing of why you need GPUs or unique Converged Systems, why you need intelligent networking with it, is because the processing alone is about 23 exa flops, just in the G 2000 Today moving to 30 some exa flops in 2024 just the process what's being created. Then when you look at aggregation here's the standing one that people have to also understand is over 178 terabytes per



NetEvents
inter@ctive

- 7 -

second of data that's being stored in the Forbes Global 2000 enterprises, per second per second. And then that's growing to over 621 terabytes per second 2024 so I posed that in interject this just to kind of level set the why, right, there's an explosion of where data gets created needs to be processed enables intelligent workflows, and all this awesome technology you just heard, has to get placed and that's why what I saw Baron in your slides for the audience is you're talking about, hey, it's about to go through another whole wave of growth in processing, more intelligent processors, and you're going to see the hyper scalars and the large enterprises really driving that. So sorry to interrupt, Doug, but I just think that the relevant of the why, to go with the what kind of ties to this topic, I hope that helps Baron.

Baron Fung, Research Director, Dell'Oro Group

Yeah, that really helps for me, and just touching on one of the slides. We see the hyperscale companies really racing to the top to build more data centers. And we're seeing the kind of the enterprise is doing the same especially some of the I think the big financial institutions are returning to growth this year, I think seeking colocation providers like yourself as partners. So yeah, just wondering how this build out relates to the data gravity index. Yes, but how is this distributed reasonably,?

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

What we see is the shift to a data centric world here, AI here, intelligent workflows right, you hear the build to unlock insights. That's why you need all this special processing and all these technologies, and we need to network together from cordage. Think of it like a proxy of demand right for this specialized supply and that proxy demand is driven by population rights driven by connected devices, it's driven by this data creation rate, and then you just start to map that globally, what we're seeing is the largest provider of hyper scalar capacity, like we build the data centers or a lot of them for the hyper scalars, and we also provide the capacity for the large enterprises, and when you see that kind of that demand supply those two big demand entity types, coming together, what you're what you're starting to see is you're seeing accelerating growth rates, but you're also seeing accelerating growth rates that's more, let's call it fit for purpose. You know we see another wave of, and I'm sure all the, all the intelligent geeks on the call here with me, would say, high performance computing is going through its next wave probably like the third generation of it, why because you have physics limitations. So when you think of the demand of data centers, it's shifting from, you know, yes the hyper scalars and some of the biggest will build their own but they're moving to third party multi-tenant data center providers that can do hyperscale to traditional retail colo and you're seeing them where, look at this industry, it's an accelerated growth rate, it's going to how do I solve this globally and Baron I think that's where it gets really exciting because it's like kind of like the hidden infrastructure place where it all meets, and what I would say to you is that the numbers that you're projecting may actually be conservative, and that is as an in the audience as a media, it's an area to track where data centers can be a proxy for how fast digital transformation is occurring, and how data is getting utilized. Food for thought.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Tony, do you also see a change in the requirements for these data centers? Especially when it comes to moving data closer or the processing closer to where the data is created?



Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

Yes, we see both Jonathan. What you see is, you see the cores of if you go to the hyperscale core traditional where their compute model in storage is that's continuing to accelerate right like they can't build it fast enough, and they're pushing how do I drive horizontal scale in a more intelligent way faster, and then leverage technologies like you all provide, to be able to say, How can I automate it, how can I provision it, etc. So that's, that's continuing to happen. But what we're also seeing is the densification of systems in non-public cloud environments, let's say cloud adjacent sitting between, let's call it the middle between the edge, whatever that religious definition is, and the core, right, of legacy and next gen, is you see data aggregation of public private data sets you see analytics, and you see AI, turning on and what ends up happening Jonathan is that we're seeing this DPU GPU like the hot conversation just had starting to drive up density within a cabinet, starting to really run very hot, so you need structured cabling. I don't know if anybody's ever looked at the back of these cabinets, but when you look at the back of these cabinets, the amount of cables coming out of these specialized processors GPUs, etc. The structured cabling that's needed is changing why because then when you start getting into mathematical calculations. If you don't have synchronous calcs occurring, your math gets off your model doesn't learn your AI is useless. So, so that's what I see Jonathan.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

No, we seen the same thing. I like to say that at the end of the day, you know whether your data center is for a medium enterprise or Amazon. Its electricity goes in, mathematics comes out, and then the byproduct is heat, and you know the real value I think for all of our mutual customers is how good are we at reducing the heat and eliminating it so you can get more mathematics and I think that's spot on and doing it in different places these days as well.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Yeah, so if you think about the challenge that developers and enterprise have with all of these systems that listen to the conversation, it's like there's lots of different environments to think about now. We have been primarily focused a lot on bringing the public cloud experience back to the on prem environments or the private clouds, but now it's just not enough right we have to think about multi cloud and okay how developers want to develop anywhere and wherever the data is they want to compute power, Whether it's an edge or public cloud or, or even out to the branch locations. So, it creates a scenario where you need to think about unifying operations. And that's where I spend a lot of time from a networking and security standpoint. Since We've abstracted everything into or virtualized everything into kind of the processing power of the server, networking and security functions that we can even deeper and deeper as we go along, like we started out with just networking, and now we're kind of moving into the advanced security functions such as IDS and IPS and load balancing and even some of the network detection and response capabilities that we can distribute everywhere and unify across lots of different environments because it's really the operational aspects become really difficult to maintain, as these environments continue to kind of grow in various different locations.

Baron Fung, Research Director, Dell'Oro Group



Yeah, just to add on top of that, with a lot of these solutions that have more servers disorientated. I mean there's been a growing trend in which a lot of vendors are delivery kind of servers or IT as a service right in terms of a consumption-based model. So, the enterprises can help them conserve capex right which is crucial nowadays. So obviously Dell has introduced the apex, recently, just wondering how that appeals to the enterprise customers from that perspective, and if Doug and Jacob can talk about how these kind of services could interoperate with multi cloud environments that would be helpful as well.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Yeah, the way to think about it, you go back to the why. You know to Tony's why, is if you're a business, if you're an organization, and your business is data, or your business is digital, you need some kind of infrastructure, you need to invest in that infrastructure. But again, if you're a business, the accounting model it much prefers, a pay as you go kind of model right and so, the ability to be able to deliver, the exact amount of predictable, whether it's compute performance or its networking or its storage etc. on demand, I think is critical, it's the fundamentally that is the appeal of the public cloud, and yet there's going to be times when enterprises need that infrastructure closer to hand with under closer control in order to, to direct to pull that innovation, out of the system. That's really where we see the value of APEX as Dell and I'm sure you know, our competitors have similar situations, too. And it's really important that those things happen. I really appreciate that you've got Tony on the line as well because we're also seeing that growth in these in these data centers in colocation providers, pushing these facilities closer to where data is created at the edge, etc.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

You know Jonathan, I've seen two wildly different approaches to consumption now in this extremely elastic environment where it's like, I have to plan for, Black Friday in a retail environment, and I think the ability to get the on demand resources without the CapEx spend is one of the leading drivers of cloud consumption, not to mention just the inherent agility associated with the cloud environment versus hey I have to wait for my vendor to ship me servers. It's just it's a huge win. On the flip side though you've got a server storage network, not picking on anybody, have supply chain constraints in the current market. The one thing I've noticed is slightly different, it depends on the source of capital for the company. Tony and I have a shared background working together, (inaudible) banks in the past. And when your cost of capital is that of a global financial institution, sometimes it's a lot easier for them to buy and capitalize, as opposed to pay recurring OPEX costs, and so I'd say it's just a function of your cost of capital and your source of capital. But yeah, I mean the general non-financial services enterprise, I think you're spot and this is a conversation we have with most of the procurement teams, they certainly want to see as much determinism in the pricing models as possible, which does sometimes belie the cloud first message. Yep, right, like hey I'm gonna go cloud first, oh wait I didn't expect to you know my developers use it a lot faster than I expected and who's the idiot who left 100 VMs running for three days, and that happens. I'm three XLS or something, it's like, oh, that hurt. So, yeah, getting the controls aspect is something that a lot of companies don't take into account when adopting a cloud first message, and they learned the hard way that consumption controls procurement controls and some of the things that they've developed really strong operating models for on the enterprise side both from acquisitions to



whether it's ITIL or whatever on the day to day operations are still very relevant in some of these cloud environments.

Baron Fung, Research Director, Dell'Oro Group

Yeah, it's great to hear that, vendors have a wide range of innovative solutions, 250 different type customers, some providing them with sort of like cloud like experience, like on premise, which is really important nowadays to give enterprises extra flex building. So, let's move on to the next topic. Shifting gears a bit on this the second discussion topic, it's related to the proliferation of accelerated computing. I touched upon this in some of my earlier slides, particularly with coprocessor so it's like GPU FPGAs custom ASICs that are designed for more specialized workloads, and how they can become more efficient than just using general purpose CPUs and pointers, you cannot really just put in more of these accelerator servers in a data center, without really addressing the whole ecosystem right from compute, storage networks, and also enabling developers with the application frameworks to produce valid use cases, and so on. So, just want to take a few moments for each of the gentlemen, to talk about some of the potential bottlenecks right or architectural changes that are needed before we see wider adoption of accelerated computing. First of all, assuming that you know, the demand for AI applications, continues to be really strong, as well. So, any, any thoughts there on addressing system or ecosystem bottlenecks.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Again, I see a kick off out there for sport and yeah that this might be one that's designed to get Jonathan and Tony to just pile on and either argue with you which might just be fun for the entire audience who knows but I mean, let's throw a rock and see what happens, right. Yeah, I used to joke that the entire AI movement was actually built off the back of crypto mining. Everybody's like hey let's do more crypto mining, let's spill and you literally saw, AntMiner asix rolling off the leading edge of the node off of TSMC, driving the latest and, tight geometries coming out of there so you then watched a lot of GPUs get consumed and then applications get developed that utilize those I think there's always a pull/push there. You know in finding that unique application that generates enough volumes to get the chipsets made that provide these accelerations and I think you've correctly identified whether it's GPUs or AI specific processing or FPGAs are a little more accelerated but general purpose, or these GPU environments, you know, whether it's a smart NIC or whatnot that offload some of the processing from the core. There are applications and use cases for all of those today. And like we said earlier, it's a race between the creation of value in the accelerated piece and the system on chip or their core CPU being able to absorb enough of that function for it to be practical, cost effective and efficient and execution and that's the ongoing race. The other last thing I'll throw is a lot of this is always in my mind, being a networking company been somewhat networking centric I've always felt that every time the network advanced in capability, whether it's the next order of magnitude speed performance the next 50% reduction in latency, the next evolution, the ability to have guaranteed or reliable or self-correcting transport, every time that's happened you've taken one of these peripherals and you've extracted it from being bound to the CPU, and again a good example, if you look at the Grace processor architecture coming out from arm interconnected with the NVIDIA GPUs using the NV link for technology, you're seeing an example of a coherent transport being used to allow for the intelligent

- 11 -



decoupling of these outside of the sheet metal of the server, and I've kind of personally always envisioned that there's this day we get to where I have racks of memory, racks of CPU, racks of storage, racks of dpus, racks of crypto accelerate, whatever it is, and I have a capable network architecture fabric transport whatever you know term you want to throw at it, that interconnects these which allows for the dynamic assembly of resources necessary to process a given set of workloads at that time. And then most importantly the releasing of those when they're done executing so they can be repurposed and re consumed by others, whether this is literally GPUs being interconnected over IP, I don't know. But, I've kind of always envisioned that's the natural progression we're on as CPUs get faster, accelerators happen and networks become eminently more capable generation by generation.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Yeah, when I was in the 1990's we used to argue, you know, back in the good old days, so to speak but which constraint would we alleviate first bandwidth or storage? Who would have guessed that today, the answer is, you would have guessed, the power and distance are the real constraints right now. physics sucks, exactly. And I think you're calling it a lot of ways, we'll get to the future of technology in a moment, I think, as we discussed this but the advancements and the innovations around things like Gen Z Consortium, and these other alternative AI interconnects, going beyond PCI Express are critical to enabling like I said a lot of this accelerated computing because the physics of the physics of power and the physics of distance just don't lend itself to, you know keeping everything in this monolithic device and. And so, I see that, that being a big, big deal. I think the other ultra interesting part of this is that we as an industry talk about the killer application. Like what's the applications going to drive this. And the irony I think is it's actually, it's the hardware drives the applications in so many ways that the software development model has to follow what the architecture delivers to them. You can't write AI code for GPUs until there's a GPU in a device and being sold. And, now we're in this really interesting world where there's so much choice, and so much innovation that ISVs are gorgeous, they should be, you know, it should be the best, should be the funnest time in the world for them, they can they can develop anything for anything.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

I think, yeah as these constraints we saw even back in like the, like as big data was first coming about right, we needed to move the processing power to where the data is because we're not going to move it. We're seeing that kind of same thing we'll take a different view at it than just the applications, AI and ML into the infrastructure. So, we look at the same thing when we just moved towards like an MDR solution was distributed. Before we used to have to kind of tap it on, these taps on the wire and try to pull data off the wire back centrally to do security controls, but now we've taken that same notion and said okay can we do this on the server in line to when the packets are going through in in a distributed way. So, I think what I see from an AI and ML perspective is that it also gives us far greater opportunities so instead of having just these services that the trust boundaries we can do it everywhere, as we implemented our MDR solution within the actual east west traffic in a data center we can get to almost a 90%, better false positive ratio, because we see everything. You can look at every single hop along the way and not just the trust boundaries where you are, you start to just look at aggregate data



instead of actual packet by packet flows. I think it unlocks a lot of interesting opportunities from the security world as well.

Baron Fung, Research Director, Dell'Oro Group

So Jacob, do you see like kind of a smart mix sort of enabler for these accelerated computing because they have the ability to kind of look at the packets, do some automation network, and so on?

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Yeah, I think it helps it makes things better. I mean you can do it today without, but I think it unlocks other opportunities. So, what I think mentioned before, is that okay so we say we virtualized environments, but how do we extend that into a high performance database system that maybe wasn't virtualized? We have moved down the path of agents, the first, but then that shares the same memory space of the High Performance database and maybe it's built on a custom operating system that just doesn't really support it. So we can air gap it out into a smart NIC and have that same distributed functionality to those bare metal high performance systems as well. But then, if you take that one step further and this was interesting with four eight cores, but as we're at 64 cores now it becomes an interesting kind of management, and cost. Take out of like, well, yes we are taking up some CPU power from those bores to do networking and security functions so on a large scale this one server now looks like a data center of the past. So, can we start thinking about costs, as well as we move forward.

Baron Fung, Research Director, Dell'Oro Group

Sure, sure, with extra computing there's obviously a lot of implications in terms of the mechanical, thermal aspects of it. So, from Tony's perspective, what kind of solutions in terms of, you know, form factor of rack, architecture and or Innovative. Innovative. Innovative solutions like immersion or liquid cooling, are the some of the solutions you're considering implemented into your future datacenter belts?

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

Yeah, I think that's a good transition there and, you know, if you kind of pull it back of, you know causation. You know we really kind of go, hey the technology continues to innovate. Right, it's moving to more composable. If I recap kind of what I'm hearing from Doug and Jacob and team, it's more dynamic right, can be assembled, reassembled. Those environments then still tie back to what services what applications,, what use cases, which is what Doug started the brilliant evolution of this last dialogue of accelerated computing, and when you think of it accelerated computing is early innings and that early innings is because if you don't have all your data in the right place or you don't have all the applications exchanging with that data, you don't have it networked right, you don't have your securities and trolls like Jacob just said. You come back down to that data centers, you know what you do in the physical data center, almost becomes a second thought, and it needs to become part of the first thought, As you're kind of architecting and what, what we're seeing is, is that we're seeing that when you start running, whether it's short, or wide, right, analysis, and you start running that on a more constant continuous spaces. That's where you're seeing the scale up again, It's almost like remember the old scale up, scale out right team, you're kind of like seeing that scale up again and in that scale



up, it's translating to physical attributes of hyperconvergence and density of power, cooling space along with them what you're putting into the you know the modular kits and boxes and switches and so on. And you'll see that with each of the vendors that you've talked to, to hear what we're doing is we're already dealing with the fact that you have the same schizophrenia, as you cut across organizations that there's inconsistent standards and approaches. So how do you, you got to be able to accommodate those different ones so it's a very fit for purpose world, even in the physical right bridge to the digital, and in so immersion right containment, you've got different containment strategies to extrapolate heat as close as possible to where the heats being created from this infrastructure, as I talked about earlier, you have cabling Baron that is as important cable management people it's a hidden thing that's as important as the cooling along within the power distribution, and then how you tie that power distribution, not just to the rack or to the pod, but how do you time that power distribution that is sustainable, right, that is available, and you're dealing with power destroy disruptions in the grid, you're dealing with alternative energy sources coming into the buildings right it's a, it's a whole picture architecture which is what I think you heard, I've talked about in the others that you got to be thinking about it, right, so from the business to the infrastructure, right down into the data center, and then the only other part, I'd say Baron is what you've heard today is the data center is of the future is shifting to being much more right distributed, it's going to be much more decentralized, where you're going to have smaller versions of it or tailored versions based on the volume densities of workloads and data flows. And I think you got to think about it from a physical and virtual security sense because I think that's what Jake was hitting on earlier is, it's a zero trust world, right, you got to have security controls that ingress egress. You got to have it physically, you've got to have it virtually, so I think that's how to think of it Baron and that's how for the audience to think is that you got to be looking at, boy this is a very codependent business architecture technology architecture to physical deployment, architecture, and all of it has to link together.

Baron Fung, Research Director, Dell'Oro Group

Right, yeah, that's a great segway to our last topic which is what the future of data centers look like. So ,we touched upon this topic earlier about how data centers are becoming more and more decentralized. Enterprise data centers, I think, will continue to stay like these, adopt a hybrid multi cloud strategy going forward. So just wondering, as these data centers become more centralized, what's the implication for cloud enterprises, companies, in terms of, you know what changes are needed for infrastructure networking?

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMware

All right. I was gonna pile on what Tony was saying beforehand about kind of the whole world is already changed a bit. This past year from kind of the perimeters dissolving. Right, from the data center in general with everyone working from home, and so on right it's interesting. I've talked to various different security teams, they actually thought it was easier when there was a hard perimeter, even though maybe it may have been less secure, who knows. But it was easier. Their job was easier because they're like oh I had this hard perimeter and that's what I deal with on a day in and day out basis but now the perimeter has dissolved hasn't gone away it's just turned into hundreds of micro perimeters, whether it's your user connected to a SaaS application, your users connecting to an application to public cloud



NetEvents
inter@ctive

data center connecting to private cloud data centers, and so on. So there's obviously definitely a new edge, I think, Baron you showed it on one of the slides of this new edge compute that's popping up and that's the whole Secure Access Services edge market that is saying okay well if there's no perimeter, how do we combine some of these micro perimeters into maybe a pop base location, so we can start serving, not just the security functions and networking functions to interconnect everything but also when we start thinking about pushing that compute out to the edge as well from those architectures so yeah I mean data center, the future is not just one location as many locations. I think we have to get the operations right and make this successful because as much as I know Doug, you talked about the different racks of maybe, processors are in a rack of storage and a rack of whatever, if it's not in a way that that's operationally feasible then we're enterprises, probably won't adopt it in the private world as much.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

I was gonna say almost the same thing, which is the location stops being relevant right, but the location of your data storage, compute networking whatever can only stop being relevant inversely if you're centralizing something which is control. If the business doesn't feel they're in control of their data applications and what's running where and they don't feel they have visibility into it that they can trust it to Tony and your point on zero trust architectures, shout out the (inaudible) on creating that ages ago. You know what you find is if I can't trust it, I won't let it go, yet but if I can trust it, then location becomes less relevant. And I'm capable of putting compute, storage, whatever, where I need it to be to most effectively supported by end users and over the last year we have seen a shift in end user population. The whole zero trust thing, big fan of it and the idea of creating micro perimeters around the Protect surface of an asset and shifting it to as close to the asset as possible, greatly reduces the blast radius impact of change. And, the thing I like about networking is I generally have configuration is fairly for most devices somewhat immutable, I mean I make a change that I could run on it for hopefully weeks or months that I'm having to make another change quickly. Whereas in the world of load balancers and firewall devices and Security Policy and controls implementations, you find a much more mutable, highly dynamic rapidly changing set of controls based on changes in threat model changes in user location. So, the again the D parameterization and creating these micro or PICO perimeters where we want to call them, or as close to the asset as possible, allows for me to adjust those policies quickly and not have the risk associated with screwing something up and blocking my entire data center from accessing its DNS server or putting 100 gigs of traffic on a one gig link or any of the numerous examples we've all lived through of galactically cosmically stupid outages caused by human error.

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

I think Arista talks about a cognitive network. I think if I build off of that for the audience just to bring this back is the datacenter, the features both logical, physical. It's decentralized right, and everywhere, but you have, been tight like a data center become centers of data, if I recap going a different way, the network needs to be intelligent. You need VMware types to be composable, you need Apex type delivery of the kid as a service, and you got different instances tailored where, that data creation, but it goes back to why I'm pushing all of us, and pushing the audience that the digitization of processes, whether it's the business process, or what you're hearing a lot of us talking about is the infrastructure



process just to enable that business process, that's where people, the perspective I think gets lost, because they think of it on a point only, not on end to end workflows and processes of business and infrastructure, now have to be decentralized and be ubiquitous, yet have to be intelligent, on an (inaudible) basis, which is Doug say that's kind of a fair statement. Jacob does that kind of tie back to what you're both saying?

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Yeah, I mean you bring up an interesting point too, It doesn't just have to be end to end and this is what we're seeing, I mean we haven't really talked about data center, but technology a lot but it's also people in process as well. We have a lot of those conversations with many organizations around okay well how do you, especially with end user kind of being in a silo from app, an app being in a silo from others right, so I mean it. You have to start breaking down some of these silos to start to push forward with some of these architectures as well.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Yeah, it's a good reminder that that it's really all about the people, that the process is there to enable the people, the hardware, the software, we're all there to enable the people who are doing the innovating in these data centers. I would say that my top three, where's the data center future thing step one, it is autonomous, it's worth, it manages itself because we really don't want people spending time on that, on the rote work of keeping things running. I think that there's a really nuanced conversation that comes out of the distributed idea which is that when we talk about edge it's not a place in the network, the edge really is the edge of it. Right, it's where the edge of where people go to make innovation happen. And in order to manage the risk and all these other things you have this transition away from the idea of being safe and more towards the idea of being agile right ,and the people, in order to live in this zero trust environment, in order to maintain the quality of the innovation that they're doing, they're going to have to focus, It has focused on being agile, as opposed to keeping just keeping the trains running on time,

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Jonathan, kind of riffing off something you and Tony said, with the word cognitive, you with automation, I think one of the lessons we've learned is the amount of data to your point of exabytes of data getting created earlier Tony, the speed at which is smooth, the speed at which things are processed. We are currently beyond a human's ability to operate, troubleshoot, it's like, good luck. Hey, I'm going to look at, a billion packets per second of sniffer traces and figure out what's happening. Right. So it's time for us to build the machines and train the machines to look at the models we want that our network implemented on, and then ensure adherence to the model, using that to identify discrepancies and identify positive behaviors as well. I mean, I hate to go to look AI and ML is gonna solve everything but there's proven use cases of supervised learning models for automatically selecting which Runbooks were likely to correct a particular set of issues that I have experienced in a networking or full stack computing environment. And I'm not saying that it's perfect every time, but gosh even 80 to 90% accuracy on selecting, we've seen this collection of events, it might be this root cause and here's a runbook you use the last 27 times it fixed it points a new opera it gives that new operator a heck of a lot

more confidence in making that adjustment or change to keep the environment running while also targeting them at the real problem, as opposed to letting them sit there and chase noise and chase their own tails for days. So, I think this concept of systems level observability network level observability feeding into how do I monitor and manage learn from the environment I'm operating to then make relevant changes to keep that environment as operational as possible is another key area of continued innovation and investment.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Yeah, totally see the same thing from a security world to where it's you're shifting from like chasing bad to, how do I keep a yes, I guess. Yeah, right. And that's the model where I see a lot of efficiencies, zero trust as well it's like it's not about trying to kind of block people, I want to enable them to connect and keep them connected as long as possible.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Yeah, there was an interesting research paper done out of MIT that looked at analyzed all the controls positions to keep a nuclear reactor performing the way you want it to. And instead what you're actually need to monitor is when the controls are out of the positions that are known to be good and work, and you can actually apply a lot of that control theory back to say network and distributed and compute systems and say, here's what good looks like, we should really only start worrying when we start seeing things controls positioned being moved outside of this good, that's when we need to investigate and apply deep learning or ml models or forward looking observability into it.

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

There you go Baron, now that's a soundbite, augmenting, Right, if you can kind of distill it back, augmenting decisions and actions right and making it instead of chasing bad chasing good enabling new. I mean this is all good geek stuff right soundbite.

Baron Fung, Research Director, Dell'Oro Group

Right, we'll keep that in mind for sure. I know where the top dollar so now. Should we move wrap things up so. So just maybe each of you take a quick moment to give the audience's some key takeaways that will help them. So, we'll start with Doug.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

I'll tell you just one quick funny story. There was a name change that Arista went through in 2008, and the one of the key criteria of naming was it was required to start with an A so we'd be listed alphabetically ahead of our competitors now. So, we covered a broad swath of things ranging from today still the primary limiters, of data centers scale performance execution and so on seem to be more physics bound in many cases than technology bound on capacity. I think we covered a really good area on the limiting factor in addition to physics is going to be the human capacity to operate and manage at scale, which we see continuous innovation in observability machine learning and systems managing systems as opposed to people managing packets or things like that. I think that's going to be a key evolution that allows us to move to this D parameterization of the network as we increase the

- 17 -



centralization of the control to enable the decentralization of the asset which gives the enterprise operators the comfort necessary to be confident that data and applications can be processed remotely, and wherever they need to be to offer the best end user operator experience without compromising the integrity of the enterprise or integrity of the environment. I think that's one of the big ones. (inaudible) the whole network evolves things, deconstructors, I wrote a paper on that, in 2005, and it still seems to be happening so cool. I'm sure it will reach a point where it stops happening, or the thing that has to happen becomes back to physics limited and the distance of a one meter cable is a problem, we will get there. And you know when that happens, you'll see this interesting shift as long as processors continue to evolve to more things being shoved into system on chip designs because that's as close to the actual processing as you can get, we've always seen that with level one, level two caches, we're now beginning to see, like we said with the ice lake processors with, you know, network function enhancements in the form of crypto library support quad and things like that being driven further into the CPU itself and then like you said the AI ml evolution is driving one of the primary accelerator implementations in the form of GPU adoption and dedicated co processing adoption for that as well as we see smart NICs and dp is starting to appear but I think there is going to be a constant push pull between the CPU manufacturers and that, especially as we see the ARM processors continue to take off, because a lot of those are designed with some of those SOC designs that implement those functions natively like we saw in the Annapurna labs implementation at Amazon, Jonathan.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Yeah, so I'm gonna double down on the on the people idea. Pat Gelsinger this year said that sometimes it can take 10 years. To move a week, and sometimes it seems like in a week we can advance a decade, and we just have come out of one of the biggest changes in our collective industries that you will probably ever see. Within a week. It stood up and delivered in ways that all across organizations that we would never have thought possible. And you know it's a huge kudos to it out there and that the future of the Data Center really is tied back to what these organizations are looking for. I believe that that they're going to pivot from what they from the reaction to this past year toward the toward innovation and transformational activities, they need a platform, and an environment that they can innovate on an innovation engine, if you will, I think that it'll be one that is focused on, making their lives easier, through the autonomous operations, but also being able to deliver on this idea of enabling him to be agile and create innovations, regardless of where the data is, but also regardless of multi cloud choice model that they have built.

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

Yep, if I kind of build off of that. So, if I sum up and takeaways for everyone. We talked about the why being a perfect data storm right, we talked about in shared that it's really data centers, two centers of data that are more decentralized and more ubiquitous. It's because the physical and digital divide, right needs to converge and that's because of physics, I would throw in compliance and regulations are also going and globalization, that I would tie with it. And then we talked about the innovation and capabilities of all these new accelerated technologies are going to unlock, when you think of it on an end to end basis and remember our quotable soundbite augmenting decisions and actions to chase the good instead of chasing the bad. That's what I'd leave you with.



Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Yeah, so, we had been focusing a lot on the public cloud experience on private cloud and now we're moving to a ubiquitous multi cloud experience, which includes edge and everything in between. And what I'm really excited about is this is unlocked. A lot of great technologies that we can do more in a distributed fashion software, right, like, whether it's processor capabilities or smart NICs or anything in between, or moving up the stack right especially from a networking and security standpoint, like we started, and a down packet by packet, looking at how do we recreate networking, from a virtualized world. And now we're all the way up into side cars and Kubernetes land right as we think about service meshes and API calls so like we're networking, security is kind of distributed now, in a way that we can start kind of thinking about higher level technologies and security initiatives. So for me it's like I want to get applications in production faster, wherever they may be, where I want to enable users to connect to those applications and keep connected on any device anywhere and any location. For me it's definitely about the outlines.

Baron Fung, Research Director, Dell'Oro Group

Thanks, gentlemen, so I guess we'll move on to the next section. Q&A

George Rickman, NetEvents

So George Rickman here from NetEvents. I just wanted to say before we start the media Q&A If you are a member of the media. Just a reminder, feel free to raise your hand and I'll unmute you and you can ask your question that way. Or if you don't want to ask verbally, feel free to use the Q and A function on your menu bar, and also just specify who you're directing the question to that'd be great. So let's move on to the Q&A I do actually have a journalist, Hector Pizarro from Diario TI. He's actually asked to speak so I'm just going to unmute his microphone and I'll let him ask this question. Hi Hector, you should be able to unmute yourself now. There we go.

Hector Pizarro, Editor in Chief, Diario TI

Thank you George.

What is your opinion about the increasing need for electrical power for data centers in the coming years, that is resulting from the growing computing demand versus the expectations around the green and renewable energy sources?

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

Certainly, Hector. Great question, thank you for engaging. The answer is the demand is going to continue to increase. The challenge on all of us in the industry is how do we drive efficiency per watt per square foot and do it in a sustainable way. That's why you see folks like ourselves and the other major players, I'm driving a significant portion of our capital raise for our businesses to build data centers and infrastructure, under green bonds, So I would track the financing of this industry, I would track also how you take and utilize renewables in alternative energies, but do it where you're also dealing with the batteries and the storage of the power, because you have to build a deal with various scenarios various consumption levels right it's not just a state a steady state, even streamline and then you have to deal with disasters as you track, like we had a very interesting one over here in Texas where I'm based



NetEvents
inter@ctive

- 19 -

recently, where we had records of freeze temperature etc. that was never dealt with before, so I would pose that it's a big issue. The demand is going to continue to increase, innovations are occurring, and it's going to be something that we all have to track and make sure that we create a future for our kids and so on.

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

Kind of piling on Tony, one trend I've seen happen, as well as with advances in WAN technology, the ability to go 100 gig 400 Gig over a long haul and the wide area network. We've seen large smart enterprises, say hey you know, I'm designing my next generation data centers, I want to co locate those data centers, next to their most critical asset which is green renewably generated power so you see the Columbia River Gorge, you see, just outside Las Vegas or whatever the location is, but the idea being, get close to cheap renewable energy sources, you then see it's something different we used to always embed these meet me rooms and big telco facilities into those data centers, we might only see two to four links coming out of the data center, it goes back to a transit hub, which is located where the major fiber paths across the US intersect. So, you see transit hubs appearing in Atlanta, New York, Chicago, Dallas Vegas whatever, and the data centers are located in your power, the telco communication rooms are being built near where the facilities are and your long haul back hauling them together. This allows the location of the asset close to its most critical resource whether that resource is connectivity, or that resources power because the two are not the same in most environments.

George Rickman, NetEvents

Great, thank you. So I've got another question from Guy Matthews who is an editor at AI Business. He wants to actually verbally ask this question so I'm just going to unmute Guy. Guy you should be able to speak now.

Guy Matthews, Editor, AI Business

Yeah, it's my feeling that the typical data center of 10 years from today is going to be more or less unrecognizable to us, we get sort of traveled through time. I don't know if the panel agrees with us, If so, in which case, in what fundamental ways will the data center change?

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

I'll take the first cut at that Guy. I think you're right, time dimensions will change. I would say go back up a level and look at where is the digitization occurring of physical buildings physical locations cars etc. So the operational technology, and then the information technology, and then you heard what Doug was saying, is you're going to see, you know, you could do processing, you could do a transit or connectivity, you can do by power source you can also do by where data needs to be aggregated so if you look at it from a workflow lens, you kind of go, there's going to be a data center in a box that's going to be in a building or tower. As soon as that starts to evolve and the world goes to 5G and things become more connected there's going to be a data center within most major metros, or if you follow kind of the study of urbanization and you look at like there's a book by Parag Khan on Connect Geography, he's talking about these mega cities. So within the mega cities which would be like a San Francisco to Los Angeles as an example, you would probably see a mesh of data centers or a

- 20 -



collection of data centers, because you're going to deal still with that data aggregation massive amounts are that needs to be enriched and shared, and it kind of goes back to that physics thing. So physics will still dictate density will dictate, and then you have the extreme edges of either I'm trying to get from region to region, you know through submarine cables or I'm trying to process intelligently in a building as an example. I think I would pose to look at it that way Guy. And understand that it's ubiquitous, it's distributed, It's decentralized and localized, and it's going to be very fit for purpose across the value chain.

Jonathan Seckler, Senior Director, ISG Solutions, Servers & Partners, Dell Technologies

Thanks. All that being said, I think that the data center of 10 years from now will still be recognizable as the data center, to Tony's point, there's this physics and power required, will the data centers be bigger or larger? The answer is yes, right. Will they will they be in remote locations or centralized locations? The answer will be yes. Will they be at the edge or are centralized? The answer will be yes. I think you'll still see all of those trends continue, but you'll see more of that fit for purpose.

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

I will add to that, you're going to see it in a server closet. Server closet been around a long time, right, like trading floors right they've been in the retail store, they've been in the warehouse, the manufacturing floor, they're not going away, they're going to evolve, right and then it needs to be why you mean intelligent management and software and networking, you know, to bring it all together.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

I think as we move these forwards if the data center still looks like a data center, great. In the end, I'm hoping that the user experiences has improved. Developers are able to develop wherever they want, however they want with the tools that they need at the time, be able to push to production, quickly without a bunch of roadblocks in the way. But then security teams having the tools distributed and automation in place all at once and users being able to work from wherever however with whatever devices so as long as we continue to push those forward I think we're going to be in a better place.

Baron Fung, Research Director, Dell'Oro Group

Yeah, I think we'll see more new classes of applications driving those data centers like autonomous driving, you know, real time sensing industrial automation, potentially, we're gonna see a wider range variety of data centers large to small different form factors, different designs, a greater variety of what we're seeing today in this mega data centers. So, I think it will stay tuned for that.

George Rickman, NetEvents

So I've got one final question, this will be our last question as we've run over a little bit. It's from Steve Broadhead, who writes for Computer Weekly, and he says, forecasting the future in IT has never been easy, but the events of the past 12 months seem to have made a hard job even harder. Notable, especially in the early days of the pandemic hitting and the WFH initiative being rapidly enforced, were some of the big name casualties who suffered downtime – for example, MS Teams, Citrix and Zscaler as ones reported to me by family and friends – i.e. real users! Having been involved in product testing all

- 21 -



the technology within a DC – and all the tech that delivers the data and applications to the endpoint – for decades now, I know that technology IS capable of standing up to extreme spikes and ramp-up in usage, so what is going wrong and why?

Doug Gourlay, Vice President/General Manager, Cloud Networking Software, Arista

So, I'll kick a small comment on this off because it's just sporting some of the most complicated systems designed in the world our, you know, multi billion transistor, semiconductors and multi 100 million line of code, operating systems and distributed systems and applications. We've developed models that make them reliable. The ADA tooling that you use in semiconductor development, the CI CD test driven development environments that we use in software development, that allow us to test before things are deployed test after things are deployed, validate that every single line of code does what it's supposed to do, that we have input validation and so on in these systems. We're watching an evolution in networking and infrastructure management that goes under a variety of monitors. Networking is code infrastructure is code net DevOps, whatever you want to call it, the reality of these is we're saying, what worked in this SDLC the software development lifecycle model that we can apply to infrastructure, how I mean, when's the last time you heard of a networking person testing a configuration before it was deployed outside of, I stood up a little mock-up of a lab and I let my buddy look at it. The majority of the changes were fingers on keyboards affecting changes in production environments. And this is common across, I'd say greater than 90% of the Fortune 500 systems today. And what we're seeing is a sea change evolution of that where people are coming announcing I went network automation or I went systems automation what they're really saying if you unpack that a bit more I want repeatable outcomes. I want programmatic inputs, I want a source of truth. I want a self-documenting system. I want something to test before it rolls out like I do it every single other production system on the planet. And I need to now do that in my infrastructure, whether it's server network system security firewall, whatever. And every single change touches so many systems that building these model based environments is incredibly complicated. But I think this is one of the big evolutions we're seeing not just because of the change that is, you know, was pushed because of the pandemic environment which shifted users to new locations and might have reduced some of the operations teams abilities to sit on top of the system, but I think it started, You know, even further back where enterprises are saying how do I keep up with the cloud, how to deliver the similar levels of reliability, how do I deliver similar, similar levels of agility, how do I help my business have more deterministic costing and Bytown the risk and the, this is requiring fundamental changes in the way we not just design, build, but most importantly operate the infrastructures we offer to our users and our clients on behalf of the businesses we work in.

Jacob Rapp, Director and Lead Technologist for the Networking & Security Office of the CTO, VMWare

Yeah, totally, I think, to echo your point. So many times I go ask okay where's your network diagram or what is this, how is this connected over here? Okay, say well let me open my spreadsheet, or let me pull up my wiki page right, those things are out of date the second you add information to it. So yeah, as soon as we start automating it to infrastructures code or whatever buzzword, you want to use as part of it we can kind of create those meaningful processes that can be repeatable, but in the end, you're only as good as your high availability scheme. Right so outages happen errors happen code, bugs, take place. So how do we create an infrastructure that is based off of maybe an SLO, and not just an SLO

- 22 -



from social service level objective from a networking standpoint uptime point standpoint but also from a security standpoint. So, I can tie back in business criticality for what I'm running into my decisions is like what part of the application is crucial, right, and how do I make sure that that is, is up no matter what. And the correct processes and procedures and tooling are put in place.

Tony Bishop, Senior Vice President, Platform, Growth & Marketing, Digital Realty

If you sum it back with Doug and Jacob, would you say that the part of the answer to the question is it's the shift to industrialized execution of workflows, right across systems, kind of like what you used I think Doug earlier the nuclear reactor if you think of what that industrialized controls are now you're talking about doing that on a very broad basis, And that's what's keeping it from getting there, and then kind of maybe a proxy of demand if I go back to that part of the question is, you know, use data and the growth of data and the exchange of data as a proxy of where demand is going to go, and then measure the industrialized quality of execution across industries across companies and then you can kind of decide what's going to happen where, and what's limiting those firms or not, maybe that's a way to think about it.

George Rickman, NetEvents

Great, thanks. So, Baron it's over to you, that was the last question. I'll leave you to say thanks to the panel and then we'll pass over to Mark.

Baron Fung, Research Director, Dell'Oro Group

Yeah, great. Yeah, thanks panelists for a very lively discussion debate. Hopefully the audience took a lot of valuable information out of this. And thanks audience for attending, thanks NetEvents for organizing this great event.

Mark Fox, NetEvents

Yeah, thanks everybody. Thanks, Baron for chairing the great session. And thanks to the panel. Nice lively session. Sorry about running over, but the content was great, so thanks, everybody, and look forward to welcoming you to the next event. Cheers for now. Thank you.

