Intel

Edge Infrastructure Orchestration and Management

CORPORATE PARTICIPANTS

Alice Puckett
Moderator

Hein Matthee
NearbyComputing SL – Senior Pre-Sales Engineer

PRESENTATION

Alice Puckett

Welcome, everyone, to the Intel Network Builders Webinar Program. Thank you for taking the time to join us today for a presentation titled “Edge Infrastructure Orchestration and Management”.

Before we get started, I want to point out some of the features of the BrightTALK tool that may improve your experience. There's a Questions tab below your viewer. I encourage our live audience to please ask questions at any time. Our presenters will hold answering them until the end of the presentation. Below your viewing screen, you will also find an Attachments tab with additional documentation, reference materials, and websites and documents mentioned in the presentation. Finally, at the end of the presentation, please take the time to provide feedback using the Ratings tab. We value your thoughts, and we'll use the information to improve our future webinars.

Intel Network Builders Webinar Series takes place live twice a month, so check the channel to see what is upcoming and access our growing library of recorded content. In addition to the resources you see here from our partners, we also offer comprehensive NFV and SDN training through the Intel Network Builders University. You can find a link to this program in the Attachments tab, as well as a link to the Intel Network Builders newsletter.

Today we're pleased to welcome Hein Matthee from NearbyComputing. Hein has 21 years of experience in the IT field, focusing on technical pre-sales and solutions architecture, in both Software as a Service as well as edge orchestration. There was also a three-year stretch of working with Intel, helping customers leverage Intel technology. He has extensive history in working with Infrastructure as a Service solutions, and is now applying this expertise to help companies realize the power of edge computing, leveraging infrastructure to solve challenges in unique ways.

Welcome, Hein, and thank you for your time to join us today.

Hein Matthee

Thanks, Alice. Hi, everyone, welcome to this webinar. Let us get started.

So, the purpose of this webinar is going to be talking about edge orchestration and infrastructure management, but the really interesting part is that there might be a few questions here, and my first question is, what does infrastructure management have to do with edge orchestration, and how do you see edge orchestration? And I will hopefully endeavor to answer most of these questions as we go through, and hopefully, I might clear up some misunderstandings about specifically what the edge is, because the reality is that when you start talking about infrastructure and the edge, they are very closely intertwined.

But before I continue on with that, let me talk a little bit about NearbyComputing, and who we are, and why we feel that we are in a good place to talk a little bit about infrastructure management when it comes to the edge. So, a little history about us.

So, we were incorporated in 2018 as a spinoff of the Barcelona Supercomputing Center. Most of you might know of it, a leading research center here in Europe, definitely the prettiest data center when it comes to supercomputing centers. We're also-- NearbyOne is a product that we position our orchestrator, and it's based on our IP. Some of the important shareholders about our solution, as well as partners, I think this is important to share, is from a strategic perspective. So, we obviously have Cellnex, which is one of the largest telcos in Europe;

Lenovo, which I think needs no introduction, as well as Red Eléctrica de España, which is one of the-- which is the sole transmission network for electricity here in Spain. And then obviously backed by some research institutions, and very importantly, we have a longstanding and very fruitful collaboration, obviously, with Lenovo, with some SIs like Minsait at Indra, and then also, obviously, Intel, and I have to again thank the Intel team for the opportunity to show this webinar. We've done extensive work, and we're part of the Intel Network Builders program, and hopefully we can expand that even further. But this is very briefly about us.

What's the objectives of this webinar? And I will also state that I understand everybody's time is short, so I'll try not to do a death by presentation and hopefully we can get these concepts through quite quickly. But the two main objectives here is understanding of basics of edge computing. How does it relate to infrastructure? Secondly, how does NearbyOne help when it comes to challenges around provisioning and managing? And to understand these specific things, we need to move away from infrastructure a little bit, and we need to move to applications. And you might say, “But, hold on, we want to talk about infrastructure. Why are we talking about applications?”

And the real reason is that the reason why the edge is such an innovative concept and why everybody's talking about the edge is the fact that you can start leveraging applications. So, even though you have a potential requirement for managing the infrastructure when it comes to disparate sites or remote sites, you might still have an end goal of leveraging applications because that's the main reason why the concept of edge was born. That being said, the edge is also not necessarily a new concept.

For those of us that might have been around for the last 20 years, we know that this whole concept of on-site or on-premise, moving into the cloud, and then moving back again, it’s a never-ending cycle. What makes the edge computing and the concept of the edge a little bit different is because of all the innovative concepts and applications that you can run, and these are some of the examples. We're talking about traffic control, connected vehicle, mission critical sensor-based analytics, video analytics around security, private connectivity. I mean, everybody's talking about 5G, and 5G is one of the great catalysts for positioning the edge, but there's also obviously the ability, content delivery. I mean, we all consume, especially in the current pandemic times, a lot of us have consumed content like Netflix, etc. So, content delivery is growing strength to strength purely because of-- and this is another potential concept of being able to get your content at the edge, which means that you could consume it easier. And then obviously people detection and counting. And lastly, but not least, online gaming, which is an interesting concept when we start to talk about the edge cloud, the telco cloud specifically.

Now, here's a real question. How do users consume these applications? Now, some of these questions might seem very obvious, because we live in a cloud world, and we're used to leveraging the cloud. So, the first one is through the network seamlessly. The second one is QoS, or quality of service, or SLA-based, so it’s about reliability, security, being responsive. And then as a service, Software as a Service, subscription-based.

Now, edge computing is important because it allows management of IT systems and access networks solving a complex set of challenges so that users can enjoy edge applications like cloud applications. And apologies for the slide. I see it's a little bit messed up, but the concept here is that edge applications can be consumed the same way that you consume cloud applications, because I think that's the new normal. Everybody is used to consuming data at a cloud level. I mean, we all have cloud accounts. Everybody uses Office 365, Google Cloud, etc. The idea being that you want to be able to leverage the resources with the benefit of edge computing, and we'll get to what those benefits are, but consume it in a similar way to how you would consume cloud applications.

Now, this is where it becomes interesting. What is the edge? Up to now I've been talking about everything other than infrastructure. So, hopefully, this is going to be the first point where we start talking about infrastructure specifically. But before we do that, let's talk about the three pillars. We say that edge computing cannot be defined by a single point of view. We say that there's three very important pillars that all form part of the edge, and depending on who you talk to, they’ll focus on one of the three pillars.

So, the first one is around virtual network functions or containerized network functions. Now, you might wonder what are these? These are classified like SD-WAN, SDN firewalls, and then the one that's the most obvious, LTE and 5G, whether it's a core, entire 5G core, that you can run at the edge, and provide connectivity in the remote site, or even more powerful, the ability to run just the user plane on-site. So, for example, providing extended connectivity and extending your 5G network by way of leveraging edge sites. So, that's the first one, so both running virtual machines as well as containers, that would be the first pillar around the edge, and for some people, that is all that

they see, especially nowadays, when you talk about 5G, and you talk about edge, a lot of people are immediately jumping to private 5G, and that's a valid use case, and that's a very powerful use case.

The second one is application and service operations, I already briefly spoke about that, and we say that edge computing is driven by applications. Video analytics, augmented reality, virtual reality, asset tracking, IoT analytics, all these functions, you can now actually do at the edge with very low latency, and we'll talk about that a little bit more as we go on.

And then the important one that we want to talk about in a little more detail today, which is around infrastructure operation. So, we're saying that edge computing must be hardware agnostic, but it also requires infrastructure awareness, and we're saying it's a little bit more than that. We're saying that you need to be able to provision and manage that infrastructure, and when you talk about infrastructure, we'll go into what the definition of that infrastructure is, because it can be a lot of different things. It doesn't necessarily have to be a customer premise equipment, or user equipment, or end device. It could be a server, it could be a lot of different things, and we'll expand that a little bit.

So, we're saying that our solution, NearbyOne, and unfortunately, I see that the slides are not showing very clearly, but I will talk you through it so that everybody can be clear around what the specifics are. So, I'm going to focus purely on the infrastructure side. So, the first important pillar around infrastructure is around unattended lifecycle management.

That is near zero touch provisioning, that is around the ability to manage cloudlets, and cloudlets would be like, for example... For example, if you... I just want to try something here. There it is, cool. It was just me not clicking fast enough. So, near zero touch provisioning, you also have the ability to leverage specific cloudlets. So, for example, you can have a private cloud, a customer might have provisioned his own private cloud, and you might want to manage an aspect of that private cloud, and that aspect of a private cloud could also be you wanted to manage a service in Azure, or you might want to manage a service sitting in AWS, which is obviously quite commonly used. So, we're saying that when you have infrastructure, from an orchestration perspective, it shouldn't just be the edge service, but you need to be able to extend it to various cloud services as well, and the near zero touch provisioning is important because that is, how do you deploy devices?

Now, when you start talking about remote sites, I'm going to use the example and the vertical of oil and gas or mining. So, now I've got the situation of I sit with a site that, let's say, is 200 kilometers away from everything else, it takes me a bit of time to get there. Now, I need to get a server on-site, I need to provision it, I need to get everything up and running. The only thing I know about that site is that it's got some connectivity that's decent enough. Wouldn't it be great if you've got the ability to actually ship that server completely bare metal on-site and have the ability to provision the services that you want on-site without any intervention, and that's the concept of near zero touch provisioning. The near zero touch comes down to the fact that we preload that server with a small little ISO that allows the server to connect to the controller that can sit anywhere on the cloud, and we can then provision whichever services we want on that device, and it could be pure infrastructure services, but it could also expand to services such as VNFs or CNFs like firewalls or SD-WANs, etc, and it can even extend to services such as applications, whether it's IoT analytics, etc. So, all these functionalities can be leveraged from the infrastructure operation. So, that's the first one, lifecycle management and being unattended.

The second one is multi-site management and I briefly spoke about this, but the fact is that you need to have the ability to manage on-site, on-prem sites. So, for example, devices sitting inside your environment remotely, that 200-kilometer site I was talking about, and the second one is the ability to manage cloud sites. So, again, having a single interface that can talk to both Azure, or please note not Azure specifically, but to your own site inside Azure, because we're talking inside those specific sites, as well as private clouds, as well as physical hardware, as well as virtual machines that can be provisioned on some OpenStack environment somewhere, we have the ability to give you a single pane of glass view of all of that.

And then thirdly, virtual as well as physical devices, x86 and obviously expands more than that. We can do lots of different devices. We can do IP cameras; we can do all kinds of different devices that we can manage. So, that's the concept of infrastructure operation.

Now, there are others, and I'm not going to spend too much time on these because the main focus is around infrastructure, but these are, for example, VNF/CNF. It's important to note that the NearbyOne solution is fully containerized, and cloud native, which means that we leverage the power of Kubernetes to actually do a lot of this functionality. That allows us to do containerized functions, as well as virtual functions. It also allows us to do traffic steering, and quality of service, and then the telco function interplay is where we start looking at being able to orchestrate 5G cores, we've worked with quite a few different partners, we're completely agnostic, as well as the ability to manage SD-WANs, and then the application side is around bootstrapping, health, leveraging specific devices like SR-OV, the ability to look at what's on a physical device from a NIC or a network card specification, and being able to provision activity applications, the ability to use accelerators such as Movidius from Intel, or even others, resource-driven, location-driven, being advised wherever the best location is to put a specific node, and the ability to do KPIs and performance models through automation, automated control loops.

But as you can see, this creates a holistic approach where even if you focus on the infrastructure, it's already important, and the idea being that we're saying is that if you manage an infrastructure, and then you slowly want to decide maybe I can actually get more out of this infrastructure than just managing the infrastructure, which is already important, then we can actually start extending it to running specific virtual network functions, and then further, starting to run specific applications that give you the full power of the edge.

Cool. Now, this is all done through a single pane of glass, and that's a drag and drop approach, and it's important. So, you can see a geographic layout of where every site is. We even have virtual probes that we distribute when we roll out our solution so that you can determine latency between various nodes, because this is important, especially if you start running specific applications. And then also, again, when it comes to infrastructure, the ability to see some telemetry about the servers themselves, what's the performance of the memory, what's the utilization, is the server running hot, is the server running cold, all these types of things, and this can all be given to you either via APIs or a single pane of glass.

Now, let's back the edge a little bit. Now, the edge is different things to different people, but the one constant when it comes to the edge is that wherever you are in the edge, there are devices being deployed. So, when it comes to the far edge, you've got specific end devices or sensors, and that's mainly what you find here are applications, devices, compute storage and data processing. On the premise edge, you've got hubs and gateways. On-premise infrastructure, it could be UCPEs, universal customer premise equipment. Then you get the access edge. All of this is obviously classified as the near edge, so near to the network. So, access edge, network access points, self-sites, access points of presence. You'll see at the bottom; we talk about the type of connectivity. So, for the far edge we normally talk about wireless, so Wi-Fi and things like that. For the near edge, you start talking about 5G. Potentially, you've got an environment where you can run fiber, etc. And then a concept that we started to explore a little bit is around the metro age, which is we're finding with smart cities, a lot of cities are starting to put down regional aggregation points, where you might have a network point of presence. You start following a bit of a hub-and-spoke model where you've got lots of remote sites connecting into a central site that then connects back into a more regional site, and I suspect that the architecture of the edge is going to evolve as more functionalities become available. This all connects into the cloud and data centers, and it's important to note that I suspect, in the foreseeable future, this is going to have another split, because you're going to start finding edge data centers and edge clouds, where you actually have a data center that has more edge than a proper data center connecting to a more core data center, but these are just-- as you can see, the one constant across all these various sites is the fact that there's hardware sitting at all these different sites, and the hardware at all these various sites going from a far edge to the core needs to be managed, and if you've got the ability to manage it through a single interface, so much better.

Now, why do you want to go to the edge, and this sounds like the unspoken question. We talk about how great the edge is, but why. So, when we started talking about the edge, the thing is that technology improved to the extent that you could have things like the Internet of Things, which everybody talks about, IoT. But the problem is that when you start talking about that, you need some connectivity, and not just connectivity, you need some processing, and we quickly realized if we start doing that processing in the cloud, there's certain challenges with that. The fact is the cloud might be very far away from where you're doing the processing, which means with every hop to connect to the cloud, you add more latency to it.

So, when a computation happens in the cloud, you might have a scenario where you might lose packets, or potential loss of data, longer outage time for you to actually get the data back on-site, so when it comes to data recovery, it becomes a challenge, data vulnerability, and very importantly, latency. So, the fact is that if I ask a question, I do some processing in the cloud, it takes me a certain amount of milliseconds, to potentially even seconds, to actually get a response back, and that's not ideal for some functions. So, with a cloud environment and us relying on the cloud, there were certain functionalities that we could leverage purely because of the fact that there was just too much latency. So, video analytics was a really interesting use case where you couldn't get facial recognition in time because maybe people would move along, or you just couldn't capture the face fast enough because of the fact that it's being processed in the cloud. That's one.

The other side when you talk about pure infrastructure is if you have infrastructure that relies on very low latency, so for example, I have a server that powers a transmission network, as an example, and I need to pick up if there's any issues like fluctuations in their transmission network, because if that's the case, I need to be able to shut that system down. Then suddenly, latency becomes a massive issue, because that could be the difference between a safe network or a safe transmission network and a potential surge that damages equipment, and it's just a random example.

Now, when we started moving to the edge, you started to find the fact that you could leverage... instead of having to rely on the cloud data set-- and you can still aggregate the data in the cloud, there's nothing stopping you, but the idea being is if you've got a server sitting very close to where you need to do the compute, you could actually leverage that server to do the compute locally, or to even action things like, for example, I need to shut service down very quickly, which allows us to do very low latency, the computation happens at the source data, low latency, reduced bandwidth, better user experience, less data transmitted, which might not be an issue for some, but in some regions, it still is, much higher security, no chance of vulnerabilities, and at the end of the day, faster response times. So, that's one of the reasons why we moved from the cloud to the edge.

So, to just change pace a little bit. So, end-to-end cross-domain automates specific orchestration, this just talks about what we see as-- so we call every-- every device or environment that we manage, we encapsulate it in a block. So, you'll see that we referenced there a lot of blocks. And what we do is, we have different blocks, like public cloud blocks, that's an infrastructure, that's a specific virtual machine or a pod or a container inside AWS, Google Cloud, etc, and we've got the ability to manage and interface with that. You could have servers or equipment sitting out in the back office in your environment. We have the ability-- might be a private cloud, we've got the ability to interface with that, and manage that environment, whether there's virtual functions running on that, or whether it's literally just making sure those servers are running fine, and it doesn't fall over.

Then MEC blocks, and this is when we start talking about cameras and routers and devices. Let's be honest, these devices, that's still infrastructure, you still need to be able to manage it. It might not be the traditional thing about-- it's a server, but it's still a device that needs to be managed, and what we do is we create an abstract layer, present that in our solution so we have the ability to manipulate it, manage it, and shut it down. So, that's a very important understanding of all these various devices that come with mobile edge or multi-access edge computing.

Then, obviously, the edge nodes themselves. So, you can have servers sitting in all your remote offices, or mining sites, or oil and gas sites, routers, etc, to be fair, any device that has the ability to communicate via NETCONF, we have the ability to talk to, and then various different IoT devices.

So, as you can see, the scope of orchestration, when it starts looking at infrastructure, is quite wide, and then, very importantly, on top of this, as you can see, there is extra functionality that you can run.

How do you onboard edge nodes? So, we use some very interesting technology. There's some interesting technologies like Redfish, etc, but I'm not going to go into the specific details, and you're more than welcome to contact us, and we're happy to talk you through it and see if we can help you. But I'm going to use Lenovo as an example, because we've done some extensive work with Lenovo on this. So, it doesn't mean that we have to just use Lenovo, but we are very comfortable with leveraging Lenovo, and we've done this for quite a few customers and sites.

So, the idea being is that this whole near zero touch provisioning that I spoke about is where you can actually take a server, a small little ISO that's shipped out with the server, we can take this server, we can take it out on-site, and we can start the server up as long as there's some form of connectivity, and through a single interface, we can connect to this, we can decide what services we want to drop or literally just making sure the server is visible so we can provision a server correctly. We can get platform telemetry, chassis telemetry, network telemetry, make sure that all the services are running perfectly fine. Through our orchestration platform, we make sure that it's registered in inventory, and we have the ability to manage it, and we've got a little database that allows us to track where the services are.

Now, we support lots of different devices, but specifically with regards to Intel, we obviously support not just Xeon D and 3rd Generation Xeon scalable. We also support Atom devices, i3s, i5s, i7s. We support both Intel X710 Fortville NICs, as well as the new generation of Intel NIC. So, we support pretty much all the Intel architecture. We support other architecture as well, but specifically for purposes of this discussion, we're very comfortable working with Intel. We've also done some integration with Intel. It used to be called Open SSL, Intel Smart Edge Open. So, we have the ability to leverage some of those services that provides extra performance when it comes to managing remote sites, when it comes to specific applications.

So, as you can see, that's just the infrastructure, and then on top of that, you've got the ability to start accessing a catalog that we have onboarded of functions and applications if you want to run that on top of it. So, if your only requirement is to run that infrastructure, make sure that it runs fine. We can do it, and then on top of that, we have the ability that if you decide, hold on, I need to start leveraging some functionality here, I can actually start running up virtual functions and applications as well.

Now, who are the key players when it comes to edge computing? So, we all know at the end of it, it's the end user and you've got a customer. And then you've got service providers and you've got system integrators, and system integrators, I think, would add the most value because system integrators normally understand both the enterprise space, and they also understand the potential telco space because as an orchestrator playing in the edge, we realize that we have to play in both sides. So, we sit at the other side of this, so as a standard, our software, pure software, we sit and the system integrator and the service provider can leverage it.

Then you've got some optional network functions or edge applications, or optional MEC functions, or even 5G or UPF functions or LTE or cores that can be leveraged, and then very importantly, the virtual infrastructure management or the controller's infrastructure management, which is we run on top of various different Kubernetes platforms, but we provide much more. So, the Kubernetes platform such as your OpenShift, or your Rancher, or your Vanilla Kubernetes, it supports a certain function. We extend that function even further to the hardware layer where we can actually reboot servers, we can connect into specific devices, we can start them up, we can shut them down, etc. And obviously, optionally, you can actually even manage your connectivity, your radio systems, and then very importantly, the specific infrastructure, whether it's routers, Cisco. We have discussions with quite a few of our vendors, whether it's firewall devices, everything else, we can manage both the virtual infrastructure as well as the physical infrastructure.

Now, what are the value propositions of leveraging NearbyOne? So, the first one is unattended provisioning and monitoring of remote sites. It's an obvious one, and what does it do for you? Improving operations to ensure operational excellence. The fact is, if you don't have to send people on-site, especially if it's a remote site, you can actually get the sites up and running a lot faster. You can leverage your resources much better.

Reduced time to deployment, fast operational efficiency, and increased return of investment. The same principle. Shipper's server, you don't have to prep it beforehand, you can ship it on-site, you can decide what services you want to roll out, and you get the server up and running much faster.

Seamless and transparent management of edge infrastructure. That allows you for a reduction in cost spent on management resources.

Easy KPI and operational metrics allows for early warning, ensures minimal downtime, and that's, at the end of the day, the crux of it. If we have the ability to detect potential issues, or see signs of issues, we can actually do something about it, which means that we reduce downtime, which is for customers absolutely perfect.

Centralized management through a single pane of glass that allows for scalability, and that allows for reduction of time and cost.

And then application ecosystem for added value around Computer Vision, AR, and others, and it's optimized efficiency and performance at edge locations.

I want to go back to the second last one, which I think we need to just talk about very briefly, and that's a centralized management. But the fact that you can manage a very heterogeneous environment through a single pane of glass, wherever it is in Azure, Amazon, whether it is sitting in a private cloud, or it's virtual, physical, all through a single interface, that allows your customers so much more flexibility in regards to being able not just to manage sites, but potentially even migrate sites.

Cool, and then let's talk about some of the use cases. I'm just going to very briefly talk about them. Maybe this is just to get the audience thinking a little bit about some of the potential applications. So, the first one is traditional IoT networks, where you have IP cameras or sensors, and things like that, so the ability to manage those very diverse environments. Then you also have the ability to manage remote offices. So, we're not even talking about leveraging edge. We're just talking about, you've got maybe lots of remote offices, and you need to be able to manage servers with devices sitting on-site premise, etc. It's a potential way for you to manage those environments. Distributed sites, I guess it's where you've got lots of mines with lots of distances between then, and you'd be able to manage all these various sites. And then this is an interesting one, managed services. The ability to leverage the edge to, firstly, generate extra revenue for customers, and the other thing is to also manage the customer's environment much better, because the extra revenue will obviously come in the form of edge devices, and being able to provide edge services, but being able to manage lots of different devices through a single interface allows you to manage services centrally, which, again, comes back to the value proposition.

Thank you very much. Alice, back to you. I'm sure there's some questions from the audience, and hopefully, we can try and answer them.

Alice Puckett

Excellent. Thank you so much for sharing such great information. We do have a few questions coming in from the audience, so let's get started. Let's see, the first question asks, what is the advantage of using this type of solution over traditional network management tools?

Hein Matthee

So, I think that's a good question. I think the real problem with-- I mean, traditional network management tools, if you speak to somebody that manages them, they'll obviously like using them, and a lot of these have been functioning for a lot longer than the edge has been around. What I will say that is very important about that is this gives you much more flexibility because it's not just managing your environment, and your devices. It's also about being able to provide extra services to solve potential challenges on-site. So, the infrastructure management is the gateway for you to start leveraging other services, whether it is specific connectivity services, like 5G, managing your connectivity through some kind of firewall, or router, or even being able to provision specific applications on-site such as camera, like Computer Vision, video analytics, facial recognition, etc.

Next question, Alice.

Alice Puckett

OK, perfect. Our next question is, what verticals would benefit from this type of solution, whether for infrastructure only, or for applications?

Hein Matthee

So, the verticals that we see a lot of value in is Industry 4.0, because as you can imagine, some organizations have a very extensive network of factories, and they need to be able to manage them centrally, especially for visibility. Even though you might have resources on-site, it's good for management at a central location to be able to keep track of these things. Then I would also say oil and gas, mining, marine ports, where you have the potential for lots of spread-out services in a specific area. For example, in ports, you've got lots of containers, you might need to be able to manage lots of resources from one side to the other side. It could be quite a distance, so it's maybe easier for

you to manage all these resources centrally. So, that would be one. I would say mining, oil and gas, ports, and then obviously, Industry 4.0, those are verticals that could definitely leverage, whether it's just managing infrastructure or obviously providing applications on top of it.

Alice Puckett

Great, thank you. OK, our next one asks, could you expand around managing infrastructure at the edge, and what you would need for adding connectivity or applications to the solution?

Hein Matthee

So, it's... So, the biggest question with these type of solutions is normally you think, but hold on, I need to manage infrastructure, but I need to provide a whole bunch of extra services, and I need to provide a whole bunch of extra devices on top of it, and the real easy answer without solution is that, specifically for applications, if you, for example, have a camera on-site, like IP camera on-site, that you already have another device using that IP camera, because let's face it, some of these services are already running on-site, we're just not getting the full value out of them. You don't need extra services other than the edge node that we would roll out on-site to manage the environment. You could literally through a single central console, you could provision that specific application on top of it.

The same goes for connectivity. Now, with connectivity it's a little more interesting because you might need to roll out, obviously, your radio devices, but the server itself could run the 5G core as well. So, the only thing that you would potentially need is for radio for you to be able to connect to that site. So, this is the important thing. It comes down to service consolidation, the ability to run multiple functions, all running on the same server, because I think we're sitting in the paradigm of everybody trying to run specific devices, appliances for specific things, and the fact is that with modern x86 architecture, we can actually consolidate a lot of this, and that would be definitely something.

So, if you've got a device that you already provisioned for managing infrastructure, and it doesn't have to be a new device, so you can repurpose a server that you have on-site to do this, you could very easily, if a server is underutilized, roll out specific functions through your central management interface to be able to start managing, whether it's connectivity or applications.

Alice Puckett

OK, great. I have one more question from the audience. In the presentation, you have Lenovo listed. Is there any limit on what hardware could be orchestrated?

Hein Matthee

Not at all. So, we are perfectly comfortable to work with any hardware provider. We have done some work with Dell, we've done some work with HP, we've done some work with Advantech. We've been in talks-- we've been working with Lenovo, but the fact is that we could potentially work with any hardware, which means that if you've got servers on-site, and you want to leverage them, or you have a specific vendor that you prefer, like a Super Micro for example, and I'm mentioning all the vendors here, I'm not choosing one over the other. We're completely agnostic and we're perfectly happy to work with any devices. At the end of the day, there might be a certain amount of interfaces for us to understand, but we could very easily manage and leverage those different hardware providers.

Alice Puckett

Very good. Well, that concludes our webcast today. Thank you again for joining us.