



Managing Today's VDI Infrastructure – Part 2

Christian Chavez, Aspire Solutions Architect, with Brad TerEick, Cisco's Global Lead in Cloud Infrastructure, and Chuck Foley, NetApp Senior Director.

Christian Chavez: Hi everyone and welcome back, this is Christian Chavez, Aspire Solutions Architect, sitting in for Doug Stevens for PART TWO of this two-part series on Managing Today's VDI Infrastructure. If you missed PART ONE, you can head over to aspiretransforms.com, under the Resources tab select Digital Aspirations Podcasts. With me again - we were having some great discussions - is Cisco's Global Lead for Cloud Infrastructure Software, Brad TerEick, and NetApp Senior Director, Chuck Foley. Welcome back, guys. Appreciate the time that you're putting in.

Let's pick up right where we left off on Part 1. So, in Part 1, we were talking about some of the management tools for consideration. We were talking about deploy, manage, and optimization. We didn't get a chance to finish, but I do want to pick up on optimization. I know there's a lot of things that are out there in the industry for optimization. We've done it a lot for the VSI industry or the VSI workload as well as for VDI. So, specifically around VDI, Chuck, I'll put this in your lap first. Optimization, what's NetApp doing to help VDI optimization?

Optimizing Your VDI

Chuck Foley: I think we're taking it from a couple of different layers. The underlying infrastructure is all the data that your VDI users are using. And NetApp's ONTAP technology has a number of optimization techniques such as dedupe, data compression, and thin provisioning, et cetera. Also, it optimizes the management of it in terms of the backups. You can literally just go point-and-click from a web interface and ensure that all your data is backed up to inexpensive object stores.

You can also turn on tiering so that your high-performance data is on high performance tiers of disk. But using artificial intelligence, machine logic, your less often, less used data can be moved off to very

inexpensive object store, so that stuff that's not used very often doesn't take a lot of cost out of your pocket.

Then, you move up from the storage layer to the compute layer, and there's some really interesting things we're doing there. Our VDS or Virtual Desktop Service solution which is a SaaS delivered global control plane that will deploy, manage, and optimize your VDI environment for you will do the things you wish you had an army of VDI admins doing. Which is, as VDI hits enterprise scale, you can have hundreds or thousands of VMs. They get spun up and used, and users load on them, and users come and go. And you can end up with a lot of lightly used VMs, a lot of orphaned VMs, if you don't know how to go turn them off, and that's just wasted money.

So, what VDS does is, VDS takes your policies where you've said a certain number of users should fit on a certain size VM, or a certain number of users of a different class may sit on a bigger size VM. And VDS will watch who's on and who's using the resource. And it will kind of manage using what we call live scaling. It will manage the largest number of users to a given VM size. So, your biggest bang for the buck that you've decided is the right service level agreement. What that prevents is you having dozens or even hundreds of 10, 20 percent utilized VMs lying around that you're paying for, especially in the cloud, pennies per minute.

And one of the interesting things you're going to see is, last year, we acquired a company called Spot. Spot has fascinating AI driven technology that can reach out and acquire reserved instances, which you can normally get from hyper scalars in the cloud world at 25, 30 percent off. Or even Spot Instances, which you can reach out and get for up to 90 percent off. It can find and apply those to your workloads. And the really neat thing is Spot Instance is a great example. You get them for ten cents on the dollar. But most of the hyperscalers have an agreement with you, if you use that and save that money, we can take it back in minutes. We can take it back with no more than three minutes notice. It's really hard to tell you to use your machine.

So, what Spot does is, Spot has years of algorithms to say what types of machines are most likely to be grabbed back by the hyper scalar. And if they've provisioned one of those for you and they think, "Hmm. I have a VM sitting there that's running 20 users." That's going to get grabbed in the next few minutes. They will go find another machine for you on the market. Provision it for you and move your workload over so that you don't face that disruption. So, when the hyperscaler comes and grabs the one back that you had, your users don't even know.

Digital Aspirations in Business Podcast – S2E2

What that does is, that uses machine logic automation and AI learning to be able to provision compute resources in this highly dynamic world at a fraction of the cost you would normally pay, unless you have an army of people watching. So, when you look at those layers between the storage efficiency layer, the resource utilization layer, and then the acquisition and provisioning component of compute, you can end up with a really cost effective VDI environment.

CC: That's amazing. Even just hearing about Spot, I think that it speaks volumes to the experience that the end-user has. So, not only are we automating the infrastructure and building cost optimization, but we're not impacting the user while doing so, and that's pretty solid.

I know I've worked a lot with workload optimization manager, specifically when it was on-prem, when it was CWOM and when it was Turbonomic. Intersight now has integrated that with IWO, who also optimizes the infrastructure. So, Brad, tell us a little bit about IWO and how it manages, or it optimizes infrastructure today.

Brad TerEick: Sure. And we touched a little bit on this in the first episode of this podcast. But I'll say that, if we look at the goal and the role of Intersight Workload Optimizer, it is a SaaS delivered real-time decision engine that works with both public and private cloud environments. And so, for the private cloud environments, obviously, we mean on-prem infrastructure. And its goal is to assure workload performance and it provides workloads the resources that they need when they need them. So, it analyzes all of the resources that are in play for a given application, and it helps the customer understand what they can do to better improve the overall experience for the customer of the application.

And that overall experience can be related to specific high-water marks that we set with the Intersight Workload Optimizer application, or they can simply be cost constraints that we put in there. So then, this delivers a platform, or it gives them the resources that they need in a more optimal fashion.

CC: So, one of the things I specifically wanted to touch upon, because even though we're having a VDI conversation, I still think there's infrastructure that is providing services to the desktop. And that could be your VSI infrastructure. I know IWO is very driven. It's not doing what NetApp is doing with Spot, but it's at least optimizing the infrastructure, either VDI is sitting on top or the server infrastructure that's providing services to the desktops. I think there's definitely a component that we need to consider when talking about management. There could be developers that are on these desktops. There could be just users that are consuming applications that the business is either delivering to them or delivering to their clients that they need to interact with.

I know Intersight, aside from what you had mentioned, whether it's IWO or whether it's the peace certs or the TAC connected, there are some integrations that are happening for some of these desktop users specifically around the development as well. I know there's an integration now happening. The user of desktop services could be people like developers, right? I know there's a lot of remote working developers today.

There are set of tools that Intersight provides to those desktop users, and in this is really just in thinking outside of just desktop itself, but also the infrastructure that we're providing to those users and some of the services. And some of the services that are coming out of Intersight today, specifically around integrations with Terraform, integrations with some of the workflow automation with hybrid clouds. What can you speak specifically around what it's doing in that area?

Simplifying and Streamlining IT Operations with the Cloud

BT: Sure. Well, when you think about Intersight itself, it really simplifies and streamlines a customer's IT operations and provides this consistent cloud operations platform for them, no matter what it is within the Intersight environment that you want. So, if you think about the various resources or the various aspects of IT, you have an operating model for the way you manage virtualized infrastructure. You have an operating model for the way you manage Kubernetes infrastructure. There's an operating model that a line of business user may use for developing applications that run on, say, a Kubernetes environment. There is an operating model for, really, any aspect of the IT shop.

One of the things that Intersight does is, it brings all those operating models together onto one platform. So, now, if you are, say, a line of business developer and you are wanting to take advantage of Kubernetes-based resources that, maybe, you're used to just writing your own check to Amazon and requesting resources living in the cloud. Well, what an IT shop really wants to do in that scenario is, an IT shop wants to be the broker of those resources because it's following corporate standards and corporate guidelines.

So, with Intersight, we provide a vehicle by which a line of business developer can request the resources that they need to support their application. And those resources are coming through the Intersight platform themselves. So, if they are looking to, say, provision a series of virtual machines to support their applications or Kubernetes cluster to support their applications, Intersight can be that vehicle for

providing those resources. And those resources can live in on-premises infrastructure, or those resources can live in the cloud. It doesn't matter.

CC: Excellent. So, guys, honestly, I kind of want to break from script here and I don't think we can end the podcast without having this conversation specifically. We've seen it happen in the IT infrastructure, meaning just your standardized virtual server infrastructure. We went from an on-premise solution where everything was just hardware-based, and people were building data center to now cloud. So, this aspect of cloud is to being able to consume resources without having to build data centers and then just pay as you go kind of mentality. "I got a credit card. I get signed up. I can pull resources and just start working."

There has to be an aspect of desktop that is actually moving to something similar, where we're now consuming functions and different aspects of cloud that aren't even relative to servers anymore. So, we're paying for like a managed service or managed solution that's out there. So, specifically, I have to imagine that we're headed towards at least desktops where you're just paying for desktops out there. You're not even building the infrastructure anymore. Not even in cloud. It's just you're now consuming a desktop and you're paying for those desktops and letting somebody else manage it and run it. Chuck, is this something that you see coming down the road here? Is this something you can speak to?

Maintain Productivity through Workplace Changes

CF: Absolutely. I think you're going to see some really exciting things from major industry players over the next few months, specifically in this world which some people refer to as DaaS or Desktops as a Service for the reasons you mentioned. The global pandemic had so many large enterprises realize, "Okay. Maybe it's not the smartest thing to have employees running around with 2,000 bricks. And then, I have a corporate event like a pandemic. And, literally, my enterprise is spread around 2,000 bricks in the world. And then, they're not all another working right. So, let's get all my employees working in the cloud."

And I got to tell you, I think that many enterprises, large companies were talking to, are saying, "Yeah. I'm not sure that I want to manage all of that resource in the cloud. If I'm going to put it in the cloud anyway, do some of you other players in the enterprise infrastructure landscape want to manage that for me?" And instead of me managing my 2,000 users on VMs and storage networking, maybe you just want to get 2,000 licenses to a desktop of Type A, Type B, Type C. So, I think you're going to see an accelerated movement on a business model trend that way.

Now, the reason that's happened is interesting. It's happening because to run hundreds or thousands of users on a common infrastructure prevents a really significant scale problem, as we've talked about before, with performance, availability, data integrity, compliance. And a lot of enterprises are saying, "I don't want to take those 5,000 points of light, turn them into a bonfire, and be responsible for it. I want to take those 5,000 points of light, put them up there, and let somebody that knows how to manage large bonfires manage that for me." And so, you're going to see some pretty significant announcements.

I'll give an example. Here at NetApp, we have an offering called VDMS or Virtual Desktop Managed Service. It has been NetApp using our enterprise technologies, like Azure NetApp Files in the world of Azure, combined with Microsoft's WVD or Windows Virtual Desktop Infrastructure, combined with NetApp's Virtual Desktop Service to use automation and machine logic to manage desktops. We wrapped that all together. And we have customers that say, "Hey, just give me a thousand end-user accesses. You don't even need a cloud account. NetApp, you're running it all. You're billing us just X amount per user per month. All I want to bring is my users and my apps."

That has really taken off as a business model. And we're not the only ones. And that's why I say I think you're going to see some pretty significant announcements through the rest of this year. And it's, I think not only a technology trend, but a business model trend that our customers are saying, "We're demanding so everybody ought to keep their eye on."

Securing VDI with a Shared Security Model

CC: So, what does that do? So, a lot of things that we talked about, the management aspect. We didn't even get into the securing aspect of it. The deployment, the management, optimization, is that something all that's going to be at the hands of the provider at this point? Do we need to consider that? Is there anything that we should consider? Is it just, "Hey, we're just going to buy the service and we can then add some security, obviously, to the desktops that we would traditionally do inside of our IT infrastructure." But how does that change a user's perspective in terms of how they go about delivering desktops now with the type of service like that?

CF: So, I think that what we're seeing is a shared security model. And that whoever is providing that Desktop as a Service is responsible for providing secure infrastructure, data security, and secure connection. And this is where kind of multi-tenancy in the cloud as it's grown up and evolved comes in. Because if I am an infrastructure provider, I'm using my infrastructure to let you, as a user, tap into that

for a hundred or a thousand in place. But you still need to do that with your tenancy control. So, it's your active directory that's touching on my infrastructure. And so, we're applying these multi-tenant environments to take your particular security policies and apply that to the infrastructure that are being managed for you. And it's what the industry calls a shared secure model.

CC: Got you. Okay. So, in thinking about that, I think that's definitely the way things will be moving. We saw it, again, with some of the virtual server infrastructure or at least applications in this case. Server infrastructure still means servers, whether you're using an AWS or not. But in terms of applications, these things are monolithic and we're just consuming services from some provider that's providing the application back to us without having to deal with infrastructure.

Specifically, I know Cisco is sort of moving to that model. And I just want to touch upon this because I think it's still relative enough in the conversation that we're having, is that there's this concept of moving the management plane from infrastructure, from on-prem, out into Intersight. And you had touched on that, Brad. So, think of UCS manager, which today lives inside of the Fabric Interconnect sys, now moving out to Intersight. And, again, we're consuming services that we're no longer responsible for and we're putting it at the hands of a provider.

What we've seen Chuck talking about, where we're now just paying for consumption of an application, in this case, Chuck had talked about desktops. We're now actually talking about platforms that are managed infrastructure. I know IMM is something that's being talked about at the Intersight level, where we'll be taking on configurations or at least the overall management aspect of infrastructure today. Can you at least touch upon that?

Managing the Infrastructure from a Centralized Place

BT: Sure. A couple of things, one of the really nifty things that Cisco brought to the table in 2009 when you UCS came to the forefront, is that we moved the management plane to the fabric interconnects. Then, with Intersight, we wanted to bring that management plane or that control plane up one more layer all the way to the cloud. The nice thing about moving it up to the cloud is the purview of what we were managing increased tremendously. It didn't just have to be servers that were part of that fabric interconnects domain. It could be literally anything that we could manage with the platform. So, things like Kubernetes clusters, things like the actual fabric interconnects themselves. So, now, we can build domain policies and domain profiles and assign those to the fabric interconnects themselves. Not just assign those to resources that are managed by the fabric interconnect.

So, it brought everything up a layer. By bringing it up a layer, it didn't only increase the scope from a hardware management perspective. It allowed us to really look beyond the hardware. Look at resources that are consumed elsewhere. Even resources that are consumed in the cloud. So, that's where we talk about Intersight being this hybrid cloud operations platform. So, it really doesn't matter where it lives. And as Chuck said, if we are moving these virtual desktops to more of an as a service model, this plays very nicely into where Intersight is teed up and headed. And Intersight doesn't care if those virtual desktop resources live on-premises or if they live in the cloud. It can see those and consume those whether their virtual machines or containers. It can see those and deal with those no matter where they live.

CC: Definitely. And that's where I'm seeing where a lot of this is heading, and I think that's something that I wanted to speak to. In just being able to consume either application, in this case, we're talking desktops, we're talking about management platforms, we're even talking now Fabric Interconnects, which if you look at some of the things that Cisco's doing and moving things out more centralized, so it gives, as you had mentioned, an overall purview of all of infrastructure from one location in terms of management.

So, this was a great conversation. I appreciate you two joining me. We're going to close this for Part 2. I want to thank everyone for joining us. And, again, thank you, Brad and Chuck. This was a great discussion.


And for Aspire Technology Partners, this is Christian Chavez. And until the next one. Thank you so much.

Outro: Digital Aspirations is brought to you by Aspire Technology Partners, a premier technology solutions and services provider designing, implementing, delivering, and managing digital infrastructure and IoT solutions to enable transformational business outcomes, creating more agile and efficient IT environments that deliver differentiated customer experiences for your organization. To learn more about Aspire Technology Partners, visit aspiretransforms.com or email us at podcast@aspiretransforms.com.

Digital Aspirations in Business Podcast – S2E2



Aspire Technology Partners is a Cisco Gold Certified Partner engrained in solution pillars that set us apart as a true Cisco solutions provider. We are committed to the continuous improvement of expertise and skillsets around Cisco initiatives that enable us to help and guide customers in the adoption and management of technology architectures designed to transform their organization. We hold Cisco Master Specializations in Collaboration, Security, Cloud & Managed Services and is one of only 25 partners in the US to receive the Cisco Advanced Customer Experience Specialization.



Aspire Technology Partners is a NetApp Gold Partner and is closely aligned with NetApp's data driven solutions which help customers achieve digital transformation and IT modernization. As a Gold Partner, Aspire is committed to the continuous refinement of our people skills and expertise in the solution areas of hyperconverged infrastructure, cloud services, flash storage, and data management. The Aspire and NetApp partnership and collaboration allow us to guide customers throughout their digital journey, as they use data at the core, in the edge, in their Data Center or in the Cloud, to innovate platforms and services for their customers, employees, and partners.