

Lee Caswell: 750 million new applications in the next three years, more than the past 40 years of computing. How do you assimilate that quickly without overrunning your staff?

Jason Burns: You're able to roll out these applications and not have to worry about how much time it might take to stand up new servers, stand up the network. All of that piece just can be more seamlessly provisioned so that you have more time to deploy applications.

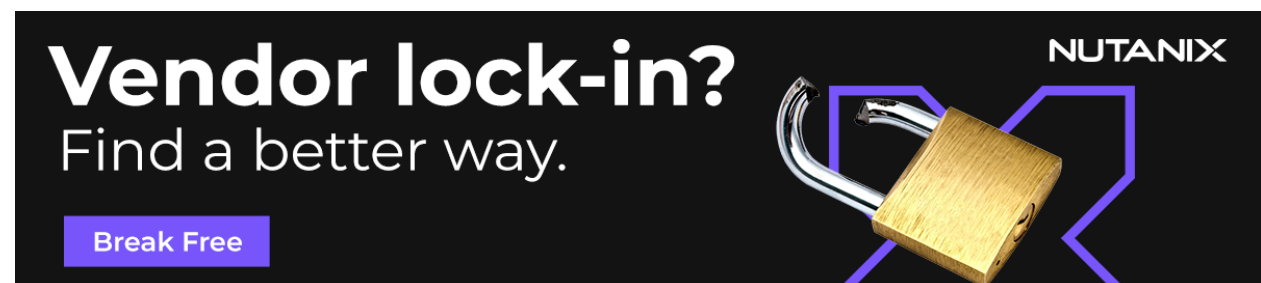
Lee Caswell: Let's say it a different way. Infrastructure doesn't get in the way of supporting new applications. The best partnerships are customer-driven.

Jason Lopez: Those are the voices of Lee Caswell, vice president of Product and Solutions Marketing and Jason Burns, Director of Technical Marketing, both of Nutanix. This is the Tech Barometer Podcast Here, they're in conversation about the evolving nature of network infrastructure and application deployment in the context of Nutanix's partnership with Cisco. They chat about different perspectives and understandings that network, storage, and compute teams bring to the table, unified communications and Cisco UCS, rapid application deployment and many other facets of the partnership

[Related: [Focus Shifts to Migration in Wake of Broadcom's VMware Acquisition](#)]

Lee Caswell: Today, for example, I was talking to a customer who's the first Cisco customer we have in the Americas, not the first we have worldwide. And one of the things that really stuck out for them was the relationship that they have with Cisco, you know, expanding from networking into compute, has really meant that they have a level of trust there and confidence. And now that Cisco's reselling the Nutanix offer, it's really bringing together the networking and the, call it the infrastructure teams, in a way that we brought in the past, the storage and the compute teams together. So I'm curious on your perspective on that, Jason.

Jason Burns: From my perspective, I've seen the same thing. When customers are deploying a solution from Cisco, it is exactly that. Oftentimes you'll see something with the unified communications on top of Cisco UCS and the Cisco networking piece as well. And so the combination there of being able to bring in the entire application on top of the computing system and then the storage as well is really appealing to the customers. And then having that delivered as a solution rather than something that they're piecing together part by part.

A graphic with a black background. On the left, the text "Vendor lock-in?" is written in large white font, with "Find a better way." below it in a smaller white font. A purple button with the text "Break Free" is positioned below the text. On the right, there is a gold padlock with a silver shackle, set against a purple outline of a Nutanix logo. The word "NUTANIX" is written in white capital letters in the top right corner of the graphic.

Vendor lock-in?
Find a better way.

Break Free

NUTANIX

Lee Caswell: You know, as I'm watching and talking to some of the Cisco folks, there's some really interesting things that Nutanix does that they didn't have access to before and now make it just easier to sell an [HCI-based solution](#). One is, you know, the ability to have mixed nodes, for example, to be able to have nodes with different capabilities, whether it's disk and flash or different CPU capabilities. That level of flexibility actually makes it a lot easier to sell this. And now integration, you know, support for ACI, for example, really helps from an integrated

management solution. So I'm seeing those two things really set the bar as to how we're now we're bridging the network stack along with compute and storage.

Jason Burns: Yeah, the integration of [Nutanix AHV](#) and Cisco ACI is something that I'm personally excited about. As a former network engineer, I love to see the automation point where we can make networks inside of the network fabric and have that pushed into the compute fabric. What you get from that action is a system where you're able to really easily deploy applications because the network substrate is already there for them. That's part of what Cisco has done to integrate with Nutanix AHV. That's really complemented by what Nutanix is doing on the virtualization side, both with our flow virtual networking and flow network security for virtual networks and micro-segmentation inside the hypervisor itself. So you have this physical network layer that Cisco is providing, and then you have the virtual network layer that Nutanix provides there, and that's all being tied together.

Lee Caswell: Yeah, it's kind of fun. I don't know if you're familiar with the concept of parallax, which is basically things, lines move as you move. It looks like things are moving because you're looking at it from a different vantage point. And so in many cases, the network community is looking at things from one vantage point, where you've got now the storage and compute teams coming in and looking at it from a different view. For example, things like capacity utilization and performance from a storage standpoint are around IOPS, consistent low latency during failure conditions. Those are the things that a storage person cares about. And at the same time, now you've got all this based on the underlying network infrastructure. So the ability to have both of those lenses, I think, is really an interesting way to think about how to deploy this more seamlessly over time.

[Related: [Why Hybrid Multicloud Matters](#)]

Jason Burns: The lenses is a great point of it. The network team, rightly so, controls the network because they need to have the best performing network for their applications. They hold some of that responsibility, and therefore they hold the provisioning control as well. Same thing for the compute and storage side. But going back to that comment around lens, you get visibility from the storage side of what the network is doing. And then from the network side, you get visibility into what the compute is doing, thanks to the integration of ACI and AHV. So the solution works a bit better now that these entities are aware of each other.

Lee Caswell: You've also got inter-site integration, looking at how you go and view, manage, and monitor the holistic system. And so, I'm particularly excited about this because there's over 60,000 Cisco resellers, for example, who are familiar with inter-site. And now to be able to say that inter-site can go and monitor Nutanix environments today, manage shortly, is really an exciting way to go and leverage the experience and expertise of that Cisco reseller base.

Jason Burns: The inter-site piece is also very interesting because I came from a world of deploying Cisco UCS with local UCS managers. So for me, I'm learning more about Cisco through this and learning that they have inter-site as a tool. It's a tool that I wish I had when I was deploying UCS servers years ago. The ability to look at a fleet of servers and manage them all from that point would have been really helpful to me as an administrator back in the day.

[Related: [The Spirit and Hybrid Multicloud Mindset of a Healthcare CIO](#)]

Lee Caswell: Yeah, I was looking at some of the roadmap items we've been talking about. One of the things that a lot of customers have been looking for and are excited about what is coming with Nutanix and Cisco together is the ability to run on the latest blade servers that are storage-rich to be able to take that new form factor. Because certainly Cisco is known in the UCS world as having performance-oriented systems with a lot of memory access. That was the early design of UCS. And now the ability to bring storage-rich capabilities and Nutanix's value on our HCI element makes it possible now to go and run the highest performance applications on this new joint offering.

Jason Burns: Yeah, absolutely. And I'm looking forward to seeing expansion into that blade ecosystem because that's what I had seen a lot of deployed myself, these high-performance applications running on a blade backed by some sort of storage that might have resided outside the blade. And so looking at what we have as possible future options with Nutanix compute-only accessing the storage remotely, that to me just seems like a win-win for configuration in these blade server deployments.

Lee Caswell: Especially important, by the way, for the highest performance applications because of the licensing model, it turns out. So the ability to have a contained, let's call it compute-only node, for example, where you can now go and have, let's say Oracle licensing to pick one running dedicated on that node means that now you've got a predictable licensing expense going forward. And so that's important because, you know, there was a risk in the early days, you could basically have licenses spread across all of the CPUs and the cores. And of course that would be prohibitively expensive. So the idea that you can contain licensing to this actually means now you can take the highest performance applications that are licensed by core and make sure that financial exposure is guarded there.

Jason Burns: Anecdotally, one of our customers who's a service provider, they deploy a lot of different solutions. I was helping them out with a technical design I said, you know, this is how I would do it. And he said, well, Jason, about 40% of the way I design things is technical and the other 60% is driven by the licensing requirements of all the solutions I'm deploying. One solution, while I might from a technical side think it's the right answer for licensing or financial reasons, might not be a viable solution for them to deliver that service.

[Related: [Ransomware Resilience: Evolving Data Protection](#)]

Lee Caswell: Super interesting. Even as you look at what's next after databases right now, the number one workload on Nutanix. Analytics now from our own survey data is 49% of our customers are running analytics. And a big part of that, of course, is Splunk. So now acquired by Cisco, I see a pretty interesting opportunity to go bring together these high-performance analytics along with our Cisco partnership. Bring that all together and make that really a solution from a customer standpoint.

Jason Burns: Yeah. I'm also interested to get hands on that in the lab myself. Part of what we do is we, in the tech marketing team, we develop these lab environments for people and we try to demonstrate our own software and try to pull in realistic use cases where we can. One of those use cases is you have this Nutanix software stack. It's generating alert data, event data. How can you send that into a tool that's very similar to what the customer has? A lot of our customers ask us, how do I send all this data into Splunk? If I were able to, in the future, have a system that can send all that data into Splunk and then show our customers how to do that, to

me, that seems like a great offering and something that customers would benefit from. So I think that's how they deploy it in the real world.

Lee Caswell: Of course, no conversation would be complete without a discussion of generative AI. Thinking for a moment about how we're taking GPU-enabled Cisco nodes and now being able to do our GPT in a box, for example, together with Cisco UCS. Pretty interesting, right? As a way to go and say, how can I get started? You could be training a model, probably not a large language model, which is happening in the public cloud, say, but you could certainly be modifying or tuning a model with your Cisco-enabled, GPU-enabled nodes. And then being able to take some of the value that we offer on the software side to help customers get started on their AI journey.

Jason Burns: Yeah, Lee, I see people ask that all the time, which is, how do I get started? How do I deploy this? What's your best practice guide? And working with Cisco, what they have on their side is the Cisco validated designs. And so what I would love to see is a Cisco validated design for something like generative AI, kind of complementary to our GPT in a box solution for, okay, here is how you get from end to end a deployment with Cisco UCS servers, the Nutanix AHV hypervisor, your GPUs of choice, and how do you build that out from end to end? For me, the validated design is a great starting point because it shows customers that we've tested this out in the lab. And if you follow this recipe, you also can have success.

[Related: [Will AI Workloads Overload IT?](#)]

Lee Caswell: I guess it strips out both the real risk by doing predetermined testing and then also any perceived risk. So you can say, all right, if I can take the infrastructure piece out, which by the way, is what every data scientist wants. It's like, how do I make sure the infrastructure is there responsive and fast? So we do that. And by the way, reliable and day-to operations and compliance, all the things that IT brings to the party. And then at the same time, be able to deploy new applications really quickly, whether they're libraries like PyTorch, for example, or bringing in new applications, bringing in our support for the NVIDIA and VAI libraries. I mean, these are just ways to say, hey, we've got all the building blocks and can help you with putting them together.

Jason Burns: And having the validated design or whether that's a Nutanix validated design or a Cisco validated design means that you're really speeding up the deployment. You're cutting out a lot of those investigatory steps and just following along a cookbook rather than inventing the wheel on your own.

Lee Caswell: I think for a lot of customers, particularly, they're looking now at these new server-based architectures as a way to build out an architecture for the next three to five years, where you'd put all of your new applications. It's rarely just like rip and replace, take something out. That's why we still have Unix systems in there. It's additive in the sense that now you're going to take these server-based architectures, add that in, but put all the new workloads on there. Certainly that's what happened with x86 because the economics and the simplicity of moving to a x86 architecture based on Ethernet, which is what we're doing here. And then deprecating over time, fiber channel, SANs, switches, and even the expertise. So you can basically focus on the new stuff is how I'm viewing this market.

Jason Burns: As a practitioner, what I like is that that gives me the time to focus on the new things. If you'd asked me a few years ago if I was going to be learning about generative AI

models, I would have said, no, that's not really something. But now that the infrastructure side is more taken care of, that's a place where I can look at these new applications and maybe learn more about how Kubernetes works and how to deploy a modern app in Kubernetes or maybe learn more about an AI application rather than spending all my time just getting the compute and storage and networking infrastructure up and running.

Lee Caswell: So this infrastructure as code transforms into everything as code. If you think about over time, how do we make this easier for you to take and basically treat AI as just the next workload? Many customers from Nutanix started off in [VDI](#), then went to [databases](#), then went to analytics, and now are looking at AI, generative AI, as the next thing. And if you can assimilate that easily with the same team instead of having a separate architecture, boy, it just makes it a lot easier to get some operational leverage, particularly when you're having trouble finding talent.

Jason Burns: That is something that I see every day, even in some of the new designs that I'm rolling out. You might do something where you have an engineer walk through and manually set something up as a proof of concept. That's pretty time-consuming to do. Now, if you have something like infrastructure as code or your application deployment defined as code, the subsequent deployment of that can happen really fast. And so you just cut out a lot of time there for deploying that. That really is moving toward everything as code with the solution we're showing here. So you have the Cisco ACI portion, where you have these APIs for configuring your network. You have the Nutanix compute and storage part, again, with lots of open APIs to configure. And then these modern applications on top of it, which might be deployed as a series of config files against a [Kubernetes](#) cluster. You can just keep all of that code for the entire infrastructure and even application configuration in a code repository where someone can just check that in to make changes to it.

Lee Caswell: Yeah, it's always struck me that ACI was probably one of the best acronyms in the industry, application-centric infrastructure. The idea that you're thinking, hey, how do I make sure that applications are optimally located? And there's a really interesting way to think of how we introduce value jointly here, because when you bring these together, these technologies, now all of a sudden you've got a full complement of how to make a seamless outcome for optimally locating applications and their associated data over the hybrid multi-cloud universe.

Jason Burns: And as I was seeing adoption of modern applications and containerized applications, that was the one thing that struck me the most, is that all of these apps are going to generate and use incredible amounts of data, and they're going to need highly available storage that's also high performance. I didn't see on the market a dearth of those solutions. I saw that Nutanix really had a great offering in that space, that people were going to be deploying these modern apps and using Nutanix storage to back that app deployment.

Lee Caswell: Both of our companies, I think, have terrific brand loyalty, by the way, in part because of the support experience that customers have gotten. And one of the things that impressed me about the relationship that's more than just your average partnership is how we've integrated the support path so that our support systems and support tickets can be followed up and make sure that the customer experience is singular in nature. Curious on your thoughts on how you've seen that develop, and why is that important?

[Related: [How Flash Memory is Reinventing Hybrid Cloud Storage](#)]

Jason Burns: Working at Nutanix these last nine years, I've heard consistently again and again that customers value Nutanix support. And in my previous life, I was actually a Cisco support engineer long, long ago, and heard the same thing, that customers valued the support that they got from Cisco. So we have a history of two great support organizations, and being able to support our customers really is what I know that they're looking for. It keeps them coming back. It keeps them expanding solutions on Nutanix or expanding their solutions on Cisco, because they know that when something goes wrong, they can give us a call and get routed to someone who can help them.

Lee Caswell: I have a theory on why our companies are so well suited together. So my theory is this, that architecturally, from a design point standpoint, network engineers tend to make things fast first and then reliable later, because you can always retransmit. And storage companies come at problems with exactly the opposite lens, which is that they try and make things reliable first, and only then try and make them fast. And you can certainly see that through the enhancements we've made over the latest versions of AOS, for example, and moving into newer flash technologies, you know, NVMe and new CPUs and everything. So I'm curious to say, how do you think about it, right, architecturally, isn't that an interesting way to think about why the engineering teams have so much appreciation for each other?

Jason Burns: That definitely rings true to me. And I've seen that growth at Nutanix over time, where all of our subsequent software releases have basically worked on performance over all others, but reliability was always there, because from a storage perspective, you had to keep that data. Data corruption wasn't really a possibility for consideration, but we really worked hard on increasing performance over time. And it seems like we've kind of met in the middle here between Cisco and Nutanix.

Lee Caswell: I think storage people, right, preserving the data is number one and all the data services that we have. And I think that's a really interesting value also, because for most customers now, this is a really unpredictable environment. Trying to forecast, for example, your mix of blocks, files, objects, data, trying to forecast whether you have VMs or containers three years from now, and start to think about, you know, what you have on-premises versus the cloud. The simplifying design that Nutanix brings is the fact that you've got the same data services across all of those endpoints. If you go and buy different unique architectures for each, you're going to have a different way to go and protect that data, different snapshots, different recovery mechanisms, different disaster recovery. And so it's a really easy simplifying, you know, way to basically make this easier for customers to get operational leverage and deploy things in an unpredictable environment.

Jason Burns: Nutanix has a lot of that flexibility on the data services side, and I see Cisco having that flexibility on the compute side. The longevity of some of these Cisco UCS chassis, for example, or how long you might have a UCS deployment is pretty incredible. And that's just something that you can keep as a long-running deployment and take out parts of it, maybe blades, maybe certain rack servers while keeping the whole of the deployment there. You get the ability to upgrade in place and really extend the possible lifetime of a deployment. And for me, that was always the beauty of the Nutanix cluster, because your AOS storage cluster could live for quite a long time, even across multiple generations of Nutanix hardware. If you are adding newer-generation servers and removing older-generation servers, you have a pretty long-lived deployment without a lift and shift.

Lee Caswell: So one way to think of this is in a highly unpredictable environment where applications are changing really fast. You don't know necessarily if there are going to be block interfaces or file interfaces or object interfaces. By being able to buy one thing, you can actually get started. And then by having seamless licensing that allows you to go and basically deploy at will over time, you've got this flexibility now, along with the fact that we're the only company that gives you a consistent way to do data protection across all three of those protocols. And by the way, across VMs, traditional VMs, and now newer container systems orchestrated by Kubernetes. So it's a really fantastic way, I think, to go and give that flexibility for customers who actually don't want to be storage experts, it turns out.

Jason Burns: That would describe me as well, a former network engineer who has come to the storage world, but Nutanix has made that a lot easier. Storage is not as intimidating as I had thought it was after working with Nutanix. Let me say it this way. You know, what I'm hearing from a lot of CIOs now is that they're interested in consolidating architectures so that they can address the fact that they don't, let me say that in a different way. Yeah, one of the things I'm hearing from CIOs recently is about trying to consolidate the number of different architectures that they have. So for example, a simple one would be like separate filers and separate SAMs, for example, bringing those together, and a different object store, bringing that all together. And so for CIOs thinking about how they do it, they've already done that with Cisco on the networking front. They have a single vendor on the network front that gives them everything across different architectures. And so now one of the opportunities here is by working with Cisco and Nutanix together, now you've got the opportunity to basically address any of the degrees of freedom you're likely to anticipate in the future.

Jason Lopez: Lee Caswell is vice president of Product and Solutions Marketing at Nutanix. Jason Burns is director of Technical Marketing at Nutanix. This is the Tech Barometer Podcast, I'm Jason Lopez. Tech Barometer is produced by the Forecast. We have a treasure trove of tech stories at the forecast from all across the technology landscape. You can find more at www.theforecastbynutanix.com.