

Rajiv Ramaswami: As customers here, all of you expect to have solutions that meet your business outcomes. We solve some portions of that, but we don't claim to solve everything. So clearly that's why we are investing in the ecosystem here around us, the ecosystem at every level, the ecosystem in terms of our OEM partners, our cloud partnerships, our best-of-breed partners in areas that we don't play in like security for example, in cloud native [Kubernetes](#) and public cloud providers. We all realize that we operate in this broad ecosystem and we have to come together to help you get that solution. So that's what we are committed to doing, we'll continue down that path and we hope to create deeper and deeper relationships with our ecosystem.

Jason Lopez: What you're hearing is a talk Rajiv Ramaswami, CEO of Nutanix, gave to Nutanix customers at the Nutanix dot NEXT conference in 2023. He highlighted several partners, such as OEMs like HP, Dell, Lenovo, and Super Micro, a security partnership with Palo Alto Networks, cloud-native services with Red Hat, collaboration with major cloud providers like Azure and AWS, and with Citrix, particularly in VDI workloads which represents about 20% of [Nutanix's platform](#) usage. This is the Tech Barometer podcast, I'm Jason Lopez. In this talk, Rajiv addresses hybrid multicloud, public cloud, edge technology, platform services, and a brief picture of the Nutanix roadmap. But to get back into the talk, here Rajiv touches on the company's support of AI application development and its integration into services.

Rajiv Ramaswami: There are two parts to that role. The first part is many of you are developing new applications that [do use AI](#), and we want to be the platform that you can run those applications on. We are also committed to open source on this. There's a lot of open source work being done on AI. We are also working closely with Nvidia because a lot of these use GPUs and we expose DPU through our platform and help you use them in an efficient way. Clearly one whole approach around making these AI workloads work well on our platform. And some of these tend to be very industry-specific workloads like retail and manufacturing as examples that I mentioned. So that's part one. Part two is using AI internally within our company's offerings in terms of the products and services that you consume from us. So in fact, our [Prism Pro](#), by the way, the R & D team for Prism Pro is actually called AI Ops because we truly believe that we can actually make operations more efficient by using machine learning and deep learning. That's just one example of where we are starting to use some ai. Another example is we all, at least a good chunk of you actually provide telemetry data for us. We'd like to do more with the telemetry data. Of course, protecting identities, protecting privacy, of course, is very important there. But being able to look at the telemetry data, we could potentially be doing things like predictive analytics and helping you understand what needs to be changed. And so there are potentially new use cases there. And then of course when it comes to support, I want to make sure we retain our NPS there and we continue to do a great job supporting you, but there's also AI that can be used inside to help make that process even more efficient.

Jason Lopez: Now imagine, if you will, a digital landscape, with applications and data freely roaming across multiple domains. In the middle of this, there's the home base, the on-prem setup. There's the edge, on the outskirts. Then there are the public clouds. In this next part of Rajiv's talk, he hones in on creating a single, unified platform that bridges these diverse places..

Rajiv Ramaswami: The reality is yes, most of you are actually hybrid multi-cloud in the sense that you have these applications and data in multiple places. You have your on-prem, you might have your edge, you have public cloud, a, public cloud B for the most part. Each of these is different. They operate as silos. You have a different stack, you have different processes, tooling, and security to manage each of these. That's actually what hybrid multi-cloud is today. And what we aim to do is to simplify that for you with a single platform at the infrastructure level that cuts across all of these and gives you that same experience, same tools, same processes, and complete flexibility to run your applications wherever you'd like across these environments and manage them all with a single team. That's really, I think what we are aspiring to do, and we've got a lot of proof points along the way. We've got our offerings today on AWS on Azure with a number of managed service providers, and we are seeing an increasing number of edge opportunities in very specific verticals. Manufacturing, and retail, for example, just to name a few. And of course, defense.

Jason Lopez: Attitudes toward using public cloud services are changing. Nutanix initially embraced the cloud for certain services but eventually had to reconsider and move some of their operations back to a private data center to control costs more effectively. Here, Rajiv addresses this.

Rajiv Ramaswami: We are certainly seeing people being a lot more careful about what to do in the public cloud or what not to do because if you had asked this question three, four years ago, a lot of people in the room would just say, I'm going to the public cloud unless I'm a dark site. I've got some specific requirements. A lot of people would say, I can go to the public cloud. And then a couple of things happened over the last few years for those of you who have gone, I think you've realized that it isn't as easy to take everything you want to the public cloud. It can be very painful for lots of applications and will need refactoring and potentially re-platforming. So not easy to take your existing apps. And for those new apps that you're building, it's great. It's an easy on-ramp. But once you start running it at scale, and once you start running it continuously, the cost starts up. And we've experienced that even at Nutanix. My cloud bill keeps going up every year and we keep trying to look at what we can do to optimize this. And I'll give you one example of what we have done, on a much smaller scale. One, during COVID, we started a program called test drive. Test Drive ran on Google Cloud. Actually, it's a nested hypervisor for those of you who want to be technical, but they didn't have a bare metal offering at that time. And it proved to be quite popular. We didn't charge for it because it was just an online proof of concept. And the use has continued to grow like crazy for us. And our cloud bill, therefore from Google, keeps going up every year. We said, okay, well we've got to now it's starting to get serious. This is here with us to stay. We have to treat this properly and figure out how to optimize it. And so what did we do? Well, we moved back a chunk of what we call the steady state workloads back to our data center. We run that very cost-efficiently. Then, for burst capacity, we use the cloud, and it's working out well. It's certainly saving us some dollars when it comes to cloud costs. And what you're going to see is a lot of companies realizing going through the same journey that we did, which is there's an easy on-ramp to the public cloud. It's an easy way for you to get a new app out there quickly, and if you don't do it right, you're going to be locked into the public cloud. At least in our case, it was portable. We could move it. It's tempting to go use all the services, get there quickly, and build this app and run it. But then you realize that as you start scaling, at the end of the day, they have to make their margin. So you end up spending more and you can certainly, I always say a well-run private cloud infrastructure can certainly beat the public cloud in terms of pure cost. There's no doubt about that. And I think lots of you probably have experience with that and you've done that work already. So we are certainly seeing that as, I don't know if it's maybe a little too early to call a trend, but I think what certainly we're seeing as a trend is people being much more circumspect about what's going to go in the public cloud. And I can't say I've seen a trend in terms of everybody moving back from the public cloud.

Jason Lopez: Edge means different things to different people. But Rajiv points out its evolution has paved the way for a new generation of edge applications. New use cases demand advanced computing power, containerized applications and centralized automated management systems.

Rajiv Ramaswami: The way I would think of edges would be, let's call it reasonably high compute edges. Some people call it near edges. And then other people say, I call it far edges where you have much lower cost, very simple solutions that are more OT type deployments. So for us, we are more focused on the compute edges. In fact, to put this in perspective, what you are deploying as a server in the data center three years ago, you can actually get more computing in a single server sitting in the edge today than you probably got in your data center a few years ago. So imagine tons and tons of compute potential available in the edge by deploying very few numbers of servers and a small number of nodes, one-node deployments, and three-node deployments. And so there's a lot you can do with it. But more importantly for many of you, there's a whole new set of applications that are emerging at the edge. If you're in retail, we are talking to people about fraud detection. We are talking to people about automatic checkout manufacturing. We talked to car manufacturers who are saying, I want to have completely automated visual inspection of defects through machine learning and ai. So just to name a few here. And so there's just a range of these new edge applications coming in. There's a lot of data being generated at the edge as well, and you might be able to do training for some of these AI applications in the cloud, but inferencing will likely be done locally. And as you look at this, these edge use cases actually starting to get fairly compute-heavy, and we are engaged with many customers today in terms of helping our solution optimize for those specific use cases. Now, what's common in addition to meeting these specific application needs is that most of the new applications are containerized. There's a lot of AI being used and data management being needed, and there's also a need for clearly centralized management, provisioning, upgrades, everything because you're not going to be able to go out there and have people manage your edges. It has to be

automated. So that's what we are seeing today, and I'm sure this is going to continue to evolve over the next few years.

Jason Lopez: Rajiv's talk moved on to a vision for simplifying the management of applications across platforms. He says it comes down to a platform service vision.

Rajiv Ramaswami: If you look at applications today, modern applications are being built with containers and Kubernetes, and that Kubernetes substrate is available for you everywhere. I would almost go so far as to say it's starting to be commoditized. It's everywhere. You can get it on-prem, you can get it in AWS, they have EKS, and you can get it in Azure. They have AKS. You can get it everywhere. Google has GKE. So that the compute substrate is available everywhere. In addition to the compute substrate, all apps also need a set of data services. Most apps, almost all apps will need databases. They'll need messaging and streaming. They'll need caching, they'll need search. And this is a set of services, most of them related to data that's available today in the public clouds. The only issue there are the public clouds, so that they tend to be siloed. It's not easy for you to go from one to the other. It's not easy for you to have flexibility in terms of avoiding locking in. And so what we are looking to do is to provide you with a consistent set of these data services across everywhere, right? Across all native substrates. You can run it on top of a Nutanix infrastructure, of course, wherever that's available. But we are also going to enable these services to be available natively on AWS, Azure, and other native cloud substrates, so that you can truly then think about building an app once using this set of services. You could see how easy it is to be able to deploy that app anywhere and also not be locked in.

Jason Lopez: In this last section of Rajiv's chat with customers, he touches on the company's plan for enhancing their data services, especially in the context of Kubernetes and containerized applications.

Rajiv Ramaswami: So really you'll see a range of data services for Kubernetes and containers that will essentially help you to run these modern applications and deal with containers much the same way as you deal with a VM. We always do snapshots, but now we are going to be doing snapshots and being able to store the snapshots in low-cost public cloud object stores or any SC-compatible object store for that matter that has a lot of use cases and implications. They're going to provide you with a single console for you to look at your entire estate and manage your entire estate across on-prem public clouds and edges. And the fourth is more of this foundational vision around our future, how we can help you build these applications, modern applications in a portable way, using a consistent set of data services and being able to run them anywhere. This vision is for us to help you build portable applications. This is a five, 10 year journey for us and we are just getting started with our Nutanix database service. So I'm excited about the future of the company. I'm excited about how we can continue to meet your needs, not just for today, but also as you go forward.

Jason Lopez: Rajiv Ramaswami is the CEO of Nutanix. This was a talk he gave before customers at the Nutanix dot NEXT conference in Chicago in 2023. In May of 2024 the conference moves to Barcelona, Spain. This is the Tech Barometer podcast, I'm Jason Lopez. Tech Barometer is a production of the forecast. If you like what we do here check out more stories at theforecastbynutanix.com.