

WEBINAR TRANSCRIPT

Data Protection in Hyperconverged Environments

Simplifying Data Storage with Nutanix and Cloudian

Host:

Aaron Friedman, strategic alliance marketing, Nutanix

Speakers:

Rohip Goyal, product marketing manager, Nutanix

Grant Jacobson, director of technology alliances, Cloudian

Aaron:

Hi there. Good morning, good afternoon, good evening. Thank you for joining Nutanix and Cloudian webinar. With the explosive growth of data has made with harder to manage and scale your data storage while ensuring data protection. It makes it even harder when you add public cloud in the mix.

Today we have our presenters from Nutanix, Rohip Goyal, product marketing manager, and Grant Jacobson, director of technology alliances at Cloudian. If you have any questions during the webinar, feel free to ask them in the chat on the right, and feel free to check out our joint solution brief on the right as well, as well as links and attachments. At the end of the webinar, you will receive an email with a recording of the webinar, and at the end, we will also have a Q&A. So right now, I'm going to be passing it off to Rohip. Rohip, take it away.

Rohip:

Thank you for that, Aaron. Hi, and welcome to this webinar, "Simplifying Storage for Data Protection in Hyperconverged Environments." My name is Rohip Goyal, principal product marketing manager with Nutanix, and I'm joined by Grant Jacobson, director of technology alliances with Cloudian. This presentation will provide an introduction to Nutanix and Cloudian and how you can benefit from the joint solution to solve common IT challenges with infrastructure and the explosion of data.

We include this slide to remind everyone that our presentations and demos may include some features that are not yet released, and that you should work with your Nutanix and

Cloudian sales and partner teams before making purchase decisions based on these upcoming capabilities.

We have a fair amount of content for you with key challenges. What's happening with infrastructure today? How is the public cloud impacting your IT environments? What are the different clouds everyone's talking about? How does data growth impact it? What's new and what are those challenges and what do they bring? After that, we'll discuss the Nutanix Cloudian solution and deep dive into both of them. Then we'll wrap it up and take some questions.

Let's first talk about some critical IT challenges. So today, IT leaders struggle to manage the unrelenting complexity across their IT environment. This includes not only core data center infrastructures but remote and branch offices, disaster recovery sites and more. This complexity makes it difficult for businesses to pursue new markets, develop example products and services and tackle strategic initiatives such as digital transformation projects. Much of the challenges stem from legacy infrastructure constructed with discrete resource silos such as storage, server, virtualization, networking and security. This fragmented design makes it difficult and time-consuming to provision new IT environments to run critical business applications. On top of that, these architectures have multiple points of failure, creating a fragile IT system that could have unplanned downtime. Managing this IT complexity demands significant processes, process overhead for executing even basic tasks such as adding additional capacity or upgrading software, adding unwanted friction to IT operations and slowing the business.

Legacy architectures also impose financial burdens and large up-front capital purchases required for even small-scale deployments. This is undeniable, and it brings this complexity that has a profound impact on already stressed IT teams. It requires the tight coordination of multiple IT specialists to handle simple operations, and most importantly, leaves little time for true innovation.

The adoption of public cloud services has reset expectations of what IT enterprises must deliver. Internal IT consumers need infrastructure provision in minutes, not weeks or months. Otherwise the business will circumvent IT teams in favor of public cloud services, creating shadow IT challenges. Public cloud has changed how organizations pay for and consume IT services. Expenses, overprovisions, infrastructure is no longer possible. The business expects pay as you grow economics, where IT services are provided only when you need them, and you've paid for them, solutions that are simple to deploy, manage and scale with one-click operations for common workflows. And continuous innovation, new features and capabilities incorporated into a production environment that continuously non-disruptively adds features and functionality without [00:05:28].

To meet a diverse and expanding set of IT requirements, organizations are leveraging multiple cloud environments to power their workflows and manage valuable customer and business data. These multi-cloud designs include public cloud services, including AWS or Amazon Web Services and Google Cloud Platform, or GCP. Private cloud environments running corporate data centers or co-lo facilities, and then you have the

distributed or edged deployment, including remote office and branch offices and field-based IT sites.

While individual cloud environments may deliver particular capabilities and have a unique set of operational requirements, managing a multi-cloud IT infrastructure brings a clear set of imperatives. For example, lower the long-term total cost of ownership for IT services, accelerate overall IT agility to keep pace with businesses, bring real-time visibility of cloud costs, public and private, enabling informed decisions for application placement and data storage, simplifying IT in-cloud operations, leaving more time for IT teams to add business value, and meet stringent requirements imposed by the business while maintaining necessary security and IT controls.

Grant, maybe you can shed some light on the data growth and the problems that it's bringing to these IT teams.

Grant:

Yes, thanks Rohip – a really good foundation introduction here. And for context, on this slide, we're showing the dramatic exploding growth of data. There's more data than ever, and it's really accelerating. It's doubling every couple of years. This shows, really, the dramatic amount of data across different data types. I list here on the right everything from enterprise data, IoT, you've heard of those. We have increasing amounts of data archiving for compliance and long-term retention. There's an explosion of life science data, bioinformatics, genome and so forth. Your healthcare data is increasingly important for medical imaging to patient records and so on.

The challenge here for organizations is how do you manage that and how do you protect that? And let me suggest, as additional context, what's driving that are these data-intensive applicational workloads, everything from manufacturing, financial, insurance and so forth. Analytics increasingly a big deal with big data, so-called big data. So these application workloads run very nicely on the Nutanix platform, and what we're showing today here is looking at that, looking at the implications of the data and how that can be stored efficiently, effectively, cost-effectively, actually, and how to manage and protect that.

Another challenge here from the storage side is where do you put the data? How do you keep it managed and accessible? in the old days, it was pretty simple because the amount of data was fairly light, but over time, it becomes more challenging with more silos of storage or silos of data, and it really doesn't scale. The technology, the IT technology and the infrastructure, the legacy infrastructure doesn't scale and it's really not possible to keep up in terms of budget, in terms of staffing, even in terms of real estate, where does it go. And companies are now looking at this and realizing they need a modern optimized solution such as Nutanix and Cloudian.

Rohip mentioned the multi-cloud world, the imperative to take advantage of new features and new ideas here. That is the same as well with respect to storage. We have

enterprise IT working the stations already deploying hybrid cloud for their storage – that is the data center private cloud as well as one or more of the public clouds, and they're also increasingly committed to a multi-cloud architecture where that could be private cloud, multiple public clouds, it could be the distributed cloud Rohip mentioned, but how does that all that work together? How do you have data everywhere in multiple locations have a unified view of that, have access to it in fairly real time and make that work cost-effectively? How do you manage that complexity?

That's what our next chapter here is to talk about how we address those challenges, all of them, with the Nutanix Cloudian solution. So let me move to the next slide and make the point here that we're talking about a Nutanix offering as well as a Cloudian offering, and these are complementary. Rohip, perhaps you could kick that off and talk about Nutanix Buckets.

Rohip:

Sure. Thanks for that overview and those challenges, it actually plays well into the joint solution that we're talking about, and Nutanix is working on a new object storage service offering that's going to be available soon, and it basically will be part of our enterprise cloud software and available as a service, so you'll be able to enable it without actually having to set up a new set of infrastructure.

But think about some of these cases where you're running these VMs on top of Nutanix, and you want to extend and have some capabilities of utilizing a local object store, and it's great that we're going to have that capability. We're going to expose that through an S3 API and then you'll be able to consolidate your enterprising cloud work while doing that. And you'll be able to centrally manage all that.

And the great thing about partnering with Cloudian here is that when you want to be able to tier off, we're working towards doing some work on that, being able to extend the capabilities of our object storage offering into what Cloudian has, and Grant, maybe you can elaborate on that.

Grant:

Yes, absolutely. Cloudian Hyperstore is really complementary with the Nutanix Buckets. They're both object storage. Cloudian has perhaps the industry's most compliant S3 API, so all the S3 compatible applications work seamlessly here. We are at Cloudian a software-only company, but we do make that available in our own appliances, which is shown here in the slide, but we do work as well with Nutanix, which we'll talk about soon here, as APM for example, to deploy inside the Nutanix cluster.

A little more about Cloudian – it's enterprise software defined object storage. Our customers can start small with just three nodes, and this scales seamlessly and also

non-disruptively to petabytes and exabytes. It's very simple. It's a single-storage space across that, and it is hyper-cloud and multi-cloud ready. So, as I say, it's a complementary solution, works very nicely together with Nutanix, even seamlessly for that. And let's go to the next slide just to show you three deployment options or solutions with this combination. Let's see if that slide will come up. There you go. So Rohip, perhaps you can make a few comments on this one before we talk about the other two.

Rohip:

Yeah, perfect. So hopefully you're getting an idea of how the two solutions are working together to provide you the best breed of the infrastructure stack and being able to extend out your capabilities for our structure data out to Cloudian.

With the first portion of this, what we're saying is, again, think about the Nutanix solution and how we basically allowed you to consolidate your three-tier environment into a single hyper-converged solution to be able to run your VMs on top our enterprise cloud OS. And having the new capability of Nutanix Buckets, which is going to be released later, which is going to be that object store capability, and think of all the features that are available to Nutanix today for all the users where we have enterprised great data services like compression, [00:16:34] coding all built into the platform, and this new service will be able to take advantage of all that, and then being able to eventually tier that off to Cloudian.

Grant:

We can talk about the other two. So like the first example where Cloudian is positioned for extended capacity for the Nutanix environment, the second one is a variation on that. So here, as I mentioned earlier, we're showing Cloudian running as a VM in a hyper-cloud provider inside the Nutanix stack, talking directly via S3 to an external Cloudian cluster for that extended capacity. This could be a second location. This could be multiple locations in the single picture here.

And then we have a third scenario where Cloudian is actually backing up the VMs and the application workload data from Nutanix as a storage archive or just for backup or recovery using your favorite backup company between the two. I think we're all familiar with the different options here, but this is just a generic way of saying back up and archive is increasingly important. Companies want to lock their data, have that available. There's also a reference here to public cloud, so we cater seamlessly to Amazon to Azur as well as to Google Cloud platforms. And so any of these scenarios can fit many use cases. I will mention, I forgot to mention in the second one that Cloudian running in the VM is actually a certified solution, a Nutanix-ready offering today. All of these solutions are available for consideration. So let's move onto the next one, next slide, and I'll turn it back to Rohip.

Rohip:

Perfect. Yeah, thank you for that. Hopefully you're starting to understand there are a lot of options, there's a lot of choice for you as you're looking at these solutions. And there's a lot of benefit as well. So Nutanix has approached infrastructure offers with these real, tangible, material benefits to customers in terms of greater agility, faster time to value and lower total cost of ownership. A study done by IDC quantified the benefits of Nutanix infrastructure found that customers get over 500 percent return on investment over a five-year period, and 60 percent lower TCO within a seven-month payback period.

Let me talk to you a little bit about the solution benefits. So the first one is more about overcoming the management and scaling complexities and challenges, and the idea there is that instead of just replacing your existing infrastructure and trying to solve it with some sort of point solution, think about it holistically. Think about how you're going to manage this environment and how you're going to handle it. And really start to understand how the joint solution is going to solve your problems and make it easier for you to follow that data, so as that application is riding that data out from a VM that's running on Nutanix to an object store locally, and then eventually maybe tiering it or maybe it's running a VM on Nutanix and then Cloudian's got a VM running on Nutanix that's going use a policy to push it off to the exabyte's worth of space that Cloudian offers. But think about how that becomes easier to manage and scale.

And then in terms of seamless across multiple locations, it's really about looking at how the solutions scale beyond a single site and being able to handle multiple sites, and both solutions across the board will allow you to do that. And they're going to be simple and efficient for a data-intensive workload, so you can run your VMs, which could be like your OLAP or OOPP database workloads on a Nutanix cluster and run your Cloudian virtual machine on there as well. And if you'd like, you can also have that data tiered off to the Cloudian solution for longer-term archival as well.

In terms of flexibility and deployment, there are a lot of options. Don't feel that you're kind of constrained by a one-size-fits all solution. Think about the workload that you're running and then think about that workload and how the backend platform has to adapt to it in terms of memory configuration and CP configurations, in terms of capacity configurations, in terms of locations, in terms of reliability, making sure that you've set it up for all tolerance. You want to have choice, and at the end of the day, between the two solutions, that's the idea that you're going to have that choice.

And in terms of working together with these backup partners, many of them like the [00:22:53] of the world, all write to that S3 API, so what that does is it allows you to utilize things like Nutanix Buckets or utilize Cloudian's solution as well because that S3 API has become the de facto standard and a lot of these backup and recovery options are utilizing that natively in their software, you don't have to change anything. You can automatically start covering those today.

And once you've actually made a decision, you've made a choice on a platform, think about how you're going to actually go and set those up, which one of those solutions in terms of capabilities and features are going to make it easy for you to set up and deploy that, and eventually manage and expand that solution, so both together, the Nutanix and Cloudian solution scale, and they can do that independently. So if you need to add more compute resources or more storage resources or memory type resources on your Nutanix side, you have that capability. If you want to add that to the Cloudian side in terms of expanding out the storage, so really think about how that makes it flexible for you.

Again, we want to make sure that you realize the benefits, and I covered some of that on the previous slides in terms of cost effectiveness, and know that TCO and ROI over that five-year period, you're going to get that investment back. And the same on the Cloudian side, they're the best at making that dollar-per-gig price, what's available on the market, competitive.

I'm going to dive in a little bit on the Nutanix solution here and give you a little bit of a deep dive. Nutanix enterprise cloud addresses the cost and complexity associated with traditional three-tier data infrastructure constructed with separate storage, compute and storage networking resources. Building data-centered infrastructure with legacy architecture requires partnering with multiple vendors, [00:25:26] servers and the storage networking virtualization and more. And then you have multiple weeks and months that go into designing, procuring, deploying and provisioning that infrastructure, and it's difficult and time-consuming. So those processes to increase capacity to meet those business needs is not always easy, right? So there's always multiple management consoles staffed by specialists of IT teams, and then it's challenging and troubleshooting the environments when you don't have the performance that you need.

The Nutanix enterprise cloud simplifies that design and construction and deployment of IT infrastructure. It converges all IT resources into a single, scaled-up solution that can run any application. A Nutanix-powered infrastructure basically can be deployed in 60 minutes or less. IT teams can start small with typically just three server nodes and scale capacity, storage and compute without limit. Scaling is 100 percent linear. Each node adds a proportionate amount of storage and server capacity. If you need to double the overall capacity of a Nutanix deployment, you simply add 2X new nodes. No more guess work – 100 percent software solution that can run off the shelf servers from all leading server manufacturers, and the entire software stack can be updated in a single click and without disruption to the live environment.

In terms of enterprise-grade storage and security, two of the fundamental requirements to replace proven IT infrastructure is the ability to provide all enterprise-class storage features and to ensure the best overall security. Nutanix enterprise cloud solutions routinely replaces all [00:27:31] storage solutions from all leading storage vendors and provide a rich set of enterprise-class storage capabilities also shown on this slide.

For maximum flexibility, Nutanix provides and delivers software-defined storage with the granularity of a singular virtual machine. Additionally, Nutanix solutions are built with a

security by design methodology and include the richest set of security features available in the AT/I industry today.

I'm going to deep dive a little bit about what happens under the hood. One reason why Nutanix has become the most popular hyper-converged solution in the market is that distributed systems architecture that provides unlimited scalability, high performance and unmatched resiliency.

Here are some of the basics. Nutanix software is deployed on a cluster of servers or nodes starting with at least three or four in a typical data center. One or two node deployments are available for robo-installations. Each node has CPU memory storage, DSF or ACD or all-Flash and hypervisor. The nodes are connected via an industry-standard ethernet switch. The core Nutanix Acropolis AT/I software runs in a user-mode VM called a controller VM on each node in the cluster. The controller VM creates the single-storage pool using the direct attached storage in each node in the cluster. This is a distributive storage fabric, or DSF. It completely eliminates the need for a standalone storage gateway like SAND and NAS products. The DSF distributes data and metadata across all nodes in the cluster and provides enterprise storage features such as de-duplication, compression, erasure coding, cloning and snapshotting in a truly distributive system with no single points of failure or hotspots. This fully distributed scale-out architecture allows you to expand capacity by adding nodes when needed without limits. You get linear predictable performance and capacity when you need it.

The Nutanix solution delivers data locality to provide the best VM and cluster performance. All IO requests from a VM are handled by a local controller VM on that node. When a VM writes data, one copy of that data is always written to the local node. That way read requests can be handled immediately using data on the same node without having to go over the network. If the VM moves to a different node, then over time, the data also migrates to the remote node to restore locality.

A quick snapshot of our native hypervisor. So Nutanix offers this native virtualization layer. It's called Nutanix AHV, and this is what Grant was alluding to in the beginning when he was talking about the Cloudian solution running on a VM on top of AHV. So basically AHV is to leverage popular KPM hypervisor as the foundation, and what we've done is we've added enhanced security, faster performance, and most importantly, integrated management of the virtual environment in all the VM. It's basically integrated a hypervisor into the single cloud stack. It also enables Nutanix customers to automatically intelligently place every VM to optimize cluster CPU, RAM and storage resources for the best performance and greatest efficiency.

It's got built-in security with STIGs, which are security technical implementation guides, and self-based, self-healing, those resources, and the state of the hypervisor if it drifts from one group configuration for any reason. And it provides the capability of being managed through Prism, which is our management plane.

In terms of providing a single platform for all apps and data, Nutanix provides capabilities for storage services through block file and eventually we're going to add

object to that mix. Then we've got lots of virtual workloads running on top of it, so think about any of the workloads that you have today, enterprise mission critical workloads, so a lot of our customers are running those on Nutanix, and then cloud native.

And then this is a little bit more on what the Nutanix offering is going to be around the object storage service. It's called Nutanix buckets and it's really tailored around the DevOps, long-term retention and backups. And in terms of DevOps, there are a lot of developers out there and sometimes they're working on these applications that are running on Nutanix, and they want the capability of riding out to a single global main space, and maybe they don't need that much space, but they want that cross-functional area where they can work with teams that are local or they can work with teams that are remote, and they're looking for that global main space. That's one area that will be interesting to be able to utilize the Nutanix offering.

In terms of long-term retention, you have capabilities of doing the active archive or WORM capabilities that are set on the bucket level, and that'll help you meet those regulatory requirements. And then backup, where we've hinted on it before, being able to leverage existing backup software and writing natively to the S3 API that we'll expose. And again, think about how this works and how it works with the Cloudian solution that we talked about earlier and being able to leverage this Nutanix Buckets running alongside VMs, and then also having the option of running it as a standalone solution as well.

At this point, I'm going to hand it back over to Grant to dive a little deeper into the Cloudian solution, because now you've got an understanding of the infrastructure stack, how it works, and what values Nutanix brings to the table, and I'll let Grant do the same for Cloudian.

Grant:

Yes, thanks, Rohip. So let's look at the other side of the solution here, looking at Cloudian. As I mentioned earlier, Cloudian hyperstorage actually is software-defined storage. We have our appliance that we're showing here in the slide, but it does run as software as well on Nutanix as well as other servers. This is really On-prem S3, the industry's most compliant S3 object store. The customers we have may use this with three nodes. We have some using it with hundreds of nodes, multiple petabytes scalability, same architecture, the same platform for all of that, and the same feature set.

As I mentioned earlier, a lot of the benefit includes having this unified view of data across locations, including public cloud. This is really managed centrally that way. The access is available. There's a single global main space that makes this happen, and what we're doing here with this is the idea that you have a single solution here with Nutanix that is kind of ready for today but also ready for the future.

Of course, being very cost-effective is an important point here. How do you scale all of your data without breaking the bank? We've got very good economics here as well. There's economics on the Nutanix side and economics working for you on the Cloudbian side.

Let's talk a little bit more here. It's limitless scalability, but it's non-disruptive. You can add disks or add nodes or add locations non-disruptively. The performance is linear. It's the same management as I mentioned throughout. Zero downtime scaling is a key feature as well as full data protection, and I'll talk more in a minute about data protection.

The storage set to me is one of the ways we achieve the great economics here. In one of our 4U appliances, you can get up to 840 terabytes – that's about five petabytes – in just a single rack. So you can get as much as you need and only when you need it and it's very easy to scale there.

As you deploy this, this really appears as a single storage pool across your locations. The data protection is actually one of the unique advantages here. There is a very flexible bucket level storage policy for how you protect your data, and we make it possible to use erasure coding as well as replication, not just across the complete deployment across locations in the same way, but you can actually mix and match replication and erasure coding even within the same rack or within the same location or give this data a different level of protection. So it's extremely flexible, but it also makes it very cost efficient, because not every EC and RF scheme has the same cost profile.

And on that, on this slide here, I'm showing you both the replication option as well as the erasure coding option on how we do that. I think as you go larger in terms of the Cloudbian cluster, the erasure coding is going to be slightly more cost-efficient, but replication, especially for applications or use cases where those copies are actually needed, may be appropriate. So we can work with you on what's the best way to set that up.

One last slide here on Cloudbian, other unique capabilities, we have many customers who are service providers, providing storages and service or backup as a service, so we have a multi-tenant QoS and billing capability that's built into Hyperstore. Any enterprises also use Hyperstore internally with the multi-tenant capability because they have various business units that they then provide storage for that then get billed back, so that's a nice feature as well. We talked about the policy-based tiering, so mention that again.

Another key one that is increasingly becoming valuable is you have the data, as we've been talking about, which is valuable in and of itself, but then you have that goes with that the metadata that goes with that object. And being able to do integrated search and data analytics on that metadata is very invaluable. So there's a way that our customers are extracting value from their data by leveraging the metadata search and also adding their own custom metadata tags to that data. And so you have the data and then you

have the metadata, and the metadata is becoming increasingly useful as a tool not just to find something but also to look inside the storage.

A few other things I'll just mention very quickly – we've been talking about Hyperstore, which is the object storage. Cloudian also offers hyperfile, which is a NAS controller connecting via NFS and SMB, so the applications that may not be S3 come in as a file stored as an object all in the same platform, all managed the same way. And one last point I'll just mention is our smart support, so having that level of service and support that goes with the technology is also very important and a key part of the offering.

Use cases here, as I close this chapter out, this is as we said, local S3 storage within Nutanix has the certified offering in a VM on Acropolis hypervisor. Cloudian is also nicely available for capacity expansion or to extend the local capacity as a storage tier. Lots of customers use this for their backup and archiving, and that's a way they achieve the compliance and retention that's needed. And then of course, the hybrid multi-cloud tiering, that's another important use case as we mentioned in the beginning.

And then I'll just click through this very quickly here, if my buttons cooperate, here we go, yeah, this is just a nice summary of what we've been talking about – S3 API, this is on premises solution that adheres to the cloud, limitless, non-disruptive scaling, and with the erasure coding and replication options, we're getting unmatched data durability, which is important for people who are wanting to protect that data. And all of this, both Cloudian and Nutanix, are enterprise-grade, reliable – Rohip mentioned no single point of failure and so forth – that's as well on this slide.

So to just refresh again, these are the three deployment options or the three different solutions we've been talking about. Rohip covered a lot of the details around Nutanix and I introduced the Cloudian part, but this is, again, what it looks like altogether. Let me pass it back to Rohip for any additional comments before we take questions.

Rohip:

No, that was great, Grant. Hopefully the audience is starting to understand all the options available, so really think about how you can extend your capabilities using Nutanix Buckets and Cloudian, and think about how you can run Cloudian on top of Nutanix and extend out that way or as we've said as a standalone solution. So there are a lot of options for you based on the workload that you have, so think about the workloads first and think about how you can leverage the solution to match the needs that the business is driving.

I'm going to pass it over back to Aaron at this point.

Aaron:

Great, thank you, Rohip. We're now at the Q&A portion of our webinar, and before we answer the questions that we have lined up – looks like we have a few coming in – if you do have any questions feel free to ask them in the chat. Also check out our solutions brief, and if you want any more information on our partnership, go to our website as well.

Here's the first question: We have a hard time adding capacity without disruption. Can I add capacity to Cloudian non-disruptively without impacting Nutanix? It's probably a question for Grant.

Grant:

Yeah, I'll take that. Yes, we get that a lot. It's very easy to add capacity, and it is, as I mentioned earlier, non-disruptive in the sense that unlike some solutions where you have to start and stop or take out and put back, this is fairly seamless and fairly quick. So you can add capacity within the nodes. You can add nodes and locations and you can do that as you go, kind of when you need it. So you don't have to overprovision as some solutions require. You can just do it as needed and no other changes are required.

Aaron:

Thanks, Grant. And what is the biggest benefit in Hyperstore with Nutanix?

Grant:

Right, I think it's what we've been talking about, probably the simplicity in terms of management is one thing that I would summarize based on everything we've been talking about, simplicity of management. Economics is another thing. People care about the cost of infrastructure and so on, and both of these solutions, and especially together, kind of make it easy to deploy and expand in the future because a lot of economics are very predictable. So simplicity, cost-effectiveness are two key benefits.

Aaron:

Awesome. And what is needed to get started with this?

Grant: Rohip?

Rohip:

So in terms of getting started, I'll tackle the Nutanix piece and maybe Grant, you can tackle the Clouidian portion. So the Nutanix side, it's easy. I think we covered it in-depth during the presentation, but just a quick recap. It's simple – use three nodes to get started and you can actually deploy that in scale to meet the requirements that you have, so if you're a smaller company, you have the capability of having these smaller types of setups, and if you're expanding out and your business is growing and you're an enterprise-size company, you have the ability of leveraging that and expanding the clusters out. We've got customers who are running over 2,000 nodes, so it's pretty impressive. That means that they're running multiple workloads on that platform, and with the features that we talked about around data locality, you have that capability of really scaling this linearly.

Grant:

That's great, yeah, and I'll just chime in here with respect to Clouidian. As an external deployment, we start with three nodes, and that's really for data protection, and that can be as low as 50 terabytes. Not all data is big, but 50 terabytes, depending on the application, and that expands and scales up to petabytes and beyond. The other way to get started is, as was provided, presented here, is to deploy and run Clouidian Hyperstore as a VM within Nutanix, and that can be expanded as needed, and it can expand within the cluster, the Nutanix cluster, or it can also expand with an external Clouidian Hyperstore cluster to support that extra capacity need. So pretty simple, Aaron. Back to you.

Aaron:

Great, thanks. And it looks like we have a few questions around object storage, specifically Buckets. So does Buckets have multi-cluster support, Rohip?

Rohip:

Yeah, so we are working towards the capability of adding multi-cluster option, but at the first iteration of it, it's going to be a single cluster, especially around thinking about workloads where you're running your VMs and maybe you're running files on Nutanix and now you want to expand that out to running objects, and having that capability of having object on that Nutanix cluster, and then know that when we add features and functionality, we actually think about how they're scaled. So the object storage solution is designed that way. It's been designed from the ground up, but also leveraging all the goodness that Nutanix has on the platform so what I meant about the ground up piece is that it's actually designed for scale, so it can handle all the metadata information. It can handle all the capacity management, and it's going to actually enable that controller VM, which is that storage controller I talked about, to be able to direct it on instead of

putting all that metadata into the controller VM. It'll actually keep it separate in microservices location where you can actually have a complete scale for that object storage solution. And then eventually, working with Clouidian and being able to push that off to longer-term, like that exabyte location size would be the next step.

Aaron: Great. And can Buckets support terabytes to petabytes of data?

Rohip:

Yes, so in terms of Buckets, it definitely can. You can start small, very small indeed, and we can handle a few kilobytes to terabytes to petabytes opportunities. So think about how the Nutanix solution can scale and then think about when your needs go beyond that and you want to expand into Clouidian, you have that capability as well.

Aaron: And what about supporting multi-part uploads?

Rohip:

Yes, so in the release, multi-part upload is definitely going to be there. We looked at what we need to build into the initial release, and things like WORM capabilities to meet those technical regulatory requirements, object tagging, object versioning, multi-part upload, handling small files, we're going to build all that into the native object storage solution that we have.

Aaron:

Awesome, thank you, Rohip. Looks like we don't have any more questions, but like I said, feel free to check out our joint solution brief. Go to our website and check out our joint solution. A recording will be sent out to you via email. If you could give us a rating, that would be awesome. And with that, I think we're going to sign off, so from the Nutanix and Clouidian team, we wish you a great day, a great evening, great afternoon wherever you may be. Thank you. Bye.



CLOUDIAN, INC.

177 Bovet Road, Suite 450, San Mateo, CA 94402

Tel: 1.650.227.2380 | cloudian.com

©2019 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, and HyperStore are registered trademarks or trademarks of Cloudian, Inc. All other trademarks are property of their respective holders. T-NUTW-1218