

23-Oct-2023

NetApp, Inc. (NTAP)

Investor Meeting

CORPORATE PARTICIPANTS

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

George Kurian

Chief Executive Officer & Director, NetApp, Inc.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Anthony Lloyd

Vice President-Technology Services, OpenText

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

OTHER PARTICIPANTS

Tim Long

Analyst, Barclays Capital, Inc.

Sidney Ho

Analyst, Deutsche Bank Securities, Inc.

Victor Chiu

Analyst, Raymond James & Associates, Inc.

David Vogt

Analyst, UBS Securities LLC

Meta A. Marshall

Analyst, Morgan Stanley & Co. LLC

Irvin Liu

Analyst, Evercore Group LLC

Steven B. Fox

Analyst, Fox Advisors LLC

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Matthew John Sheerin

Analyst, Stifel, Nicolaus & Co., Inc.

Frederick Gooding

Analyst, William Blair

Jake Wilhelm

Analyst, Wells Fargo Securities LLC

MANAGEMENT DISCUSSION SECTION

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Thank you, everyone, who's here with us today and on the webcast. We appreciate your time joining the NetApp tech session here at INSIGHT.

Just as a reminder, we won't be covering any financial information today. We won't be giving you an update on the quarter. This is all about understanding our technology and our competitive advantage and how customers use our technology better.

So with that, I'm going to kick it off with the Safe Harbor. So, my apologies. It's a bit long and I need my glasses to read it. Each of the 2023 INSIGHT Financial Analyst Tech sessions may contain forward-looking statements and projections about our strategies, products, future results, performance or achievements financial and otherwise. These statements and projections reflect management's current expectations, estimates and assumptions based on the information currently available to us and are not guarantees of future performance. Actual results may differ materially from our statements or projections for a variety of reasons, including macroeconomic and market conditions, global political conditions and matters specific to the company's business, such as changes in customer demand for storage and data management solutions and acceptance of our products and services. These and other equally important factors that may affect our future results are described in reports and documents we file from time to time with the SEC, including factors described under the section titled Risk Factors and our most recent filings on Form 10-K and 10-Q available at www.sec.gov.

These forward-looking statements made in these projections are being made as of the time and date of the live presentations at the presentations are reviewed after the time and date of the live presentation, even if subsequently made available by us on our website or otherwise, the presentations may not contain current or accurate information. We disclaim any obligation to update any forward looking statement based on new information, future events or otherwise.

So with that out of the way, let me just give a quick recap of the agenda. So, George Kurian is going to kick us off with some brief statements about what you can expect here at INSIGHT. Then we're going to have a couple of guys from our AI team come in. They're going to do all open Q&A. You can ask them any questions you want. Then a couple of people to talk about enterprise storage. We'll take a short 20 minute break and reconvene at 11:05 Pacific Time. Then we'll have someone from product marketing come up. He can answer any questions you have about anything.

We'll see if you can really stress them on that. Then a couple of our cloud guys will come out, talk about the breadth of our cloud solutions, both storage and otherwise. And we'll have a customer panel with a couple of real NetApp users. And again, I encourage you both in the room and on the webcast to ask questions. We'll be fielding questions through the webcast. It really is your time to talk to the technical experts here at NetApp. Once this session concludes this afternoon, you're free. Those who are here, free to wander about the INSIGHT experience. The general session keynotes will start at 3:00 this afternoon. Please go to those if you're still here and then they'll be a reception that you can attend from 5 to 7 tonight. So hopefully, everyone who's here can do all of that and make it worth their while. Again, really appreciate everyone joining us and coming. And with that, it is my honor to introduce George Kurian.

George Kurian

Chief Executive Officer & Director, NetApp, Inc.

Thank you, Kris. Thank you and welcome to those who are on the livestream. Thank you for joining us today to those that you're here in person. Thanks for making the trip. This is our first INSIGHT in-person since 2019. And the theme of the conference is very germane to the industry context in which we operate. It's about turning disruption into opportunity.

What we've seen over the last few years is that the range impact in terms of scope and scale and pace of disruption is accelerating. Range, geopolitical, macroeconomic, supply chain, you name it, you've got it, right, as the world transitions from a relatively stable geopolitical architecture in the post-war era to the dawn of a potentially new global economic architecture and political architecture.

In terms of the pace of disruption, you're seeing the pace in terms of the ongoing impacts as well as the scale and scope of those impacts accelerating. And what we have seen, as well as others have seen is that those businesses and organizations that are data driven are better set up to understand disruption and respond and transform it into opportunity than anyone else. If you look at the research conducted by the Boston Consulting Group and Google, they have shown that organizations that are data driven were 5% better positioned to drive revenue or cost and productivity improvements or rapid response to disruptive change in 2022, that gap has widened to almost 14% in 2024.

So, the competitive advantage of being data driven is wide and accelerating. And that is, of course, before the impact of the AI. AI, organizations that are hybrid cloud and data driven are even better positioned to capitalize on the AI trend. And we will share more about that tomorrow. What you'll also hear us talking about is that being data driven is not an easy thing for customers to do. There have been two widely adopted approaches, neither of which has been entirely successful. One, stop down. Hey, build a monolithic data lake and dump all of my data in it. There was Hadoop before data lakes and there were data warehouses before Hadoop. Those have not entirely succeeded, and they struggle to become enterprise scale as the nature of data and the need to derive value from those large investments continue to get pressured.

On the other hand, the other approach, which is grassroots, say let a thousand experiments and let every department try its own approach has also not scaled. They have typically gotten stuck in proof-of-concept, where you've got a great proof-of-concept, but you try to take it from proof-of-concept to production is very hard to do. And so, what we will talk about is that you need three things to be data driven. One of them is an integrated data organization where all of the people deciding your data strategy and data engineering and data science and business analysis work together. They don't need to report in one place, but they need to work together. And you need to have a clear sighted view of what data and data projects have the most important business impact. That's number one.

Second, from an operating model, you need to treat data as a product independent of the underlying systems. Today, for example, what you find is businesses operate and manage systems and business processes, right? So, you go and talk to a customer. They have a CRM system and they have an ERP system and they have a supply chain system and they have a BI system sitting on top of those system. And then they got a big data system sitting on the side and on and on and on. What none of those systems have is a complete view of a domain, for example, a customer or an employee or a vendor. And that integrates not only data from your transactional systems, but from external data sources. And that's increasingly important because to make good data decisions, especially to power AI models, from work that's going on in the industry and from which one of our guest speakers is well associated with, having good data is a major source of competitive advantage for AI. We will, for example, make the statement that AI runs on data and data runs on NetApp, especially unstructured data,

which is 80% of the world's data and is much faster growing than structured data, right. So, that's the second. You got to think about data as an operational model, as a product independent of the underlying systems.

And then we'll talk about a modern data architecture, which is what's important there is that we are pragmatists. We say that you cannot evolve and transform every part of your data architecture. That's just not feasible from a risk, cost, speed standpoint. And so, our position is that you want to build the right balance of transformation and modernization in some parts of your data architecture with stability, evolution, integration in other parts of the architecture, because it gives you the right balance of flexibility and speed, risk, and cost. And an important foundation of that modern data architecture is what we call an intelligent data infrastructure that combines hybrid multi-cloud data storage with integrated data services and AI powered cloud operations, monitoring, optimization, and automation.

It builds on what we set out with the data fabric. In 2013, we stood up on stage at INSIGHT and said the world would be hybrid and multi-cloud, that you would need to manage data across not only your data center, but all the places that you would put your data. Meaning all the leading public clouds and Software as a Service solutions and we have delivered upon that. Today, our technology is a native service in all of the leading Public Clouds, and that position allows us to now provide much more capability for all the modern applications that run on it. The second is we recognize that the needs of clients have moved beyond just the hybrid multi-cloud infrastructure.

So for example in data services, we delivered on the concept of portability and flexibility in terms of data management. But you see increased needs for security and governance of your data. And so, we've made important steps forward to bring what we think is the world's most secure data storage, whether it's in your data center or on the leading public clouds backed up by guarantees to make your data always available, that's what we call secure by design. And we will have one of the industry's leading spokespeople for Secure by Design Director Easterly with us today.

We talked about the fact that AI is powered by data and data runs on NetApp. And so, we have an awesome demo that combines some of the world's leading Gen AI tool chains in the public cloud with our data management solutions, so that you can version data and you can integrate your on-premises data with the world's leading large language models in a way that's secure and protected that allows you to comply with increasing mandates for data lineage and protection of private data, all of those [ph] tool (00:12:54) chains, we're ready. We've got awesome stories for you.

And then, of course, one of the important things that we will tell our clients is that in a world of constrained resources for IT and talent and the need to move faster, it's even more important to build silo free architectures. We'll talk about Hadoop, for example, where everybody said, hey, let's build a silo that combines applications and infrastructure and storage. And it sounded awesome for a couple of years. And now everybody is going, oh man, I got to re-platform out of Hadoop, which means I got to rip apart my analytics landscape, my computing schedulers, my operating system environments, and of course, my storage, right. And hyper convergence is the modern version of Hadoop.

So, we'll talk about how you can build a silo-free infrastructure that's hybrid multi-cloud by design, that supports the needs of any app any data type, anywhere you want it. And so, I'm excited for our product teams to talk to you about the announcements we have. We have a rich innovation pipeline as well, and you'll hear more from us over the next few quarters as well.

So, stay tuned. Thanks for coming. I'm excited about having you all here, and certainly excited about the real customer problems that we're solving with our technology and the partnerships that we formed over the last several years.

Kris, back to you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Am I on? Okay, I'm on. All right. Thank you, George. We appreciate your time today. I'm glad we got a few karate chops out of you. It's good to see that energy and impact. All right. So, next up, I'm going to invite [ph] Russ Fishman (00:14:40) and Andy Sayare from our AI group and you guys are welcome to sit or stand, whichever you're most comfortable with.

Unverified Participant

I think we'll sit.

[indiscernible] (00:14:53)

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Okay. I'll join you up here.

Unverified Participant

Yeah, perfect.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. And so, while you guys are pulling all your AI questions together, we figured we'd get this one off first, because I know it is top of mind for everyone. I'm going to ask [ph] Russ (00:15:07) and Andy to introduce themselves and say a little bit about what they do and you guys can start queuing up your questions.

Unverified Participant

Perfect. Thanks, Kris. Well, thanks for having me today. [ph] Russell Fishman (00:15:19) is the name. I'm responsible for Product Management at NetApp for our solutions business globally. Most importantly, AI. And I'm joined here by Andy. Andy, do you want to introduce yourself.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

Hey, everybody. My name is Andy Sayare. I run our alliances for our AI go-to-market. I've been at NetApp for about ten years and lovely to be here today and nice to meet you all.

Unverified Participant

Perfect. Well, listen, I think we talked about having 5 minutes of sort of an overview and then we'll open it up to questions. So, I'll start off by saying, obviously, the market is moving very fast in AI, but it's not a new thing for NetApp, right. NetApp has been focused on the AI market for over five years. We have hundreds of customers who have invested hundreds of millions of dollars with NetApp in AI. And we're in a sort of, I would argue, a pretty unique position in that rather than just taking our products and applying them to the problem statement of AI, what we've actually been doing is we've been investing in our products to make them AI ready, and we've been doing that over the last few years.

So, what we have now is this fantastic portfolio that can be applied in lots of different ways, puts us in a really different position to others that I think will just put a sheen on their products and say everything's AI ready. We've really built our products to help our customers adopt AI. You probably heard George talk about three things that the company is focused on. I'm obviously here primarily to talk about how we help our customers adopt AI and really the focus for NetApp is exploitation of data more than anything else. So, this concept that the customers are sitting on data that has value that they are unable to derive without the help of AI, that's what we're really there to help them with.

We're obviously working to invest with AI in our products. That means that using AI to make our products better, things like anti-ransomware protection, that sort of stuff. And obviously making us much more efficient as a company when it comes to operations and development, et cetera, et cetera. But in terms of helping our customers build, so there's a few things that we do that I think no one else really does in the business.

Firstly, AI isn't really focused on any one particular deployment methodology, so it's not really about on-prem, it's not really about cloud, it's really about all of it. And NetApp has this very unique position in the market, which is that we have a leading storage operating system, which is available on all the clouds and on-prem. It makes it much easier for our customers, regardless of where their data sits, where it's being generated, where it needs to be used. We bring all those worlds together, and that's really important when we're not talking to IT because typically when we're talking about AI, we're not talking to IT. We're talking to data scientists, we're talking to lines of business, we're talking to CDOs and data owners. These are unique and different buying centers that NetApp typically has addressed in the past and we've really worked hard to craft our products and our message to go after that, which is why we have been successful.

The success, though, is not just about what NetApp does ourselves. It's also about how we work with our partners. And so, I was going to ask, Andy, would you just speak for a few minutes about what we're doing with our partners and how our partner ecosystem [ph] has helped us (00:18:38)?

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

Yeah, absolutely. As [ph] Russell (00:18:39) said, part of the success for NetApp in AI has been the co-innovation that we've done with our partners. Our leading AI partner is with NVIDIA. And NVIDIA and NetApp have together delivered five or six unique solutions that we've co-developed, co innovated and have brought to market and as [ph] Russell (00:19:00) said, have been adopted by hundreds of customers.

So, we are in addition to NVIDIA, we also have other partners in AI. We work with Domino Data Lab, for instance, Domino has a MLOps platform that is available on both cloud and on-prem and works in a hybrid manner. So, it's very complementary to NetApp's offering. In addition to that, we work with Run:AI. We've worked with several of

computer vendors, including Cisco, Lenovo and Fujitsu. And of course, we work with the hyperscalers as well. So, all of these – this ecosystem comes together to bring solutions that work for our customers.

As [ph] Russell (00:19:44) was saying, when you build AI for an enterprise, oftentimes you're drawing on data that's stored in one cloud or another, sometimes multiple clouds. Often it's on-prem, perhaps in multiple locations as well. And the key to successful AI for many companies is to be able to bring all that together, to be able to train models and be able to deliver the kinds of solutions that they're building. So, that's a quick overview. And when we get to Q&A, we'll be happy to talk more specific about those.

QUESTION AND ANSWER SECTION

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Well, thanks, guys. I appreciate that overview. I think it was a great level set for what we do in the world of AI. So, I will look out to everyone and we got some questions.

Tim Long

Analyst, Barclays Capital, Inc.

Q

Oh, hi, I'm sorry. It's Tim Long with Barclays. Could you talk a part – two part question here. Could you talk a little bit about as we move into this AI world. Maybe talk about, like, quantity of storage and maybe quality of storage. Obviously, some of the largest customers are not using NetApp [ph] for Dev (00:20:57). They're using white-box solutions. So, can you talk a little bit about how you see this transition to AI affecting vendors like NetApp and you could talk – if you could just touch on the hardware side as well as the software side, I'm sure there's a pretty good software play as well.

A

Yeah. What I would start off by saying is that, NetApp really focuses on the entirety of the lifecycle of data. So, that's the first thing. So, I think typically people will talk about this concept of speeds and feeds and pushing data to GPUs. You hear a lot of folks talk about that. That's really just one part of the lifecycle, though. So, NetApp's really thinking about where the data is being generated, how it's being organized, how it's being unified, how it's being prepared, how it's being then stuffed into GPUs for training. Obviously, very important part of it, but that's just one little bit of it. And then, of course, what happens after that? So, what happens when the results, the validation of that model happens, the iterative nature of that kind of life cycle of AI training. So, there isn't really a good answer for what you're asking. I think if you really just focused on the bit where you stuffed data into GPUs, you'd be very narrowly thinking about this market that what we're seeing is that the entirety of that lifecycle is what's driving revenue for us.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right, Sidney [ph] over here (00:22:19).

Sidney Ho

Analyst, Deutsche Bank Securities, Inc.

Q

Thanks. Sidney Ho with Deutsche Bank. Andy, you talked about this core innovation within video with five or six different products. Can you expand a little bit? How does that being adopted by customers? What's kind of [ph] attach (00:22:31) rate, can customers use other vendors during the process?

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Sure. Absolutely. So, we started out five years ago with a video with a reference architecture. Essentially what we did was we paired our storage technology with and video server technology, the DGX platform, and they're switching the Mellanox switching, which they acquired. So, together that made what we call ONTAP AI, which is essentially our operating system managing AI workloads in this converged infrastructure stack.

So, that was the basis for many of our customers initially getting started with us. Over the years, we've made some variations to that based on the way customers are adopting that technology. So for instance, some people don't want to build that infrastructure on premises. It's expensive. It requires a lot of power and cooling that most corporate data centers are not equipped to deliver. And so, customers are looking for alternatives to that.

So, one alternative is what's now called DGX Cloud that NVIDIA offers. It's currently available in OCI, but it will be moving to Azure and GCP and other clouds in coming months. And so, their customers can then consume the same kind of architecture but not have to have it on-prem themselves. They're essentially renting it. Other customers are looking for a model where they can have their own dedicated equipment and once again not put it on-prem, but put it in a colo like Equinix, which can then be connected to all the clouds and allow them to manage their own equipment. But in fact, let Equinix do the management, the daily operations of that equipment. So, there's just three variations of a delivery model for our reference architecture.

In addition to that, we work with NVIDIA on their SuperPOD, which is of course large scale AI training, often focused on large language models and other very large training situations. And between all of these, we have – as we have had said earlier, literally hundreds of customers running on those today.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Thank you. I think, [ph] Gloria (00:24:59), there's a question from the webcast.

A

[indiscernible] (00:25:03).

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Okay. Okay. Never mind then. All right. We'll get to you next, David.

Victor Chiu

Analyst, Raymond James & Associates, Inc.

Q

Hi. Victor Chiu from Raymond James. Can you remind us which platforms specifically are strategically positioned to target AI/ML workloads? Is it a combination of the A-Series and the storage grid [ph] unstructured (00:25:24) solutions? And then in general, where do you see storage [ph] falling (00:25:28) in the spending priority relative to computing, GPU, hardware acceleration when it comes to building out AI solution.

A

Yeah, you're probably going to hate the answer that it's everything, that it really is, right. And again, back to that whole concept of lifecycle. So, everything from – honestly, we still see hyperdisk, we see flash, [indiscernible] (00:25:49) and that's on the sort of [ph] unstructured (00:25:51) NFS side. We see objects. We obviously have a file-object duality on our ONTAP systems. We also have the storage grid solution as well. And then, we see a lot of take up with our 1P Cloud Solutions as well.

So, it's really all of it. And over time, I don't see any one of them really pulling ahead. I think maybe we'll see more capacity flash over time. That would be my guess at an industry level, as a focus because a lot of that, again a lot of that pipeline isn't stuffing high speed data into GPUs, that's part of it. But I kind of think about that as the last mile.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Yeah, I would offer one other piece here that's – that I find very interesting, which is, AI is all about data. I mean, that's obvious. Everybody understands that when you bring data together and be able to apply machine learning to it, then you can get insights that haven't been available before. And yet, when companies go and say, I want to do AI, what's the first thing they do? They go out and buy GPUs.

And the GPU servers are generally the first thing that customers think about. They begin to deploy those GPUs, and they realize, oh my goodness, we can't keep these machines utilized. So, we need to have a data infrastructure that's going to allow us to keep those machines utilized. So, it's worth the \$0.5 million per box that it's going to cost them to acquire those.

So, what happens very quickly in many of our sales cycles is we get brought in almost immediately following the purchase of GPUs and GPU servers, and then we help them build out a data strategy that works for AI, for those companies.

A

This is worth mentioning, I think someone said, oh, you know, white-boxes and that sort of stuff. The reason that people come to NetApp is for the data management.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Yeah.

A

That's why we win and we win consistently because of that.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

And just as a quick translator, [ph] Russell (00:27:49) mentioned 1P Cloud Services, for those of you who aren't fully versed in the NetApp lingo, that is the First Party Cloud Services. So, the Azure NetApp files, AWS, FSx for NetApp ONTAP and then most recently Google...

A

Cloud...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

...Cloud.

A

A

...NetApp Volumes.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

...NetApp Volumes. All right. And...

[indiscernible] (00:28:12)

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Let's get David here. [indiscernible] (00:28:14).

A

A

Yeah.

David Vogt

Analyst, UBS Securities LLC

Thanks, guys, for doing this. David Vogt at UBS. Just maybe as a follow-up, can you kind of explain sort of what the revenue opportunity or the software mix looks like as we move from traditional storage use cases to multi-cloud AI use cases, higher GPU utilization, to your point about, adding on a \$0.5 million box. Should we expect a much more robust software story going forward, as this permeates to multiple different sort of end customers and use cases relative to where we were over the last 5 to 10 years?

A

Q

Well, I mean, so back to that point about data management, right. If you look at NetApp's portfolio and so you think about BlueXP, for example, as a unified control plane. If you think about our essentially a software – to

A

software defined storage products, if you think about other things like InstaClustr, for example, and what was known as Cloud Data Sense, but now known as BlueXP Data Classification, we're starting to see a lot more pickup in those areas, right. And that makes sense because, when AI started, it was very much a data science conversation, right.

So, lines of business data scientists, they're the ones we were really engaging with. What you started to see was CDOs and data owners getting very concerned about how that data is being used, right. I was like a moderator in a nuclear reactor, right. They're coming in with carbon rods, they're slowing everything down. And there's this sort of push and pull between the data scientists who want to get out there and move quickly and the data owners who want to slow everything down.

So, that – what's interesting, of course, is that NetApp has this amazing portfolio capabilities that can address both sides of that, right. So, what we're starting to see is a lot more engagement on the data owner side as well. So, my expectation is that more of our data management capabilities will be consumed at time.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Yeah. And I think the other factor to consider here is regulation. We fully anticipate that regulation is going to hit the industry and customers are going to need to comply with that regulation. And we're encouraging customers to start looking at that today and using tools, as Russell mentioned, that can automatically filter out personally identifiable information, for instance. So, that never makes it into the models or to be able to create auditable models so that if a customer or if the government, for instance, wants to go back and look at how your AI came up with a particular answer, you have a stored image of that model with its data at the time that, that was created, which we think is going to be very important moving forward.

A

And the last thing I'll just say is that regulatory compliance, I guess you hear that all the time. What we hear more from customers, anything else is, is commercial concerns about the sensitive – commercial sensitivity of data and the reputational risk of having that data leak. It wouldn't necessarily be a legal issue in some cases, but it would be obviously a reputational issue.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Absolutely.

A

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Let's go to [ph] Gloria (00:31:11) with a question from the webcast.

A

Hi. This is a question from Aaron Rakers from Wells Fargo. Does the DGX Cloud Solution utilize NetApp ONTAP AI as primary storage backend versus other alternatives?

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

So, I'll take that. So, currently, no. So, currently, DGX cloud is leveraging an alternative storage for its scratch space. So, this is the area that the GPUs use as extra space from their internal storage. What NetApp is doing is working adjacent to that to help connect the data from multiple sources, multiple clouds and on-prem, to be able to bring that data together, to be able to train those models.

A

Back to the lifecycle point again, right, which is that, scratch space in general is just that – is really the last mile, right? It's ephemeral data, which means that it's not stored long term, it's not protected. The data management capabilities are generally not included there. What we're finding from customers is that, that is not us – that's not a complete solution.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Yeah. All right. How about Meta in the back?

Meta A. Marshall

Analyst, Morgan Stanley & Co. LLC

Q

Couple of questions. Maybe first, you guys talked about customers getting a NVIDIA server and then figuring out that they need to have a data management solution behind that. But in many cases, they haven't even gotten the NVIDIA server yet. They're still waiting. So, I guess just where are [ph] like and – staying (00:32:50) with kind of the regulatory conversation. So, I guess, just where are you in your kind of customer or where are your customers kind of in their conversations of thinking about this?

And then maybe just as a follow-up like over the next five years, is the greater opportunity for you guys on the data management to prepare for training or is it on inference, eventually inference?

A

That's a good questions there. Well, I mean, I'll take the second part first [indiscernible] (00:33:18) and we'll [ph] go back (00:33:20) from there. If you think about the lifecycle of AI, inferencing is the runtime essentially, right. I think it's a little iterative, so maybe I'm oversimplifying it. But training is development. There's been a lot of focus on training last year because everyone's been working out how to exploit AI. But we're moving rapidly into a new phase, and that's the operationalization of AI, right. And operationalization is all about inferencing. It's interesting because it's a completely different [ph] buying sensor (00:33:48).

So, in some of our more mature customers, especially in financial services, for example, who have already heavily adopted AI. AI is becoming part of their core business processes. When something is part of your core

business processes, you're worried about all the same things that you would be with any other enterprise service. So, that's manageability, observability, operationalization, supportability, right. And that's a different [ph] buying sensor (00:34:12), but that's very well aligned with NetApp's traditional messaging, right. Our products are designed to be easy to manage. They are regularly accepted by IT departments. They are well-situated data centers, et cetera, et cetera.

So, we do see a huge opportunity on the inferencing side. Inferencing is going to be 85% of the runtime, if you will, the lifecycle, but not necessarily 85% of the revenue. I'm not saying that. Before Kris [ph] slaps like (00:34:43). I'm not saying that. Obviously, it's a very different set of performance characteristics and storage requirements around that. Do you want to take [indiscernible] (00:34:51)?

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

Yeah. So, the other thing I would say is that, we naturally have been doing this for five years and there are a lot of customers doing AI and machine learning across a wide variety of use cases where there are DTXs and other systems deployed. And data management has been an important piece of those deliverables.

So, we could talk about use cases in manufacturing for defect detection, we can talk about computer vision cases, we can talk about a whole bunch of different cases across health care and life sciences and financial services where customers have deployed these solutions. Remember that these large language models have really only captured our imagination here for the last ten months or so. And while there is a lot of interest and a lot of noise about that, AI has been around for a while and there's been a lot of customers that have deployed these systems.

A

But I would just say one other thing, in terms of specifically DGX, right. So, DGX is a relatively small part of the [ph] training (00:35:50) market mostly NVIDIA has been quite open that they are selling into [ph] OEMs (00:35:55) like – like the Dells and the HPs [ph] and what have you (00:35:58). We sell successfully into all of those environments. We're not really tied to NVIDIA service where we work with all of the big manufacturers.

Last thing I'll just say is that, some of and NVIDIA's most advanced [ph] training (00:36:12) GPUs have, it's well understood that they have a significant supply chain issue right now relating to the lithography that they're using, the two millimeter process, two nanometer process through TSMC. There are a bunch of other GPUs out there that are much more readily available that, NVIDIA has been openly pushing people towards include things like the L40S and what have you. So, we don't see the market gummed up at all, if that's what you're asking, not at all.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Irvin, right down here.

Irvin Liu

Analyst, Evercore Group LLC

Q

Hi thank you. This is Irvin Liu with Evercore ISI. So, do you see a share gain potential or opportunity presented by AI or is or most organizations are going to stick with their incumbent vendors and avoid a major upgrade or a major transformation prior to jumping into the AI journey?

A

I think I'll start by saying, I think we're extremely well positioned. And again, back to that comment that we've been building for AI for five years, five and a half years. So, I think we have a portfolio that makes us competitive in non-traditional NetApp and non- NetApp customers, that's how I would describe it. Again, because of our portfolio, because of our capability. So, yeah, I mean, I think there's – there's always the opportunity.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

And I would just add to that, I think, you're getting more specific. Customers tell us they want to be able to aggregate their data from multiple sources. It's difficult for many of our competitors to do that. That's something that we have put a lot of investment into and are able to bring that data together for customers. So, we think we're very well positioned to help in this hybrid world.

A

And lastly, I'll just say is [ph] it still was (00:38:02) an engagement strategy, right. So, firstly, the ability to go and converse with the data scientists, understand what their life is like and be useful to them, and understand that intrinsically data scientists, for example, don't care about infrastructure. That's not an interesting thing to them, right. So, the first question is who are you and why are you here. But we've become very good at connecting the challenges that a data scientist has in accelerating AI adoption and development to our underlying value. That's how I describe it.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Why don't we grab Steve here in the middle?

Steven B. Fox

Analyst, Fox Advisors LLC

Q

Thanks. Steve Fox from Fox Advisors. I think on the last conference call, George talked about how the real uplift in AI comes with the industrialization of the applications. And you touched on it a little bit with healthcare and defect detection, stuff like that. Can you just sort of talk about how you think that develops because if that's really when the S-curve takes off, what do you envision like over the next few years where those applications are most likely to develop and develop quickly?

A

Yes. So, I'll talk from an industry perspective primarily. So, what's really interesting about our experience is that it's not really tied to any particular industry vertical. And actually, there hasn't been – we haven't seen a lot of commonality in use cases. The engagement model with customers has typically been around, let's understand what data you have and how we can exploit it and then find the right application that would actually be useful. It's kind of how I'd describe it, that has started to change at an industry level.

So, we're starting to see a number of horizontal use cases appear, come to the fore. Probably the most obvious one is based on generative AI and chat bots, specifically customer service chat bots, which, I think what's really

happened is, is that it's moved from a – this is a way for customers to innovate to a this is required for us just to stay competitive. And so, if you talk about the S-curve, that's where I think the S-curve is hitting because it's going to hit every single customer [ph] who (00:40:07) wants to talk to us about generative AI and they want to talk to us about those sorts of use cases. And those become very repeatable, mostly because of the use of these pre-trained models, what they call foundational models. So, we're starting to see a lot of that. So, yeah, I think, yeah, from an industry perspective, I think we're just at that point, that inflection point right now.

Andy Sayare

Director-Global Strategic Alliances, NetApp, Inc.

A

I could also add that I was fortunate to attend the TED AI Conference in San Francisco last week. Did anybody get to go out for that? It was fascinating conference, speakers ranging from Andrew Yang to Stephen Wolfram. One of the things that I learned there was that there are 2,700 funded startups for large language models right now. I was blown away. About 35 of them have done foundational models, the rest of them are all building on top of OpenAI and Llama 2 and a bunch of other models and going very specific with [ph] one (00:41:07) particular area of focus they're going after. So, we're going to see an explosion of large language models over the next couple of years that are going to be very tailored to specific industries and specific use cases.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. We have time for one last question from Mehdi. It's got to be a quick one, though.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Make it quick. Thank you. Mehdi Hosseini, Susquehanna International. Just as a follow-up talking about actually the decision makers implementation, talk about data scientists. I think a lot of these points you're making, you're referencing to the folks that are involved in training the model. What I want to learn from you is in the next 15 seconds is who's actually going to deploy this at the enterprise level? Who is actually going to be working with CIOs and CFOs in deployment of these train models. It's great that there are 200 startups, but to go from a startup to actual use to actual deployment to actual realization of improved productivity is going to be the key. And to me, quite frankly, a Chatbox has been the most frustrating experience, and I compare it to Alexa, Alexa didn't really lead to significant growth in the storage, but deploying of AI for product – for realizing productivity improvement, I think could be a key. And I just want to learn from you. How are we going to go through that journey? Who are these decision makers you are going to employ to do that?

A

Yeah, there's a lot there. [indiscernible] (00:42:42) I'll try to make this quick, as I'm going to get in trouble otherwise. So firstly, IT are the folks that are operationalizing and deploying this stuff, right? Ultimately, they're the ones that are tasked with waking up at 2:00 in the morning if something goes wrong and getting it fixed. So, there's a lot of focus on getting IT ready to do this stuff. But the data scientists are helping make the decisions, but the lines of business have the [ph] checkbooks (00:43:09), right. And as I said, the CDOs and data owners are the ones that kind of holding it all back a little bit because they're concerned about regulatory compliance, commercial concerns, et cetera., et cetera. So, yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Well, I'll have to give you the hook now.

A

Yes. Okay.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Thank you guys so much.

A

A

Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

I really appreciate you coming. I know there are more AI questions. Just reach out to IR. We can help hook you up...

A

A

Yes. Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

...after we announce earnings [indiscernible] (00:43:36).

A

A

Thank you, everyone.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Now, I am super excited to introduce Octavian Tanase and Sandeep Singh, who are going to talk about Enterprise Storage. So, guys, come on up. You are welcome to sit or stand, whatever you're most comfortable with. It's entirely up to you. Whatever you want. All right. Standing. Okay. So – sitting. Okay.

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Sitting. Sitting it is.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right, well, why don't we kick it off with each of you introducing who you are and what you do at NetApp?

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Hi, everybody. I'm Sandeep Singh and I'm the Senior Vice President and General Manager for Enterprise Storage. I've been with NetApp now coming up on 11 months. And why I'm super excited about being here at NetApp is we are enabling unique outcomes for our customers across the spectrum of helping customers to be able to save money with the lowest cost of storage over the data lifecycle, helping them simplify at scale and through that lens, increase productivity, lower risk and helping them become more secure and protect against ransomware and cyber security attacks and through that lens of having the most secure and protected storage infrastructure, helping them become more sustainable as well and then ultimately harness the power of cloud and AI as and when they're ready. That's why I'm super excited to be here in at NetApp, and to be here with you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Octavian?

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Good morning. My name is Octavian Tanase, I'm the engineering guy, so I'm ready to answer your question and your question and your question again, from the perspective of somebody in engineering.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Perfect. And some of you might recognize Octavian. He's presented at these events for us in the past. He's been at NetApp, not as long as me, but a good piece.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

You guys could have bought more. I would have been retired. But here. I'm back here.

QUESTION AND ANSWER SECTION

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. So with that, we're going to open it up to questions. Otherwise, I'm going to have Octavian re answer the questions that were asked.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

A

But there was an awesome question. I'm sorry I didn't catch your name.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Mehdi Hosseini.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

A

Mehdi. Okay.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

[indiscernible] (00:46:06) in the back.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

A

Right.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Can I elaborate my question?

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

A

Please.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Sure.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

I just want to know, like, if enterprises are going to go hire consultants or if Dell is going to build out their consulting, how is that going to play out? Because I think ultimately enterprises are going to limit their IT staff and

rely on somebody else to come in and tell them how to deploy it and tell me if you disagree or how is that going to play out?

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

A

I actually don't know who hires consultants at NetApp. I think it's mostly the CEO or the board. So, let me tell you what I think it's happening in engineering, right. Everybody – in any large organization, your ROI comes from making your engineers more productive, right. And generative AI, it's a great technique for that, right. So, there's a concept of copilot. So, if you in, somehow are able to introduce in the edit, compile, debug, process that an engineer goes through [ph] a co (00:47:07) and AI assist, that engineer will be able to write better code, more secure code, more effective code, right.

So, I believe that any large enterprise that has a pool of engineers will want to use generative AI. The question is, can they – can that be done in a secure, less risky way, right. So, if you, let's say you have some proprietary information, you don't want to lose copyright, right.

Because all of a sudden there is a generative AI engine, that uses that. So, we believe that there are safe ways to do that. I think first of all, you can create a taxonomy for your products in code and say what is core versus context? Let's say for your context code where you may not necessarily care about your IP, you can use straight up a OpenAI public LLM. For something that is more sensitive, perhaps you can use, I think [ph] Russell (00:48:06) talked about the pre-trained LLM [ph] that (00:48:09) you can deploy within your enterprise and you can augment that with some interesting proprietary information. I think that process is called fine tuning, right?

So, we expect a lot of these enterprises to come in and take petabytes of data and information and augment with that data, the pre-trained LLM. So, it could be perhaps a good co-pilot for somebody in engineering. Makes sense. And I don't know if this is going to be done through consultants or not. I think just the way you've seen the explosion of consultants in the cloud space, probably there will be an explosion of consultants in the AI.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

So, what...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

On the mic, please, so the webcast can hear you.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

May I have a follow-up?

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

A

Please.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

You got a mic? So...

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

So, I have time. [ph] I mean, you can do (00:49:03).

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

So, can – is there a business opportunity for NetApp to help customers to do this? Is there, I don't want to say consultant, but is there a business to build around this in terms of deploying it, as you sell terabytes of storage but then who is going to be at the other end to deploy?

Q

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, we're not talking about future business endeavors nor financial things at the event. But I will say we are helping customers plan and decide their AI data management journey already today, right. We're working with hundreds of customers who are currently deploying AI and we're helping them figure out what that data management behind it is. So, I mean, generally, yes, there is opportunity for us to participate at a higher level with our customers. Do you agree?

A

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

100%. What you just articulated, Kris, is exactly the pattern that we're seeing out there. All of the learnings that we're getting in working with hundreds of customers on their journey to AI or GenAI, we're taking those and then sharing those best practices with the customers.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. David in the front.

A

David Vogt

Analyst, UBS Securities LLC

I just want to follow up on that point. So, as it becomes more pervasive in terms of training models, inference, data sovereignty, what is the differentiation that all of these sort of applications bring to an enterprise? Doesn't it become ultimately somewhat table stakes in the sense that there's multiple offerings? They tend to be somewhat similar and sort of you have to do it. But there's really no revenue uplift for the corporate, not for NetApp and I'm saying for the company. And it's more of a cost saving tool at this point.

Q

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

I think it's hard to think of a company that has such a complete end-to-end capability to address all the phases of machine learning and AI, right. So, you guys have talked a little bit about inferencing and it was emphasized as very important, but most people will probably have to acquire data and there is a phase where people take data in, put it in a data lake and try to cleanse it. That's a process that takes time. It's automation intensive and so forth and you need that large practically infinite storage container to do that.

A

And that storage container needs to support heterogeneous data sources, not just unstructured data, but sometimes structured as well. So, [indiscernible] (00:51:52) this data lake that has that simple interface to read and write data into? We got that. Then there is a phase where you're training, your data based on some algorithms that you've chosen. You need a lot of throughput, right. So, you want to be able to take some of this data from the data lake in a very simple, cost effective way, move it in the place where you train the data that it has to be very close to compute right to those NVIDIA DGX systems or GPUs that they've talked about, right.

And then after you train the model, then you're going into the inference phase where latency not throughput is the most important thing, right. Because you're going to ask me a question. You want a quick answer there? So, I feel that our flash systems running the ONTAP data platform are uniquely positioned to, in one system support that whole data pipeline from data acquisition in the data lake to the model training to the inference. What do you think?

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

A

I think the – what Octavian is sharing is basically we provide the customers that flexibility end-to-end with really infrastructure that is ready for AI when you – to that prior comment that was made which is basically AI works on data, data runs on NetApp. And as we look at, basically generative AI, it's tied much more to unstructured data. That is where we are looking at how we can enable unique use cases for customers there. And at that very high level in terms of what you mentioned, I see kind of three ways so it becomes important for organizations to think about. One is the productivity enhancements that Octavian was just talking about developer productivity becomes incredibly important from an AI perspective.

The second area is all of the supporting functions of how AI can bring productivity boost there. And then the third key area is really how can AI help design new customer experiences or net new overall business and revenue models or product models there. So, that's where I think AI becomes a both a table stakes as well as a opportunity for customers to innovate with their data.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. And we have a question on the webcast.

A

This is from Samik at JPMorgan. He's alluding to the move of AI to on-premise to leverage data. His question is, if that mix shift does happen, how can you increase your differentiation to some of the other large storage vendors? Is adding more products to the portfolio or simplifying the portfolio the way to go?

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

A

So, with the data, AI running on data and largely a lot of the data sitting on-prem, NetApp [ph] if gets (00:55:10) the opportunity now to help customers leverage their data. Certainly, the – lot more of the AI and generative AI is working on the unstructured datasets. So, first is basically with the data that's already resident on the NetApp storage infrastructure, we can make it seamlessly accessible to the data engineers, to the data scientists, and that's through the integration with the MLOps platforms.

So, that's one of the areas. The second area that was talked about in the prior session is really enterprises as they're looking at adopting and deploying predictive AI and generative AI. Certainly, the data privacy becomes important, the model, traceability and the associated datasets tied to those models become incredibly important. This is where NetApp's data management capabilities become pivotal in enabling our customers with a overall model, traceability model, versioning use cases. That's where we see a tremendous opportunity of helping responsible AI deployments in the enterprise for our customers.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

That's awesome. Can I [ph] ask (00:56:37) something about Kubernetes?

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

You absolutely may [ph] ask (00:56:40) something about Kubernetes.

A

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Well, because so most of these AI applications, ML applications are modern applications, right? So, the lingua franca for application scale in virtualization right now, it's Kubernetes, right? This has been born in the cloud. It's pervasive in the enterprise. So, we believe that our investments that we've made in building a Kubernetes middleware for applications to simplify the lifecycle of deployment, in protecting these applications will come handy because we believe that many customers will go back and forth between, the on-premises [ph] estates (00:57:16) and the cloud, depending on the type of services that they want to take advantage of, right.

A

Moreover, we have some interesting technology. It's called FlexCache. You're not going to remember, but it's like a caching technology that kind of helps one take the data from its source and make it available to compute. But there's more. I thought that they were going to talk about GPUDirect. GPUDirect, it's an interesting technology by NVIDIA that it's trying to kind of simplify data access from the, let's say the server or the network storage device that has the data all the way to the GPU itself.

So, the way that works is right now, let's say you would be talking, let's say, over – using NFS. So, that would be me talking to you. You're the CPU, you're the GPU. And I'm telling you something. You're telling her something, right. Then we invested in a technology called RDMA, which is now I'm going to be able to not say something, but the brain, right, the memory, will be connected to his memory. And then back into Kris' memory, GPUDirect, basically it's memory to memory between the data and Kris' brain, what is a GPU. So, that's a technology that we've implemented recently and we have tremendous performance results in like 170 gigabits a second in getting the data from network storage directly into the brain of the GPU. Did that – was that okay?

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

I think that was good, although it's frightening to think that we would have a direct memory connection. I know some people in the audience have question about QLC technology, I'm sure of it, or we can keep talking about AI.

A

Victor Chiu

Analyst, Raymond James & Associates, Inc.

Q

Just one last quick question on [indiscernible] (00:59:25). Victor Chiu from Raymond James. In the intermediate term, do you envision the growth in demand for AI solutions potentially cannibalizing traditional storage orders? Or do you see it as purely an incremental opportunity? Maybe elaborate on how that plays out.

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

A

I think that will continue to bear out in the market. The reality is that data is essential to every organization, and so data continues to grow. Data is the underlying infrastructure tied to all of the application workloads. So, customers absolutely need to continue to service the need for the underlying data infrastructure to support, whether it's [ph] their (01:00:11) high performance [ph] file (01:00:12), their virtualized environments or their containerized Kubernetes environments, right. In addition to looking at basically how do they harness the power of AI, make it a competitive advantage before it becomes unnecessary table stakes across the board. So, I think there's an opportunity and a need, more importantly, for customers to do both versus one or the other.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Okay. I'll go with Tim and then we'll get you.

Tim Long

Analyst, Barclays Capital, Inc.

Q

All right. I guess I'll ask the QLC then.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Thank you.

Tim Long

Analyst, Barclays Capital, Inc.

Q

Yeah, just outside of AI, that's obviously one of the trends is you guys have been in several quarters now with the QLC based product. Could you talk a little bit about kind of what that's doing to segmentation within the market. And how – where are we in the continuum? Is this something that's going to permeate more through different applications, use cases and in different parts of your portfolio? So, can you just kind of give us a sense as where we are and where this is going and kind of your differentiation with it as well? Thanks. Yeah.

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

A

So, I'll touch point kind of – the little just a short history of kind of what we introduced earlier this year and where we are. Earlier this year, we announced and introduced capacity flash through the lens of the NetApp AFF C-Series and that was designed to help provide a basically value proposition of near the speed of flash at hybrid economics for customers and really targeted at three key use cases. One is basically customers who have hybrid flash or hybrid storage systems with 10K based hard drives, making it affordable for them to transition over to all flash and be able to get faster, denser, more sustainable, that is continue to happen.

Secondly, for the application workloads, whether it's virtualized applications or database applications or other file environments where roughly 2 to 4 milliseconds of latency is more than enough for those application workloads, our capacity for flash provides that best price performance for the customers. And then the third use case is

targeted at the disaster recovery and secondary storage use cases. Since our announcement and the launch of C-Series, we announced it in February. We started shipping in March. We're seen with capacity flash with C-Series. It's become the fastest ramping product within NetApp's history.

So, that's where we are at this point where now in our portfolio, when you look at it, basically we have our AFF A-Series for the best performance for the performance intensive application workloads, we have capacity flash as providing that price performance proposition of mere speed of flash with hybrid economics and then we've got [ph] FAS (01:03:28) with the lowest cost overall. All of these design centers are fully interoperable.

So, when I mentioned, basically we're able to provide customers with the lowest cost is what data, data we know is dynamic. A lot of the data is cold and so customers are able to leverage C-Series and seamlessly leverage either our [ph] FAS (01:03:54) systems or our storage grid systems and have automated and granular tearing as well as tearing all the way into the public cloud for providing that lowest cost data over the data lifecycle. The other thing I'll mention is NetApp is unique in enabling all of these use cases as well as across [indiscernible] (01:04:15) Unified as well as block an object with an underlying single storage OS.

That's NetApp ONTAP. Thanks to Octavian and team and what that is enabling for customers, whether it's for serving their different application workloads or a structured or unstructured data set or across on-premises or with the first party native cloud storage services and the public cloud is the simplicity at scale, not just within silos, but simplicity at scale. And that's enabling customers to remove complexity of bespoke infrastructure silos. It's enabling them to increase productivity and lower risk by having consistent management and automation, consistent and comprehensive data security, consistent and comprehensive data protection across the board. And an overall consistent support and vendor experience.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right, thanks. I know, Sidney had a question here.

A

Sidney Ho

Analyst, Deutsche Bank Securities, Inc.

Thanks. Sidney Ho with Deutsche Bank. Just want to follow up with the C-Series, that right now sounds like you guys are targeting the 10K hard drive. If you kind of look at road map for the company, is that a road map that potentially at cannibalizing 70 to 100 as well in the future?

Q

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

I won't be able to comment on anything road map.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

But you can opine about technology trends and when flash starts to erode that.

A

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Yeah. I think, so when we look at the future, we're continuing to look at opportunities for helping customers be able to get to all flash and be able to through that lens, right accelerate, get more efficient and denser as well as more sustainable. We will continue to look for those opportunities into the future. When you think about basically

A

with all flash, certainly QLC technology has made it much more such that the storage operating systems that are flash optimized, that can write to flash in a flash friendly manner are able to make sure that the endurance levels are such that customers can get a 7 to 10 year lifecycle with those systems. And that's how what that is enabling customers to then be able to shift over from the existing 10K hard drive-based systems. With the 7.2K near-line SAS systems, that is still further out in time. And when you look at basically customers being able to get the lowest cost through that lens and still be – especially from a hybrid flash perspective of still being able to get really good performance as well as high availability of the systems, that from a flash dollar per gig just a raw perspective that's still further out in time.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

And just in case it wasn't obvious, ONTAP is one of those flash optimized...

A

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Yes.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

...operating systems.

A

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

And QLC is a cool technology. We like it even though, initially we thought, man, four electrons on the cell there, flipping them on and off there. It's going to be hard. Apparently, they have their own personality. And they're quite resilient, right. I'm making a joke because initially when that technology was introduced, people weren't really sure of the right cycles and all that stuff. And what we've proven in the last few years is that QLC technology, it's awesome, right.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yeah.

A

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Yes.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Additional questions. You guys pepper me with QLC market opportunity, customer use cases all the time. I can't believe you're so quiet right now. All right. Well, to wrap it up, I know you guys spend a lot of time talking to customers. What are some of the cool things that you're seeing customers doing with NetApp technology and the unique reasons that they tell you they're choosing us, that what NetApp can do, that no one else can do?

A

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

They're very excited about QLC.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

I shouldn't have shut down the AI conversation, apparently.

A

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

So, I spend roughly 50% of my time traveling and meeting with customers and partners in terms of the unique things that they're able to enable with NetApp. And the one pattern that I've seen is customers who are migrating and looking at harnessing the power of public cloud and being able to deploy and run their mission critical application workloads. NetApp is enabling that for them. At the same time, those customers as they're consolidating data centers and they're looking at bringing much more of that agility and a cloud like experience on-premises. We're seeing the combo of where basically the customers are leveraging NetApp Keystone for on-premises and then the NetApp first-party cloud storage services for being able to deploy their application workloads in the cloud. That is a unique area.

A

Second unique area, Octavian touched upon FlexCache. And so, customers who are, [ph] EDA (01:09:55) or high performance file use cases where they're going to have a large presence on-prem. But then they, developers or other personalities or personas within an organization want to be able to leverage subsets of data and through the lens of cloud. FlexCache becomes an integral portion of that technology that customers are leveraging to be able to get that hybrid workflow enabled for them. We're enabling that for customers.

The other area is when you think about virtualized estates and VMware environments, customers are going through a massive upgrade and tied to that a refresh cycle. And those customers as they're looking at basically how do I optimize my on-premises today and then basically get the full flexibility and future proofing for hybrid multi-cloud environments tomorrow. NetApp is enabling that unique use case for customers where NetApp is the only certified and supported enterprise storage with VMware hybrid cloud across all three public – major public clouds. So, let me pause there and invite Octavian to share some thoughts as well.

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Customers love us, but mostly beyond the technology, which I love. They love the fact that we're integrating with the ecosystem, right. So, at the end of the day, when you deploy a data management and storage system from NetApp, you want to make sure that, that works well in whatever landscape that you're deploying, right. It works well with a Cisco environment. It works well with the VMware environment. It works well with a Kubernetes environment. It has the right APIs to enable a data protection vendor to build, I don't know, forever incremental skinny replication of the data. We're learning a lot and doing a lot of development with our AWS, GCP and Azure Cloud Partners. So, I feel that the competitive advantage that we're building and why many of our customers appreciate us, it's the investment that we have in the ecosystem.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Well, thank you. Last call for questions. Anyone, anyone.

A

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Unless you have more AI questions. We'll be here for...

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yeah, yeah. [ph] We don't mean to (01:12:35) stop your AI questions.

A

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Okay.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Okay. Well, thank you guys very much. I really appreciate your time. And I'm sure we'll make sure that you're seen by this audience more often, so thank you.

A

Octavian Tanase

Senior Vice President-Cloud Engineering, NetApp, Inc.

Thank you.

A

Sandeep Singh

Senior Vice President & General Manager-Enterprise Storage, NetApp, Inc.

Thank you, Kris. Thank you, everybody.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. And so now we're going to take about a 20 minute break. We will reconvene at 11:05 Pacific Time. Thank you very much.

[Break] (01:13:00-01:33:27)

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Welcome back, everyone, to the NetApp Financial Analysts Tech Session. We've got a few more sessions this afternoon and kicking that off – or I guess, it's still morning. It just seems like afternoon. Kicking that off, we've got Jeff Baxter, who is VP of Product Marketing here at NetApp.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Jeff, have a seat.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Thanks.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

To get it started, why don't you say a little bit about who you are and what it is you do at NetApp?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Sure. So, hi, everyone, thanks again for joining all of us. I'm happy to be here. I also feel like it's afternoon somehow. So, I run Product Marketing at NetApp. I've been at NetApp for 15 years. So, long time believer and veteran here. I started out in our field organization, I worked in our product management organization for a while and I've had the privilege for the last few years to lead our Product Marketing organization.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Well, great. Okay. And again, like the other sessions, we're going to open it up for Q&A. And I'll just warn you, if you guys don't come up with questions, then I'm going to ask questions. So, that should be a good encouragement.

Jeff, as Head of Product Marketing, you've got the purview over and of all the enterprise storage platforms and really have been helping define our hybrid multi-cloud strategy.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Right.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, why don't you say a few words about what NetApp is uniquely doing around hybrid multi-cloud and how what we've done with the cloud vendors positions us and truly differentiates us from other storage players?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Sure. So, I think you heard from George this morning around intelligent data infrastructure and our positioning there, and really what's underlying it is the fact that we're the only enterprise storage vendor out there that has a single operating system, a single platform that can work for any data, any workload, any application from a technical level; right file, block, NVMe protocols, object protocols all across the board.

And I think that's interesting from a technology perspective and it makes life easier for our customers, but what that really lets us do is be sort of force multiplier as we expand that out to the cloud. So, Ronen and Pete will be

coming up here in a little bit, and Ronen runs our first-party cloud business and we'll talk about how we're able to extend that onto the cloud.

And the nice thing for us is when we operate is really the only vendor available as a native first-party cloud service, we can do that on a single operating system. So, it really helps center the R&D that all of Octavian's team does – gentlemen who was just up here – into that single sort of focus point, and then we can expand it out to reach the entire market.

And so, the really interesting opportunity for us that we've, I think, exploited so far and can continue to exploit is not just on-prem customers, not just in the cloud customers, but really customers that span both. And I think that's one of the few places that NetApp is truly uniquely positioned to be able to exploit in the market.

QUESTION AND ANSWER SECTION

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Questions? Okay. We've got some questions. David? Then we'll get to you, Meta.

A

David Vogt

Analyst, UBS Securities LLC

Great. David Vogt, UBS. Thanks, Jeff, for doing this. So, when you think of your go-to-market from an enterprise customer perspective, the fact that you can deploy across multiple different sort of platforms, how has that changed over the years as customers are looking for potential maybe other vendors that can do something at some point down the road? I know you're first to market with a lot of your offerings, but how's the competitive landscape changed over the last, call it, five, six, seven, eight years in your perspective?

Q

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. It's a great question. I think that, in the competitive landscape, obviously everyone's moved to all-flash, right, and that was – Sandeep discussed that and some of the others discussed that. I think that a lot of our competitors are still looking to try and rationalize their portfolio offerings. I think that, to their credit, they've probably gone down from some of the major ones from having seven different storage operating systems to having three or four, right? We still think we have a lead and a pretty substantial moat there, quite frankly, competitively in that regard.

A

They've also started to – and we take this as a validation of our strategy. We've seen, especially over the last year, people starting to tiptoe into having their offerings on the public cloud typically as marketplace offerings. And as you're probably aware, we first had a marketplace offering of ONTAP in 2014. And that was our first – called it ONTAP Cloud at the time, right, and put it on AWC in 2014. And so, I still think there's a very substantial moat there.

There's really two parts to the moat. Can you technically make all of your storage operating systems, especially if you have multiple ones, work on every different cloud? Is there any technical impediment to all of our competitors doing that? No. I mean, that – and most likely they've all indicated that eventually they'll get there. So, maybe that's a five-year moat that we've built, maybe a ten-year moat.

The other is from a business partnership perspective and where we're the only ones that are actually provided as services by Microsoft, by Google, by Amazon. And that is harder to put a year on, right? That's harder to say, will they cross that barrier, will they be able to say, hey, we're an important enough partner that Amazon is willing to invest in the co-engineering to build Amazon FSx for NetApp ONTAP. Same thing with Azure NetApp Files.

For a lot of those cloud vendors, they look at our 30-year track record in building what we believe to be the absolute best enterprise file system on the file side. And that's something that we don't think is replicated by any of our competitors. So, once you have us especially for a file system, do you really need to add a second or third competitor to have an enterprise class file system? I think the answer is no. And even if the answer turns out to be yes, it's several, several years worth of co-engineering to do that.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Yeah. Okay. Great. I'm glad everyone's back. So, Meta, actually was, yeah, early in the question...

Meta A. Marshall

Analyst, Morgan Stanley & Co. LLC

Q

Meta Marshall from Morgan Stanley. In the past, maybe a lot of the cloud customers were kind of new to NetApp versus kind of being customers who had been on-premise. I guess, just over time, how have you kind of adapted marketing to kind of bring customers along on that journey from kind of on-premise to cloud?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

A

Yeah. It's a great question. So, just even in my organization, right, so we unified, we used to have them separate. My product marketing organization covers both cloud and on-prem, right? Octavian that you saw before has a unified organization. So, we have a General Manager for Enterprise Storage, Sandeep; we have a General Manager for Cloud Storage, Ronen, who you'll meet shortly. Right? So, we treat them as separate businesses from that perspective. But in terms of going out to the customers, every one of our gold pitches, when we go out and do a one-on-one conversation with the customer, it's always covering what we do on-prem, what we do in the cloud and how we link them with the hybrid cloud data services. Because that's the other important piece is it's not just about, yeah, you can store your data here and you can store your data there, but being able to link them together, right? And so, increasingly, that's where we're seeing a lot of crossover from customers.

We have those cloud-first customers that in some cases they're going to be cloud-native forever. Right? They may have been cloud-first and cloud-only and they'll stay cloud, and that's 100% we're on board with that. Some of them may have discovered us first in the cloud and then when they go to do their next tech refresh of their on-prem environment, they discover that there is compelling advantages to refresh into NetApp. And so, we gain a competitive advantage on-prem as well. And we definitely try to exploit that in the market where possible, because it gives us an end to that customer, it gives us an introduction into that customer and it also gives us a compelling technical differentiator as to why they should refresh their competitive on-prem gear to our on-prem gear.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Matt and then Steve.

Matthew John Sheerin

Analyst, Stifel, Nicolaus & Co., Inc.

Q

Yeah. Thank you. Matt Sheerin from Stifel. I'm hoping you can talk about your marketing strategy by customer segment. You've got enterprise customers. You've got thousands of partners that you work with mid markets. Could you maybe differentiate those sectors and also how you leverage your partners, whether it'd be the MSPs or VARs that you work with?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

A

Yeah. So, I'll do a little bit of that. I also note I tend to be more on the product marketing side of things. So, I don't want to speak for our CMO, but I'll give you a little bit there and then we can always do some follow-up as needed. We do tend to segment – so, the larger scale enterprise customers are higher touch, as you would expect, and we work directly with – we partner with and have our marketing really on a, I would almost say, customer-by-customer basis, right? So, it's sort of a surround and go directly to where the customer is from that approach.

The commercial and the – so, you are talking marketing or go-to-market more particular? Either one. Okay. Pick my poison. So, I think the go-to-market is obviously the sort of higher touch marketing for the enterprise and our big global customers. Right? And we tend to align there with and partner with our hyperscaler partners. Right? So, we go to market directly with AWS, Azure, Google in calling on those accounts as well as handling the on-prem side of the business.

The commercial, as you mentioned, a huge part of our business is through VARs, through our partner network, and we continue to expand that. We released the new Partner Sphere program and have kind of revamped a lot of what we've done with our partner ecosystem. And so, doing a lot of that commercial business, a lot of that go-to-market is predominantly driven by the channel for us, and I don't expect to see that change. And so, I think that's sort of the basic segmentation I would say.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Yeah. Okay.

Steven B. Fox

Analyst, Fox Advisors LLC

Q

Steve Fox with Fox Advisors. In your opening remarks, you said that as a native vendor – first-party cloud service provider...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

A

Yeah.

Steven B. Fox

Analyst, Fox Advisors LLC

Q

...you can expand into the entire market. So, how does that play out? Like, what do we think about that meaning in like one to two years versus three to five years, what would be the outcomes that we should be looking for?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

So, I don't want to get into...

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

No road map ideas.

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

But you can talk big concepts.

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. Yeah. I think what I meant by that is it doesn't limit us to the existing NetApp installed base on-prem. It allows basically any Azure customer, any Amazon customer, any Google customer to take advantage of our services natively within those clouds. So, it dramatically lowers the barrier of entry for people, right?

A

The barrier of entry for on-prem is it could be measured in weeks or months or even longer just in a typical tech evaluation cycle and kicking the tires, getting something into the data center, deciding if you want to do a wholesale tech refresh and then you typically stick with that for three years, five years or longer. In the cloud, right now, someone could logging into the Azure portal, spitting up an instance of Azure NetApp Files without any interaction from NetApp whatsoever and be up and running, decide if they like it or not and make decision within 30 minutes to become a NetApp customer. And so, that what I think the immense opportunity is for us going forward is to really expand that service addressable market to every Azure customer, AWS customer and Google customer.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yeah. All right. Oh, [ph] get here and then I'll get you, Gloria (01:44:49).

A

Irvin Liu

Analyst, Evercore Group LLC

Hi. Irvin Liu with Evercore ISI. So, I wanted to ask about the Keystone bare-metal offering. The use case – I mean, is the use case of this offering meant to target the transformation and modernization of current on-prem workloads, or is this more meant to help customers repatriate certain workloads back from the public cloud?

Q

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yes. So, it's honestly [ph] – not to be tried (01:45:18). It's honestly meant for both, right? So, there's some customers who are looking to just continue their modernization journey. In some cases, they want to reduce their data center footprint and they're going to [indiscernible] (01:45:29) standard colo model, right, which is nothing

A

new in that regard. I think what Equinix Metal allows people to do is have that colo model where they're co-located next to the different clouds, right? So, it plays perfectly into our strength. If we think of our strength as being around hybrid cloud and hybrid multi-cloud, if you're replacing bare-metal NetApp storage near every major cloud, you eliminate a lot of any latency or distance sort of considerations. And so, that's really what allows the Equinix Metal with NetApp storage to have a lot of interest for those customers.

And then, obviously, for repatriation customers, I mean, from a NetApp perspective, we're willing to support customers wherever they want to be. So, we are in a lot of ways very neutral to that sort of discussion, right? So, we provide TCO calculators. We provide all the information to customers. If they are on a given cloud, we'll help them optimize in that cloud with our Spot portfolio that Pete Lilley will be up here to talk about. We'll help them optimize their storage spend by moving to Azure NetApp Files or AWS FSx. If they still find that their cloud spend is in excess of what they think they could spend if they repatriated, then we make it incredibly easy for them to repatriate. In a lot of cases, if they decommissioned their on-prem data centers, Equinix Metal would be a perfect location for them to repatriate too.

So, it really for us is about giving customers the freedom of choice to operate whatever they want and building that up. And the nice thing about Equinix Metal is that NetApp Keystone is all about removing friction from the buyer experience, right, turning it into a storage as a service. Now you remove the friction of where is it located, actually installing it on-prem. One of the challenges with storage as a service compared to the public cloud is even if you decide to do it and procure it, it still takes time to ship the box into the data center to stand it up, right? That latency is still measured in days or even weeks compared to cloud services where it's instantaneous gratification.

Using something like Equinix bare metal because it's already staged there, they can get a very cloud-like experience, but they're operating still on bare metal storage within a data center since it's pre-provisioned there and we're running it for them on Keystone. So, it's kind of a – it's a little bit of a best of both worlds there in terms of immediate access to bare metal storage, but with more of the cost economics of running it on-prem.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. I think we had a question from the webcast.

A

This is a question from Wamsi Mohan from [indiscernible] (01:47:47). In NetApp's view, does HCI matter anymore?

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

[indiscernible] (01:47:52).

A

A

Why did NetApp emphasize HCI few years ago, but doesn't talk about it anymore, hyperconverged?

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

So, I think the HCI market remains where it is. I think NetApp decided to de-invest being a part of that market I think a large part of that as we saw some of the workloads that were very common for HCI deployments moving into hybrid cloud deployments. And so, when you start to look at VMware Cloud and virtual desktops moving into cloud-delivered models, when you start to look at software being delivered as software as a service like Office 365 as opposed to being delivered on individualized desktops, it just started to become I think clear to us and clear to a large part of the industry. And I don't think this is a NetApp specific phenomenon. I think we've seen in the entire industry...

Siri: [ph] ...information (01:48:40).

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Oh, sorry...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Siri would like to say...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Siri wanted to help.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah, Siri wanted to – Siri converged infrastructure, everyone. So – a new market segment. So, I think, yes, we've generally just disengaged there. I think for most customers we have converged infrastructure stacks. We continue to invest in our FlexPod partnership with Cisco. We think that meets the needs of customers for simplified on-prem infrastructure. And then, for the most part, some of these new buying models, some of these things like Equinix Metal, like public cloud services the need for simplified infrastructure that HCI was trying to solve in a more, I think, to be honest, sort of complicated methodology.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yes. All right.

Thank you. Just going back to ONTAP, where are we with the ONTAP evolution? Is there one version of ONTAP that is close to end-of-life? And if so, would that create opportunity for NetApp to go through a upgrade cycle?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

It's a good question. So, ONTAP itself isn't end-of-life and we continue to offer generally. [ph] Our general – (01:49:47) past track record has been to every six months have a new major version of ONTAP. And so, that continues to be our vision for being able to do that. We tend to – we have a hardware road map. I'm not going to go into the details on it. Obviously, any time we release a new generation of hardware, that creates an upgrade cycle.

One of the things that NetApp does as a cultural principle is that we try not to force upgrade cycles through hardware or software obsolescence. Right? It's a choice. It's something that we think has gained us loyal customers over time by not saying, hey, there's a new version of ONTAP and you need to buy the new box that just came out this month in order to run it. I think that would be a short-term gain for us, perhaps, in terms of creating a bump, but in terms of creating customer dissatisfaction, that's not the way generally the industry has evolved. It's not the way NetApp has done business for as long as I've been at NetApp. And so, I think, from a customer satisfaction, trying to do the right thing by the customer.

There will always be things that new hardware can do for them in terms of performance, in terms of efficiency, in terms of density, and we're always going to continue to innovate there. But if we can offer software innovation on customers' existing platforms that they have under support, we're going to continue to do that. We think that's the right thing to do for customers.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Before we get to Tim, I'm going to inject a question.

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, how easy is it for a customer to upgrade from one version of ONTAP to the next?

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

It's totally nondisruptive. So, it's – and we even automate the entire rollout across an entire cluster. So, we've been doing nondisruptive upgrades. A customer – we'll have customers who've been on 10, 15 years' worth of ONTAP that'll just sit there and roll forward and not only non-disruptively upgrade their software, but non-disruptively upgrade their hardware within a cluster, right? So, we've made that incredibly easy. And we've introduced new programs, even like our Storage Lifecycle Program that we introduced over the last year that allows people to essentially subscribe to hardware replacement as a service.

A

It's a financial engineering model, right, but we've backed it up with the engineering to be able to just non-disruptively replace controllers. So, I tell people it's basically the free iPhone every two years plan, right, where you're paying in advance for the controller, but that allows you to have it be as a standardized sort of part of your ongoing OpEx or support budget as opposed to a CapEx bump every three to five years. And so, that really we think over time that will allow customers to sort of stabilize their spend as well as create more guaranteed refresh opportunities for us.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Okay. Now to Tim.

A

Tim Long

Analyst, Barclays Capital, Inc.

Thank you. I wanted to go back to kind of go-to-market. I think it was maybe close to a year ago, there was talk about increasing focus and penetration on small, midsize businesses. And I think you mentioned something in one of the other answers about kind of changing the partner program or ecosystem. So, can you talk a little bit about where you – what changes are you making there and what kind of success are you seeing so far as you try to increase the breadth of the go-to-market?

Q

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Okay. I may need to defer that question.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

I can help out with this one.

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Okay...

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Right. So, we did look at how do we diversify, right? NetApp is really strong. We sell to every Fortune 500 company. But, there is big opportunity to continue to push down-market. So, we did launch a new partner program that makes it easier for our partners to earn money by leading with NetApp. And we also introduced entry products to the product portfolio to better address that market.

A

So, definitely, it's a conscious push for us to make sure that we continue to broaden and work with a broader swath of customers. Once you sell to 500 – or the Fortune 500, where do you go? So, that's what we're focused on. I would say, so far, it's going well. Good initial feedback on the partner program and then the new entry products are also performing well.

All right. Any other questions? Yeah.

Q

Thanks. Yeah. Since you sell to every Fortune 500 company, can you kind of discuss the go-to-sales motion from cloud-native customers to enterprise on-prem solutions? How involved is the cloud partner in that sort of discussion, go to market, bringing them onboard? And what's the sort of sales motion in terms of time line effectively? Like, how long does it, from kicking the tires to – yeah, I know, obviously, someone could just spin up an instance immediately, but more complicated sale or more robust sale, how long does that generally take?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Okay. For the latter part of that question, I'd like to ask you to defer to when Ronen is up here, our GM for that business. I think he probably will have a better answer to that. Not trying to skip the question. I just want to get you the right expert for it, right?

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

And he'll be on shortly.

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. I think he's next. So, you won't have to wait long to ask that question. From how we engage from a go-to-market perspective, can you – that was...

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Let's get the mic.

A

Q

Going from a on-prem customer to a...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Okay [indiscernible] (01:55:12)...

A

Q

...cloud customer sounds pretty straightforward from a...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. Yeah.

A

Q

...extolling the virtues of moving that direction. But a customer that maybe spun up an instance that's cloud-native, working backwards to a more on-prem solution if they want a hybrid solution...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah.

A

Q

...what is that sort of motion look like and how involved are the public cloud partners in sort of that process?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. It's a good question. So, I mean, the public cloud partners, as you'd imagine, right, are the most enthusiastic about pushing stuff out of their own clouds, right? With that said, it's a symbiotic and realistic partnership, right? So, we have – especially for those large Fortune 500 companies, we have dedicated teams that work and they know from each one of these organizations which major partners they're engaged with, they engage directly with their partner in AWS, right? So, we tend to align our cloud selling organization. It's actually aligned based upon the hyperscalers' individual regions, right? So, if Google Cloud has a selling region, we'll have our regional director assigned to it or DM assigned to it. And so, we organize the same way they organize, they kind of go at the hip.

And so, the idea is if something comes up as they're both co-engaging with the customer where they say, hey, I'd like to have an on-prem premise for this as well, then that's the lead that the NetApp rep will take off. It's not – we wouldn't expect the Amazon rep to go try and help to push to close that deal. That's our responsibility. But we've built a pretty successful partnership and they're very realistic about the fact that these hybrid cloud architectures exist. I mean, that's the reason for some of the things like AWS FSx for NetApp ONTAP is not just the on cloud, but they recognize there's such a large ONTAP installed base, they recognize that the reverse will happen.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Other questions? Back to me then. Okay. So...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

[indiscernible] (01:56:57)?

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

That's right. Tell me why NetApp is so great...

[indiscernible] (01:57:00)

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

[indiscernible] (01:57:01) where did you get your shoes [indiscernible] (01:57:03)?

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, for this question, right, at the beginning – or at the end of last year, we announced a whole set of new products in the portfolio, which introduced the C-Series, which is the QLC-based technology. We also introduced the All-SAN Array.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yes.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, when Sandeep and Octavian were here, we talked a little bit about C-Series, but we haven't touched on the ASA yet.

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, maybe you could say a few words about why we introduced that product because ONTAP does blocks...

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yes.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

...so why is there ASA in the family?

A

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. I said softballs. So, I think there's actually two parts for the ASA. It's interesting. And I think one is around market opportunity and really market presence and the other is the technology answer. So, the technology answer I think is simpler. By having a block optimized simple solution with the ASA, we're able to just make it really dead easy for our block customers, simplifying the interface, simplifying the setup and, more importantly, we're able to do things like symmetric active-active technology, which allows for much faster failover times, which tend to be more important for block-critical workloads. So, that's a feature. Being able to be symmetric active-active has typically been [ph] limited sort of legacy firm arrays. It's in fact EMC Symmetrix (01:58:17). Right? If you think about where the name came from 30 years ago, something like that, right, it was all about being a symmetric active-active architecture, right? So, being able to do that at a very low affordable modular all in NVMe price point, that was where we really focused the ASA.

A

So, that's the technical answer and that's where it differentiates. And by the way, it's still the same exact to ONTAP, still able to be matched in the same way, same APIs, ability to replicate between the two. So, we're not changing the operating environment and we're not splitting off. There's not a separate code stream for Octavian to have to manage. We just have these optimized features that we're able to basically flip a switch on and turn on in a SAN-only implementation.

The market side of things is sometimes – and this is actually my job on a daily basis, right? NetApp doesn't get credit as a block storage vendor a lot of times in the industry. Right? Because we were built 30 years ago, we

introduced network attached storage. Right? And 20 years ago, we introduced unified storage. We get a lot of credit on that side of the fence, but what's not recognized is that we have 20,000 customers that run SAN storage, right, across 50,000 storage arrays. We have 5,000 of those customers, of those 20,000 customers, have NetApp storage arrays that they run nothing but SAN block workloads on.

So, that tells us, A, there's a market; tells us, B, that we perhaps do not get the credit or the coverage. And so, perhaps, we're not getting that automatic consideration opportunities. If a customer is building a short list and their customer that we don't have a touchpoint with, are we always making their short list for block storage? I think that's an open and a good question and opportunity for us.

By putting the ASA out there, it allows us to have a focal point for our marketing, for our go-to-market team to go out there and aggressively say, yes, we are in the block storage market. It also gives us some pricing flexibility to go after the block storage market in a way that doesn't necessarily arbitrage our unified storage business. So, all of those sort of three reasons I think are really why we went into the ASA market.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Okay. Great. Still no hands. All right. So, we also introduced an entry-level product in the A-Series family...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

A

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

...the A150. Someone earlier today asked me kind of why did there seem to be not only from NetApp but other vendors a big flurry of entry-level products?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

A

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Why did we introduce a lower end product into the portfolio?

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

A

I think there are a couple of reasons. One is the desire to – and just to repeat what you said, right? If you've already conquered – if you've already taken the Fortune 500, where do you go, right? The other, I think, important point is the cost of flash have gotten down finally economically enough that it doesn't make much sense to build an entry-level system if the cost of the storage on the system still blows out an average – an entry-level customer's budget. With the advances in not just bringing down the cost of TLC flash, but with QLC flash and others, it starts to get to the price point where it just makes sense for customers that are on hard drives in an entry-level space to adopt entry-level all-flash technology.

So, I think the reason that NetApp did it and probably the reason a lot of the market did it is because we finally reached that inflection point on pricing where you get down to a price point where a mid-sized business or smaller commercial business could get into all-flash technologies.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. And then, I think one of the things that NetApp offers that's probably not well understood is BlueXP.

A – [03L9C7-M Jeff Baxter]>: Okay.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

And it would be great if you could explain a little bit about what BlueXP is and how it differentiates us and helps our customers.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Yeah. So, BlueXP is our unified control plane. We rolled it out about a year ago based on a lot of the technology that we had started to build for our public cloud instantiations. It allows our customers in a single pane of glass to manage all of their on-prem storage as well as all of their cloud instantiation, not just the ones bought through the marketplace, but also things like Azure NetApp Files or AWS FSx for NetApp ONTAP so they can manage and it allows our customers really to go through two different models.

If they're primarily Azure-centric, for example, they can manage Azure NetApp Files entirely through the Azure portal, right? They're essentially a Microsoft customer. They are a Microsoft customer. They're thinking of it through that lens. Right? And everything integrates there. On the other hand, if they're hybrid multi-cloud and they're more, say, a storage customer, right, that just happens to use multiple different clouds, they can go through our unified control plane and have the same single experience on Azure NetApp Files as they do on AWS FSx for NetApp ONTAP, as they do on-prem with our AFF line.

And so, that allows us to do very cool things in terms of data services built on top of it. So, it's one thing to just say, hey, I can provision software, but if I can drag and drop from an on-prem system to a cloud system and set up replication in a couple of clicks, it's something that basically none of our competitors can do. And you can see from then on, we can add additional services. Tiering, caching, all these services we've built over the past few decades, we're now able to expose to that single pane of glass.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Well, I think our time is up where I have at least one of our next speakers. So, I will set you free...

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

And thank you very much for your time. I appreciate it.

Jeff Baxter

Vice President-Product Marketing, NetApp, Inc.

Thank you all for your time. Appreciate it.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. So, now we're going to get into the world of cloud, specifically. So, since Ronen's here, I'm going to invite him to come on up. Here's Ronen Schwartz, who's the head of our first-party cloud storage services.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Hi, everybody.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Hey, Ronen.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Good to be here. Yes.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Thanks for coming. Someone else is going to join us, but I think he's getting mic'ed outside. So, let's start. Have a seat.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

And why don't you tell us who you are and what you do at NetApp in a better way than I just did?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

So, good afternoon, everybody. I can't believe I'm only 24 hours here because based on my voice, it sounds like I've been here a little bit longer.

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Oops, sorry.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Hey, Pete. My name is Ronen Schwartz. I joined NetApp about 3.5 years ago basically to lead our first-party cloud journey. And this actually includes, from a leadership perspective, leading the engineering team, the product management team, the strategic alliance that we have with the three hyperscalers that – at this stage.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Okay.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Short description. Yes.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

That's great. And then, Pete Lilley also just joined us. Pete Lilley comes from the Instacluster acquisition. So, you might notice a bit of an accent when he talks.

Peter Lilley

Vice President & General Manager-Instacluster by NetApp, sNetApp, Inc.

That's right.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Pete, why don't you introduce yourself and what it is you do?

Peter Lilley

Vice President & General Manager-Instacluster by NetApp, sNetApp, Inc.

No worries. [indiscernible] (02:04:58) comes from Australia [indiscernible] (02:05:00) traveled a long way. So, I...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yeah, that's true.

Peter Lilley

Vice President & General Manager-Instacluster by NetApp, sNetApp, Inc.

[indiscernible] (02:05:02) in the chair. No. So, yes, thanks very much. I'm taking my glasses off. I'm Pete Lilley. I'm the VP and GM of the Instacluster business. I was actually one of the co-founders of Instacluster and I was the CEO of the business leading up to its acquisition by NetApp in May 2022. And I am responsible for all of Instacluster's as part of the CloudOps portfolio, which there are three businesses in that in my group under Haiyan Song, which is

Cloud Insights, Spot and the Instacluster business. So, all of the product development around what we do from a platform and enterprise open source is really part of my business.

QUESTION AND ANSWER SECTION

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Well, great. So, you can see we can cover all things cloud here. And I'm sure you guys have questions. Otherwise, you're going to have to listen to me ask more softballs. So, all right. We've got Meta in the back.

Meta A. Marshall

Analyst, Morgan Stanley & Co. LLC

Q

Oh, sorry. Maybe just on the CloudOps portfolio, clearly bringing together kind of all those different acquisitions has been kind of a journey for you guys. Just where do you guys – like, where has the synergies kind of come from all of those different acquisitions and kind of pulling them together and where is kind of the work still being done there?

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Certainly talk about how we are using Instacluster and attaching...

Peter Lilley

Vice President & General Manager-Instacluster by NetApp, sNetApp, Inc.

A

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

...NetApp cloud storage.

Peter Lilley

Vice President & General Manager-Instacluster by NetApp, sNetApp, Inc.

A

Yeah. Yeah, yeah. Absolutely. So, from a CloudOps BU perspective, you've got Cloud Insights, Spot and Instacluster. I think there are tremendous advantages between the three products themselves in terms of being able to leverage each other's capabilities to be able to drive more automation, more capability, and make the CloudOps portfolio effectively more intelligent and as intelligent data infrastructure.

The other part, which I think is really, really interesting from my perspective as the GM of the Instacluster cluster business is the bridge that Instacluster can help make between the cloud-native part of the business and NetApp's traditional storage business. So, we're able to leverage cloud first-party storage, which I'm sure Ronen will be very happy with. But we've just done our first integration of first-party cloud storage through a solution called Postgres on Azure and NetApp Files. 300% performance increase – up to 300% performance increase for leveraging Postgres on ANF, which is amazing compared to what's available in the hyperscaler Postgres deployments and what traditional open core and other competitors can do with that technology.

And the other fantastic part about that is it's not just about the price performance, it's about being able to bring some of the advanced features to the usage of that technology that's pretty amazing. So, data tiering, advanced replication, quick snapshots, disaster recovery. These are problems that enterprise users of Postgres have had for a long time. And the ability to one-click deploy or [ph] API call (02:08:29) deploy that technology out through – in Azure is amazing being able to leverage that.

And I think, at the same time, on top of that, with Instacluster and what we're doing with hybrid cloud customers and hybrid cloud environments is being able to offer the same as a service cloud experience in any cloud. So, any hyperscalers – so, whether it's Azure, AWS, GCP or on-prem – is offering customers a pretty unique experience in using and leveraging this very powerful enterprise open-source software.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

Maybe I'll add also. I'll add two things to what Pete was just mentioning. The first one is when you look into database optimization, there is a layer of storage that you can optimize for the databases. And this is something that the hyperscalers have done to a certain degree to some of their databases. And through this partnership, we are basically pushing the performance, the efficiencies and other capabilities basically to the best that you can. So, I would say, like, I also really, really appreciate the knowledge and the depth of implementation that the team is bringing. And this is pushing us to bring an overall better, better storage.

I would give one more example, which is Cloud Insights. Basically, today, a lot of the AWS field is using Cloud Insights as a way to demonstrate and show the customers that are looking into migration use cases, this is how your existing environment look like, this is how it will look inside AWS. And it's basically helping us or helping AWS in this case accelerate the migration to the cloud. So, I think there is many points of synergies. We gave two examples that are, I'll call it like, already implemented in a good scale.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. And [ph] Gloria (02:10:22), I think you had a question from the webcast.

A

Yes. Aaron Rakers from Wells Fargo. As we think about NetApp's multi-cloud integration, do you have any color on how many of your traditional on-prem customers are leveraging NetApp's native cloud offerings? How has this progressed?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

So, obviously, we do know.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

You can talk about what customers are doing, the deployments, the reasons why?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

Yes. So...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Sorry, I've apparently put the fear of God in everyone.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Extremely well-trained. One of the advantages of having a cloud solution is you actually do know what customers are doing and how at least on the high level – on the high level degree. But very specifically to the question, we're seeing across many, many verticals customers that are implementing in the cloud I think one of three patterns. The first one is basically expanding their data center into the cloud. They do it in a case that they are pressed on their short-term storage availability, they have long-term plans of managing data centers, et cetera. In this case, what they are doing is tiering, shifting back up [indiscernible] (02:11:36) into the cloud. This is I think one pattern definitely very well embedded and implemented.

The second one is basically migration or building the similar workloads instead of on premises basically in the cloud. We've seen a massive growth in SAP in the cloud, Epic in the cloud. We're seeing databases that are moving to the cloud. VMware, there is a lot of push in VMware moving to the cloud. Customers like that it's not that they're moving there or leaving their on-prem, but they're choosing for different workloads, which one should be on premise, which one should be in the cloud. Again, multiyear of customers adopting it in a very big way. And I think the third part of it is basically customer that are innovating new workloads and sometimes they're doing it in a cloud-first approach.

I think the most common one and if you have the chance, I recommend you'll see the demo, is a lot of the AI workloads are innovated or starting from the cloud. Not always, but that's a very common pattern. I think others are implementing Kubernetes as a platform for their applications in the cloud. I think in all three of them, we're seeing really, really good adoption and not just good adoption in the last six months, but actually good adoption in the last few years.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

But maybe I will say one thing is we are using the cloud as a way also to acquire new customers. Obviously, AWS, Azure and GCP have a very broad market reach. There is a lot of customers that NetApp have not necessarily has its customer on-premises that will have their first NetApp ONTAP experience in the cloud. I think it's true also for the rest of the portfolio, but it's not limited to the – it's definitely not limited to the customer base.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Mehdi.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Just one follow-up and I may have missed this. If you're focused on expanding business in the cloud, how do you prioritize with AI projects? And I am asking this because to me AI is more of a on-prem deployment. Do you find yourself competing with other parts of the company; question A? And question B is, if I am right that AI is more of an on-prem, then what are your thoughts about the future of NetApp's cloud business model?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

I am trying to see how do I answer whether AI is an on-prem or cloud part. I think there is – customers are making choices about where to design, where to innovate, where to deploy. And in many cases, these choices involve both on-premises and the cloud. I have customer that I know that have developed in the cloud and deployed on premises. I have customer that have developed or started development on premise and deployed it at scale in the cloud. So, I definitely don't see this one or the other.

I also don't think that NetApp should – I will call it like – that NetApp should decide for the customer where are they going to do the AI innovation. Our goal is to support the customer with the best-of-breed storage, optimizing their AI promise – their AI basically goals. I do think that there is an interesting change in the market with gen AI that is basically giving more value to unstructured data when previously machine learning and so on brought more value I think to the semi-structured and structured data. I think the patterns there are a little bit different. We are going to support our customers in both of these journeys. I think the gen AI just naturally from OpenAI to Google, Vertex AI, et cetera, there is a lot of the gen AI that is done cloud first. As I said, I think our goal is basically to support the customers wherever they are.

I do want to call out, I mean, we did demonstrate as part of Google Next together with the bunch of announcements that Google have made. You could have seen basically one storage partner, NetApp, that actually have already done the full integration to Vertex AI, that actually supporting the customers and how to augment the LLMs or the data models with proprietary data in a secured way how to bridge on-prem and cloud data. I think this is NetApp. If you're here in the event, you'll be able to see similar demos – I think one on main stage later today and I think two others through the sessions – of how tightly we're integrated into the GCP AI, how tightly we're integrated into AWS, both SageMaker as well as the gen AI technology, and the same for Microsoft. You'll see us doing that across the board. We're doing it also fabulously on-premises as well.

A

All right. Tim.

Tim Long

Analyst, Barclays Capital, Inc.

Q

Thank you. Two, if I could. The first one might be a quick one. One, do you think there's going to be, as things evolve on the hyperscaler public cloud, any hardware play for NetApp at all, or is this going to predominantly be software?

And then, second, on the software side, could you talk a little bit about talent and resources and competing with other high-profile tech companies? Because, I think with some of the acquisitions, maybe there had been some

departures. So, if you could just talk more broadly about how you're building kind of the internal engine here and be able to keep fueling it with software talent? Thanks.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

You'll start?

A

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

You want to do part one and I'll do part two, or [indiscernible] (02:17:43)? But maybe we should get – what was the – could you repeat the first question again?

A

Tim Long

Analyst, Barclays Capital, Inc.

Is there any hardware play...

Q

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Yes.

A

Tim Long

Analyst, Barclays Capital, Inc.

...that you're going to pull in on to that...

Q

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

So, I'll start with that.

A

Tim Long

Analyst, Barclays Capital, Inc.

... [ph] customer (02:17:54)?

Q

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Yeah.

A

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

You start with the second one.

A

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Yeah.

A

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

So, definitely, there are specific workloads that will be best supported by a single tenant solution [ph] to a customer that (02:18:06) that's what they want. How does the customer consume it from a subscription and so on, I think there is a lot of flexibility. There is a lot of flexibility there. So, I do think there is hardware opportunities there. There is hardware opportunities as well. I think in some places you will get full visibility because it will be basically public and other places you're just behind the service that is basically running and you remain anonymous from that perspective. So, I definitely see as larger workloads are moving, as very demanding workloads are moving, in some places, single tenant will make a lot of sense.

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

A

On your second question – and thank you. It's a really good question. The – yeah, one of the value – and I'll talk from an open source perspective, enterprise open source is one of the interesting value propositions that Instaclustr brings to customers is that actually finding resources to run these types of technologies at scale with deep open source knowledge is actually really, really difficult. It's highly competitive. There isn't a sufficient level of expertise out there for all of the industry to consume the available talent pool. And so, as a business, it's actually really, really critical to have your own internal programs to be able to develop the engineering talent that you need to have to continue to sustain the capability and the competitive advantage that you've got.

And we realized that from day one when we founded the company that Instaclustr [indiscernible] (02:19:51) story long before for NetApp is that that would be a constant challenge and we had to build a robust program to create and train, raise and sustain exceptional open source software engineers, both from a DevOps perspective and a development perspective. And we continue to do that today.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

And then, I'd just add, Tim, I think you're asking about NetApp internally, how do we attract and retain key engineering talent? I think, not to speak for Ronen, but I will. One of the key advantages we have is ONTAP, right? We are the only company with a single primary storage operating system, and that enables our R&D to leverage really broadly. Right? So, we get a lot of leverage. And maybe you want to actually add some color to that?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

Yeah. So, I think when you look at it – and I think the first part you're absolutely right is we are basically building on the foundation of ONTAP with a lot of talent and skill set that we have and we have built through the years and basically optimizing that into the three clouds. So, we're building on a very, very strong foundation, unlike you build some new storage offering from scratch, right? We're also leveraging at scale and especially in storage, it's really, really important, the testing framework, the scale framework, the resiliency and all of that we're leveraging it at scale.

I think for the cloud-specific talent that we have, I think, like any other company, we are based on people and that's the most critical asset that we have. I sometimes tell my team, if you want the front seat into the cloud journey, into the AI journey, we are actually that front seat, right? Like, you see three cloud vendors [ph] first hand is the 1P (02:21:41) working with them and you get to see a lot of innovation as it comes to the market, even

ahead of the time that it comes through the market, through these tight partnerships. So, I think it's a very, very exciting place to be in general and definitely for the developers.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Got a question in the back.

A

Frederick Gooding

Analyst, William Blair

Hi. How are you doing? Frederick Gooding with William Blair. Just wondering if you could spend a little bit on the go-to-market motion? You guys talked about a lot between on-prem and cloud, I'm just wondering a little bit more about how you're telling sales reps to think about that and, more specifically, also I believe Jeff mentioned, you made it in terms of your relationship with partners and you guys are making it easier for them to make money, wondering if you could provide a little more color on that?

Q

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So, probably not the right team to talk about the on – the VAR partner program that we just launched, but certainly talking about how we go to market with our cloud partners is a great one. And I can follow up with you on the other one.

A

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Yeah. So, I think what we have – I will call it like – aligned especially since the beginning of the year, we aligned our specialist cloud sellers, those that are focused and we have team that are focused separately on AWS, on Azure and on GCP, we basically made them aligned to the structure, to the GTM structure of the hyperscalers themselves, with the goal of basically – the group behind them to create the design wins per workload and then for the sales team to basically focus on enabling the hyperscaler sellers, but even more important, enabling the hyperscalers workload specialist. Some places they're called black belt and other places they're called workload heroes. Every hyperscaler have a little bit of a different terminology. But we're basically aligning ourselves with them so that we are empowering them in parallel to basically empowering the end customers.

A

This is actually giving you a very, very good leverage, right? If you are a hyperscaler seller, you have about – I don't know – 200, 300 things to potentially sell. We are basically helping them identify the best solution for the customer for the different workloads. And that's kind of our focus. We do it systematically, meaning that there is design wins, there is published calculators and so on, and then basically helping the customers and the hyperscaler sellers in basically proving the value leveraging these tools. So, that's basically is the main motion.

There is actually a motion of the hybrid customers. I think Jeff was just describing it toward the end, which is hybrid customers that are using the cloud for resiliency, is using the cloud to augment their on-prem. I think this is – and NetApp supports that through the direct sellers that work with the customers and basically through the regular presales organization.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

David.

A

David Vogt

Analyst, UBS Securities LLC



Great. Thanks. I just want to go back to the comment you made about working with your customers that are designing or developing applications either on-prem and then moving to a public cloud or innovating in the public cloud, moving back, and it's not your kind of position to kind of tell them how to run their business. But can you talk to how you bring CloudOps to that conversation to help them manage potentially incrementally higher costs or complexities or technical challenges that they may face going back and forth in either direction? And should that ultimately be sort of an incremental service that most customers take because you're trying to optimize and solve for potentially a complex solution?

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.



Yeah...

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.



Do you want me to start, or you're start...?

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.



[indiscernible] (02:25:39).

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.



Yeah. I'll start a little bit and then you continue. I think that you said it already, and especially in the last six to nine months, cost is an important thing for each and every customer. Right? So, as we were bringing – I mentioned just the TCO calculators and all of that, right? So, basically, part of actually giving the customer the right recommendation for a workload does include the technology best practices, but it does include the total cost of ownership recommendation.

And this is actually where the CloudOps portfolio is coming to play. Right? And it comes to play in multiple levels. It comes to play in how much is that a compute-heavy workload? Can we help you with compute optimization? Is this a database workload? Can we help you with full service of the database, delivering the database as a service to you as a customer. Can we help you with insight that help you find the bottlenecks across your entire environment so that the workload is optimized? So, I think the work – if you go from the workload, then storage is a critical component, but the entire optimization where our CloudOps is focused on is a [ph] natural – it's an actual (02:26:55) next steps.

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.



It is. It is. And we're able to bring to bear the complex team of technical account managers, the sales and account management teams and the engineers, the customer success engineers to get engaged with the customer that's looking to deploy in multiple contexts and look for and position the right kind of solutions at the right time that really bring the benefit to the customer, whatever stage of the cloud journey they're in, whether it's going one way, going the other or going both ways at the same time.

And then, talking about my part of the business, which is the Instaclustr business around enterprise open source, the experience that we offer to the customer is really and as a service experience both on-prem and in the cloud. So, it's natively the same and we'll get more the same as we continue developing those capabilities. And so, really, it does come – the customer has all the maximum flexibility, but they just need the right access to the right tools and capabilities at the right time to help them make those TCO-based decisions.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

And I think the...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Yeah. [indiscernible] (02:28:15)?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

Yeah. Just to finish one more – one last point on this. Customers that only do migration as is to the cloud, they get very limited benefits. The idea is that you migrate, then optimize and or optimize and then migrate, and this is where it's a very natural fit. And we're able to guide the customers on that on their journey.

David Vogt

Analyst, UBS Securities LLC

Q

That was going to be my follow-up. So, I know it's not your purview in terms of helping them design or innovate their applications, but when you brought in in a discussion from the CloudOps side, is it as the innovation is happening or is it a customer says to you, hey, we've got this great application, we're working on this workload, what do you think we should develop it in the cloud, do we develop it on-prem, or is it more of a secondary consideration at some point in the journey effectively from the customer's perspective?

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

A

So, we see customers, I think, at all stages. So, it can be we're doing something new help us understand what the deployment challenges might be. They may have a strategy in mind. So, I'm thinking about a customer that we're working with at the moment that is looking to make a move from a hyperscaler-deployed environment back to an in-house one. And they're giving all of that consideration around TCO. And so, there's an engaged process in that and having a discussion with the customer about all of the benefits of doing so and what that cost will be and how that falls out and what the opportunities are to optimize.

And my experience working with our customers has been that, as Ronen was saying, optimization doesn't necessarily happen out of the box either. It can often be a peak, followed by some optimization, followed by another peak, followed by some future optimization as features and capabilities evolve. But the whole goal of the trend over time is towards optimization.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

And I think that workload implemented in the cloud is not a onetime opportunity.

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Yeah.

A

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Even if you arrived in the late stage of the current workload, the next one is just around...

A

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Yeah.

A

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

...the corner. Right?

A

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Yeah. So, as you said...

A

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

You can arrive – if you show and explain the value, you'll be early in the next one.

A

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Exactly. I mean, the other types of customers that we would experience from the planned and doing a project through to my cluster is on fire, please help us now because it's a critical application is another example of a customer where you're almost in rescue-type optimization to get that customer stable and then you bring them into the optimization discussion off the back of that.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Sidney.

A

Sidney Ho

Analyst, Deutsche Bank Securities, Inc.

Thanks. Sidney Ho with Deutsche Bank. Well, this is – I'm going to [ph] (02:31:13) financials. But, at the last earnings call, you did talk about the shift towards first-party storage services versus subscription. Is that simply the cyclical nature of the subscription business going up and down, or is it a more strategic move towards first-party storage, whether it's from NetApp or the customer perspective? Thanks.

Q

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

So, I'll clarify our statements and then hand it over to Ronen. And basically, we said we believe our emphasis and our biggest opportunity is around those first-party storage services. We see that as an absolute unique differentiator for the company, a massive opportunity, and that's where we're really putting the wood behind the arrow. Ronen?

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

A

Yeah. I'll second what you said. I think if you look into our general – into any customer, at any stage, whether they were or were not NetApp customers, very high percentage of them have an existing relatively large commitment to the hyperscalers. Being a 1P or a first-party offering means that you do not need – there is no need for a new contract, there is no need for a new engagement. There is a need for the workload or the team that works on the workload to make the right choices when it comes to the infrastructure in this workload. It's a massive advantage.

This advantage translates from a financial perspective to consumption, because they don't need to sign a new agreement, they don't need to have a new commitment. They just need to start using this environment. And that usage is translating into consumption. And that consumption is what you see eventually in the financial report. I think what we're saying is that now that we have three hyperscalers with agreements like that and so on, we'll see the consumption really translating that into revenues and so on.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Well, thanks, guys. I really appreciate your time. Thanks, everyone, for your great questions. I'll set you free to go talk to customers.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Thank you. Yes.

Peter Lilley

Vice President & General Manager-Instaclustr by NetApp, sNetApp, Inc.

Thank you.

Ronen Schwartz

Senior Vice President & General Manager-Cloud Storage, NetApp, Inc.

Thanks, everybody.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Thank you. Thank you. All right. Well, now, I think, it is the part of the event that you guys have all been waiting for some actual real customers that you can hear from and what they're doing and what their big challenges are. So with that, I'd like to invite Anthony from OpenText and Phil from Lawrence Livermore National Labs to come on up stage. And I did that for memory, so I greatly apologize. It's been a long day already.

Anthony Lloyd

Vice President-Technology Services, OpenText

[indiscernible] (02:33:49).

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Hey, thank you so much.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

Hi.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Hey, Phil. Thank you so much. All right, why don't you guys have a seat?

Anthony Lloyd

Vice President-Technology Services, OpenText

Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

So we're going to hopefully get most of our questions from the audience. But I'm going to kick it off by just asking each of you to introduce yourselves, where you're from, and kind of what your IT challenges and environment look like. So that might take a while, but it'll give them some time to queue up some questions.

Anthony Lloyd

Vice President-Technology Services, OpenText

Sounds good. I'm Anthony Lloyd. I'm VP of Technology Services at OpenText. I manage all of the infrastructure and operations for corporate IT. That covers everything from data centers, cloud, network, telecom, storage and compute, end-user services, site support. Oh! Gosh. And there's more. The service desk, the operations center, and pretty much anything that touches an application or an end user. So I'm responsible for corporate IT. So we have a line of demarcation between corporate IT and the commercial side of the business just because we protect all of the back-end systems, the HR systems, financial systems, things of that nature. The commercial side of the house really handles all of the customer-facing applications that we sell for revenue.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Great. And Phil, a little bit about you.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

I'm Philip Adams. I'm CTO for the National Ignition Facility at Lawrence Livermore Lab. I'm responsible for all of the infrastructure, everything that from the underpinnings that runs the control system to how we analyze and process data to being able to make sure that we have a 30-year scientific archive that is there available for our researchers and our visiting scientists. It's quite a bit of a task and challenge to be able to manage something that broad, that vast for requirement set that is always changing.

QUESTION AND ANSWER SECTION

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Well, and you guys are also doing some really cutting-edge innovation and technology. So that's got to be a pretty data-intensive environment. How do you think about setting up your IT environment to deal with the massive quantities of data that you must face?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

You know, we tried very hard to make sure that we looked at ourselves as not as a unique entity in the environment. It's very easy, especially for a National Lab, to say, okay, we're going to go off. We're so different and varied in our needs that we're going to go build something unique. And what we ended up doing is saying, look, we've got the same Lego blocks that are available to almost everybody else.

Let's take that innovation that has been done in industry and assemble it in a unique way to be able to do low-latency operations for a control system that has a way of being able to provide lifecycle management of data over time and pin that to our databases. And, you know, being able to leverage technology as it is makes sure that my team can spend more time helping the scientists rather than trying to uniquely innovate things that industry has already figured out how to do.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Well, and Anthony, when we were talking last night, you were telling me about the massive M&A pipeline that you have to deal with and integrating all these different companies. Why don't you say a few words about some of those challenges?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

Certainly. So at OpenText, we grow by acquisition and because we go by acquisition, we have the opportunity, and I'll use that term, of trying to figure out what the best solution is to integrate things that we may not always know about. We go through a due diligence exercise, but you don't always get the full picture until you get under the covers. So part of this is having flexibility not only to be able to integrate different technologies, different solutions from different sources from around the globe and be able to do that in a seamless manner that allows us to do it in a short period of time, in a very secure manner, and not have to recreate the wheel every time we do it.

NetApp gives us a lot of those capabilities because we operate in all of the hyperscalers. We have the ability to acquire and integrate anywhere in the world. And one of the beauties that we have is no matter what that

environment is, we can have a solution from NetApp that [ph] allows (02:38:02) us to get that done successfully and quickly.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. And Phil, how do you use NetApp?

A

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

We use a lot of NetApp technologies in our environment. We leverage FlexPod and AFF in order to get low latency compute in our environment and low latency access to data. We've leveraged Fabric Pool and StorageGRID behind our Oracle databases to be able to give a transparent lifecycle – data lifecycle management on that dataset. We leveraged SAN Manager and [ph] SAN file (02:38:43) in order to be able to do a comprehensive back-up environment. So pretty much listed...

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yeah.

A

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

...a quite a bit of NetApp's portfolio of applications in the suite.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Definitely. And I see you nodding along. So it sounds like you're also using a pretty broad swath of our technologies.

A

Anthony Lloyd

Vice President-Technology Services, OpenText

Basically the same technology that add a few more. So we really rely a lot on ONTAP, Google Cloud Volumes. We use a lot of that technology because it allows us to quickly integrate different solutions and not have to go out of the box to figure out if we have to have a different way of doing things every time we encounter a different acquisition model.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

I'm loving because I'm seeing you nod, so definitely you guys are using so much of the NetApp technology...

A

Anthony Lloyd

Vice President-Technology Services, OpenText

Yeah.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

... [ph] you're forgetting (02:39:24) the different names of the products. It's kind of exciting to me. All right. Well, let me turn to our audience and see if there are any questions. If you guys have questions about how these guys are utilizing technology, the challenges they face. Your chance to ask real customers real things. Steve?

Steven B. Fox

Analyst, Fox Advisors LLC

Q

So one of the hardest things to figure out is...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Say who you are for the webcast.

Steven B. Fox

Analyst, Fox Advisors LLC

Q

I am sorry. Steve Fox with Fox Advisors. So one of the hardest things to figure out as an outsider is how you become more efficient with storage. It's generally thought to be a consumable, but it seems like every cycle you're able to consume more with less. So can you talk about maybe how you use NetApp to do that and what that means for your infrastructure purchases now versus maybe three years ago?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

Certainly. But what we've been doing is primarily if it's in the cloud, obviously, we're using ONTAP or Cloud Volumes, but if it's on-prem, we have various solutions depending on if it's a file or if it's an A700, depending upon the performance and everything else associated with it. What we're doing now is really trying to move to more of a consumption-based model so that we don't have to make those CapEx investments. So we're really looking at Keystone and looking to how we can use that as capacity on demand, whether we use it on-prem or whether we use it in the cloud that allows us to quickly deploy what we need, not just buying additional capacity, which you end up doing when you're deploying on-prem, but you buy what you need, and then you can quickly spin up additional capacity or downsize it as you need to based on your needs. That gives us a great deal of flexibility, and it really helps us to manage our financials much more efficiently.

Steven B. Fox

Analyst, Fox Advisors LLC

Q

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Phil, anything?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Yeah, I'd say that, you know, the first order, the National Ignition Facility is a research project where we're trying to study the phenomena of high energy density science. So we capture all the data. We store all the data for 30 years because of just how hard it was to get the data in the first place, you know, and so in a couple of ways that we've tried to be – I think some of that what you're asking about is how are we thinking about reducing costs and whatnot, when was the lifecycle management that we did on the data to reduce the total cost of ownership of

storing that. Once we get a better understanding of the types of data that we need to and fuel our machine learning algorithms to be able to understand exactly the types of data that you really want to store, we can get a little bit more efficient in terms of being able to do localized processing nearest to the diagnostic endpoints. Maybe then, you know, have that be a little bit more intelligent at that point and only feed the data up that we really want to store long term.

But right now, when we're in a cycle right now where we're – everything is new, everything is amazing, you don't really know where the breakthroughs are going to be. You got to keep it off right now. And I think as time goes on, we're going to be a lot more efficient with the way how we do analysis and data. But, yeah, for the time being, it's all about how to reduce the cost to maximize taxpayer money.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

And are you guys using the full suite of NetApp efficiency tools so that you're effectively storing more data than you have space for?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Yeah.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

Yeah.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Tim?

Tim Long

Analyst, Barclays Capital, Inc.

Q

Hi. Anthony, if I could just follow up. You talked about the move to kind of more consumption based as a service. Is that unique to storage for you or are you looking across some of the other silos of infrastructure? And if anything that's different when you look at storage compared to the others? And then maybe for both. Could you talk about kind of when you look at your whole storage environment, is it all NetApp? Do you have others, other competitors of NetApp? You don't have to name who they are. But maybe if you can just talk a little bit about how you choose a certain vendor for a certain application or use case? That'd be great. Thank you.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

So to answer your first question, yes, we are looking to move to a consumption model across all of the infrastructure. We've already started that process with the compute platforms, and it's worked out well. Now we're starting to go down that path with the storage platforms. One of the challenges, of course, you face anytime you change from a CapEx [ph] and OpEx (02:43:52) conversation is, what's the cost construct. And we have to make sure that it make sense financially in order to do that, not a price point, so becoming very competitive. So because of that we can start to exercise in that manner.

To answer your second question. So we go by acquisition. So we may get one of everything depending upon what the acquisition is. But at the end of the day, NetApp is our standard. And so we move from whatever that alternative storage device may be, we move to NetApp over time as we integrate.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

I assume you have a lot of stuff in your environment.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Well, you know, we really looked very long and hard at our requirements that we had. Like I said, there's a lot of varied requirements from a control system to a 30-year archive. And the needs of a storage device really change in terms of sub-second response to, okay, you know, how do I store this big chunk, you know, petabytes of data for a long term. So in that case, once we started putting everything down on the list, we said, okay, you know, how is it going to be available? How is it going to be reliable? How is it going to be manageable? And more importantly, well, how do you back it up?

You know, if something happens to the data, like ransomware issue, God forbid, how do you deal with that? And so, once we looked at all of our requirements and we analyze them, we went through a pretty long path, pretty long decision path, and we chose NetApp. And then once we did, it was very easy for us to go, okay, let's keep implementing the technologies that makes sense for us. And they built on top of each other to give us a very comprehensive method for being able to handle our data.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Question in the back.

Frederick Gooding

Analyst, William Blair

Q

I am Frederick Gooding with William Blair. Just curious, so we talked about AI earlier in this session. So just wondering, A, are you guys looking to more implement AI capabilities like [indiscernible] (02:45:57)? And B, do you see that as like a whole another budget in your guy's mind like IT infrastructure or is that a part of like your existing budget just for storage in general?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

So we are looking to implement AI where it makes sense for us to do it. Obviously, we have the concern is that everybody else has, you know, AI has tremendous capability. But you also have to weigh the security concerns as well as ensuring that you have the ability to separate your environment in the event that you are using a capacity-based solution or you're operating in a shared environment such as Office 365 where Microsoft has implemented AI in their solution. So one of the first questions I asked was, okay, how are you making sure that if somebody's tenant becomes compromised, you're protecting [ph] them (02:46:46). And so those things you have to really focus on and ensure that you've got that validation.

So, yes, we do want to take advantage of it. We are putting it into our own products. But it's also the understanding and awareness of you don't want to let the genie out of the bottle until you know what the bottle

can do. You know, you can get good things, but good things can turn bad very quickly if you don't have the right security, you don't have the right ethical controls, and you don't have the right governance model around it.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

You know, we live in this time where you always are wondering if you're getting the right news, the right information, there's big buzz around fake news. And AI is certainly victim to that, right, the data that comes in, how you train your models, and the decisions it's going to make out of that is then really key. You saw what the SolarWinds exploit, for example, that was the way of poisoning the well and it impacted a lot of companies. And so, I see AI as going to have a lot more positive benefits than negative ones, but they are things that governance is going to have to be key in terms of how we address those things, making sure that we have a very tight handle on how we're feeding that AI tool and then the decisions that are being made. And is there a check for the process that's making the ultimate call? If you just automatically go with it and there's no counter check, that may be problematic.

Livermore Labs, we've been experimenting with AI for quite some time. In our particular environment, I think we're still at the machine learning phase, again because of where we are in our phase right now, trying to train the models and making sure that things are right, as they say some models are useful. So it's an iterative process until you get to the point of something that is going to be interesting in our environment.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Yeah. I think that iterative modeling part of AI is really interesting because I think so many people just want to think, well, you train the model and done.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Yeah. Right.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Right. But you – I would imagine you're having to constantly fine tune your pre-trained models with new information as it comes in. How do you think about keeping things current and moving them along?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

You know, I'll give an example. I mean there are two areas in our program that we spend a lot of time and trying to get very deep and one is Optics Inspection, because we're looking for defects in our optics as you put fluence through glass, you can get [indiscernible] (02:49:36) inclusions and after a while we have this optical loop process that recycles the optics. So it's always a question of, well, when do you land the plan or, in the case of NIF, when you bring it down for maintenance to pull those optics out and refurbish them. That means less time for doing experiments, less time for science, so we want to make sure that we are very smart for one we do that type of event.

The other issue that we spend time looking on is the small little BBs that we use for filling with hydrogen, for actually being this – our main target. And we're pushing the edge of manufacturing as we know it. So they're always the pits and a little inclusions that we need to do to make sure we machine what's possibly the most

smooth surface that we know on the planet. And those things, you know, we're using machine learning to be able to find out is that a defect or is that a piece of dust or something that we need to deal with. And so, you know, those algorithms that we have to actually do detection is where we're spending a lot of time right now. Eventually, you know, those things now are saying, okay, is a defect, is not a defect, are going to be smarter as our scientists get more involved with it to be able to make sure that it is making more right decisions than wrong decisions.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Okay. Is it a hot dog? Is it not a hot dog?

A

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

Right.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Meta in the back.

A

Meta A. Marshall

Analyst, Morgan Stanley & Co. LLC

Maybe a question for Anthony. It sounds like you guys have a pretty standard playbook for kind of bringing a company in and then, you know, maybe starting with kind of moving to Keystone and figuring things out from there. I guess just wondering, you know, how has that playbook changed over time in terms of, I'm sure the economics of just the scale that you've gotten has kind of changed some of the [ph] waiting (02:51:33). So just wondering, you know, has it been more towards subscription? Has it been to kind of optimize on-premise just how has that playbook kind of changed over time?

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

So initially, we didn't have the technology available to really consider off-prem. And so we invested heavily in on-prem solutions. As the ability came along for us to get things in the cloud, NetApp brought those capabilities along. We started to move to those models much more expeditiously for a couple reasons. Number one, enabled us to spin up that capacity much more quickly. But the other thing was we were able to get what we needed versus having to oversubscribe and that was extremely beneficial to us.

A

Now, as we have now the on-demand capacities available from solutions like Keystone, it gives us even greater flexibility because we go by acquisition, we may close a location, we may close data centers, we may do all of these things, I cannot be reliant upon having something in a site that may go away at some point in time. So having moving everything, not just our storage, but we're moving all components of our infrastructure where it makes sense into the cloud.

So we've moved to an SD-WAN solution that gets us off-prem. We're not reliant upon those things. You know, so many other components of our environment are in the cloud when it's financially and functionalities feasible and worthwhile. But if it doesn't fall into those categories, then we remain with it on-prem and with a limited amount of footprint because we know the likelihood of those locations being impacted is very low.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. I saw [ph] someone (02:53:11) hand, Sidney.

A

Sidney Ho

Analyst, Deutsche Bank Securities, Inc.

Thanks. Sidney Ho with Deutsche Bank. I want to follow up with the questions on the AI questions earlier. One of the speakers earlier talk about when you build these AI infrastructure, you buy the GPU first and then the next thing you buy a storage. Just from your experience sounds like there is a [ph] lack between them (02:53:33), you're talking about the iterations of evaluating the need. From your experience what is that, that lead time going to be like, and is it because – the consumption model that you guys are going after, does that mitigated the need for buying storage in certain period of time?

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

So, there's pros and cons to everything. I would say that one of the biggest challenges we face is that not only do we want to make sure that the technology has the right capabilities, particularly from an AI perspective, for me and my environment is a little bit different than yours. But for me, it's how – do I have the ability to get automation, self-healing, all of those things that reduce any capabilities, any possibilities of us having downtime or us having to have manual intervention to remediate anything that may occur. So moving to that on-demand solution, we don't have to worry about doing upgrades. We don't have to worry about maintaining it. We don't have to worry about that thing that's on our partner.

A

So from that perspective, depending upon their adoption and integration of AI into their environments, those things may come faster or they may come slower. Now, if you're building it yourself and you're deploying it in your own environment, like Phil is, that's a little different. I – in my world, I am not to do that. I'm a pretty much a standard infrastructure and operations person that I got to be able to scale and get performance and reliability as cost effectively and efficiently as possible. In his world, he is doing a lot of very custom things. So his challenge is much more difficult than mine.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Well, let's hear about that difficult challenge.

A

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

Well, if I understood your question, I mean, I think the – it's a little quirky for the horse. Because for me it's all about the data first. You start collecting data on things that you find interesting that you've instrumented in your environment, and then you start having questions on terms of, okay, what do I want to learn out of that. What deeper meaning is in there? And that's where you need more compute. You start doing things like data wrangling and figuring out how you're going to make sense of this, and you start organizing that data and everything from data provenance and systems that you're going to be able to do to organize that and get deeper analytics from it.

A

And then you find out, wow, I don't have enough compute to be able to answer that question and you start limiting yourself in terms of, well, I can answer these types of questions, but not these. I may not be able to forecast fully for where I want to be. And if you're lucky, you can get more hardware, more equipment and further on and answer those things all the way up until you go, man, I really could use a GPU.

And so I don't see going to that path of wanting GPUs until you [indiscernible] (02:56:29) beginning of that. You know, as a data scientist looking through everything, saying, okay, I'm at that point now where I can't build those neural models until I have enough course to really chew through this. And it doesn't allow you to escape the hard work. The hard work is the data wrangling, the cleaning of the data, organizing the data. All that stuff is time consuming.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Yeah. David down here in the front.

A

David Vogt

Analyst, UBS Securities LLC

Sorry. Thanks again, guys, for doing this. So maybe for Phil. So I would assume given your status as a National Laboratory, nothing that you do is in a public cloud. It's all self-contained within the laboratory from an on-prem effective solution. So does that change your thought process? You mentioned at some point down the road you need a GPU. Does that change your thought process in terms of what your ultimate infrastructure looks like down the road? Is it definitely you have to go down sort of an NVIDIA GPU-type solution or is there other Ethernet-based solutions like Slingshot that are more than adequate to kind of meet the needs that you have? And so just curious about how you're thinking about that.

Q

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

It's very good question. [ph] I only think (02:57:42) about that. So, I'll say, we have available to [indiscernible] (02:45:57), so anything that is a FedRAMP Cloud we can use, where we have an ATO to be able to use one right now. And we have leveraged it for certain things. Certainly, we have outside vendors that we have as trusted partners that do things from creating diagnostics for us to recalibrating things and sometimes they send data back to us.

A

And so we leverage AWS and S3 so that they can upload data into that and we download from that, for example. So there is some footprint that the National Ignition Facility has on cloud. The greater lab, they have been using [indiscernible] (02:58:29) cloud quite a bit and pushing a lot of other things out to the institution, but I personally can't speak to that. We have from the NIF had some requirements, for example, where we did need some GPU capability and there was a certain set of work that we could not farm over to Supercomputing because they're too busy. [indiscernible] (02:58:51), sorry, and we said, okay, let's see if we can spin up some stuff on AWS, get it going and answer this question.

And I think it's very powerful for being able to meet our gaps between the time to order, the time to deliver, the time to implement in our datacenter versus okay – okay team, we've got some time to rent somewhere else and be able to do that. So those are powerful. But the other things that we look at in our environment since we have an on-prem building is some of the safety and failure modes and effects analysis, and we want to make sure – I mean, there we've always been on-prem presence because [ph] we want (02:59:31) safety interlock systems and those kinds of things to be locally controlled and those are good things.

So we are very careful about where we put things, how we compartmentalize things. I think some time back in the past, we consolidated [indiscernible] (02:59:49) and then you found out while we couldn't take certain parts of our infrastructure down for maintenance and patching and those kinds of things. So, cloud fits into our strategy in a

way of saying, well, how do we best leverage this in a way that is going to make sense. I think you and I were talking about some of this yesterday that some of the challenges of wanting to lift and shift a legacy application and just stick it out in the cloud doesn't make physical sense. You really have to collapse that whole thing and make a truly cloud native application. And so there are things that we're still working through. I wouldn't say we've got all the answers on it, but we are dabbling in it.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Mehdi?

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Thank you. Just one quick follow-up clarification. You said the consumption model is helping you with OpEx, which is CapEx and also helps with TCO. Can you clarify me, are you talking about compare apple to apple like buying NetApp in the past was transactional, now it's consumption, or are you comparing NetApp consumption model to just renting a storage in a cloud?

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

No, I'm talking about comparing apples to apples. So if I was going to go and let's say, I'm buying another on-prem flash storage solution...

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Sure.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

...from NetApp.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

And I'm going to get that same type of solution in a consumption model...

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

...from Keystone.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Sure.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

What we have been looking at previously running the numbers is what's our cost per gigabyte.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

Yeah. That cost at one point in time was not competitive. That cost now has become very competitive.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Sure.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

And it has allowed us to have the flexibility to make the decision.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Okay. And how about the depreciation schedule? Are you now able to like assume a 10-year [indiscernible] (03:01:50) especially since its OpEx, then you don't even have to...

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

We don't have to deal with that. But there's other benefits that...

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

...that we also gain. We get carbon credits. We don't have...

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

...to worry about disposing of the asset at the end of its usable life, which we end up paying to do. That goes away.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Sure.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

So we reduce our carbon footprint, which goes towards...

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

...our getting to a zero emissions model by 2035.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

So all of these things come into play as part of these decisions.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Sure.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

They give you benefits by moving away from this CapEx investment into an OpEx model.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Right. And just two quick follow up. Does this consumption model – I'm sure it includes services or any kind of upgrade, right?

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

Yes.

A

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

And that makes a flash more durable compared to in the past. Like 10 years ago, we didn't know if the flash would be good. After four, five years, now it's – there's a vendor that guarantees it in terms of a consumption model and services, right?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

Yeah. That's correct and it takes the workload off of your staff from an operational perspective.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Okay.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

You don't have to worry about upgrades. You don't have to worry about patching. You don't have...

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Sure.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

...to do any of those things. If there's any type of an issue, our partner...

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Sure.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

...takes care of.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

And then one last item. In terms of a backup cold storage, how does the consumption model impacting the cost structure associated with cold storage or for backup?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

It really doesn't unless you go to a different solution. For us at the end of the day, our backup costs are our backup costs whether we're backing up our storage that's on-prem or whether we're backing up our storage that's sitting somewhere else, it's still that cost is a set cost. So at the end of the day, it doesn't really impact us from the backup perspective. Now if we need to do something different to do that, obviously, then there is a cost impact.

But we're able to utilize our standard products and not have to have a deviation, which is another one of the advantages of using a constant solution across a vendor because you have that consistency across the tools.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Got you. Thank you.

Q

Anthony Lloyd

Vice President-Technology Services, OpenText

You're welcome.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Thank you. Questions? All right. [ph] Give it down (03:03:52) on the front.

A

Q

Thank you. So, Anthony, I know you mentioned earlier your business model or OpenText business model as one of the frequent acquisition. What have you learned in terms of how much of your infrastructure over time becomes more consumptive based versus what is stranded capital in terms of, hey, we can't shut this down? It doesn't make any sense either from a flexibility or a cost perspective. Is there kind of a good general rule of thumb in terms of for every dollar of IT spend, we try to move \$0.75 of it to some sort of consumptive solution versus stranded capital that's more fixed on-prem or however you want to define it?

Anthony Lloyd

Vice President-Technology Services, OpenText

So I can't say that we've gotten to that level of granularity yet, but I can say this as the technology has continued to become better, our goal has been to continue to move towards this consumption-based model because of the other things that I was just talking about. It reduces the workload on my operational staff. It reduces the need for us to have to be involved in maintenance and patching and upgrades which takes a great deal of time, and in addition to the time factor, there's a cost vector associated to it. Because I don't have resources doing all the work because they are focused on these activities. But over time, as technology continues to improve and the price per gigabyte gets lower, it will continue to reinforce that value proposition as well as the carbon footprint benefits that we gain.

A

Q

Okay.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Scanning for questions in the back.

A

Jake Wilhelm

Analyst, Wells Fargo Securities LLC

Q

Hi, Jake Wilhelm with Wells Fargo. Could you talk a little bit about how you see the move towards continued disaggregation in the data center with technologies like CXL affecting your storage and memory architectures over the next several years?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

I think a lot of that really depends upon the use cases. My – how do I say this? In the world that I operate in, we have a pretty well-predefined set of parameters around the performance characteristics of our applications, our databases, because we're primarily talking about HR and financial applications, those really backend systems that really run the business. So those things are a little bit different than Phil's world, where, you know, he is the mad scientist that's doing all of these out-of-the-box things that are one-offs and they're trying to figure it out as they go along. My world is a lot more predefined. So, I would say the impact to me is a lot different than there may be for other companies that have different use cases. I don't really have a lot of exotic applications and things like that, that we're supporting. So, it's a little bit different for me.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Well, I'll say that we've tried really hard to make a lot of our environment predefined.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

Yeah.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Which is why we've leveraged FlexPod and we've got stovepipes in certain environments, so where things may seem disaggregated what we have is pools of aggregation instead of one big aggregated pool. And as I mentioned before, that gave us some benefits for being able to take down different environments, being able to patch, being able to manage the unique workloads so that we're a little bit more fiscally responsible when it comes down to expanding our infrastructure to know, okay, this environment here needs more compute. This environment over here needs more [ph] RAM (03:07:20), especially with our databases and some of the analytics, they're asking for a lot more RAM, as we're loading things more in memory and doing a lot more in memory compute. So those are, like I said, we're trying very hard to not be this mad science [indiscernible] (03:07:41) here. Again, we're using the same Lego blocks that everybody else is, maybe a unique application for how we've assembled it, but yeah.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

We all have the same problems.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Yes.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Well, actually, that was going – well, Mehdi has got a question. Yeah.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Sorry. Just this is for Phil, actually, [ph] non-storage (03:08:00). Phil, as you look into the future, how do you see GPUs and ASICs path crossing? Do you see a day where there will be less of a GPU and more of a ASIC solution for AI application?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

I hope there's both. Yeah, I'm looking forward to a time when there's more code on chips so that, you know, everything from where the – you know, I live in a world where there's a lot of industrial Internet of things. So I'd love to have the data as it comes out of those things to be able to be bagged and tagged and label that at the central point. And as it's flowing through my environment to be able to get a complete manifest of where it's been, what's accessed it, because this is where, as I mentioned from an AI perspective, you need to know the provenance for where the data is flowing. You need to know if somebody poisoned that well for that information if you're going to hope that this AI brain is going to make a good decision for you.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

But is there a wish list? Is there actually a realistic target on the horizon?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Well, I do see that where somethings that are happening now and even in some releases I think that you guys are making is...

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

That might come later this week.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Yeah. You might see somethings that, that may be able to do, localize our on-prem processing.

Mehdi Hosseini

Analyst, Susquehanna Financial Group LLLP

Q

Thank you.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Irvin?

Irvin Liu

Analyst, Evercore Group LLC

Q

Thank you. Irvin Liu with Evercore ISI. So, Anthony, you talked about several of the advantages of spinning up storage in Keystone, such as OpEx versus CapEx and the consumption – the flexible consumption model. But is there a connectivity or low latency advantage from being in a colo data center or on peering – that you get from peering that's worth mentioning as well?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

So for us, we basically provide all of our storage where our compute is. So if we're using compute in the cloud, let's say, GCP, our storage will be at GCP. If we're using an on-prem and one of our colos on one of our data centers, it's going to be there. When you start trying to connect things from on-prem to the cloud is when you get into major trouble with latency and application performance and cost, because every time you take data out, dollars go up, right? So we don't do that. So it's really important that you ensure that you don't start going down those paths where, first of all, your applications are going to be severely impacted. But then every time you – every time things get chatty, it's going to cost you a lot of money.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. Questions? All right. Well, we're getting close to the end of the show. So I'll ask kind of one question to wrap it up. It'll probably take the full time. You guys are in really different industries. Sounds like you're both aiming for a more off-the-shelf approach to your technologies to make it a little bit easier for you. I was going to ask what the challenges you face are, but let's make it a bit of a positive and what are the opportunities that you see ahead of you from your – for your IT environments.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

There's so many. The ability to deploy it faster, to have greater reliability, to have self-healing to minimize the need for it having more support personnel and resources which reduces our overall cost. They are exponential. You know, the whole conversation about shift left, shift left, shift left, that's what everybody wants to do. So the better the technology is, it reduces the need for us to have Level 1, Level 2 people. The technology will take care of that. So that's nirvana, right? So you can spin it up as you need to. You can deploy it where you need to. The time is minimal. You know, all you have to worry about is getting the connectivity to where it needs to be in a way you go.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

What are you excited about as you look forward into your technology crystal ball?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

Well, it starts off, like I said, with the data and being able to access more of that data, our scientists are just ruthless when it comes down to accessing data and analyzing the data. And that's also one of the reasons why I was smiling, as you were saying, the cost of pull things back from the cloud. You know, you don't want to do that.

Anthony Lloyd

Vice President-Technology Services, OpenText

A

No, you don't.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

But, you know, so just being able to enable that, those discoveries, I think is a huge part of saying, we've done the right thing. Looking forward, as I mentioned, I'm enthused about AI or I'd say cautiously enthused about AI. There's more – a lot more good that's going to come out of it than bad. I'm looking forward to its prospects as it relates to cybersecurity, so almost every day hear some places getting hacked into. And I'd like to be able to take this on the offensive rather than the defensive. That 1,000 rocks coming at you and even if you looked at it and you bat it off 999 of them, that one comes through, but if AI can be that other layer, as I'm going through layers of my security, that's something else is there saying, okay, we're seeing something else that would be outstanding.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Are there any technologies on the horizon that you think are particularly exciting or just stuff getting cheaper and faster?

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

A

I'll start off, cheaper and faster will help, because the volumes of data are increasing. There is more [indiscernible] (03:13:57) that you want to grab in datasets, more aspects that you want to look at in order to be faster, better and more efficient going in the future. So those things are naturally going to increase data sizes. We talk about a lot in terms of how do we be – how can we be more intelligent about what we store and not store everything. So as you do that, it hopefully will help the more [indiscernible] (03:14:25) curve. And in terms of the demand for more space versus trying to be responsible when it comes to clean – everything green and sustainability. So those things kind of come to mind.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

Anything on your horizon that you think?

Anthony Lloyd

Vice President-Technology Services, OpenText

A

I won't say specific technologies, but I will say to echo Phil's points, cheaper, faster and more secure, those things are really important. The amount of data that we have to maintain and support is growing exponentially. That doesn't go away. So to your point, how you manage it more efficiently, how you reduce the amount that you have to keep and only keep what you need, but also being able to ensure that you can access it based on the use case in the manner that you need to. From a performance perspective, those things are all critical. So those are the things that for me are really going to be important as we continue to go down this road.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

A

All right. All right. Final call for questions. One in the back.

Q

Sorry.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

No, it's great. It's better that you close it out with a question from you than from me.

A

Q

Yeah. No. So both of you guys kind of laughed about egress fees, and so, I guess, I was just kind of wondering from the perspective of have you brought stuff back and pay those fees, or is it just those end up being kind of stranded islands of stuff in the cloud?

Anthony Lloyd

Vice President-Technology Services, OpenText

So the answer is, yes, which really brings to the forefront the criticality of doing the homework upfront. Everybody seems to think that everything in the cloud is free, which it is not, and you have to do the analysis to understand a few things. First of all, is the application a good candidate to go to the cloud? You have to do the homework and understand what the cost is to remediate the application, because you don't want to do a lift and shifts because what you're going to do, you're going to have a bill that you will not be able to afford and then you're either going to have to fix it when it's in the cloud or you're going to have to bring it back.

A

And so you don't want to do that, but you've got to do the homework upfront, understand the cost to remediate the application, understand the cost on your application teams because you're going to pull most folks from doing other work to remediate these applications, or you're going to pay a third party to come in and consultant do that work. Then you're going to pay the cost associated with the application running in the cloud if it's optimized. So you've got all three of those numbers you got to compare and compare that to what your on-prem costs are. Then you can make a fair analysis and assumption of exactly is this a worthy application to move to the cloud or not? Because if you don't do that, you're asking for trouble.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

That resonates with me quite a bit. The only thing I've got to add to that is, you know, and it's kind of obvious, right, most – in most cases where your datasets are is where you want to do the compute. The problem comes in to where you find out that you've got some data on-prem or some data that's out in the cloud or some data split across multiple clouds. Then what are you going to do when you need to do that deeper analysis to pull everything in and actually manipulate that data. And so at that point in time, now you've got no choice but to pay those egress fees as you point from all of those and slamming it into your machine learning algorithm.

A

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

All right. Well, Phil, Anthony, thank you guys so much for spending some time with us. I really appreciate it. I appreciate you being NetApp customers. And thanks for talking to everyone.

Anthony Lloyd

Vice President-Technology Services, OpenText

Yeah. Thank you.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

Right. Thank you.

Anthony Lloyd

Vice President-Technology Services, OpenText

Thank you for having us.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Thank you, Phil.

Anthony Lloyd

Vice President-Technology Services, OpenText

You're welcome.

Philip Adams

Chief Technology Officer-National Ignition Facility-IT, Lawrence Livermore National Laboratory

Right.

Kris Newton

Vice President-Investor Relations, NetApp, Inc.

Thank you so much. Great. I really appreciate it. All right. You are free to enjoy INSIGHT. I know you're in high demand, so thank you. All right. Well, that concludes our program today. At 3:00, the keynotes go off, so hopefully we'll see you all there. And then again, for those who are here physically, cocktail reception at 5:00 where you can mix and mingle with all the speakers that you saw here on stage today and more as well as the entire INSIGHT attendee program. So thank you very much. Have a great day. Oh! And there are box lunches outside.

Disclaimer

The information herein is based on sources we believe to be reliable but is not guaranteed by us and does not purport to be a complete or error-free statement or summary of the available data. As such, we do not warrant, endorse or guarantee the completeness, accuracy, integrity, or timeliness of the information. You must evaluate, and bear all risks associated with, the use of any information provided hereunder, including any reliance on the accuracy, completeness, safety or usefulness of such information. This information is not intended to be used as the primary basis of investment decisions. It should not be construed as advice designed to meet the particular investment needs of any investor. This report is published solely for information purposes, and is not to be construed as financial or other advice or as an offer to sell or the solicitation of an offer to buy any security in any state where such an offer or solicitation would be illegal. Any information expressed herein on this date is subject to change without notice. Any opinions or assertions contained in this information do not represent the opinions or beliefs of FactSet CallStreet, LLC. FactSet CallStreet, LLC, or one or more of its employees, including the writer of this report, may have a position in any of the securities discussed herein.

THE INFORMATION PROVIDED TO YOU HEREUNDER IS PROVIDED "AS IS," AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, FactSet CallStreet, LLC AND ITS LICENSORS, BUSINESS ASSOCIATES AND SUPPLIERS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE SAME, EXPRESS, IMPLIED AND STATUTORY, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS, AND NON-INFRINGEMENT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NEITHER FACTSET CALLSTREET, LLC NOR ITS OFFICERS, MEMBERS, DIRECTORS, PARTNERS, AFFILIATES, BUSINESS ASSOCIATES, LICENSORS OR SUPPLIERS WILL BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, INCLUDING WITHOUT LIMITATION DAMAGES FOR LOST PROFITS OR REVENUES, GOODWILL, WORK STOPPAGE, SECURITY BREACHES, VIRUSES, COMPUTER FAILURE OR MALFUNCTION, USE, DATA OR OTHER INTANGIBLE LOSSES OR COMMERCIAL DAMAGES, EVEN IF ANY OF SUCH PARTIES IS ADVISED OF THE POSSIBILITY OF SUCH LOSSES, ARISING UNDER OR IN CONNECTION WITH THE INFORMATION PROVIDED HEREIN OR ANY OTHER SUBJECT MATTER HEREOF.

The contents and appearance of this report are Copyrighted FactSet CallStreet, LLC 2023 CallStreet and FactSet CallStreet, LLC are trademarks and service marks of FactSet CallStreet, LLC. All other trademarks mentioned are trademarks of their respective companies. All rights reserved.