REFINITIV STREETEVENTS

# EDITED TRANSCRIPT

NTAP.OQ - NetApp Inc Insight Conference- Technology Session for Financial Analysts

EVENT DATE/TIME: SEPTEMBER 23, 2024 / 4:30PM GMT

**OVERVIEW:**

Company Summary





REFINITIV

## CORPORATE PARTICIPANTS

**Kris Newton** *NetApp Inc - Investor Relation*

**George Kurian** *NetApp Inc - Chief Executive Officer, Director*

**Pravjit Tiwana** *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

**Russell Fishman** *NetApp Inc - Solutions Senior Director, Product Management*

**Hoseb Dermanilian** *NetApp Inc - Director, Global Head of AI Sales and GTM*

**Sandeep Singh** *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

**Jeff Baxter** *NetApp Inc - Vice President - Product Marketing*

**Casey Shenberger** *Hyland Software - Cloud Platform Architect*

**Scott Brindamour** *Lumen Technologies, Inc. - Vice President of Product Management*

## CONFERENCE CALL PARTICIPANTS

**Steven Fox** *Steven Fox - Analyst*

**Aaron Rakers** *Wells Fargo - Analyst*

**Louis Miscioscia** *Daiwa Capital Markets - Analyst*

## PRESENTATION

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Hi, everyone, and welcome to Insight 2024. I appreciate your time and especially for everyone who made it in-person. Before we kick off, I'm going to read a Safe Harbor. Each of the 2024, Insight financial analysts, 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 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 find -- we file from time to time with the SEC, including factors described under the section titled Risk Factors in our most recent filings on Form 10-K and Form 10-K Q available at www.sec.gov.

The forward looking statements made in these presentations are being made as of the time and date of the live presentations. If the presentations are reviewed, there are no presentations. If the Q&A sessions are reviewed after the time and date of the live presentation, even if subsequently made available by us on our website or otherwise, these Q&A sessions may not contain current or accurate information. We disclaim any obligation to update or revise any forward-looking statement based on new information, future events or otherwise.

So with that, said, thank you. Let me welcome you again. As a reminder, today's sessions are all technically focused. We won't be covering any financial information. So this is your opportunity to ask our tech leaders, questions about our products, our services, our value proposition, the competitive environment.

I know those are always top of mind questions for all of you. So you get to hear it straight from the horse's mouth. And after this session, we'll -- for those of you who are here. You have some free time, the keynote start at 4:30, you're welcome to go to the keynote.

There's a small section of reserve seats that you're welcome to sit out, but you're also equally welcome to go find a customer and sit next to them and get their live reaction. After the main session, the showroom floor is open and you can wander around if you need anything, find one of the IR team and we're happy to help you.

With that, it's my pleasure to invite George Kurian to the stage. George?

---

**George Kurian** - *NetApp Inc - Chief Executive Officer, Director*

Thanks Kris. Good morning, welcome to NetApp Insight. Thank you for taking the time to be here. We are super excited. We have a lot of innovation payload that we will be sharing with our clients that our teams have been working hard on over the past year.

We will talk about with our clients that we are at the start of the era of data and intelligence. What we've seen over the last several years that we've shared with our clients is that data driven organizations are outpacing those that don't have their data well organized over many years.

But now the tools that are being made available are even more powerful than they ever have been before on two dimensions. The first is, you are able to analyze a set of data that is frankly for any organization, the preponderant majority. 85% to 90% of an organization's data is unstructured data, meaning files and videos and documents of various sorts, it's conversations with customers and whiteboard sessions and so on.

It's about 85% to 90%, and we are the unquestioned leader in that part of customers' data landscape. The second is not only are the tools more powerful to be able to normalize and understand that data, but they have almost human-like intelligence to be able to understand the domain in which they operate without human involvement and be able to switch domains.

So go from one domain to another, go from one modality of operation to another, right? So going from text to video or video to image and so on. And so we talk about the fact that we are at the junction of data and intelligence and what's required for success is that you need to have your data and data strategy well organized.

We said that to our clients last year, and we are reinforcing that to clients this year. I was just out on the road over the last few weeks with clients, and I'll draw two banks, one whose data is well organized. They are already using Gen-AI tools. They've got a hybrid architecture and they're making good progress with those tools.

Another bank that had a classic siloed custom built architecture with lots of different landscapes. They are a year away from buying an AI computer and then starting to put their data on it, right. So profound differences, you will see, for example, that in industries like life sciences, where there has been regulation and requirements to have high quality data catalog to the right way with clinical data codes and with procedure codes that they are making much more rapid progress than those that have not had their data well organized.

And surprisingly, some of the more regulated industries are making more progress than the unregulated ones, so data and data strategy super important. The second is when everybody has really powerful tools and all of the world's data supporting those tools, your domain knowledge and insight becomes even more important.

And then the third is the ability to take your domain knowledge, your data tools and then apply it in an iterative test, learn and adapt loop so that you can graduate some of these projects from proof of concept to production and then close the loop back to make your data science environment even more strong.

And then finally, a data ecosystem that enriches your data, just like your business ecosystem supports your business. Those are the four key things that we'll talk to clients about. And then we'll talk about the two challenges that organizations face in using their data effectively to support advanced use cases like AI and there are two challenges there.

One of the familiar challenge, which we are skilled at helping clients with it's a data management challenge. And so what does that mean? How do I find that data across my landscape that I might want to use for my AI project. How do I govern sensitive data so that I can bring forward? My security and access control model into my AI landscape, I can ensure privacy and so on.

And then the third is how do I actually keep my model environment and my data environment in sync as the data gets fresher and fresher or data progresses through the lifecycle as it always does. So there's a data management challenge.

The second challenge, which we have observed from interviews with around 800 clients over the last year. And I've personally participated in probably 50 discussions with clients around the exact same problem is that AI is being built as a silo. It's got custom networking, custom chip, and it's not integrated into your data landscape.

So there's so many clients that we've met that said, hey, I've got my AI computer stood up, but I can get data to it or we had a semiconductor vendor that was copying 300 terabytes of data a week to tried to keep their supercomputer moving and I asked the gentleman who led that project, I said what's your lifelike?

And he said one word [Hell], and you said it will not scale the way it is. Now if you think about that gap, the chasm between your AI environment and your data and operational environment for us, that looks just like how cloud ex cloud was many, many years ago, we stood up here in 2013, and said, If cloud was to become useful and seamless, you needed to have a bridge between your on-premises environment and your cloud environment.

And we call that bridge the Data Fabric and we innovated over many years to make cloud and enterprise data work seamlessly together. And today, you will have many of those cloud partners talking about the AI journey that they are taking on with us.

And so for us is a familiar problem, and we are innovating to bring capabilities to the world of AI that don't exist today and those capabilities exist along three fronts. First is a set of tools and applications that make it easy to find the data. All of your data estate so that you can quickly explore it and then choose what data you want for your AI landscape.

The second is to have bring AI to your data, which is much easier than trying to bring your data to AI, you see data is the gravity part of that equation. And so we know how to bring AI to your data and you see two or three transformative capabilities that we bring there.

The first is the best infrastructure for AI and we will talk about the third architecture that third generation distributed system architecture. The first generation were shared nothing architectures like Pure FlashBlade or Isilon or Cumulus.

The second was distributed architectures with centralized transaction management like what vast data has and the third is actually a truly distributed architecture where the transactions and file operations are distributed across the system.

And so that's the second area. And we combine that with capabilities that allow you to have model versioning and data versioning synchronize to have traceability of your data together with the models and to have highly efficient patented technologies for data retrieval, retrieval is the first step of what you call rag in your inferencing world.

And then the last is the set of capabilities that we have built over many years that we've enhanced for the world of AI where you can bring all of your security policies and privacy and controls across your AI life cycle and be able to detect changes in the data. So that you can efficiently apply those changes to your models.

And so super excited. We think that we have step function improvements in our capabilities for unstructured data, reinforcing our leadership there. We also have really good announcements in block storage, cloud storage. That Pravjit will talk about the security capabilities in our portfolio. And all of that falls under the umbrella of intelligent data infrastructure.

So I'm excited for you to hear from our technologists today, but equally from our clients and the partners that we co-innovate with. Thank you for coming. Look forward to a great conference.

---

## QUESTIONS AND ANSWERS

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Thank you, George. Appreciate it. So now we're going to start the first of our Q&A session. I'm happy to introduce a new face to you all Pravjit Tiwana, get that right? All right, who heads up our cloud storage group.

So one of the questions I get from you all a lot is what is the value proposition of your cloud storage services. Why do people choose to use NetApp in the cloud so you can ask directly, why don't you come on up?

Thank you. And before we have you start asking him questions. Pravjit why don't you introduce yourself say a little bit about what you do and then we'll open it up to the audience and I'm going to switch places with you.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Hey, everyone. My name is Pravjit Tiwana. I lead the cloud storage group and SVP and GM for the cloud storage encompassing all three hyperscaler services from us and AWS, Microsoft and Google.

I joined NetApp earlier this year in March, although I have been in cloud world for pretty much all my career. What really excited to be here is the world is seeing really very interesting times with the growth of data, which we are seeing zettabytes of data is being produced every week.

And if you intersect with what is happening in cloud, right, like the migration in cloud, the momentum in cloud and overlay it with the advent of AI It creates a perfect storm right now in terms of right like, the intelligence, the capabilities, which you can build on top of your data.

So from that perspective, right, like since there is so much of data growth, there's so much of every organization is thinking of how to use AI to further accelerate their businesses and menu mix, cloud, this thing. The complexity and scale at which organizations have to deal with it is becoming more and more complex.

And that's where if we see pretty much every business or organization have some cloud strategy, either they are all in or they are exploring it or they have one or two mission-critical workloads running in the cloud. And that's where we play that role because we have our first-party cloud storage services available in all three hyperscalers today.

And we have a Deep, thanks to Deep partnership at all three of them, we continue to innovate on behalf of the feedback which we get from our customers. So I can talk about it or we can do a Q&A, whatever --

---

**Unidentified Participant**

So a question we often get from investors we talk to is, sort of how do we think about your competitors playing catch up to the first-party services that you already offer on the public cloud? You obviously go through the marketplace.

But in terms of capabilities, how do you think about where the roadmap is to sort of stay ahead of the competition on that front? Do you at all see a risk of down the line, sort of some competitors moving to becoming sort of the first party services on those public clouds. How do you sort of think about that risk?

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

No. Yeah sure, right. Like if you see the landscape today, we are the only first-party cloud storage available in the cloud with all three hyperscalers, right?

Yes, the marketplace offerings might be there, but marketplace and first-party is kind of like an apples to oranges comparison because the deep integration which we have with the entire stack there, the whole sales motion to billing to specialist support and operations which you get in first party is not the same thing with you get in marketplace kind of offerings, but that said, right, like having multiple players are having competition is usually not a bad idea for customers right? Like we don't have any mechanism to say who will do it. And we don't need to speculate that either right like who else will be in the first-party cloud stores versus down the road.

But where we are focusing on is that we are being relentless in terms of the innovation we can do on behalf of our customers. The number of capabilities which we are building working with our hyperscalers is phenomenal. And you will hear a lot of those capabilities either we have announced in the earlier part of the year and some of those we are announcing later this week.

So you'll see a broad array of, the one thing which I have seen turning cloud services at a very large scale for a very long period of time. There is no compression algorithm when it comes to things like scale and building the capabilities over and over again, like iterating on those capabilities over the years, you cannot like if somebody storage compute network, these are kind of parameters, right?

Like where it takes time to understand how to build scale, how to build operational strength. If you see in my organization, we spent so much of time on things like security reliability availability performance. Those are the things which are not just easily replicable from anyone because it takes years to build those kind of capabilities and we are investing those -- in those a lot.

But when it comes to marketplace offering that we don't think that's Apple-to-Apple comparison today. But we cannot also say that, hey, what's the future for our competitors down the road, but in short customers having wide choices is not a bad idea, it raises the bar for everyone.

---

**Unidentified Participant**

Thanks. So many of the customers who end up buying the cloud storage are kind of new customers to NetApp. And so just how does the value proposition to those customers kind of different than those who would have traditionally been NetApp customers? And what's kind of the most effective way to get those customers on board?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Yeah, you're right. Three out of five customers, which we get today in our cloud storage they are net new to NetApp and the remaining part of those five are using us in hybrid mode. So from that perspective, yes, there is a -- The good thing about being available in the cloud is the -- is being available in front.

And especially when you are integrated into hyperscalers, it's being available to all the customers and they can themselves play and figure out right, like what's the value which we provide. We see lots and lots of that developer based ecosystem also happening in our first-party cloud services where customers drive start with the small workload, learn that how much goodness we bring in terms of performance or about multi-protocol support or the data protection capabilities which we have or the security investments, which we have done.

Like this year, working with Microsoft for almost six-eight months, we have been working on the SFI initiative from Microsoft likely so that customers are getting things like not just our security goodness, but they're also getting all the bars and the security controls, which are hyperscaler partners are defining.

So combination of those excites, these new customers to start using end. We have been fairly focused on price performance optimization in our hyperscaler clouds, and that also resonates to the customers and the -- but the final thing is right like the AI. integration, which we are able to provide.

If you see most of the science behind AI is being driven by that and also hyperscalers like all three hyperscalers are very, very heavily invested into building the AI stack along with obviously and videos in the metals of the world, but being using our sources and hyperscaler setups makes it almost seamless to use the AI stacks with hyperscalers they have built.

Without the need to like replicate data, or clone data or copy data and create silos, they can do it wherever the data resides and can use that AI.

So the combination of our security, our AI, our price performance, our operational strength and the capabilities which you have built in ONTAP for almost two decades. All those combinations resonate with the new customers as well as customers who are using us in a hybrid mode.

---

**Unidentified Participant**

Hi, it's [Tim Longo] at Barclays. You mentioned AI maybe a two-part. Talk a little bit about kind of how your business is affected by a lot of these large language models that are going on now. And I'm sure it's not major at this point and then maybe walk us into when we get into inferencing and there's a lot more on-prem and bursting to and from the cloud. So how do you see that dynamic impacting the cloud storage business for NetApp? Thanks.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Thank you. This unstructured data and file, like most of the AI is going to run on data and data runs on NetApp right? So means it's a perfect combination for us. So we are really excited that AI growth in AI and the improvements in AI, which we are seeing, at least in the last 18 months, it has accelerated a lot. It is really good for our customers to be able to.

Our mission there is that we want to provide the AI capabilities irrespective of whatever LLM models you want to choose right there on your data, right, like you don't have to copy clone and those kind of things that are so we don't want to create silos.

So if you look into earlier this year, we launched workload for Gen-AI, that is basically a capability in which you can integrate LLM models from Amazon Bedrock into the data, which we are storing in FSx for ONTAP, just with few clicks right? Like you don't have to copy it to say at three ObjectStore or any other storage system, you can do that.

So we see, that is the -- And once you have integrated some capabilities like you're better off condition models, then you can also start overlaying it with other capabilities which you might choose from marketplace or you choose from hyperscalers or you choose from us also, right? Like so that provides a very rich ecosystem for customers to use that AI. And there was a second part of your question. I forgot that part.

---

**Unidentified Participant**

(inaudible - microphone inaccessible)

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

So let's talk about the hybrid cloud, right? Like because we see very significant growth in our hybrid cloud setup year-over-year. So taking that same case of what we have done with Amazon Bedrock to use LLM models in case the data which you have stored in FSx for ONTAP. You can also combine it with their on-prem storage.

So things like I'd like with Snap mirror. You don't have to copy the whole data, but you whatever is the data set which used to want to use for inferencing or your chat bot or whatever. You can combine that data into FSx for ONTAP. And you will see these kind of pattern also in other hyperscaler partners.

So you will -- from that perspective, we want to provide and working with our hyperscaler partners. We want to provide a mechanism in which you don't have today. So here is AI for cloud, here is AI for your on-prem. We want to provide a seamless integration, and that's the path, which we are on.

On inferencing. We do support some of those capabilities today for your data and with our hyperscaler partners, we earlier this year launched Gen-AI toolkit for Azure NetApp files as well as Google Cloud NetApp volumes, which provides -- you can use it against a proprietary data, but few clicks, you are able to build your drag interfaces on top of that and you can build applications like Chatbot, knowledge base, all those kind of capabilities with a few clicks only.

So inferencing part and the RAG, and Gen-AI toolkit -- the workload factory part of our also support building your whole RAG and infrastructure with just few clicks. I highly encourage you to look into. We are doing these demos this week and different sessions and also on the floor. So please do -- if you have time, please do look into those because these are really, really exciting and step forward capabilities which are coming.

---

**Steven Fox** - *Steven Fox - Analyst*

Steve Fox with Fox Advisors. Maybe just you mentioned a couple of things on the roadmap, but can you just sort of step back and talk big picture of how you envision the roadmap and without, I guess, giving away the announcements you have this week, what are the big things we should think about that you're focused on, say, over the next 12 to 18 months for improving on?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

In AI or in general?

---

**Steven Fox** - *Steven Fox - Analyst*

In cloud storage.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Overall in cloud storage, right. I think it's -- majority of our roadmap is controlled by our customers, right? Like whatever we are building is based on the feedback from customers. So if I have to look into what dimensions we are working on, right, like without going into right, like what land ran and all but the dimensions I can talk about, right?

So that from a dimensions respective right, like we know that customers have trusted us to land a lot of their mission critical workloads right? We read about their ERP systems of their databases like Sequel Server, Oracle or high-performance computing or streaming content. The list goes on in terms of the kind of workloads they run with us.

Our roadmap in from that perspective is highly focused on making sure that if customers have to run their mission-critical workloads, we are the best destination for us then. So that's one dimension of what we are working on in terms of focus.

The second dimension for us is what we discussed about AI. Our mission is that we want to bring the best of AI stack to the data, which you, as trusted with us, be it on cloud or on-prem. And we provide those capabilities so that you can build instances, RAG interfaces, whatever you want, paired with almost like in a frictionless manner. That's our second dimension.

That third is the cost optimizations, we have like if you see in the recent past in this year, we have shipped a lot of capabilities, our customers have told us that they love capabilities like autosteering, which is now available in all three cloud for some good period of time.

Which basically move the data from [Part two card called tier] or a different tier based on the data access patterns without the customer have to do something about it, all the things which we are the all the ONTAP averageness like compression, compaction, deduping, thin- provisioning, all those.

So we'll continue to keep our focus on making price performance optimizations. Our goal is that a cost should never be the reason why you don't select our first-party cloud services. And the fourth dimension, which we will continue to focus on is in our hybrid.

We understand that a different customers that are at a different stage of their cloud migration. But irrespective of whatever stage they are in. We want to provide the best capabilities so that the not just migration, but deployment and operations of those workloads and the data which they bring to us is highly, highly optimized.

And that's why it I'd like we'll continue to invest in things like disaster recovery, data mobility and couple it with the elasticity that you get in cloud in terms of instantaneous, cloud capacity or bursting or those kind of capabilities. So it's a four things, right, like the workloads, AI price, performance optimization and hybrid.

---

**Unidentified Participant**

Pravjit, not sure if this is sort of the focus, but cloud ops, the cloud ops part of the portfolio. Can you talk about sort of how do you think about the value add relative to the cloud storage part of the portfolio that you have sort of more focusing on? And then when you think about inferencing as well in terms of AI workloads. Do you see the value proposition of cloud ops changing on that front eventually?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

I'm pretty sure Ian, is also talking today on cloud ops. I can give a little bit high level.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Yeah so we don't have Ian today. But if you could be very high level and I can always follow up with you of semicon, any cloud ops you want.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

We have a different leader who runs the cloud ops portfolio. I'm sure you know her. We have like a bunch of places where we have synergies in terms of right, like making it into the hyperscaler cloud, especially around things like cost optimization, goodness, which comes from out of the portfolio or the observability stack, which we can integrate in those.

So we are looking into it. I would not say we are just looking into it, but many of our customers are looking and using in that fashion, especially around observability stack and cost optimizations aspects and so on. And we will continue to work with cloud ops in that sense to build a comprehensive portfolio. That's where all the aspects of our unified storage or data protection or cloud ops portfolio. They all work together to provide frictionless way to manage your workloads. I think that is the part and I'm not super deep into the cloud ops portfolio, but this is what I know what we are doing.

---

**Aaron Rakers** - *Wells Fargo - Analyst*

Yeah, this is Aaron Rakers, at Wells Fargo, how is that share gains evolve for block storage? I'm not sure this is a class specific question.

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

No, -- what's the question?

**Kris Newton** - *NetApp Inc - Investor Relation*

So Aaron, I'm going to have you hold that until Sandeep comes on. But I do think you can talk about block storage in the cloud because I think that's the thing that people don't expect to hear from us.

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Actually, block storage is a part in our cloud portfolio in our first-party services, which is really which personally caught me by surprise when I joined. And last one year, we have seen 140% growth in our block storage and first-party cloud services.

So it is pretty significant growth for us. And we are really, really proud and excited about that. And the good thing is right, like most of these customers, they go to say AWS console or some other console and the Central region, the block capacity because they are familiar and they know that how much ONTAP richness we have built over the years into a block storage portfolio and they can use that in the cloud also. So yes, block is an integral part of our strategy as well as consumption from our customers.

That was the question about block, right?

**Unidentified Participant**

Thank you, [Wamsi] at Bank of America. I want to ask you a little bit about when you look across your implementation with different cloud providers, it's kind of different the way the technology is implemented. So as customers are deploying, say AI workloads. Do you think that one architecture is more favored versus other architectures? And what are maybe some of the benefits of one versus the other if you see any

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Yeah, I think the way customers use us as from the interfaces, right, be it from the console or from the APIs or the SDKs right? In general, like the nitty-gritty of the architecture are kind of like opaque to put from a consumption perspective from the customer. So right.

And then from that perspective, our goal is to bring uniformity. So we have uniformity in the same terms or price like the protocols that you use to access data, irrespective of what cloud you are they have common rightly so.

And then we have a Blue-XP our manageability, which you can use to manage multi-cloud setup. So it gives you a single interface and all. So matter point being that, we are trying and based on customer feedback, we are trying to make it uniform so that you don't have to figure out -- customer doesn't have to worry about, if it is the one implementation behind the scenes, second, they get the same performance, same capabilities as they will get in one hyperscaler versus the other.

I think the second part of your question, right, like hey, is it good to do. I don't think like in computer science like right, like there is a perfect answer for these things right, like I think both have their mechanism. I think GCNV and ANF side, we pretty much do everything from the service delivery engine capability which we have built. So that is a lot of the property.

FSx shows unique way to do it. And that is also very, very the kind of scale growth, which we say, we see with that is humongous, but customers are not bothered about that. We have a lot of customers who use us in multi-cloud setup using FSx (inaudible) with ANF, or something like that. And for that, we make it seamless for them.

**Unidentified Participant**

And if I could just follow up for the customers, I think to [Meta's] question that was you've got three out of five customers are new to NetApp on the cloud. So for the ones that are not is the AI deployment currently, whatever they're doing with that, is that more on-prem centric versus cloud centric? And are you seeing any signs of early migration of AI workloads on NetApp to the cloud?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

We are seeing science, but you also have to look into the landscape of the businesses, right? Like a lot of businesses today, our clients are in the learning, more fine tuning their models and those kind of things, right? So for that's why I was talking about, Gen-AI toolkit and those kind of things.

So customers are showing a lot of interest. And I'll say a little more than like just POCs at this point in terms of that. But in fullness of time, I truly believe that, yes, it will be one of the core players in terms of right, like especially using the AI seamlessly between cloud and on-prem.

---

**Louis Miscioscia** - *Daiwa Capital Markets - Analyst*

Okay. Lou Miscioscia, Daiwa Capital Markets.

Similar to [Monty's] question, I'm just wondering because this was all deployed at various different stages and some of the cloud, you are customers of yours first. Where are you with all of them in the sense of you just talked about trying to be uniform, but some more advanced and you're trying to take those more advanced features and sharing with the other one, some sort of wondering if there's some that you're just doing a lot better with and if so, how would you take that to the other one?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

You mean like from (multiple speakers) Azure versus capabilities in AWS versus Google that -- in that sense.....

---

**Louis Miscioscia** - *Daiwa Capital Markets - Analyst*

Yes, in total revenue and growth and all the things we care about.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

So just a reminder, we're not doing any financial updates. So Pravjit will not talk about the financial ramifications, but you can talk about the high-level capabilities and partnerships across the clouds.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Yeah. I think you have to look into a broad spectrum of things here, right? Like if you see on the AI side, right, like it's a relatively new space, right? So that standardization and uniformity will take a little bit of time to land between not just hyperscalers, but in the industry in general, right.

But when you come to the other things, right, like how you access your data, be it by different protocols, right, be provide capabilities like, hey, you use the same protocol to access. We have windows or Linux, right without needing to duplicate the data or the unified storage capabilities.

From that perspective, there is a lot of uniformity which we are bringing in terms of access as well as from our Blue-XP framework or the capability right? But in fullness of time, right, like we do expect our customers to start using more integrated services from these hyperscaler sevices.

I think that's a natural progression. We don't want to control or be in the middle of it, where they find value, they should absolutely use that and yes, there will be little bit like maybe the object store and Amazon is different than Microsoft or the security, if they want to use something other capability from them that might be different. But yes, there will be a little bit of that. But that is the part which we're trying to make it seamless to our Blue-XP arrangement for multi-cloud setup.

---

**Louis Miscioscia** - *Daiwa Capital Markets - Analyst*

Okay, thank you.

---

**Unidentified Participant**

Thank. Understanding from the customer standpoint, they're most interested in kind of being able to not have to replicate all this campaign from various sources, I guess from a cloud customer standpoint, is there anything that they're asking you to kind of work on developing other than just accelerating kind of the back and forth that they think can kind of help optimize their services?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

I assume your question isn't reference to AI right?

---

**Unidentified Participant**

Yes.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Okay. Yes, we are deeply integrated with our customers that's -- there are a lot of capabilities which we are building in terms of right, like how to manage your Rag based infrastructure. That's the part which I was talking about with our Blue XP workload factory. Now you can a few clicks can enable that.

So what customers are asking us to simplify the complexity of the AI. Like, there is a lot of complexity involved in that today and how you set up your infrastructure. How do you bring your foundation models? How do you build inference on top of it.

So what they are telling us is right like, hey, we are trusting to store our data with you. We are running workloads on top of you. So make the whole AI ecosystem also simplify. So our goal is to make it frictionless and simple, and that's what we are doing, and that's what our customers have told us to do.

---

**Unidentified Participant**

Like, I guess, I mean, what are Azure, Google, Amazon as like, are they asking you for anything different than your customers are asking (inaudible) for?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

I wouldn't say it has an ask, but yes, we are deeply partnering and how to simplify it for our customer. So we do a lot of engineering and technical conversations with our hyperscaler partners to make it frictionless. They do like our hyperscaler partners and fully understand that there is lot of worlds data which are non NetApp, right? Like and they want to utilize that with AI.

So from a simplicity perspective, yes, we have a lot of integration discussions, a lot of discussions with our hyperscale partners and that defines the roadmap or the capabilities which we build in cloud.

---

**Unidentified Participant**

Just a clarification. So when you do all this work to simplify, the sort of, let's say, Rag for your customers and the cloud customers want you to do that? Like does this come in as a feature that you offer to enterprises? Like how do you get monetized for it? Or are you just more dependent on eventually more data finding its way to the public cloud? Or do you then have a feature that you can actually monetize at a premium with the customer for putting that work?

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

I think monetization part, we can discuss a little differently, but this is a core capability which customers needed today and especially in the future monetization part is I'm sure like, if you bring goodness to what customers want, monetization part is separate.

Right now, our main focus is to make sure that we make AI seamless to work with our -- yes, there are some monetization aspects of that. We haven't fully built that model. I don't think anyone has been the full model yet on like, right, the AI monetization part. But in fullness of time, we do expect it to grow our business. And I think Hoseb Dermanilian a lot of going to talk more about AI. Is it in the next one? Yeah.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

And we have time for one final question on cloud storage.

---

**Unidentified Participant**

Great, thanks. Pravjit, can you talk about what is your understanding of how the hyperscalers? And obviously, each one is different, but they are key considerations for build versus partners has changed over time. And because they do the hyperscalers, do source and write a lot of the raw components, they could do a lot of themselves, but you need to talk about complexity about scale and you hit on a lot of these topics.

But what I guess, if you really want to narrow down, like are we, the crux of the question really is, are we at a tipping point in terms of like there's only so much that they can do themselves that building more to third party services and native first-party services like of NetApp, therefore, the floodgates will have opened or are about to open in terms of give me sort of --

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

No, I get sense of your question. I don't think I want to say on behalf of a hyperscaler, if they are on tipping point or not. But what I can say is that right, like one thing, common across all three hyperscalers is they're very good at listening to the customers, right? Their customers are telling us that, hey, we want -- say, take the case of NetApp storage right? Like they are telling them that they want those capabilities in the cloud and they are responding to it in a right manner.

Now there are two parts side like we have been a lot of innovation, a lot of capabilities over the last two decades as I was talking about earlier, right, like there is no compression algorithm to basically say that, hey, we can build the same capabilities in next and months or something that's part A.

And part B is, like we ourselves are on behalf of our customers are not also stopped doing innovation. So we are also accelerating that. So combination of what we have built the richness over the years and our laser focus on the dimensions which we've talked about. This is a combination of this

will continue to keep all of us relevant, and it's the right thing for customers perspective also to have these capabilities. So our hyperscaler partners understand that. It's not like we versus and kind of a situation.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, thank you, Pravjit. I appreciate you coming up and taking all the questions and handling all the questions, I really appreciate it. So thank you for your time today, (multiple speakers) next meeting.

---

**Pravjit Tiwana** - *NetApp Inc - General Manager and Senior Vice President of Cloud Storage*

Thank you so much.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Thank you. And now because AI. is such a hot topic, we have two presenters who are going to come and handle all your questions. I'll happy to invite Russell and Hoseb up. Some of you have heard from them before. They've done a lot of work for us. So before we kick it off with what I'm sure is an endless stream of questions from the audience. Why don't you each introduce yourself and talk a little bit about what you do related to AI for NetApp?

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

(inaudible - microphone inaccessible) So Russell Fishman. So I lead the Product Management for solutions, and that happens AI solutions in particular. My responsibility, partnered, very closely with Hosed and our product and engineering teams. It is on making AI real for our customers by taking products from our portfolio and combining them with third parties primarily to create complete use cases that help our customers adopt and accelerate their use of AI.

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Hoseb Dermanilian, good morning, everyone. I run the AI sales and go-to-market for NetApp, been with NetApp for 10 years, and I've been doing this for six years. So I am very fortunate to be on this journey for almost six years now with NetApp. So I'm happy to answer any questions you have about what customers are using NetApp for? Or how do we see the market growing from a customer standpoint.

---

**Unidentified Participant**

Very big picture question. So George talked about the journey from cloud to the data fabric that started in 2013. So now you're starting on a similar journey, any mistakes that you would warn us about the you're going to have to overcome as AI evolves because I think a lot of Wall Street's concerned about the hiccups, not the endpoint?

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

So you want me to cover on the journey or --

---

**Unidentified Participant**

Just like, what you're looking for -- and just like what you're potential problems and challenges?

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Potential problems and challenges. So I think what we are seeing is customers today are trying to use this large language models, but the first trial is not is not as good as they want to be that they can deploy it in their own enterprises, right? And we're talking if I'm a customer, if I'm a large retail customer and I want to provide a robot that will answer support questions and that cannot be trained on a whole basis that was available to the entire globe and it might answer questions differently.

Some of them, they are not even going to resell them and they can identify the name of their CEO. So I think that is the biggest challenge them understanding that they need to bring that data to these models to make these models more specific.

I think the overall expectation today that they can just go and use dislodging win models and it will apply to every problem that they have. What the biggest challenge we see is that expectation becoming a frustration, which then you people either stop using the technology or they adapt and Okay, if we want to use this and we're going to do this one, two, three, four, bring the data find 200 model, rag, the model, et cetera, et cetera.

So I think that's from my perspective, is one of the challenge. You could tell the other challenges, cooling power data center, but I think for the wider enterprise, the hyperscalers will be a big player in this. So even talking to customers, those who doesn't have the power and the cooling and data center capabilities, they will go to the hyperscalers, but then they will face this LLMs not being trained specific for their own models, right?

So it's a mix of both. So either they fixed or cooling and power for their data centers and train these models up, ground up in their data centers or that data that exists today on premise is bringing to this large and which models and open-source models or the hyperscalers, but then that becomes the data challenge. Do we move the data, the data is sovereign and all that so that's what we see.

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

I might add to science at what Hoseb said. So if I look back at that transition, as you said from a data fabric. One of the big challenges we had was we had the technology, but actually getting customers to adopt it was complex. It was difficult a required knowledge and sophistication of those customers also required knowledge and sophistication on behalf of the partners that make it real for most of the customers.

So I think one of the things that we're really focused on this journey, a journey is to accelerate adoption. And so that really means having a fantastic set of partners that we work with. So Hoseb, mentioned the hyperscalers, but our go-to-market partners, we already have a partner square program that focuses on getting AI adopted by our partners with our technology portfolio and also the fact that we have focused on integrations with third parties.

So we already have this amazing the rich set of ecosystem partners because we don't believe we can do it ourselves or that we don't believe that we're the only people that people need to work with to deliver AI we think we're an accelerator.

But what we're doing is by focusing on that complete picture is helping people adopt. And that's really this kind of inflection point of democratization where Hoseb been brilliant helping customers who want to get ahead who are willing to go build things themselves, maybe have the data sets have that sophistication. #

What we're starting to see, of course, is, many more customers who want to adopt without having to go through the development was like just buying off the shelf. So that's the inflection point that we're starting to see and all those integrations that we already do or the way that we're already helping that.

---

**Unidentified Participant**

Hi, thank you. Maybe you could talk a little bit about the most common use cases, applications that you guys are seeing NetApp getting involved in. How broad is the that set that you highlight two or three? How broad is the set and how do you see the overall set growing?

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Yes, absolutely. So I'll do it on over the past six years. The journey, right? So we started actually six years ago when we certified our storage with Nvidia DGX-1. That's what we saw, customers -- that time, it was deep learning and neural network. So there was no Gen-AI much of language models, et cetera. It was mostly customers who are developing their own in-house models to, for example, we have one healthcare customer. This is six years ago, right? I'm not talking six years ago, we deployed Nvidia DGX with NetApp ONTAP to train models to detect anomalies in the x-ray images and et cetera, et cetera.

So I think image processing and computer vision was a very big use case and is still very big use case before and after lunch larger language models. And then over the time, we started seeing customers build AI center of excellences. These are large customers, we have some of them here at the conference. They going to speak later on where it's not five users. It's basically they have multiple data scientists scattered across the nation as well as across the globe, where one of them will do a model training to detect fires in the forest.

The other ones will detect, something in the war zone, et cetera. Because what I'll give you an example of a system integrator who does multiple different things and they wanted to build AI center of excellences. Again, this is based on the GPUs they purchased in-house with the data that needed to feed those GPUs.

Obviously, when large language models and whole chat-GPT and Gen-AI boom. Now we are starting seeing those customers who, as I mentioned, want to leverage the goodness of the cloud, but then their data has gravity and their data is on premises.

Those use cases is mostly we're seeing AI agents. Customers want to build the AI agents to either respond to support to either respond to customer service. Now we're seeing people writing software through AI. So that's another thing that we're seeing as a use case. And this is now where the cloud and on-premise becomes together. And NetApp is actually the connective tissue from the data perspective. So that's the evolution of the life cycle. We've seen.

Now aside of this, we also have seeing a handful of customers who probably are witnessing is those who are building their own large language models. It's not probably as big as the hyperscalers, but it is big enough to call it a super cluster Nvidia likes of term SuperPOD, they call it SuperPODs.

They are -- I wouldn't say it's the majority of enterprise because again, back to the cooling energy power, all that requirements. But we have seen customers actually build that large supercomputers to specifically train foundational models and this is because they either want to provide this as a service to some other folks, they want to compete with the hyperscalers or they just don't want to leverage someone else's model. So we have seen that type of customers as well. I hope that answer for question..

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

I had a couple of pieces of that. So yeah, you said you talked about personal productivity, chat bots and copilot enterprise knowledge management, obviously pretty, pretty broad a use cases that have from a vertical perspective, broad applicability.

But what we're starting to see absolutely is more interest in highly verticalized solutions. So we talked about LLMs and talking about SLMs and excellence and those that have been specifically augmented with the knowledge that's necessary in each particular industry.

So there's going to be an explosion of that. And we're absolutely we're seeing that in the market, and that's really probably closing the gap between folks trying out these generative AI solutions and actually making a very useful in a more specific context beyond just general productivity.

So what you see, for example, with Microsoft and their copilot for example, but that they're very broad. And then moving into industries like legal, finance, healthcare, in particular, life sciences, we see these very specific versions of LLMs, SLMs is coming into play. So I mean, that's going to be definitely where things are moving.

**Unidentified Participant**

Can I just ask on when you talk about go-to-market on this front, how is that different from the traditional infrastructure go-to-market in the sense that if customer today wants to sign up and sort of get more educated from you in terms of what that technology should look like they want a full understanding of the complete road map, including going up the inferencing before they even sign up in terms of what a four-step looks like and how would you compare those aspects?

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

I'll answer, then Russell you can. So it is different than a typical, hey, I want a tech refresh my storage right. I have two petabyte requirement and they put an RFP out differently to different DNA, a cycle. We're seeing cycle of six to nine months. Sometimes we're seeing given faster ones because if they know what they want to do.

But definitely the conversation starts from, let's look at the ROI, let's test this and do a PC, -- a proof of concept POC and then move on into integrating this with the stack. Usually, as Russell said, it is not only storage conversation. So it is a conversation where we are sitting in the room usually we have Nvidia, usually, we have our partners like Domino Data Labs and others because, again, this is not -- it's the cloud probably nine years ago, right? They want to understand what's the value of them doing this first of all. So that's what it starts from an ROI perspective. And then it comes on.

Okay. What's the stack is the ML Ops provider, the GPU providers, the storage provider, and then the people who are doing services on top of that. And that's the reason we and have we put a specialist team, that are highly specialized AI sales personnel, including technical folks as well, who are really going to engage in these conversations in a more deep dive conversations rather than just talking about storage because as you said, it is a conversation that is going from an ROI -- as I mentioned, from ROI to POC to kind of purchase, if you would like in the end.

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

It is an evolution here for sure, though. I mean, I think it is definitely been a highly specialized sale. It's generally being driven by the line of business and specific AI practitioners inside our customers. What we are seeing is that, of course, as it becomes more mainstream, more of the components inside a company necessary for turning a POC into production, getting involved in that sales process.

So as well as these new buying centers that Hoseb talked about us focusing on even our traditional buying centers are much more involved in those purchasing decisions. Again, mostly because these other folks that have been involved in making a real shift, Germany quite good at getting it to a POC. But then when it comes to production raising and scaling it, that has been a massive challenge, so that the more traditional folks that are involved in making those systems production ready are much more involved now in much earlier on in the sales process because they know that eventually it's going to come.

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

And IT definitely getting involved more and more. I mean, six years ago, one of our first sales was pure to cardiologist and one of the hospitals in UK. And we didn't even talk to IT folks. It was the budget was staying in the cardiology department. They were buying everything and they were just asking a question how to manage this data.

And now we're seeing IT being engaged more and more, especially that with the Gen-AI, what we are seeing is data being copied all over the place data being replicated all over the place, the security of the government is also going to come into the IT in the end and people will ask why is this data sitting somewhere that is not certified why is this?

So that's why actually starting to gain control over this project more and more because of all the now the challenges that comes with the data. In the end listen, I think you can have the best models out there. But if you don't have your data grounded to those models, you're not going to get anything out of those models, unless you're doing just basic stuff, which we all do know. There is a right new job description right now, like if you want to do those type of AI. You can it's available out there. But if you want to now put this into a real ROI, that's where now I keep saying, okay, you need access to the data. This data is sitting here this data as privacy, Intuit. I cannot move it to the cloud, et cetera, et cetera.

---

**Louis Miscioscia** - *Daiwa Capital Markets - Analyst*

Lou Miscioscia with Daiwa Capital Markets. Creating IT applications is very difficult and takes a lot of time and you guys are sort of on the front row to that. Just curious as to where you think things are right now or maybe the quantity of proof of concepts and when would they actually possibly be deployed into real applications we're talking about months quarters years, we're just trying to understand that even though AI is transformative which is still three or five years out or something a lot sooner than?

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

Well, I'd start by saying that we've been doing this for a number years, and we have a lot of customers already very much in production. So that there's nothing stopping people going from to production. I think one of the things I was talking about earlier though, is just how broad it's going to get in terms of customers that don't -- haven't necessarily invested in for the sophistication necessary to go build it themselves versus those that are more like just to move to what I'll call it value phase, just going straight past development strain to value.

So I think as an industry have to on an industry basis, we're definitely seeing a move towards more turnkey solutions that will shorten those sales periods. And those moves are immuno POC directly into production significantly, but that's an industry wide thing. And so that someone mentioned ahead of sort of more industry-specific SLMs, for example, would be a good example of what would accelerate that.

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

You're right. And I think one of the things that needs to happen to accelerate this is those who already have things in production, they need to start sharing it somehow and what the value is they have achieved by doing this because a lot of customers right now is like, hey, who else has done it, how they have achieved that. They got anything out of it. And we know a lot of customers have done it, but they don't want to share it, right, so I think that's the biggest challenge as well.

In terms of I mean, we heard AWS, CEO a couple of weeks ago, he mentioned, and I think it was on the earnings side, a I kind of save a lot of hours of work that it was used to be in the past. I think those will accelerate the adoption, but it depends who you're talking to if you're talking to the small to medium-size businesses is still very much early on if you're talking to know Fortune 50, they are already advanced, I would say. But then you take that middle pieces. They are in the POC stage at right now, if you'd like to categorize it that way.

---

**Unidentified Participant**

Thank you. Maybe building on the last question, I guess, where do you see kind of the biggest bottleneck right now? Is it okay. Getting my arms around data governance security, is that what data, do I even have that I could train it? Or is that kind of the product -- like I just don't have those capabilities in house and I need a product that I can just buy that can help me with -- where is (inaudible)

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

Interesting thing here, it depends who you ask. So typically, when we talk to customers that haven't gone on this journey yet. So naturally, we published a study with IDC a few months ago that talks about this particularly. When we go to customers that haven't organizations that haven't really gone down this path yet.

They typically don't see it as a data problem. They think, oh, you know, this is an infrastructure problem it's a this that, but all the ones that have already started down this path to realize very quickly is the data problem, data is the number one issue that's actually holding people back.

And so that's, of course, where we can step in and really help. But it's been -- that journey has been super interesting to see, customers will go ahead and think how well we can just move forward and tried this out and then they get they hit this data, right? And the issues that they have are Yes.

Where is the data? What type of data that we have do? I have enough to have the right sort of data to go do this in the first place that sort of assessment often trials, the decision to move forward because they haven't even thought about it and they often don't have a good handle on their data estate to start with to be absolutely frank with you, but the governance piece should not be overlooked.

The reality is, is that AI to a certain extent, we've seen it as a little bit like the Wild West money has started at the Board has been given to have a line of business or a bunch of our practitioners to go and just do stuff. And in fact, if anything, to go to break some plates and cause a mess on purpose because the feeling is that the traditional structures inside customers are set up to slow down innovation.

So the idea of, hey, can I get ahead of my competition? This is going to, as you say, this is transformative. So that's going to do that. Of course, the reality is, is that the regular, -- the regulations are starting to come into place. The AI act obviously came into force beginning of last month and it stipulates a bunch of things that customers need to be thinking about legally, and it puts them on the line for some significant penalties, not unlike whilst in the GDPR. Of course, they're not enforcing it yet.

And that enforcement period will start to ramp up over the next six to 12 months. But this is really bring into focus the need to understand and control data and manage data and ensure the right types of data using the right way. And so that Wild West mentality that was pervading this whole industry that's going to move away very, very quickly. And listen, just like we saw with GDPR, we expect other regulatory environments to pick up similar rules as well. So that's just going to be the start of it.

---

**Steven Fox** - *Steven Fox - Analyst*

I was wondering if you could maybe just contextualize for us at the high level, are you seeing signs of incremental either data movement from, let's say, tape or something like that, where data has just been sitting over there now or train using this data? Are you seeing any movement like that? Is there any reason to think that the rate and pace of data growth is actually changing with AI and also at your customers?

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Yeah. As George mentioned on the last earnings call, we are definitely seeing data lake modernization projects happening more than often than it was before. And we are subjecting that to the fact that these people are trying to build or make their data AI ready basically to start.

Now moving to the next step of using these tools. Now how much that's going to change or I can't quantify that for you, but I think we are seeing that D-Link modernization. We're also what's happening onesies, also, people who have built their data lakes on old workloads.

Now that Gen-AI is requesting more access to that data because it used to be called right now, they need to bring that back to life and reaccelerate. Some of the technologies that were available in the past or today are not really either creating more cost or it's not operating at the speed they need.

So we believe that's why the data lake modernization projects are bubbling up more now. I don't know if they're moving from tape or not. Historically, these data lakes have been sitting on servers with bunch of drives in them. And then I think that was paid to the servers.

And now we're seeing that becoming more unified because also a lot of these technologies didn't have much of a cloud connectivity, especially your heavy on-prem user. So that cloud piece is now coming back into like, hey, I need to use the tools in the cloud, but the data is sitting in a call. So we are seeing data like modernization for sure.

---

**Unidentified Participant**

When you get the lead AI person, and the lead stores buyer in the room together, and they're evaluating the AFF A series. And I guess more recently, the C series. Are you -- can you confidently say at this point that because when you look at those two offerings where you offer, it's built in tiered storage, you got data replication on-prem to the cloud.

Security governance capabilities, privacy, right that the last few questions were asking about get this point. Can you say that this is the -- this checks all the boxes from a your data.

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

A or the C or both.

---

**Unidentified Participant**

If you could talk about, but both have in how do you feel like those compare it like if the answer is yes, to date that we can confidently say that. Could you prior, -- could you say that prior to these two?

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

I would start by saying the AI practitioners wouldn't couldn't care less about AOC series, right? I mean, I mean sort of storage guys, certainly interesting to talk about it, but the way that we would actually engage with the IT practitioner has much to much more to do with understanding the AI data life cycle and how our portfolio of products are able to actually track that data throughout that lifecycle, right, make their lives easier.

So that means all about things like productivity, but also creating an environment where the guardrails are automatically in place from a data governance perspective that and it gives them the freedom to go do what they want to do without putting their companies at risk from a regulatory perspective.

So that would be more of the conversation we would have with sort of AI practitioner. And up to this point, I would say that the infrastructure people have kind of been on the back end of that conversation. So they've kind of has been thrown at them and then they've had to pick it up. And then, of course, that realization that picking up stuff that was never built for production becomes very, very challenging when you tried to scale and move and move to the value phase of an AI project.

So you are starting to see more of these folks come together into those conversations. I'm sure Hoseb will talk about in a second, but the way that we talk to these different percentages is different. I mean, they again, AI practitioner, it does not care about storage, but they absolutely care about the value we deliver.

And we have this very broad ecosystem of both commercial partners, open source partners and cloud partners that enable us to take our value and expose it up in a way that's meaningful to those personas. And that's how we really engage with them.

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Back to the portfolio. I think the good news here is all of them run the same software, right? And it all comes down on what customer wants from a workload perspective. So if it is a data lake, it's different than if they're training a model, right? So then that comes down to a question of performance requirements, cooling, energy power and all that.

But again, I think one thing that differentiates us in this market and puts us up in the front is that we are our portfolio is well positioned to capture all these different requirements, while at the same time keeping the same operating system running, whether it's in the cloud or on-premises and so I hope that answered the question.

---

**Unidentified Participant**

[I just wanted to go at Wedbush Securities]. Just wanted to -- expand a bit on your initial comments about converting a prospective customer into an existing customer, and I'll go through the proof of concept and why it might take about six to nine months or even shorter in some cases.

Can you kind of walk us through what an upsell might look like once a new customer as it has been on the platform has been using solution. How does that look like and do they come to you and say, well, we really like the value. Do you have more offer? Can we explore other products or do you kind of go to them? Can you kind of walk us through what that might look like?

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Yeah. So the reality is a lot of these customers, their heavy budgets are going to the GPUs and then when it comes to store is like, hey, let's start with something that works the hit because we all know that we're the expensive pieces.

The upsell over there is like now that they start training these models, obviously, the more data, the good data model gets better, right? And that is the upsell where we'll say, okay. Now we want to train a bigger model. We want to engage more parameters into this model. We need to bring a bigger dataset that is storage right, and this is where we see it expanding, the other thing that we're back to the data lake pieces again, now that we have this system up and running.

My data sits on either different storage, different platform, different architectures, how can we make it easier to feed? If I don't want to expand that compute cluster, can we bring this now into the umbrella ONTAP, so that the upsell of now you can even modernize their data lakes to beyond that up, right?

So and sometimes we even start as largest, they needed to not that every project start small because there's no, there's no size difference here in terms of capacity. But even if we start small, let's say there is an opportunity where the model is growing, they need more data to put in there or they need to modernize the data lake so that they engage the data gateway feeding to distributors, if you would like. And that could be an opportunity on the flash and object and different parts of the business.

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

There's another aspect I'll just add to what Hoseb said, which is when we win the centers of excellence. So when you win those deals as more workloads come on. We just we get the capacity expansion. That's the reality, right? Customers are attracted to us because of our singular control plane, but we have both on-prem and cloud. We believe that AI is probably the most hybrid workloads we the industry has ever seen.

We expect customers to continue to consume resources in various places. GPU accessibility is one good example of that. But also the reality that data gravity exist in different places. And the data sources are not always not completely in one particular area.

So that's what attracts them to us. And so we're in a really good position to attract more of those workloads wherever they may end up being. So if a customer wants to, for example, have something on-prem, but then consume first party, a hyperscaler pass services around AI we're still the right partner to do that.

So when we establish ourselves as the standard, we tend to get all the workloads as opposed to some of our competition where they get one area and like on-prem and they're having to fight again to go win somewhere else.

---

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

And I'll give you an example, and I can quantify how many of this will happen as well, where we were not present at a certain customers, large healthcare where they purchased the DGX's with NetApp because of what we have showcased for them with capabilities in AI space now because they like all the goodness of what they have seen from a content perspective and back to Russell point. Now we are kind of replacing some other competitors who have been running old, SAP, Oracle and other workloads.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

We have time for one final question. And I think now everyone wants to raise their hand, last question Tim to you.

---

**Unidentified Participant**

Just on the customer front. A lot of talk about the hyperscalers. You get the partners there. Obviously, there's a lot of attention in the AI industry to these cloud AI companies, Korea, Lambda, et cetera. Can you talk a little bit about just from a high level, what is going forward at those companies do better? Is it better for NetApp? Is it worse for NetApp? If they go away, is it better? Is it worse?

Maybe just talk a little bit about how the interaction is different? Do they need more technical capability, so maybe it's better anything you can parse out on just kind of that different customer base because they have gotten pretty meaningful, at least on the GPU side.

(multiple speakers)

---

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

I would first start off by saying that, of course, these are service providers and the way that we tackled those providers is very different. So if you talk to customers, it becomes about joint service creation. So how do we help the service providers create differentiated value through their services?

And that's where a lot of our value add comes into play, right? So if we are -- if it's commoditized storage services that the partner is offering, then it becomes more difficult for a company like NetApp to differentiate our AI value services.

So that still but we're really good at that, right. So I mean, this is not this is not a concern to us and what we tend to see, interestingly, with some of these big hosted a providers is that the customer base they're going after tends to be larger enterprise.

So large enterprises are the ones that actually really appreciate the value of the data manageability features that we bring to the table. So they're not just looking for scratch space, they're looking for rich data environments that protect their data, classify their data, et cetera, et cetera.

So what I would say is I think we are very well positioned to go after those areas as they continue to mature. I think in their current state, they're very much just price performance in overall horsepower GPU. But as they get consumed more by enterprises, they're going to want those enterprise features, I think we're in an extremely good position to take advantage of that.

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Yeah. And we're already in talks in a lot of them. We already have customers in APAC as well. We're doing a GPU service provider type as well as in North America, some of them. Now we are talking more because they are starting realizing that if they want to offer an enterprise level services, it needs to have the multi-tenancy and the security features and all the sovereignty that you talked about.

And this is where our values kick in, right us being in the three hyperscalers. It is -- it didn't come without network. And I think that our partnership with the hyperscalers, first of all, puts us upfront and then this new GPU cloud providers, we are in talks with them so that we are also building similar features from multi-tenancy and all the security features that they need, don't also forget that us being in the industry for 30 plus years means there's a ton of customers out there who store the data on NetApp.

And if this service providers would love their GPUs to be consumed. That data needs to come from somewhere. And if the same way we did with the hyperscalers where we provided that hybrid data fabric, we have a great value to add over there because no one is going to bring that data easily today.

I think that is one of the biggest you asked about the roadblocks right everyone is holding very tight to their data that the goodness of AI is not showing up. And this is where we want to on tighter the gravity. We call it with strip of data from gravity.

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, Billy tells me we have time for one more question. So Billy in charge of everything. All right, so [Andre], I'll give it to you.

**Unidentified Participant**

Thanks, Kris. Yeah, I was just curious about the journey over the last six years, you said AI was deep learning and machine learning, which initially got really people very excited. But then the deployment challenges, everything kind of just didn't pan out maybe quite the way people wanted it to.

And I'm just wondering when you think about what the selling motion at that time was who are you selling to who when the business was driving that versus now in Gen-AI who in the business is driving that? And how is the sales motion today different ?

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

Good question. As I mentioned, six years ago, we were selling to cardiologists say it was majority of it is line of business owner the IT didn't have budget for AI at that time, as we move forward, IT started getting more control of that because one of the customers just woke up and found a whole bunch of IT, computers and storage sitting somewhere in the room.

I think that I started getting with the Gen-AI, especially cloud playing a big role. Ones IT is getting more engaged. Now the line of business owner are the ones who are pushing down the agenda and the IT is trying to figure out how to implement it, whereas in the past, it was line of business line of business and IT just watching and seeing what's happening over the time.

That's why actually, Nvidia appreciates our partnership is because of our huge experience in the data center. One of the things Nvidia wants to open up to the enterprise and with our expertise in data center and IT, I think we that is where the partnership becomes very important because of our expertise in that.

And now the line of business, yes, they have the idea. They have the budget, the board is pushing down, but it is still lending now and IT to deliver the infrastructure. So that's the evolution happening. And cloud, actually, the Gen-AI and the cloud made that part more real for IT.

**Russell Fishman** - *NetApp Inc - Solutions Senior Director, Product Management*

I'll just last thing I'll add is that what we are going to see more customers that just want to consume AI as a value rather than from a development perspective. And again, as much to do with the fact they don't have the data necessary to do things from scratch anyway or the sophistication.

So for those more turnkey solutions, that will be things that are kind of regularly things that are IT, you're going to be the ones responsible for delivering like copilot agents, enterprise knowledge management, et cetera, et cetera.

And then the it will be the lines of business that come in and want the highly verticalized solutions. But again, they're going to be more commercial off-the-shelf. I mean, I think that's the mass market here.

**Hoseb Dermanilian** - *NetApp Inc - Director, Global Head of AI Sales and GTM*

I mean, give you an example where we didn't have we use our own chat GTP version right? I'm a line of business within NetApp. If you think about me, I didn't develop that tool myself. It came through our IT and the data science teams and said, hey, here's a tool you guys can go and use.

So if you take our example, it is a great way of -- now if you're just yourself and you want to use AI to write a PDF paragraph for you. You don't need your IT. So it depends on the use case and what we're talking about here. That's why I said small to medium businesses are still like not that margin. Then you come into the wider enterprise in the top 50. That's where we see the curve growing.

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, thank you, guys. I really appreciate it. We're going to take a 18 minute break now and then we'll be back at 11:10.

All right, welcome back, everyone. I'm glad you made it back from the break. So now we're going to switch gears and talk a little bit more about core storage and all things ONTAP. So I'm happy to introduce Sanjit, who you guys saw last year if you were here. And so before we get started, why don't you introduce yourself, say a little bit about what you do at NetApp, and then we're going to open it up for questions.

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

Hi, everybody. I'm Sandeep Singh. I'm the Senior Vice President and General Manager for Enterprise Storage at NetApp. I look after our portfolio of on-premises products. I have had the pleasure of being here at NetApp for almost two years now. And over the course of this time period, I've probably spent at this point a little bit less than 50% of my time traveling, meeting with customers, meeting with partners, and also remaining focused and continue to rapidly expand our portfolio to have this ability to have a unmatched simplicity at scale as well as transformational flexibility across our unified data storage portfolio.

With that, I will open it up for questions.

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, I know we have a question from the webcast that came in earlier.

**Unidentified Participant**

Yeah, it was earlier on block. So I'll read it out. I think I can give a little more color to it. How has NetApp ShareGain views evolved for block storage as we've released new block products? And then can you talk about the install-based leverage? So the fact that we have an existing install-based area, has that improved our block traction? Or can you talk about success in block only environments that result in net new customers in NetApp?

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

Yeah. So first of all, we used to already offer block storage capabilities with our unified storage portfolio. Last year, we expanded our product offering with an all flash SAN array or ASA series of products. When we initially looked at overall block storage, we roughly have about 20,000 customers who rely on NetApp and trust us in some shape or form with their block storage workloads.

We introduced a series, we announced it with ASA A-series last year, and then ASA C-series. We've continued to see great adoption across the board from the use cases standpoint across our customers and across net new customers. The way to think about the adoption of ASA series for the SAN workloads is kind of threefold.

First, in terms of our overall install-based accounts, where we've got thousands of customers that trust NetApp, they love their on-tap experience. With ASA, we're helping them bring that on-tap experience to their block storage workloads. That represents then an expansion in terms of adoption of the block storage workload there.

And the importance and the value for customer becomes what I call the simplicity at scale. What does that mean? If you think about customers, they are dealing with complexity that gets compounded through bespoke infrastructure silos. When you think about the infrastructure silos across their overall NAS and file environments, and when you start to then expand that into silos of VMware or database application workloads, that is where -- when it's a bespoke infrastructure silo, they're having to deal with separate inconsistent management, inconsistent automation, inconsistent data security models, inconsistent operational recovery workflows, inconsistent overall experience.

We're enabling them to be able to bring that on-tap experience to their block workloads and be able to then get consistent management and automation that overcomes the need for talent and skills gap shortages. It gives them one consistent data security model they can trust in. It gives them the comprehensive yet consistent operational recovery workflows, and it gives them an overall consistent experience. So that's for our install-based customers. That's an incredible value for them.

Secondly, for a lot of the customers that are continuing to work through and figure out what's the right feature in terms of their overall VMware deployment, we're helping customers to be able to come over, offload data management, and through that get upfront savings in their VMware environment, all the way up to 25% savings there, and then be able to have this unmatched flexibility for the future in terms of the hypervisor or the container environment and or the cloud storage options there. So that's the second key area where we see customers adopting NetApp in their block storage environment.

The third key area is essentially as customers are looking at, they've got all the hybrid disk-based block storage systems. There's lots and lots of them, and being able to modernize them to all flash and being able to do that affordably. This is where finance department CFOs are still requiring and asking IT to continue to lower the IT budget.

We're enabling them to go modernize to all flash and be able to do that affordably. As part of this leverage and the flexibility that we enable customers with, they'll start talking to us, for example, in file environments and then recognize, well, there's this flexibility that exists of bringing a consistent experience to the block storage environments that gives us this optionality to expand in to these options.

I was just learning about a recent win that was exactly that use case for customers in other environments. Once they recognize, well, I've got this standalone block environment, I can get simplicity at scale, that's helping us go out and win. We now give them just an overall flexibility of having an end-to-end ASA portfolio that they can leverage us with in their block storage environment.

**Unidentified Participant**

Maybe on that front, when you talk to customers today in terms of refreshing their storage infrastructure, do you still feel like you need to convince some of these customers to move from disk drives to flash, or is that a decision they've already taken? Really, what you're doing in terms of your conversation is really going within an NetApp portfolio and convincing them prior to competition where you stand, and that's the conversation.

The second part to that, when you talk about block, how do we get comfort that your customers on the block side are not really just moving over from the unified storage they were using previously, and these are actually incremental opportunities that you're capitalizing on?

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

Yeah, great set of questions. In terms of disk to flash, the larger scale, there is a shift that's happening where customers are optimizing for flash at a broader scale. When you look across our portfolio, we've just got a fantastic set of options available to customers in terms of the all-flash offerings.

With that said, what I will say is that what matters to customers, especially with the unpredictable macro and tight IT budgets and IT budget scrutiny, is this notion of high-quality, lower-cost solutions. When you look across the data lifecycle, having the ability to have lowest cost of data over its lifecycle is important for customers. Where we continue to see great adoption across the board is when you look at primary storage, that's going to be typically all-flash.

But when you look at the data, anywhere from 60% or more of that overall data is going to be cold data. Being able to automatically, based on policy, be able to tier that data to lower cost storage options gives us the advantage of enabling an end-to-end portfolio that is underpinned through ONTAP, of being able to seamlessly enable customers to go and adopt those for their primary storage workloads -- yet also ensure that we're cost-optimizing the cost of data with hybrid flash options as well as the overall object with storage-grade options available to them.

---

**Unidentified Participant**

How do we get comfort that the block customers are not unified customers moving over?

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

In terms of the block customers, you will see the combination of the following. There are many customers who have their file and block environments in the same environment. They may have a large file environment and a smaller block environment over there, or it could be vice versa. This is where unified storage provides them with the best option of essentially being able to have file or block on that same shared underlying infrastructure. We provide that that's available to customers, and lots of customers are leveraging that capability.

On the other hand, if customers have a separate file environment and a separate block environment, this is where AFF for that file environment and then ASA for that block environment is the perfect fit for them. This is where it becomes a net new expansion opportunity or a net logo opportunity for us. Unified doesn't necessarily substitute for the ASA use case, where it's that standalone overall block environment that they're going for.

---

**Unidentified Participant**

Understanding that you've been competing against folks who were more upstarts on for your C-series product, you were competing against storage startups that didn't have the full portfolio, but a number of your competitors have announced that they intend to offer QLC products. How do you envision that that changes the landscape? Do you envision that that might change evaluation times, or do you feel like the argument has been made or the inroads have been made?

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

In terms of looking at capacity flash underpin through QLC technology, first of all, we have just continued to see fantastic adoption of C-series across the board. Yes, some competitors have gone in and announced products. We have not seen that make a difference.

Why? First of all, when you think about the customer's needs, when they're putting, let's say, their tier two capacity-focused workloads onto C-series, data management is still critical. We are bringing this notion of comprehensive data management and still making it available even in the capacity flash series.

Secondly, and similarly in terms of our data management, when you think about ransomware and cybersecurity protection and detection and recovery, that is critical for customers across all of their data. This is where we're also really changing the game in terms of enabling what's typically a post-process detection to that real-time detection in a matter of seconds to minutes.

Being able to do that, which is designed for that 99% plus accuracy that minimizes the false positives, as well as has the ability to be able to provide the accuracy of detection. Earlier this year, we became the first storage vendor with SE Labs validation to get that AAA rating with this ARP AI technology to get that 99% plus overall rating in that detection.

That becomes critically important for customers as an overall priority, and then couple that with the ability to rapidly recover from ransomware attacks. That's another area. Then I talked about the complexity challenge and how can customers just simplify at scale. Rather than going and putting and continue to propagate bespoke infrastructure silos, customers are increasingly looking at, well -- when I start to think end-to-end, not just a point solution, what matters most and continue to simplify and get this notion of simplicity at scale that only NetApp is able to deliver across a fully interoperable portfolio that is then underpinned through the lens of ONTAP.

Those become some of the prime motivations as customers are thinking through capacity flash options. The last part I'll touch on that is in terms of the economics, even being able to then cost-effectively tier the data over the data lifecycle and still get the lowest cost of data. This is another one where in the competitive products you end up getting bespoke silos versus a fully interoperable portfolio.

---

**Unidentified Participant**

Hi, thank you. Wondering if you could discuss a little bit, this has kind of been this tug of war over the years of movement to cloud and repatriation of some workloads, data sovereignty and whatnot. If you could just give your perspective on kind of where we are now, because we hear a lot of examples of both, and then related to that as we move further along in AI, do you think the calculus around that tug of war changes at all? Thank you.

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

So in that tug of war, I think across the spectrum of thousands and thousands of customers, we will see all of the above. We continue to see customers who are adopting the public cloud storage options. We've seen a lot of customers who are adopting the hybrid cloud storage options. We're also hearing about customers who, as they have scaled workloads and increasingly with the overall cost pressures, we're hearing also about some customers who are repatriating their workloads.

What is important and what we're focused on is ensuring that our customers have that complete flexibility to be able to have market leading offerings for the on-premises options available to them, whether in a CapEx form factor or storage as a service side of it, have the leading public cloud storage offerings, the first party native cloud offerings that are underpinned through ONTAP, where NetApp is the only one with the first party native cloud offerings across each of the hyperscalers, have the necessary technologies that are available to them for the hybrid cloud use cases, which then is inclusive of secure and efficient data mobility that becomes critical for customers across the board.

So we are ensuring that we have the necessary choice points as and when, whether they are going to go and be able to leverage the agility of public cloud for being able to seamlessly move their workloads and be able to still get the enterprise resiliency, the data management capabilities that they value, be available to them, or if they need the hybrid use cases, we're enabling those, as well as if they're looking at repatriating, we're providing the right options available to them.

We feel incredibly privileged to be in a position of providing this unmatched flexibility for customers and being that partner, no matter what their use case there. You asked a follow-on question, which was tied to AI, of how does that change the calculus there. When we think about AI, hopefully this is building on the prior session there -- when we think about AI, we're going to see the customers leverage, first of all, a lot of their enterprise data that sits on-premises in the enterprise AI use cases.

Ultimately, customers are looking at how can they truly unleash the power of AI and Gen AI with the context of their enterprise data. We're going to see a lot of that happening. We are enabling customers to have their AI or their data to be AI ready, as well as bringing AI to their data, to what I call this notion of an AI data gap. That's the first area.

Secondly, you will see, obviously, there's lots and lots of investment in terms of AI, Gen AI tool sets that are available in the public cloud. We want to ensure that we are providing the flexibility of enabling a seamless experience for customers to be able to leverage their data and making it available to their AI, Gen AI tools in the public cloud.

That's building on some of the announcements that we have made with the, for example, the AWS workload or BlueXP workload factory and various integrations that we've demonstrated in terms of the AI, Gen AI tool sets in the cloud.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, Sandeep, since the audience is reticent, one question I get from investors all the time is, do disk drives still have a place in a modern data center? Maybe you could talk a little bit about that.

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

Great question. Look, in terms of overall media underlying it, it becomes important when customers are thinking about what is that cost of storage? And as you think through the life cycle of data in terms of the disk drive, that hybrid flash storage, what you're going to fundamentally find is it is still that lowest cost of storage that you can get from an on-prem standpoint.

So the short answer is yes, disk drives still have a specific use case that they serve in the customer's data centers. When you look at basically from a raw overall cost standpoint, flash is still not at a point where it can substitute for that disk-based storage perspective. And so what becomes important is when you start to think about, the optimizing cost data through tiering backup target use cases, when you think increasingly about cyber vault use cases as well, ultimately you need the economics to be there. And that's where the disk-based options are still important.

---

**Unidentified Participant**

Can you talk a little bit maybe about where you think we are in the cycle of customers sweating their assets, on-prem assets? Feels like, storage spending has been relatively muted now for a fairly long period of time. Are you seeing any green shoots about things kind of starting to pick up in that regard? Or how much longer, given that you have a pretty good view into where your customer's utilization rates are and the capabilities are, how much longer could that extend out to?

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

I would say, look, overall, when we look at just the overall adoption of systems, whether it's on the AFF A-series side or C-series side, we're continuing to see just an overall strong adoption. And utilization rates, when you think about sweating out the assets, one would assume the utilization rates would either increase significantly or the customers are sweating out the assets.

I would say, generically, we're seeing that across the board. That's pointing to customers have a need. Their data fundamentally, continues to grow and they have a need to continue to go and, purchase more underlying storage capacity and systems tied to it.

---

**Unidentified Participant**

All right, Sandeep. You mentioned as-a-service and consumption models. Can you talk a little bit about some of your competitors, harp on this a lot, but it doesn't seem like at scale it's really happening. So what's your sense of demand or appetite for the different consumption models as a service versus CapEx across, more broadly? I'm sure there's some of each in your broader customer base, but if you just give us a little sense on how we should think about that potential transition over time.

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

Yeah, absolutely. Look, storage as a service, various customers are transitioning or looking at storage as a service across their spectrum, especially when, for example, refresh or net new purchases come up. The Keystone is our storage as a service offering. We're continuing to see just an overall fantastic adoption across Keystone.

I know some of the financial, information, we released that as part of the earnings announcement there, but we're focused on the enterprise segment. In the enterprise segment, we our position is we want to provide those leading options to customers, whether they are looking at CapEx, whether they are looking at storage as a service on premises, or whether they are looking at essentially cloud storage as a service options for them. We're unique in that way of being able to provide that complete flexibility with the leading options to customers.

We want to ensure that we're not pushing customers one way or the other. We're providing them the flexibility to adopt as they are ready for the right options for them, right? With that said, we are seeing overall great momentum with Keystone, and it's important to not only have that momentum, to also provide that overall service level that they're continuing to then go and build out the adoption with Keystone.

So, we're incredibly happy of what we're continuing to see there. The other thing I'll say is the customers who are looking at essentially the they want to get that agility of cloud, and they're on a journey, we also see a cohort of customers who then are leveraging and transitioning to storage as a service on-premises as well.

And this is another one of those unique capabilities that we're bringing and making available to our customers.

---

**Unidentified Participant**

Thanks, We just had the AI session, which was very good, but if we could go back to that a second. If you look broadly at your customer base, could you just break it down how many are in proof-of-concepts right now, how many are maybe not doing anything, and how many have actually gone in size from proof-of-concepts to deploying? And if you want to throw time frames into that over here, looking forward, that'd be great too.

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

That would have been a great question actually, for the last session,

---

**Unidentified Participant**

We did, but you figure you're here. So

---

**Sandeep Singh** - *NetApp Inc - Senior Vice President and General Manager - Enterprise Storage*

The way I'll talk about it more generally, right? The is that we see an incredible opportunity forward looking in terms of enterprise AI, right. The enterprises across the board have been looking at how did they go and take AI, Gen AI and use that for driving up productivity for delivering that new customer experiences and or enabling net new areas of innovation.

Many have been looking at holistically and in concepts and as well as the specific use cases that are going to be incredibly important within their context. What becomes important is essentially this notion of an AI data gap that I touched upon, where customers are challenged with how do I get my, I've got multiple data science teams, and how do I ensure my data is AI ready?

How do I ensure the right data sets are available to the right data science teams? How do I ensure that the right levels of privacy and security are assured there? And then how do I ensure that overall efficiency is available and overall model versioning and the data sets associated with those are available? Those become incredibly important for customers. This is where this notion of helping customers bridge that gap in being able to discover their data and prep their data and make that data AI ready becomes important.

Secondly, in terms of essentially being able to bring AI to their data also becomes important in bringing overall efficiency to that end-to-end AI life cycle. So when you think about the spectrum of how are we helping customers, when you think about the, one set of customers has been the overall AI as a service or GPUs as a service and a lot more of that use case is the overall foundational large language model training.

This is where if you think about in terms of our portfolio, we're discovering where a lot of those set of customers, they want the overall scalable levels of very high levels of performance. This is where we have a super pod with -- the NVIDIA super pod with the BGFS plus our E-series solutions that customers are adopting. That's one use case.

When you think about it in terms of the overall enterprise AI use cases, if they're not a lot of enterprise use, AI is going to focus on as much of the large language model training, yet there's still customers that are going to go do large language model training or small language model training. So we've got whether the overall the super pod that's there or the overall base pod solutions that are in place for them to be able to leverage it.

And then when you, we also see when customers are on that beginning of that journey in that data prep phase, they're modernizing their data lakes. That usually takes them to object and that's where our storage guard offering is seeing great adoption.

And then when you start to shift to the retrieval augmenter generation rag use case or inferencing or fine tuning with their enterprise data, this is where our overall data management, a lot of the data management challenges that I was touching upon becomes critical.

We've also put a overall converged solution that we call the AI pod with Lenovo, with the NVIDIA L40s GPUs with OVX that we announced back in May. That also provides them this converged infrastructure stack for them to be able to use for their rag or inferencing type of use cases. And then our flex pod solution is also being used in those instances.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Actually, we're up on time. So thank you, Sandeep. I really appreciate you coming here. -- Sorry, I just know he has somewhere else to go and I don't want to be the inconsiderate one to make him late for his next meeting. So I really appreciate it. Thank you.

But I think whatever questions you didn't get to for Sandeep and still have outstanding, Jeff can probably handle. So I'd like to introduce Jeff Baxter. He's our VP of product marketing. He can kind of handle a broad swath of all your questions. I know every session ended a little bit early for your taste, so you'll be able to pose all the questions to him.

But before we open it up, Jeff, why don't you introduce yourself, say a little bit about what you do, and your time at NetApp, which I think helps provide the context for why you can handle such a broad set of questions.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

So thanks, everyone. So my name is Jeff Baxter. I run product marketing here at NetApp. Before that, I ran product management for ONTAP, which is the operating system that powers a lot of our on-prem systems as well as our in-cloud systems. I've been with NetApp now for 16 years, so gone through several transitions at NetApp. When I started NetApp, I was actually in sales. I was an SE. So I was out selling our actual individual storage systems, and we were a storage company. And the evolution since then, in terms of going out into the cloud, and now the evolution that we're embarking on with AI, the company has transformed two or three times over, as I think all of you have seen. And it's been just a remarkable ride.

So I've been privileged to be a part of it. I've been in all different parts of our business, gone from sales to product management. Now running our product marketing as we continue to reinvent ourselves as the Intelligent Data Infrastructure Company. So that's a little bit about me.

---

**Unidentified Participant**

I don't know if Chris is going to like this question. It's about pricing. Obviously, you guys don't price per gigabyte or terabyte or anything like that, but there's now a lot of storage as a service. There's a lot of new applications. There's more value being placed ROI for AI. So how, as we enter this new phase with your customers, how does pricing change for the better? And I'm sure there's some reasons it changes for the worse. How's your pricing power look over the next few years? Can you go on that one?

---

**Kris Newton** - *NetApp Inc - Investor Relation*

So I think you can talk about big trends, right? What you see in the industry ONTAP would be a great thing to talk about.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. So I think in general, we're trying to be cost competitive in the industry, right? We definitely are an option where we add a tremendous amount of value-added services directly into the operating system. So ONTAP one is our ability to add all of the data protection you would need directly into our storage systems.

More recently, adding anti-ransomware capabilities directly into our enterprise system. So I think we see customers terrifically valuing what we're seeing. I don't know if any of the secular trends like AI or others, I'm not going to comment on what the pricing will do there, right? I do think we continue to be cost competitive in the market, right? We continue to see wins against competitors.

A lot of those wins are based on value and based on cost. I think for our customers, when they look at our price tag, the important discussion for them is really around TCO, right? So as you said, it's not about dollars per gig. When we have storage as a service customers, they're actually buying a given service level through our Keystone program. When we have customers still buying through traditional CapEx model or leasing model, they're really buying more based on outcomes rather than just raw dollar per gig.

And when you build in best-in-class storage efficiency, build in the data protection, build in the simplicity operating system that I assume Sandeep talked about in terms of making it simple at any scale, all of that tends to really drive down TCO. So one of our best sales tools with customers is we'll sit down and walk with them through a TCO calculator and put in all of their assumptions, all their power, right?

Being far more power efficient in their data centers, for example, can drive cost savings. So when we get to the end of that, we almost always find the NetApp is a more efficient solution over the long run from a TCO perspective.

**Unidentified Participant**

Great. Thanks. Hi, Jeff. I also have a pricing and packaging question that I actually wanted to ask Sandeep, but I think, as Chris said, this might be a better question for you.

So if you think about the storage market in terms of P times Q, I mean, we're all investors or the investment community here, right? If you think about P times Q, it's been pretty fascinating because (technical difficulty) the P, right, the unit price, right, it's on a per gig or per terabyte, well, it doesn't matter, right? The P has come down significantly, but it seems like, and I'm talking actually on the specifically on the all-flash side, right?

But it seems like it's kind of hitting, it's like, there's only so much more to go (technical difficulty) perhaps in the near term. Meanwhile, the Q has increased dramatically, right? And that's opened up, obviously off the smaller base. And if you think along those lines, how does the, so if you were to apply it to, the question for you is like, if you were to apply it to the enterprise side and then also the multi-cloud hyperscaler side, how does this play out? That's kind of one part.

And then the second part is, if you're like, just about there in terms of the P, then are you thinking about programs that you can work with customers? So maybe because their wallets are fixed or they're not growing nearly as fast as the volume growth is. So are there programs you're considering where you can maybe help them bridge the gap maybe for near term so that you can once you get past that hurdle, it unlocks this massive opportunity.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. So that's a, that's a multi-part question. I think you're right that the data growth is outstanding, right? Especially going into the AI era with the massive growth of data there. So I won't comment too much on the future price or where we see that affecting, revenues because Chris will tackle me off the stage. But what I will say is I think for a lot of those customers where they're concerned about, that's really where they're looking at storage as a service offerings, right?

And being able to go out and apply certain service levels which tend to insulate them a little bit from variations in pricing and other things like that, right? So they're buying a service level. And the other thing for a lot of our customers is as they're not sure about what the rate of data growth will be as they're entering sort of an AI space, pre-purchasing CapEx multi-years out, an alternative option can always be going with storage as a service and that lets them basically grow as they need to and basically match their cost curve more carefully to what their actual consumption is. And so I think they're addressing some of the variability through the continued growth of our storage as a service. I think that's the main thing I would answer out of that.

---

**Unidentified Participant**

Over the past couple of years, there's been a lot of changes in go-to-market, maybe different incentives for kind of cloud sales. And I just wonder how that's changed kind of product marketing or has the marketing to the customer changed or has it really been kind of more go-to-market side?

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah, that's a good question. No, I think the marketing has definitely changed, over the last 10 years. We actually went through a shift probably six years ago. I actually remember being at Insight, I want to say 2014, when we first talked about putting ONTAP on the cloud, right? And we first did a demonstration of ONTAP running on the cloud. So we're now literally at a decade.

And it was five years ago when we unveiled Azure NetApp files and we were a first-party native cloud service. I think at the time, everyone saw us as a storage company, right? And so from a marketing perspective, we went incredibly hard at being a cloud-centric company, right? And really

getting that message across and getting that ramp started with cloud. Where we started to balance over the last couple of years is bringing it back to talking about that balance that Sandeep talked about, that basically all of our customers are in some way hybrid multi-cloud, right?

They're either starting their journey well along their journey. And so we've been able to really embrace that. And you've seen that change towards where we talk about ourselves as the intelligent data infrastructure company, right? There's nothing in there about cloud or on-prem storage, right? It's about data infrastructure regardless of where it lives and about applying intelligence to it. So that really is the message we've taken out to market. And I can honestly say, all of our, like even within my team, for example, our solutions team, our launch team, all that other stuff is integrated across cloud and on-prem. It's not separated anymore, right?

So we certainly have subject matter experts, like I'll have a marketing person on Azure NetApp files. I'll have another marketing person on Sandeep's AFF systems or ASA systems. But when it comes to how we bring those solutions to market for our customer in terms of a VMware solution or Kubernetes solution or database solution, all of that is horizontal, right? All of that is what's the right solution for you.

And so when we go talk to customers now, it's not about we've got a great on-prem solution for you and let me bring in a sales specialist to talk to you about a cloud solution. It's we've got a great data infrastructure solution for you. Let's work with you on where your choice of data needs to belong. And we can help guide you along that path and use AIOps to look at your workloads. We actually have cloud advisors that will look at workloads and say, this one's cloud ready. This one you could put in the cloud, but might be more of a lift.

And we can actually guide customers in taking that journey. So I think that's the short version. It's how it's changed the marketing is it's really talking to customers about data infrastructure as opposed to either being a storage company or being a cloud company. We've elevated really the entire message.

---

**Unidentified Participant**

More on-prem question, just in terms of how you think about mix within Flash changing with your customers, particularly when customers sort of used to use more high-end Flash products now moving more towards mid-end products like C-series, et cetera. How do you think about how material that makeshift could be either because you didn't have those products earlier in the portfolio or for a variety of reasons that they were more performance conscious, they were using more high-end products.

How do you think about that makeshift and how material would that be and what would be the counterbalance?

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

So I can't talk directly about how material that shift would be. I think that would be more, unfortunately, a Sandeep question or others. I can say, I think we're achieving a healthy balance in customers, right? So customers are adopting the whole breadth of our portfolio, performance Flash, capacity Flash, even hybrid Flash. And so we're seeing, I think -- I guess I'll just say, I think we're seeing a good balance there and I can't really go too much farther than that as to how we would shift.

---

**Louis Miscioscia** - *Daiwa Capital Markets - Analyst*

Thanks. Louis Miscioscia at Daiwa again. If we think of AI before the NVIDIA moment last May and afterwards, maybe if you could talk about your product marketing budget, however you want to define it, people, whatever. I guess before what it was, which I assume wasn't very high to just where it is now in proportion to whatever else you're doing.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. I think like everyone else, we've seen the tremendous opportunity for AI, both in our customers and in our own business. There's been significant shift and interest and investment and resources in marketing AI, for example, in my own team. I won't get into specific numbers, but it's been a significant burst in resources moving towards that as an investment engine, adding additional specialists.

I think one key thing with AI is we have to be able to market to multiple different segments, including the infrastructure buyer, the data scientist buyers, some of whom don't know each other within the same customer, right? And being able to build messaging out to the market that embraces all of those different personas. So that's definitely required additional investment.

I think without previewing it at all, I think from the keynotes today and tomorrow, you will see just how much focus we've placed on AI and how much in a lot of ways over the next day or two, we'll have announcements that will obviously in some way concern AI. And you'll see a lot of the marketing messages that will come with those and how we unveil those on stage. And that's going to be a key focus, I think, of our selling and our campaigns over the rest of this fiscal year and well beyond.

---

**Unidentified Participant**

Hey, Jeff. In the beginning, when you guys were working with the hyperscalers, there were some fits and starts. It was a different motion. And so can you talk a little about the learnings and are all three now where you want them to be? Maybe talk a little about specific programs and how you've been able to get through the friction and kind of jointly sell these solutions to the end customers.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. So I'll give you what I know. Unfortunately, I'm not the subject matter expert. And I think you had Pravjit up a little bit earlier today, who maybe would be a little bit more inclined on those sort of questions. I think it took a while for us to learn how to sell an emotion where it was not a direct skew on a NetApp price sheet. I think we saw initial tremendous growth from NetApp customers who just want to run ONTAP in the cloud. And that was the very early versions.

But when we moved to this first party native service and co-selling with Microsoft and Amazon and Google more recently, that took a lot of work on the go to market side to figure out how to do that as a co-selling motion, where we could add appropriate effort in. And I think we've resolved most all those issues. I think the team that the go to market team between Pravjit working with Ashish and working with Dallas and Cesar and really the whole go to market team there have really figured out how to make that moving forward. So I won't talk about the future, but I think we're really satisfied with the progress to date of where the cloud business is gone.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, I have a question for you, and I'm going to build a little bit on [Sameek']s question. And I get this a lot from investors is right. We have the A-series, C-series and FAS products. They're all ONTAP, all unified. How do customers think and choose between them? How's the product positioning work across those different families?

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. So, you know, at a high level, the product positioning is all around mission critical, incredibly performance sensitive applications. So, you know, from a technical perspective, it's sub millisecond, typically sub 500 microsecond latency on those A-series high performance systems. That happy middle ground with the C-series, the capacity flash systems where you're getting low single digit milliseconds of latency.

So two to five milliseconds of latency is really where a lot of customers are finding the vast majority of their workloads. And I think that's is one learning, I think, over the last decade of the move to flash is we went from hard drives and we went straight to incredibly performant flash. And in

some ways that was as an industry, I'm almost say gold plated, right? You almost overshot the performance market in some ways because that was the only alternative.

There were no middle grounds along the radiation. But capacity flashes allowed us and certain parts of the industry. I think there's a lot of our competitors who have not caught up at all in capacity flash, but it's allowed us to go in and capture that middle ground where customers were actually being over served on performance, right? They were being over served on latency.

And in doing so, that's one of the ways we've been able to achieve cost savings for them. They were literally over provisioned in terms of their latency. So that's where the C-series stands. And then FAS and hybrid flashes increasingly and not it hasn't 100% moved off of primary workloads, but it's pretty much moving in that direction.

And we tell customers in general with only a few exceptions, you're not going to want to put primary workloads on disk any longer, but they're a fabulous place to replicate to. They're a fabulous place for disaster recovery. I think one place we're putting a lot of focus on is cyber vault solutions. So we announced one of those earlier this year. I think you'll continue to see things at Insight about them.

You'll continue to see us talk about it because we increasingly have customers ask us about how can we get logically air gap solutions that just keep the data cold and completely isolated from attack. And that's a perfect application of disk. And that's something we can do within our architecture because we have our disk running on the exact same architectures we have all of our other products.

The final point I'll make to Kris's question is, well, two points. One is we have a whole bunch of tools some of the customers have access to from an AIOps perspective. They're an existing install-based customer that will guide them towards what sort of system would be right for a refresh.

And then we have a wide variety of tools available for our partners and for our internal NetApp sellers who will specifically say workload by workload, performance by performance. This is the exact right platform to position for this one. And the good news is the capabilities are all the same. So they can still sell the value proposition of intelligent data infrastructure. And the idea of which platform to host it on the back end or which system to host it on the back end is more of a speeds and feeds and architecture discussion as opposed to a front start of the conversation discussion.

And finally, we have customers who have all of them, all of the above. Most of our large enterprise customers are going to have a mix of A-series, C-series, and FAS, sometimes mixing together in the same cluster, even seamlessly and using AIOps, tearing data between all of them. And so the idea is the data can flow to wherever it needs to at a moment's notice.

So data can start incredibly hot on the A-series. It can sit there for user-defined cooling periods, a week, a month, whatever, automatically flow down to capacity flash. You can sit there for another quarter, wait till your past quarterly earnings or whatever reporting period you're in, and then tear that data off onto FAS and keep it there for the long run.

And then if that data ever needs to be accessed again, it can seamlessly flow right back up that river. And so by building that sort of combined infrastructure, that's where customers are really able to optimize their service level directly to their cost.

---

**Unidentified Participant**

Can I follow up to your question there? Because you have an interesting perspective on this. So like, I think there's a perception in the investment community that our all-flash growth is a transition of business from our FAS hybrid or HED-centric business to the C-series. And so the all-flash is just a transition in revenue.

Could you speak to that C-series business, whether or not, you mentioned at the start of your answer that you thought it was a hole in terms of the latency that that product would have kind of addressed. Do you think that is then bringing net new customers to NetApp, or do you think that it kind of smooths out that kind of product kind of offering from low end to high end?

And so that's like really what's going on. And so we're just kind of transitioning people that were in a product that we didn't originally have, I guess.

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. So I think it's more than just a transition. I think that we do have customers from competitors who are sitting on legacy disk arrays, who the C-series opens up new opportunities to transition, right? They weren't able to transition to the A-series at a price point. They didn't need that performance point. And so we've been very successful going out and penetrating sort of the mid-range of some of our competitors on their refresh and being able to move over. So it's not just a refresh of our own internal disk.

The other thing I'll say is, based on my previous answer, it's not like disk is going away entirely, right? So it's not a one-for-one replacement. A lot of customers are, and a gentleman asked earlier about the data growth. Some of that data growth is in primary data. A huge part of that data growth is also in secondary data, right?

Because secondary data now can be used for building data lakes, for building data analysis. You need a third or even a fourth copy of data for some of the regulations and to have it locked in place in a cyber vault. So all of that growth and data can be served by the FAS market. So to the extent that even there's some of the FAS market moving up to capacity flash, which there certainly is, there are plenty of workflows flowing in the back end to keep the FAS market moving as well.

So I wouldn't say without going too far into the future, I don't think it's a one-for-one refresh. I think that there's legs on both of those for many years to come.

---

**Unidentified Participant**

So you guys have talked about the siloed people that you compete with. Is there a way to think about like, and you mentioned like the ability to supply customers so they can go up and down depending on data use. Is there a way to think about like where that evolution is?

Like how many customers out there, whether large or small, are still very much siloed and have to overcome that, especially in AI, versus how many have already made that transition to being more flexible, like your vision fits.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

I mean, there's a joke, just look at our market share, but that's probably not what you're asking. You know, that's a good question. And I don't think we fully know the answer. I do honestly think, all joking aside, none of our competitors have a unified stack, right? That has continued to be the NetApp differentiation. There are a couple of competitors that are, without going into names, or they just focus on one niche, so they just have one offering, and that's fine, right?

I wouldn't call it unified in terms of being able to do structured, unstructured data, object, in the cloud, on-prem. And so almost by definition, any customer we deal with who is not a NetApp customer does have some degree of fragmentation within them. There's NetApp customers who have fragmentation, right, that we're working with.

That's one of the reasons we have our BlueXP unified control plan and are bringing all of it into our common orchestration set. But just about every customer we run into who isn't a NetApp customer has some degree of fragmentation, unless they're a very small or sort of niche-focused customer. It's very rare for me to run into a customer these days that doesn't have multiple storage arrays from multiple competitors lying around in some state or the other, that they just have not been able to unify because each one is addressing a very specific need that they haven't been able to find anyone, until they met NetApp, that could meet all of those needs out of a single operating system, out of a single set of appliances.

---

**Unidentified Participant**

Hi, thank you. Kris may get mad at this one as well. I just want to go back a little bit. Before you had the QLC, you characterized it as if someone was refreshing a hybrid or a disk-based system and they wanted to flash, they would be over-served, more or less. On latency, yeah. But my guess is they wouldn't want to pay for that. So what was the NetApp strategy then before you had QLC? Was it more discounting? Was it more share losses?

And then if we take it to now, is there a risk that we start to see more workloads that are more on the hot side go to QLC because of macro, or just deal with the latency or the shortcomings? Is there a chance that there is a little bit of a price change or whatever, because it might be good enough for a little bit more of the workloads than what we're seeing today?

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

I don't know about price changes and I won't go into those, I will say, going back several years, I don't think there was significant discounting in regards to it being over served because people didn't realize necessarily that there was any middle ground. And so the, the market kind of found the price point. That was, right as it's want to do.

I do think that the fact that NetApp had hybrid flash offerings that were still incredibly performant and we had that FAS business that continued to perform meant that there were options for customers, right? Whereas if they went to some of the all flash competitors, they really didn't have an option. So those all flash competitors, they either had to pay the higher price tag or those all flash competitors were pricing. I'll just say significantly below market at that point, right.

So no, I think you saw the success story with our ramp to all flash over several years and how fast we grew in all flash and I think we did it without going into detail. I mean, you can go back and look at bearing statements from there. I think we did so quite successfully. And so I don't think it depressed earnings there to the second part of your question.

Will there be as we introduce and as we continue to rapidly grow C-Series, will there be customers that refresh their A-Series over to C-Series? Probably? But that would have been probably a natural transition regardless. And I think we're continuing to see them growing more and more high performance workloads. I mean, if you look at these high performance LLM trading models, if you look at high performance inferencing, different things like that, there's a whole new generation of workloads that are no longer underserved by A-Series.

So as much as there's the potential for some moderate shift of those, over served workloads down to capacity flash, right? There's as much potential if not more potential for new workloads coming in that are incredibly latency sensitive and want incredibly high throughput. You -- if you're hooked up to an NVIDIA DGX system, for example, the cost of keeping it waiting is far in excess of the price differential between a C-Series and A-Series system.

So customers who are spending millions of dollars to build out that sort of training environment for NVIDIA are more than happy to pay that premi between the types of flash to not have any cycles of their GPU sitting there in an I/O State in an I/O tech state.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

I guess it's back to me. So you mentioned BlueXP and I think that's probably an understood offering from NetApp, at least amongst this community. So maybe you could talk a little bit about what it is and how customers leverage it and the capabilities it brings.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Yeah. So BlueXP is our unified multicloud hybrid control plan. And in some ways, I think we get less credit for it because it's been this very gradual evolution, right? It started back in 2015, 2016 is like cloud manager. It was just a way to deploy some instances of ONTAP into the cloud.

Over the last three or four years. It's really evolved into having complete management of basically the entire NetApp estate. So a customer's entire data estate sitting there, not just on prem but in all the major clouds. So we have customers that have chosen NetApp because the whole is more than the sof its parts, right? They're able to literally see their entire data state at a glance. And then the nice thing we've started to do over the last couple of years is be able to add intelligence into it right.

So the whole intelligent data infrastructure, yes, it's a marketing slogan, but each word actually means something, right? So the data infrastructure is what we built up, not just as infrastructures on prem, but as data across all of it. So we can actually see your data, we can see metadata, we can see what's changed about the data and then adding the intelligence into it.

So for example, we added a Ransomware Protection service that can actually understand not just individual lungs or individual files but actually understand workloads. And so it'll know an entire workload. It'll know where that workload is scattered, some on cloud, some on prem. And it will be able to monitor it in real time for ransom or attacks and be able to respond, allow you to protect an entire workload in real time and recover an entire workload in real time, regardless of where that infrastructure is.

So it allows customers to not necessarily have to have an Azure specialist, an Amazon specialist, a Google specialist, an on-prem SAN specialist, an on-prem NAS specialist and then a security team to come in into the restores. You can literally have an IT generalist perform the job of so many of those specialists by being able to orchestrate entire restores directly out of that operating system. You can have -- you can drag and drop to tier from on-prem to the cloud. You can drag and drop to tier from one type of on prem to another type of on prem. In another data center.

You can set up your disaster recovery in a couple of clicks directly within that interface. And it's all included, essentially no additional charge to customers directly within an interface. So once they get in there, they're able to see the entire state of their fleet. We've added AI operations directly in there. So there's the digital advisor that will tell them if they have any risks, it will tell them about things they could ameliorate, it'll tell them about the health of their entire environment and it puts it all there at the -- the tip of their fingertips. And that has been, I say honestly, one of the more hidden but one of the more sticky features of NetApp because once you can get a complete unified view of your data state, it gets very hard to want to fragment it out after that.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right time for one last question and if it's not from the audience, you guys are going to take it from me.

---

**Unidentified Participant**

I was wondering, the point in time when that I was talking a fair amount about hyper conversion and we don't really hear that all that much anymore. But if you -- if we just step back and say there is a role for hyper conversion transition to cloud, maybe for a lot of customers. How do you think about that? Where does it fit in your discussions anymore or your messaging?

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

So I think in my own personal discussions, hyper converged has almost disappeared. which is fascinating, right? Given how much prominence it had just a few years ago. There's a couple of reasons, I think a lot of the people who were initially looking at hyper convergence realized that they could go to cloud and do it that way as a service because the move towards hybrid convergence was really about trying to simplify the stack.

And if you're trying to simplify the stack, it's a good sort of slippery slope to saying I just want someone to provide it for me as a service, right? So doing that as a service on the cloud has replaced hybrid convergence in a lot of places.

The other part is I think customers started to move down the hyper convergence path and all of a sudden became very afraid of lock in, right. So if you move down the hyper convergence path, for example, and you were on a hyper convergence platform that used VSAN or use VMware, you got an interesting surprise over the last couple of years, right?

And so customers have been, even if they're not on a VMware based platform, they've been very conscious of the fact that if they lock themselves into a hyper converged platform. They're very dependent upon a single vendor for all aspects of the pricing, all aspects of their stack. And as they're looking at how to adopt the cloud, how to adopt AI, I think they've all become very hesitant for good reason about lock in. And so we want to continue to partner and be best of breed, right?

And we build these converged infrastructure stack. So we have AIPod within NVIDIA, we have FlexPod with Cisco, right? Sandeep talked about the OVX partnership we built out with Lenovo, right? So we build these converged infrastructure stacks that are just about as easy for a customer to consume. But if they decide they don't like NetApp tomorrow, it's all open standards based. If they decide they want to switch to another server vendor, they can do that.

And I think that flexibility, especially with uncertainty about workloads continuing to move to cloud, uncertainty about how AI is going to change everything is for the most part, I think scaring people away from hyper convergence with maybe the exception of the very, very low end, sitting in some small shop somewhere. I think that's really where HCI is getting relegated these days.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, that wraps up our session with Jeff. Thank you so much.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Thank you.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

I really appreciate you as always.

---

**Jeff Baxter** - *NetApp Inc - Vice President - Product Marketing*

Thank you very much, Kris.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Okay. So now, we actually have a couple of customers that you can ask questions of. So I'm going to invite Scott and Casey up to the stage.

Hi guys. Thanks again so much for doing this. We really appreciate it. This community doesn't always get to hear from actual customers. So I think it's really important to be able to get some real-world practitioner input in there. So I'm going to just start by asking you guys to introduce yourselves, your company and what your experience with NetApp has been. So, why don't we start with you, Casey?

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

I'm Casey Shenberger. I'm a Cloud Platform Architect at Hyland Software. I've been there a long time and our company, we use NetApp to host our products as well as internally. I work in the hosting side where people host our products in our own private cloud. We also do hosting in public cloud.

All right. So we use NetApp products. We use all flash products for certain workloads. We have started using the C-Series NetApp products for other workloads that don't need it as reduced latency. We use FAS systems in, in data centers where we don't have the performance requirements of A

or C-Series currently, we also use a storage grid products to tear our, data off too. And we are moving some of our workloads into the cloud with AWS. So there we use NetApp products, cloud volumes ONTAP and FSx for NetApp ONTAP as well. So, we kind of use all of NetApp products across the board.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. And Scott,

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Hello. Wow, little volume. All right. So you can hear me but a little bit too much. All right, I'll try to whisper I'm good. All right. So, Scott Brindamour, I work for Lumen. I'm the Vice President of Product Management. I have four different areas I cover. So our edge compute and cloud business that we provide, I can explain some of the solutions that we do there. NetApp is part of that business.

I also have our data center strategy from connectivity as well as solutions ONTAP of data centers. AI is a big piece of that we mentioned earlier. Hyperscale strategy, so how we're going to market with co-innovating and co-selling with the big cloud providers. And they also have a wholesale business. So the legacy large pipes all over serving and working with joint network customers and partners going forward.

So I've been with Lumen for 18-years. So if you've heard of savage communications, that's where I started got acquired by CenturyLink. And then when Level 3 and CenturyLink merged together, we became Lumen. So NetApp has been an amazing partner. So you heard Jeff, earlier, talking about hybrid, Lumen is all about hybrid. Really what we're really trying to work with NetApp on is we're being coined through our CEO and our executive team as the network for AI.

There's been a lot of buzz around our stock. Who's gone a little bit crazy lately with a number of the custom private fabric deals that have been sold. So to a lot of the cloud providers, social media platforms, et cetera. So in the billions with billions coming that they're looking for. So this real rush for AI is real. We want to be that network that central nervous system that connects the people, data and applications together across wherever customers want to execute.

One of the reasons why we're doing it with NetApp. We have a platform that's built on our network which we call Lumen Network Storage that combines the solution combines the power of the Lumen Network to connect the data to its destination that it needs to go to or where that data is originated. So great use case for AI, we're in the midst of creating a solution with NetApp around. How we go to market together to connect the network and the data which I think are the two most important parts to build AI platforms to continue to invest in AI platforms and how you actually harness the data to use it and get customers on that journey.

So that's why I'm here today. Happy to answer questions and talk a little bit more our relationship with that.

---

**Unidentified Participant**

Yeah, I'd be curious to hear from you sort of, as you did your evaluation work and you looked at across vendors, what were some of the key points that primarily led you to choose NetApp?

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

So I think flexibility. So we're -- we standardize on, their ONTAP platform as well as object storage, storage grid. So the combination of those products to be able to meet, so we want to create a solution by which customers don't have to think about what performance storage they need to match it to the application and have a platform that's versatile that can support, whatever hybrid infrastructure and applications they want to put on it as well as connecting to the cloud on the premise and integrate with that seamlessly.

So regardless of what we're Switzerland, we work with every provider and platform provider out there. We're trying to make our product, so it supports whatever the customer journey wants to be, and that's it. So we've, done from the highest, financial trading applications, real time, data access and oil and gas to a general file system customer down in mid market.

The platform scales to the size and the capabilities and the ability to automate the delivery of storage to a customer. And even down to the we can provide for our larger enterprise customers. A go to market by which we -- we do dedicated storage with them as well. So I think the flexibility of the hybrid approach to work with everybody. It was talked about earlier when I walked in, that's the message that we had so very good alignment of how we serve the market, how we work with the cloud providers, how we work with the data center providers together. So there's a lot of -- very symbiotic relationship there.

They focus on what they know which is storage and data and apps. We focus on the network and we combine together to create solutions together. And that's really what we've been successful with NetApp as well.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

We had a similar situation except for. I obviously work on the, the technical side. And it was a similar thing we chose NetApp because of the unified interfaces that Jeff was talking about earlier, we have a very small staff to manage all the equipment that we have. So by using NetApp, we had, high performance tiers, mid performance tiers, archival tiers, all of that was in the same platform, the same operating systems, the same expertise is required.

So we looked at other platforms and it meant, we had to have one vendor for high performance and we might have to have a different vendor for an archive tier. Well, that means that you have to have people who have knowledge in both or separate people with knowledge in one platform versus the other. NetApp allowed us to just have one group and one set of knowledge. So that was a big help for us.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, I would say. I'm always having a storage vendor knock on my door trying to get a piece of our state and our customers. And it's usually an easy conversation that we've got a platform that can pretty much do what we need to.

And the other thing I'd add to that, in listening is the innovation, the ability to go to market and innovate and try new things and go to market together. We've done a lot of proof-of-concepts, a lot of prototypes, some of them do. Well, I used to run an innovation team before I moved into this role. They're very willing to jump in with us and learn together and understand where the market's going. So that's been big for us as well as we're trying to create the next big thing in the market AI is a good example of that as well.

---

**Unidentified Participant**

Maybe a bit more specific product level question in terms of NetApp's, recent launch of Block Storage products, does it materially change your buying pattern with the company itself? And then maybe a second one just in terms of utilizing the public cloud more over time, how do you see that engagement sort of changing in terms of are you using NetApp, just because that's sort of the on-prem versus evaluating other sort of opportunities as well there?

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

Well, we've been using NetApp's, object storage, storage grid for quite some time. And as we move to the cloud, we still needed a way to do fabric pooling, or some way to offload that. So we do use S3 there. But we chose to use NetApp, cloud volumes ONTAP or FSx for NetApp ONTAP in AWS both because we still got deduplication, we still get compression, we still get the performance and we have the same knowledge set like I talked about.

So we stuck with NetApp because that knowledge transferred maybe a little bit in the back end was different but ultimately the storage knowledge required to function there and we still got all of our performance, we still got all of our compression and compaction and those things that help us to reduce our data footprint. So that's why we chose NetApp, right? They're leading that field too. So we stuck with them in the cloud as well.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Any do either of you use the ASA products or do you use ONTAP in a block? like unified block?

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, we use a unified block today. So yeah, I, I don't see any, I think any changes. I mean, we try to wrap it up as a solution to abstract the technology. It's flexible enough and hybrid enough, right? Combined with object and storage grid, I actually see that the with now with we see all the data that's being created for not just AI but IoT and data analytics that's being created it. We're hearing a lot from customers that distributed is more what customers really looking for.

Moving, obviously the cloud providers want you to move all of it into the cloud providers, more and more enterprises are taking a more distributed approach. So the ability to have a platform that can work with the cloud provider and the data and the apps you have there but also connect to on-prem dedicated and anywhere in between, especially with an AI model with the model is going to be distributed, the data that you're feeding the model is going to be distributed. That's really where we think. And NetApp is a perfect opportunity for us to scale, to deliver a footprint where we need it and scale up from there, without having to change the platform.

Add capabilities as we go, BlueXP and that whole control plan and portal or ransomware and data classification has been gigantic for our customers as well that have access that app. We're, working with them to product that as a solution that we can give and provide to kind of as a multi-tenant access to customers as well. So they continue to add value ONTAP of it, right?

I see that they have the same vision that we do that distributed storage as a service with solution capabilities. We're not talking about the storage and the technology underlying. Obviously, the techno geeks are speeds and feeds and are thinking about that but ultimately enable customers to solve a business problem with a combined solution on top of it, right? That's where I think NetApp and Lumen are really working very well together. Their vision is similar that we need to add and talk about the data and the value of the data and the outcomes you're trying to get for your data, reducing the size, reducing your risk, your security profile, the ransomware where features are a great capabilities there and accelerating what you're trying to do with your data to get value out of it for time to market AI is just the newest version of that, right?

---

**Kris Newton** - *NetApp Inc - Investor Relation*

And just to clarify when you say distributed. You're talking about like what, NetApp would call hybrid cloud, right? Like on-prem (multiple speakers) cloud,

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Hybrid cloud, we bring our -- what we consider to be edge is the network edge or the metro edge, which is the edge of our network where NetApp lives, it's in the network, it's part of the network. So any customer that's using our network can use NetApp as part of that, as well as to serve. We, talk about it all the time if you're a retail giant with 3,000 locations to put storage on all those locations in a footprint to do, automated checkout using computer vision cameras and image recognition can get pretty onerous and expensive to actually put all that gear, all that storage on every look -- every premise location.

And then by the time you finish doing it, which is usually like a three year exercise, you have to upgrade the other ones right to be able to offer an alternative to that. Some of it on-prem makes sense. But there's ocean regulations. If you have a small footprint for sound, there's complexity of reporting that across all those locations.

So we believe that the metro location is a, is what we call the third execution venue, but it talks about cloud and premise that's somewhere in the middle, that's the metro edge capability. That's that NetApp lives and we enable all those abilities to do distributed, applications and data workloads with that as well. So very aligned to your hybrid cloud and I put the end. And the metro edge to add to that as well.

---

**Unidentified Participant**

Hi, maybe for both of you built different types of businesses. But when you think about, the importance of storage in your budgeting or your, your CapEx or however you want to look at it. Can you talk a little bit about where that sits? I imagine it's a little bit more for, for Casey than for Scott. And then just, there's obviously a lot of demand for data. How do you see, your spend on data solutions over time tracking with the data increase in data usage? You don't have infinite budget. So any insights you can offer on those? Thanks.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

It's key to our budget, right? We still, I mean by hosting enterprise content. We ultimately, we store people's data and we have to manage it and maintain it. So it's a key in our budgeting. The key there being, these features we get from NetApp, help us reduce that right, compaction compression, deduplication, they reduce that footprint to help us get down to more budget friendly answer to how that works.

But it's always front of mind for me, obviously, I manage the data in general, but it's very, we sit down and we work very closely with our NetApp team to make sure that we're taking advantage of all the features that allow us to reduce that footprint, but still maintain resiliency, reliability, availability. That's, the key to us is don't reduce anything like that, don't reduce reliability or resiliency, allow to have ransomware protection, things like that, but continue to drive the cost down for, per gigabyte of storage or per terabyte of storage. However, we're calculating it.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, I would say that the data storage and what we're continue to deploy and purchase is not slowing down anytime soon, it's escalating. But at the same time, I think the opportunity to optimize what you have in working with our customers, there's a huge optimization opportunity. Lots of customers are keeping hold of lots of data that they should be archiving out. So having a solution like NetApp with object store to move it to a lower cost option.

To optimize what they have on their dedicated arrays with NetApp as part of the solution, but also be able to have that tiered storage option that you can use what you need rather than having, an all expensive high performance all flash array that is going to cost you a lot more per terabyte than a tiered solution. By which you can lower your cost footprint as well.

So us internally, I mean, with all the data assets and AI that we're doing in house to understand inventory and customers and where we're deployed, The amount of data is going to continue to grow tremendously. We see it. We see the investments. We're moving a lot of app applications and services to the cloud. We see that it's exactly what our customers are doing as well.

So tremendous opportunity to reduce. So they can keep their budgets while they're expanding their storage as they go forward, you need to do a little bit of both. As you start to, I think you'll -- you're going to see it as customers start to adopt AI, they're going to try to make it take advantage of more of their data and more of their assets that they don't. They're not collecting and they're not using today. It's going to continue that increase that growth as you go forward.

So the optimization piece to optimize as much as you can and modernize that infrastructure going forward is huge as well as protecting that data as it grows. And as you start to move that key data that gives you those insights in an AI model around you to make sure it's protected and not available to the bad guys. Right.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, you get, okay. Well, I'm see.

---

**Unidentified Participant**

Kris might not like this question. Maybe she would. But, I guess, if you were to give feedback to NetApp, like what would be things that, from a product perspective you would say are things that maybe you're running into pushing the limits, so to speak. Where, maybe they could be more helpful or maybe they already are and you're in discussions. But what would you say are some of the key things that you would give feedback to them?

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah. For me is as a solution guy and a product guy. I mean, I, but I've spent most of my career on the sales side too. So I kind of get multiple perspectives but the biggest thing, and what I mentioned before is less speeds and feeds and technology about storage arrays and their capabilities and adding and incrementing a lot of what I've, I used to hear was an update on the next array, the next technology, the next capability, which is great, I think for my engineering team.

But from a product and solution perspective and sales perspective, I have a bunch of network sellers that trying to understand everything that goes on in that space and then layer on top of the storage conversation. And the hardest part in go to market has been melding those two expertise together to go after a customer jointly where you both get a value prop and a benefit back to each other.

So changing the conversation to a solution focused conversation where you're focused on joint customers problems that we both see in the market and then building solutions together to go after that problem where we're both adding a piece of the solution and then integrating it together as a conversation where we're ha we're changing the conversation with customers.

So now you're not just selling to the, architect who's managing and developing the architecture and keeping it up to speed. But you're selling to the business owner who wants to take advantage of how I get value out of my data, right, going forward. And what do I need to enable that? What do I need to experiment with AI and Gen AI models? I don't have the infrastructure. Can you partner with Lumen to give them the network and the infrastructure ONTAP of NetApp that you can deliver a solution to.

So that would be the big thing is it -- it's, the company has definitely changed its positioning his approach and is becoming more solution orientated and they may have done it to the market but didn't go as quickly with the partners as you would see. So as a customer of NetApp and a partner that were going after the same customers. I've really seen that shift and like to see it more as a person who's looking to -- to create the next innovative solution in the AI space. They're leaning in, but it took a while before they got there. So that would be the big thing that I see.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

From the technical side, because that's where I am, right? But I think the biggest thing that we struggle with or that we continue to provide feedback in NetApp is. They need to continue to provide better ways for automation and scale. So, right, as a hosting provider, we get lots of customers coming on, we're doing more and more and we want to have all that performance, reliability, resiliency I've been talking about.

But also, we don't want to -- We want to maintain our staffing levels, we don't want to have more staff to do that work. So we rely on automation and as that -- as NetApp is, makes these shifts and they start to transition. we need them to continue to give us the ability to do that in an automated way. There's some things there, here and there that are tough to do that for. So we continue to work directly with NetApp and ask them, let's make it. So this is much more scalable and much more automatable. That's our biggest feedback to NetApp, I think.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Well, let me ask the flip side of that question, which is right. You've both used NetApp for a long time. What capabilities did we bring that surprised you or that you learned through actually using the product?

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

I think that the, probably the one thing that was the most. I guess that really caught us and it wasn't necessarily a surprise but the ARP capabilities, right? We were...

---

**Kris Newton** - *NetApp Inc - Investor Relation*

ARPs is automated ransomware protection.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

Sorry. So, that functionality, kind of came around, we had been talking about needing ransomware protection and how we were going to do that. There were some, all kinds of different methods for that and then NetApp started adding that directly on box and now you can -- we can do that on box automatically. It's all Gen AI based. So it's very accurate. So I think that was one thing that kind of, although not really a surprise, it was a big help to take away. We didn't have to go. Now look for a ransomware vendor that we could partner with. Right. Our vendor just brought it along.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Not a AAP right?

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

Right.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, I would say the ransomware protection. So back to what I was saying that the shift to a solution mindset security in ransomware is huge. The ability to actually sell that as an add on service to our customer. And the security market, which we do a lot in the security market ourselves was huge. Like that is an added value that I can monetize with my customers going forward. I think cloud blue in general, the data classification capabilities that's built in the control plane to see all of your storage across your hybrid infrastructure.

All of that the want that was mentioned earlier of automation and visibility and control is gigantic. So that's as I said, that solution shift has happened recently and I've been surprised at how quickly they've adopted new capabilities as well that are benefiting not just Lumen to another capability that we can sell and add value to our customers. But for our customers themselves that they made the right choice in NetApp, they continue to get value out of the platform as they go forward. So that's been gigantic.

**Kris Newton** - *NetApp Inc - Investor Relation*

So I know you both talked about the value of the unified storage approach, but the fact that we have kind of ONTAP underpinning everything is what enables us to bring those incremental features and have them broadly scaled. Audience questions, otherwise you're going to keep hearing from me, Lou?

**Unidentified Participant**

So the, AI session from earlier said a lot of small, medium businesses aren't really yet using AI internally. So I'm not sure how big your companies are. But, are you all using AI internally yet? And if, actually you think that, your firms are just, aren't big enough yet, where do you see possibly that happening in the future?

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, we're, pretty big.

**Unidentified Participant**

Thank you.

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, I'd say 50,000, 53,000 employees as a telecommunications company that's trying to be a technology company, we're adopting. We're huge users on the Microsoft Copilot side. So we're probably the poster child for Microsoft around copilot. Our CEO, Kate Johnson came from Microsoft. So there's a connection there, but it's tremendously valuable in having data at your foot -- at your fingertips, transcribing and summarizing meetings for me where I'm triple booked. Like I'm here at a conference for the next few days, all the meetings that I'm missing, I can get access to it.

But just the data at deluge of data that's available, trying to find that content across multiple systems and email and all the things that go on is huge. So that's one example that we put in that we started small and then it's pretty much available to anybody.

We thought the price tag originally was pretty high and we limited it. But the value that we got out of that brought out of productivity from that solution is gigantic, living in a world of tremendous meetings that's valuable to us of how you can optimize people's time. And there's a whole regular course internally of enablement and adoption of that, which has been huge. We've been actually trying as a company who has infrastructure everywhere, every data center, every cloud, millions of buildings across the US. And in AsiaPac as well, we sold off our European assets recently.

The ability to understand those assets and using AI to not just trust the tech that installs something or turns up a circuit, the ability to have the used cameras to understand what's there. What the capacity is, what was installed was installed the right place. Do I got the nice blinky lights on like inventory management and actually putting that against demand and bringing that data against, sales force, forecasts and demands on a particular that may say, a tech that's installing something for one particular opportunity or customer can look at the demand and say I'm not just going to bring one chassis, I'm going to bring two and I'm going to install that.

So the ability to use the data for real time to train people of how to do particular jobs and capabilities, that's all automated and to check their work after the fact that's been gigantic. So I think throughout the process of understanding and data and how to use it has been tremendous across the business.

**Unidentified Participant**

And did you just develop that application, the last one yourselves internally with part..

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Partner and you name it, we're working with pretty much everybody in NVIDIA and Intel as well as Microsoft and Oracle. And some of the SIS well are helping us to adopt it. So it's been pretty much across the board. I mean, it's pretty incestuous relationship. We sell to them, they sell to us. Balanced trade is huge. So there's always someone looking for us to do more consulting or do other things with them as well. But the opposite of is, that partners that are willing to invest a lot with us to use this as a use case, as well going forward and we do the same with them. So it's been really amazing.

I think, for a big company that has the capital to do it, but also capital is tight in these times. It's, we think we've got a tremendous return, but we're learning what we need to do and how we need to go about it. I think it's very similar to what enterprises are. They're getting their feet wet and they're learning how hard it is and how they need to focus and where they're going to maximize the value as well.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

Casey, any comments form you?

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

We use -- I mean, we have AI like you said about meetings and, some copilot stuff. But for really, for me, I'm not super involved because I'm hosting customer data. I'm managing that underlying infrastructure. So it would be more on like the software that we write and that we host. And that's, Hyland is. We actually have our community live this week. So we're going to have some new announcements, about what we're doing with AI and where we're going there for to allow our customers to have like a, more seamless integration with AI to get access to the data that I store and manage. So, today, I don't really use it a lot, but it's coming in our products and it's coming in our cloud and we will definitely, embrace that.

---

**Unidentified Participant**

Thank you. Maybe we could get both the technical and business perspective on this question a little bit. We do hear from some of NetApp peers about how power has become an increasingly important consideration around storage. And so as you think about it, could you help us think through maybe in from you specifically. How much of your budget is actually power budget is being consumed by storage as a practitioner from a tech perspective and where do you see that going? Is it different between all flash versus disk? And from a business perspective, how important is it? I know you guys are, I think [bot] committed 10% of the fiber from corning or something recently like that, right? Like that's a lot of fiber capacity, building a ton of capacity out there. So as you're building these large data centers is your storage strategy going to need to change. And also as power in that context. Well, how do you think about that? Thank you.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

Yeah, power. I mean, power is a big part of our storage budget. And, and like you said, it's directly related to the type of storage that we're using. So as we do more performance and SSD is huge, right? That gives us more I/O and less power.

(technical difficulty) we get it in less space. It's why we've embraced using C-Series, we have plenty of workloads, we don't need sub millisecond latency for but we do need still, high I/O access and late in their latency tolerance needs to be low. That allows us to get a lot of capacity with a lot less power. And so, and then because we're able to do that. And we, and we use that with tiring that same unified approach we've been talking about all day that allows us to, you know, that archive tier maybe uses some more power, but it uses less because it's disk, that's kind of just there

and not access near as much. So, in all these new things, right we choose, sometimes we choose the type of platform with NetApp, we put our workloads on, based on, power consumption or what kind of space they're going to use in a data center. So, that's how we handle that piece.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah, power is always a big cost and consideration, especially when you're distributing infrastructure everywhere. So I think I mentioned earlier about having a platform that's scalable and modular that can meet all the use cases without putting a -- to build it and they will come kind of days at Lumen are gone. So how we optimize what we think we need now and then how we can easily scale up has been a big piece of what we do and optimizing the same footprint everywhere to make sure we have a kind of scalable, supportable, power infrastructure, cooling infrastructure as well.

Now we're, you know, getting into GPUs things of that nature they're at more order of magnitudes over with the storage system which is, unified and supporting all across all workloads. So I think from the storage perspective. We're pretty comfortable with the model, we have dedicating it where customers need, more power for customers. We've had, big financial services, enterprise customers that have deployed dedicated capacity in areas and making sure we're starting small and building up as well so that we're not, consuming a bunch of power and not using.

it's been kind of what we think about. But then going forward as you distribute, network capacity and compute and storage and security capabilities across the platform. It's becoming harder and harder that a lot of these destinations are old central offices where we had all the telephone gear, which is now in the network gear, which is now virtualized. So thinking about how you put it in a low power, we're always looking to optimize the power of every device that we put in as well.

As well as how can we contain the power goes straight with the heat as well, right? And the heat in some of those facilities to be able to cool them and do liquid cooling and some of those innovations as well. But a lot of vendors are bringing that. So we've been talking to NetApp about how we optimize that as much as possible as well. But the modular systems approach, helps with that tremendously so we don't have to deploy something gigantic and then be consuming a lot of power. It's a cost we can't get back. Right, because we're not optimizing that. So that's been the big piece of it going forward.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, since the questions seem to have died down, I'll ask one final question, not necessarily specific to NetApp, but more around (technical difficulty) how you're thinking about the future of storage and your data infrastructures. What are your -- what when you look on the horizon, what do you think the biggest things coming for both of you are in terms of opportunities or challenges?

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

I think opportunities that we've been really thinking about, composability of systems and composability of data in storage that I don't have to buy, I can assess and -- basically build and have a customer compose a system on demand that they need. Whether it be compute, whether it be memory, whether it be storage or network back plane. So we've been in my innovation time, we've been looked at a number of vendors for around down that route as well.

So how can I actually get integrated systems that have all of that together where I could build that modular on the fly? So we've talked to, NetApp a little bit about that as well. Not giving up any research in R&D that we've done. But I think that's a piece's. How do you deliver just enough storage and just enough compute enough memory to support the requirements of an application without, pre building and, building it and they won't wait until they come.

But the capability, I think that's already started is that the AI capability on top of that to understand the data categorizing that data, understanding that data and making you aware or even automating some of the -- the maintenance capabilities around that data segregating the data off trying

to work across the infrastructure layers and the stack including storage that our vision or my team's vision on the edge side is to create an environment where the application dictates the infrastructure that you need.

So if you have composable components in a solution that you're delivering, including the storage, how can the application understand and pre-determine how much storage, what type of storage where it needs? It understands the network and understands the compute and can build the system that it needs to support what it's doing now and then adjust as it goes forward.

So kind of that autonomous system that's composable that you can build up. So how do we support that in a business model going forward with NetApp, as our primary storage partner as well. So that's one of the things that, that we think about that we've been trying to do. We're probably early on in the technology but that would really optimize where we're not putting a tremendous amount of infrastructure in one place or building lots of racks in a data center that we can really compose it just have just in time inventory where I can swap out that the latest greatest performance of all flash arrays with something else going forward as the technology lead product.

It's big with GPUs that I don't want to GPU that cost me, $60,000 sitting there not being used. If I can afford it, it's locked into a server in a system, same thing with storage. How do I unlock that and create systems that are modular and snap together as you need it like lego blocks. So that's what I think about.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

I think for the future for us, it's probably very similar, right? We would storage -- storage is growing at a unbelievable rate, right? It's not getting any smaller. We need more analysis, we need more ability to know what that data is and where it's going. So for the future for us, it's just going to be to continue to optimize as we're storing that data. How do we optimize the space, the power, its resiliency? How do we optimize the intelligence that we have about the storage so that we can classify it properly and move it to the right location.

And those are all things that, we've continued, we work with NetApp on, we work with other vendors, even for the compute side, the same thing, right? We need to optimize all the capacity that we're using. So that's the future for us, we just continue optimizing that as much as possible to reduce those footprints.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right, [Mita].

---

**Unidentified Participant**

Just maybe on that, I know probably over the past couple of years, you went through a bunch of cloud optimization type of initiatives just is how you look at that optimization, different kind of or is there any way in which it's different pre or post kind of some of that evaluation you might have gone over in the past couple of years?

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

I don't know if it's any different -- I mean, it is somewhat different because the technology changed, right? Like some before in the past, it wasn't necessarily like can I put it on, SSD versus spinning disk? That's all that was the only two options, right? But now we have, single level cell, we have multilevel cell, we have quad level cell, right? We have lots of different SSD options.

So that optimization now NetApp has grown that way too, right? They originally only had all flash and that was on, single level cell now they have the C-Series, right? So as we continue to work down that optimization as the vendors bring multiple options then that gives us the ability to kind

of, to see. So it's changed in that way. But it really hasn't changed from we're still always looking a way to optimize that and put workloads where they belong.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Yeah. I think more from a solution and product perspective that the cloud providers have almost abstracted a lot of the technology and gone to our services. What's the application approach? What's the service you actually need and kind of abstracted the infrastructure way. So, but it's also they made it easy. But at the same time, a lot of your eggs are in that basket, right?

So, I think that's where NetApp has helped us tremendously to kind of be that neutral platform, but I don't see it any different optimizing cost. I always see the [pendulum] people moving a lot to the cloud trying to move it away from the cloud. I don't think that's ever going to change. It depends on where you are and your maturity as a company that you go through those patterns as you go forward.

But again, simplifying and the ability to offer, storage as a service as a solution. That's kind of our vision of what we try to do. NetApp has been helping drive in that direction as well, but also be hybrid that they can participate with the cloud, they can participate on-prem they can participate with us on the edge wherever it may be going forward and just abstracting that capability of the way and supporting the application and what it needs going forward.

So that's the thing I have seen change a little bit with the cloud providers have been leading in that generation. But there's another angle that you need as well as private and close to where your users are, et cetera that the clouds never going to get to. Maybe it will be one day, but we're not there yet. It's always going to be enough data and apps and performance that you need locally is but having that hybrid approach that supports whatever you need going forward is huge benefit, right? But making it easy, like you don't have to be a storage expert to buy it and use it and automate it right? So it can run itself is, is going to be gigantic going forward.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Last opportunity for a question. Otherwise I'll release Casey and Scott back out to talk with other customers going once, going twice. Okay. Well, Scott and Casey, thank you guys. So much for coming. I cannot. Thank you enough. I probably take that from you.

---

**Casey Shenberger** - *Hyland Software - Cloud Platform Architect*

Thank you so much.

---

**Scott Brindamour** - *Lumen Technologies, Inc. - Vice President of Product Management*

Thank you.

---

**Kris Newton** - *NetApp Inc - Investor Relation*

All right. Well, we're done. So thank you to everyone on the webcast. If anything peach your interest here, please don't hesitate to reach out to the IR team. We're happy to get any follow on questions or connect after this event for those of you who are here, lunch is outside. So thanks for staying for a late lunch with us.

Again, there will be the key notes later today at 4:30. And then the show floor is open after that. As Jeff mentioned, there's some announcements that you could expect to see in the coming days. Once we kick off insight officially with the keynote today.

So thank you again, all for coming and always don't hesitate to reach out to the IR team if you have any follow up questions. Thanks.