



Pure Storage, Inc.

Investor Day

June 13, 2017

C O R P O R A T E P A R T I C I P A N T S

Matt Danziger, *Head of Investor Relations*

Michael Pak, *Chief of Staff, Finance and Investor Relations*

Scott Dietzen, *Chief Executive Officer*

David Hatfield, *President*

Matt Burr, *Vice President, Platform Solutions*

Kevin Delane, *Vice President, Global Sales*

Matt Kixmoeller, *Vice President, Products*

Rob Lee, *Chief Architect, FlashBlade*

Tim Riitters, *Chief Financial Officer*

C O N F E R E N C E C A L L P A R T I C I P A N T S

Wamsi Mohan, *Bank of America Merrill Lynch*

Katy Huberty, *Morgan Stanley*

Aaron Rakers, *Stifel*

Alex Kurtz, *Pacific Crest Securities*

Jason Nolan, *Baird*

RK Raghunathan, *JP Morgan*

Tim Long, *BMO Capital Markets*

Dave Ryzhik, *Susquehanna International Group*

Kirk Materne, *Evercore ISI*

Simon Leopold, *Raymond James*

Eric Martinuzzi, *Lake Street Capital Markets*

P R E S E N T A T I O N

Male Speaker:

We've got a good lineup for you today but before we start, I'm going to just do the Safe Harbor that you all are very accustomed to. Before we begin we remind you that Management will make forward-looking statements today, which are subject to risks and uncertainties. These forward-looking statements are covered by the detailed disclaimer at the beginning of our presentation slides and in our public filings with the SEC, which we encourage you to review.

Also, we will discuss non-GAAP measures in talking about the company's performance. Reconciliations to the most directly comparable GAAP measures are provided in our presentation slides which are available at investor.purestorage.com.

Now I'm going to hand the mic over to Mike Pak. He's going to go through some logistics and then we'll get going with our program. Thank you.

Mike Pak:

Thank you. Tough act to follow. We've got a great program today. We're going to start with Scott's opening remarks. He's going to go through our long term vision and our strategic viewpoints, and we're going to have David Hatfield's team come up here to double-click on our go-to-market engine, and we're going to end that session with a customer panel featuring our customers from Intuit, Broadcom and Mercedes. We'll have Matt Kixmoeller and his team come up and go through our innovation roadmap. At that point we'll take a short break and we'll save the best for last, which is our long term financial updated presented by our CFO Tim Riitters, and we'll naturally lead that to the Q&A.

Just a couple of logistics. There are restrooms outside and inside, as you guys like. The WiFi password is Pure Accelerate Attendee, no password. There are water, coffees, juices to keep you hydrated, and it should be a great program.

With that, I'll turn the mic over to Scott Dietzen.

Scott Dietzen:

I think I'm all set. Thank you Pac-Man. Apologize for the background noise. We actually were going to use a different room but it sounds like the inside of a ship because there's a big generator behind making noise, so hopefully this accommodates everyone.

Can I get a quick—I want to just get a quick read of the audience. Who was at the keynote this morning? Who was not at the keynote this morning? All right, so we're going to bridge the gap and hit a couple of highlights from the keynote but we'll try not to bore those that weren't there.

One of the fundamental changes that's happening that we motivated in the keynote is the value of data is growing. In fact, this was a cover article off of *The Economist* in May declaring that data was now the world's most valuable resource, ahead of oil. We did a survey maybe a year ago where 80% of the customers we talked to said they could raise their top lines by 20 points just by mining insights they already had inside their data but couldn't get to, right? If that's true that's a phenomenal amount of economic growth that we can unleash by just getting access to value inside of the data, and this is what's so exciting by some of these new paradigms like deep learning. Deep learning doesn't require customers to build huge teams of data scientists and they don't have to go off and write programs themselves to go get at the data; they just need a really big data set and they need to train the deep learning neural nets on what they want to find and what they don't want to find, and the systems will learn, they'll learn to differentiate between the good and the bad, and it allows the customers to extract a lot of value across a great many different use cases, right?

We're involved in this for facial recognition, text analysis, supporting applications like autonomous vehicles with the technology. In fact, Pure is powering several of the big data rigs that are shown on the

screen as well as north of 500 state of the art tech clouds, right? When we talk about clouds we mean software as a service, infrastructure as a service or consumer internet capabilities. So, an incredibly exciting trend that we're on the forefront of, and it really leads to our playing across three different markets.

So, we're very much involved, as I mentioned, as a technology supplier to the cloud. We've more recently seen dramatic growth in what we're calling edge use cases. Customers are generating data from high def cameras, other sensors, that they need to be able to capture and ultimately want to employ data driven analytics including machine learning on those data sets, and as we'll see in a minute, it's too big to send them to public cloud.

This market is the market in particular, the storage-related market for AI is growing at a 78% CAGR and is slated to be a \$5 billion market in 2020. We don't see our traditional competitors in this market or this market. For AI, massive parallel AI demands massively parallel storage underneath, and the legacy technologies' designs cannot deliver on that, right? We won one of the premier AI platforms in the industry, the top three, top five AI platform. It's 100 gigabytes per second and millions of IOPS, all going on in parallel. It's just not feasible on traditional storage designs.

The enterprise market continues to be huge and is informed by these two markets, right? They want the same ease of use. They want the performance. They want the low cost of operations and the low cost of scaling and evergreen business model that is provided for these other two markets, so this is what gives us the opportunity to play so broadly around this wheel. In traditional enterprises, in clouds and on edge data centers, and I think edge may be the biggest of all; here's why.

In 2020, humankind will generate 50 zettabytes of data. Most of that is Internet of Things, cameras and sensors that we're capturing this data. The turquoise is the global Internet bandwidth forecast for the year 2020. So if the entire Internet is devoted to just uploading data from these sensors, this is how much of it we can get into the public cloud; all of the rest of the this has got to stay on the edge. That means in an edge data center close to where the data is generated or on the devices themselves. Data that's generated in the cloud stays in the cloud, no question, but data that's generated on the edge from Internet of Things applications, at least the larger data sets can't possibly go in the public cloud, unless they can put the data on trucks and they don't need to know what's going on for weeks.

So this is driving a fundamental sea change in the market. We have literally the situation that the vast majority of storage still used in production today was designed more than 25 years ago. For mainframe and client server technology, it can't make the leap to this data-driven AI savvy cloud-centric future that everyone wants to play in, even the traditional enterprises.

The reason Pure is having success is we got this right early on, right? We build a product that has the bandwidth to capture really big data. I mentioned north of 100 gigabytes per second sustained in this large AI platform. We have the performance to drive massively parallel deep learning against those data sets where our competitors don't. We've got six 9s of uptime, that's less than 31 seconds of downtime per year but unlike our competitors we don't have maintenance windows and we don't have data migration, so that data is always available to our customers' business.

The user experience is very much like the cloud in simplicity: self-driving, infrastructure I just provision and then leave it alone and it takes care of itself. We give our customers a subscription not just to software but future hardware innovation as well, so the transition to these new protocols like NVMe for those that have heard about it, that's included for Pure customers. They don't have to sweat. They don't have to throw away the last generation of their technology and buy all new.

Lastly, it's driving a significant change in data center architecture. I should mention, being software-driven and hardware accelerated is exactly the model that's used in the large Big 3 public clouds and that's the model we're employing and delivering to our customers.

One of the leading SaaS customers we have calls this density lust. They were able to take a state-of-the-art rack and double the amount of compute they could put in that rack because FlashArray//X and FlashBlade shrink the footprint, and for a particular application they drove a five-fold performance increase even as they drove up that density. That's versus in a state-of-the-art SaaS deployment. The gains in enterprise can be dramatically larger still. We had a customer that drove a 100-fold consolidation, north of 20 racks down to a 4-rack unit FlashBlade and performance went up.

This is changing the way people think about data center architecture, right? These fast networks. With these fast networks, storage across the LAN is actually as close to the CPU as storage in the same chassis. So the whole cloud model of DaaS, server local storage, is getting turned inside out because now all of those 1,300 cores in the rack are the same distance away from the storage media. In fact, with these modern technologies like NVMe over fabrics, it's actually lower cost to access remote storage than it is local.

The big clouds are embracing this model. We are embracing it. Our competitors are saying it's three years out and we are here today.

So, that's the end of the start up, right? We've got a really huge available market to address. On the enterprise side, a \$35 billion TAM and we're uniquely playing in the cloud end and the deep learning edge Internet of Things, and in both of those markets we don't see so much of our traditional competitors.

I think I am passing the baton to our president, Dave Hatfield, and KD and Burr.

David Hatfield:

All right. Thank you, Dietz. Nice to see many familiar faces. We're going to go through go-to-market execution as one of the secondary drivers for the market. We've got a huge market to go attack. We've got great innovation, more innovation that we launched this morning than we have in six years of the Company, and you combine that with a great go-to-market engine, you have an ability to get to billions of dollars of revenue in a pretty reasonable timeframe; so, we're going to walk through that.

A couple of folks are going to help me do this. Kevin Delane, here. Stand up, Kev, if you don't mind. Kevin runs Global Sales for us. He was with the company for—in his third year. Ran our International business over the last couple of years; prior to this 20-years plus at EMC, ran the Isilon business. Been doing this a long time. Great to have your help here.

Then Matt Burr is Vice President of our emerging platforms. So, one of the tricks here is getting and leveraging the established footprint that we have with FlashArray through distribution and then leveraging that with our data platform framework to sell all the product, so FlashStack and FlashBlade, and the fourth and the fifth product that come out. So, creating a leverageable engine that allows us to be able to capitalize on the innovation that's coming through the pipe is very important. We want to make sure that's repeatable and make sure that we're taking advantage of the unique assets that we've got.

Matt was the 11th employee at Pure. We've worked together in multiple companies. He and his team have sold over a billion dollars or Pure Storage since the inception, so able to kind of get it out of the blocks but also scale it as well. Great to have you in the new role.

The three of us will work together, and more importantly, once we stop talking we'll actually hear from our customers. We've got three very distinguished customers. For those of you who were in the Main Stage, you heard Matt from Mercedes Petronas. I'll just say again congratulations on getting the wheels back on the boat this weekend. I was in Monaco together and it was a tough, tough weekend, but to see that team execute the way that it has and deliver a one-two finish in Montreal is awesome. It's great to be a part of it and thank you for being here.

Justin Stottlemeyer is here, who is a distinguished engineer and architect at Intuit. So, been a long time customer for Pure, advocate and a great technologist as he sees things going forward. Andy, welcome, the CIO of Broadcom, as well. A great customer and a great partner.

I think we have a really great cross-section. Matt and Kev will go through the customers for our customers as well.

Without further ado, let me hit some of the highlights. The Company since we've started—I've been with the Company now five years. We had I think four or five selling teams when we started. We have almost 1,000 today in our go-to-market engines directly across 30 countries globally, and 3,800 channel AEs and SEs. This is a really important point, all right? Because we didn't want to be constrained at the rate in which we hired personnel in order to get our products to market as efficiently as we want to. We're doing that as fast as we can and maintaining quality, but marrying that with a great leveraged model with our channel. We'll go into that in more detail.

The three growth markets that Dietz hit on are really, really critical, right? There's a very big market just going out and cloudifying—I know it's not a word; it's Dietz's fault—but the enterprise. Going and giving the cloud attributions of simplicity, the efficiency, the performance that people are after in going to the cloud in a solution that they could take advantage of for all their existing applications, while moving where the puck is going, which is where the cloud is driving architectural decisions and the next generation data sets that we spent a lot of time this week talking about and we're very excited to see the early results on with FlashBlade.

Then we've built out a complete infrastructure. In order to do a B2B enterprise selling motion, you need to have account managers and sellers, channel account managers, alliance managers, technical sellers and SEs, inside sales, field marketing resources, and so over the last several years we've been building that out. We've got a nice foundation that we can scale upon going forward.

Focusing and prioritizing is one of the most difficult things to do when you've got a \$35 billion TAM, arguably one of the largest TAMs in B2B tech. Figuring out what the sales plays are going to be, how you're going to organize and segment your sales teams, and then ultimately deliver messages and value propositions that are compelling to the industry vertical that you're selling to, you need to do all of those things to be able to maintain focus.

So, the key plays that we have here are really centered around new stack applications, workloads. This is where the Bank of America that I—I don't know if the analyst is in the room who sort of did that, but as we look at the impact of AI, this is a 78% CAGR, multi-billion dollar, \$5 billion market opportunity that we're going after and we're just in the early innings of. So, focusing on those applications, not just in the autonomous driving cars use cases and the sexy stuff, but also agriculture, right? You know, 10% of the lettuce that's manufactured out there is actually done through robots, and a bunch of other examples that are in mainstream USA are happening every single day. So, we want to go find those applications and find the application developers that are looking to transform the businesses.

Analytics is obviously high on the list. Being able to get access to more data fast, to be able to transform your business is obviously a core part of what we do, and then the basics around Oracle, SAP and SQL. There's a very big business; you saw that wheel that Dietz went through? Half of it are applications that are going to stay on-prem. There was an article from Cisco presented yesterday and they said there was a Forbes article that's out there that 93% of CIOs said they were expanding their private clouds this year. It's just too difficult to pick, lift and shift these applications to the cloud, and especially if we can deliver dramatically lower total operating costs and efficiencies and simplicity that they were to get otherwise; there's no real reason to go there. We want to go cloudify those folks in the business applications, virtualization and BDI use cases.

So, our sales teams are segmented very standard around B2B software companies, cloud enterprise and the largest federal agencies. We have named account managers that are focused on those. Same thing

in the commercial and healthcare space, so you have dedicated capabilities, particularly with PACS imaging and other use cases around the healthcare segment.

Then the corporate and sled space, this is smaller organizations, sub-\$500 million where you want to get efficiency. So, we're working on an inside-out channel model where we can sell the products over the phone, expand them and fulfill them through out channel partners. They get the renewals and continue to upsell and work from there. So, driving efficiencies and leverage, and in this market continuing to push our resources further up the chain to manage those relationships.

As you look at the repeat purchase metrics, which I'll say here in a second, you've got the 4:1 and 5:1 and even 10+:1 up in this area, and so we want to make sure we go hunt in those areas, delight those customers and take advantage of the repeat purchase metrics and the land and expand model. Then, across the segments, it's very horizontal. Everybody had storage for block. Everybody has storage for unstructured. They need a FlashArray, a FlashDeck and a FlashBlade. We've just got to be able to create the infrastructure and place to go qualify those and speak to them in their language with their business challenges.

One of the key messages that we delivered today was that we can't do this alone. Storage is the bottleneck; it's a physics problem. It's at the core of applications not being—and analytics not getting the insights that they want, so unlocking that's key. But you need a whole ecosystem of partners to be able to integrate with and partner with to be able to go make that happen. So, 100% channel through our solution providers, global systems integrators, distributors globally to be able to get that reach. Application alliances and infrastructure alliances allow us to provide more comprehensive solutions for our customers.

We announced expanded partnership with Cisco this morning in terms of doing deeper integration with their product lines, deliver that NVMe fabric to the host to be able to enable this massive amount of parallelism that's going forward. We announced a deeper integration with VM, allowing customers to have high availability for archive and backup capabilities that are out there with their 100,000-plus customers and nearly \$1 billion in revenue projected. So, we've got great partners here that we can provide a best-in-class solution for and we're going to continue to do that.

By virtue of being 100% channel, and this is different than other storage enterprise companies where they take the top and you have direct relationships; our partners have an ability with Pure to go deliver them something disruptive and better from a technology and business model and customer experience perspective, but they also don't have the constraints of not being able to access those based upon sort of the go-to-market distribution models at the larger players. So, that in turn gets us a lot of loyalty and net new customer acquisitions. So, 70% of our new customers that come through are identified by our channel partners and brought in; that's great leverage.

The selling partner AEs, these are the actual number of partner AEs, so Pure-certified salespeople that have transacted with Pure over the years. So as you look at our 1,000 people +/- in our own go-to-market organization, we want to have a multiplier effect, and so to have 1,600 sellers that are beyond that and 1,200 accredited partner SEs gives us that 3,800 person sort of salesforce that's distributed. We're going to keep adding to that. Our partner model is not to go wide; it's to be more of a scarcity versus and abundance model but within those channel partners is to get loyalty in check. We've managed partner ecosystems for many years and the long tail of the partner eco system is expensive and you don't get a lot of return on it, right? So you want to focus your resources on the front and focus on loyalty and focus on programs and focus on accreditation versus trying to have 10,000 partners. We don't want 10,000 partners. We want 1,000 or 2,000 of the very best.

That's translated into some nice trajectory for the business. You all know this but 124% CAGR on customer growth, right, Q1 2015. Over the last four Q1s continuing to grow nicely. This is not about having hundreds of thousands of customers, candidly. We think we can build a very, very big business and we'll walk you through the math—Tim will later—with 10,000/20,000 customers that are enterprise,

that are cloud, that are the next generation organization. We're going to be thoughtful on where we go acquire customers and go get those leverage customers that repeat, but nice trajectory there.

A lot of what happened and what Coz—he founded the company—he wanted to deliver the Apple of the data center, if you will, right? We wanted to deliver that cloud-like simplicity, zero-touch admin, that's out there. As you talk to our customers this is generally one of the things that really surprises them. It's why they buy it on the front end and if it's half as good as we say it is, it's still better than what they've got, and I think that's reflective of what the high tech average is for net promoter. A 16 is what we're competing with. People aren't in love with their storage, aren't in love with their infrastructure. That's one of the reasons why they're moving to the cloud. We think you should move to the cloud for the things that make sense but we think there's a lot of things that we can do to be able to make that cloud-like experience on-prem and then provide you a platform that helps you manage it, so you can get the best of both worlds as you go forward.

So, this customer satisfaction rating we just actually got this certified again. We go through an annual process where you talk to hundreds of net promoter—pardon me. Matrix as a third party audited the company, that goes out and talks to hundreds of customers. Generally you'd expect at scale this would start to come down, especially when you're operating at such high levels. You know, holding it with the thousands of customers that we have is something that we're very committed to as we continue to build and scale the business, it's a key differentiator and it's the reason why we get the repeat purchase rates. So, great customer strength in the enterprise: 25% of the Fortune 500, 35% of the top 20 SaaS companies, 500-plus companies that we've shared with you in the past, and 25% of the revenue that comes in from the cloud. This land and expand metric are holding, right? So we're data people. More beyond just sort of enabling it but actually looking at it from a go-to-market perspective and these expand numbers are holding on the top end and across the board. So, very encouraging.

These are the logos and industry verticals that we can share publicly, but obviously with 3,350 customers that are out there—I don't know if you guys caught it but that first slide that was *The Economist* cover, talking about the next generation companies where data is the new oil, we have the majority of those customers; modern companies architecting for the future, leveraging Pure Storage. Right?

As we look at this mix, we're really proud of the companies that we can go help enable across multiple verticals, but I think as we bring these stories to light and our customers are willing to share them, I think it becomes more clear the kind of company and trajectory and why we're growing as quickly as we are.

With that said, we have a bunch of white space to go attack, right? As we look at the Global 2000, we've done a nice job. These are slower sales cycles, take a little bit longer to get in, more entrenched customers, the financial levers that they have to play, but one of the assumptions that we have as you look statistically is the disk business that's still out there, over the next three or four years that's going to turn over. There's 70% of that install base that's going to turn over, and when you come up to a depreciation cycle, you come up to a specific business problem that the legacy technology can't solve, those are buying opportunities. Once we have that buying opportunity, then we rely on our TCO, our simplicity and our customer experience to expand from there. So, a lot of focus; Kev will talk about it where we're placing on the G2000 the global companies and those applications that are next gen. So, 12% penetrated gives us a whole bunch of room to go hunt from here.

Sales productivity, this is a B2B business, so sales productivity is another key metric that we look at. This is a pretty interesting stat. We shared it a bit in previous calls but let me walk you through this because the key takeaways are here. The longest serving cohorts, number one, are well above our target model. Target model means when we share with you what we're going to go do and how much profit each sales rep is going to be able to generate, we've already proven that we can get there and beyond.

When you look at the newest cohorts that have joined us, which is number two, they're actually ramping faster than the early cohorts, which are already at the target model, and it's the majority of our capacity.

So, our 2016/2017 classes represent 60% of the salesforce that's been here for less than 15 months. So you have a double benefit in terms of sales capacity that's coming online to go drive continued growth.

We're very excited about these trends. If we saw any sort of hesitation here, we'd be looking and guiding differently than what we are, but when you see sales capacity coming on—now, this is fairly obvious, right? You've got a brand. You've got 3,350 customers. You've got a market that's maturing. You've got competitors that are going through massive transformation through acquisitions, divestitures, et cetera. All of that creates a window with our differentiated technology and value proposition to go get in the door. Once we're in the door, with all the innovation that we have coming out, we've got a repeat land and expand business. So, the sales capacity and the cohort analysis that we're doing, trends are incredibly positive.

With that, I'll hand it over to Kev to go through our top three priorities and go-to-market for 2017.

Kevin Delane:

Thanks, Hat.

David Hatfield:

You're welcome.

Kevin Delane:

As Hat has mentioned, I have been in the industry for over 20 years. I've been at—not to correct my boss, but I've been at Pure now three and a half years. I have been part of, I have built and I have led some of the best field technology sales teams that you have ever seen in the market and what we've assembled here at Pure is truly one of the best teams I have ever seen. In fact, to use a boxing term, I would take our field sales team and put them up against any field sales team pound for pound in the market today.

Now, drive and execution starts with common goals and objectives, and our three objectives this year are to build the pipeline necessary to get to a billion dollars in sales; to build and grow and innovate our partner ecosystem to be best-in-class across the whole market and different from the way it's been done in the past; and to transform our sales team into a platform selling execution company.

So, when I look at pipeline, we start driving pipeline from Day 1 at our kickoff at the beginning of the year. We drive pipeline cross-functionally and we do it on a quarterly, weekly and monthly basis. Every single week our teams review the pipeline and we are focused on growing it and hitting our goals and objectives. We do that from our sales team at the field level, our inside sales team, our partners in the market in our channel, and also our Marketing department, all focused on one common goal. I'm confident that we have built not only the team to do it but the metrics that will drive us to a billion dollars this year.

From a partnership point of view, we have a unique approach. Hat mentioned we're 100% of all of our business is transacted through the channel. Our team from a field level, the number one thing that our partners say back to me is you're very easy to work with. We allow partners to build their business. We don't control them on what they can do, can and can't do. We allow them to do what they're great at at a field level, and we allow them to build a business around our overall platform and technology.

Our partners are not only our VARS and our distributors, right? But we have a huge business in our overall global systems integrator, and then we have built out technology alliances that are different from the market in the past, and we leverage that to get reach to the market and allow us to get to our goal of getting to a billion dollars this year.

Lastly from a platform point of view, our FlashArray product, as we have seen from the growth, is best-in-class and it's the best block product on the market today. It has helped the overall market cloudify, as we use from Dietz's word, but then also Hat mentioned the legacy environments that are out there today. So if you were there this morning, we talked about our Tier 1 and what we have done in Tier 1 for functionality. Seventy percent of what's out there in Tier 1 over the next four years will actually be refreshed, and our FlashArray gives us a great opportunity not only to have the cloudified applications on us but also go after the refresh and help transform that old environment to the new way of doing it.

Our FlashBlade product, which Matt will walk through, we have a huge advantage in, one, we've built a channel, 3,800 people that we have out there between us and our partners, that we can now have educated, ready and they could sell our FlashBlade product.

Secondly is we have 3,300 customers now. So you'll see today from our customers we have up there, two out of three are using both platforms. So we have a huge opportunity to scale and to grow FlashBlade faster than we have on FlashArray, and to achieve our goal to get to a billion dollars.

Lastly, converged infrastructure is a big market for us overall. Having a best-in-breed converged infrastructure with Cisco and Pure allows us to meet the needs of all our customers from a converged side.

So, again, I've been in the market a long time. I know what our goals are and we have to accomplish this year. We have to get to a billion dollars. I'm confident in our team's ability to get there, to not only build the pipeline, to continue to the partnership, and we will transform and grow across all three platforms.

Thank you. Matt, I'll turn it over to you.

Matt Burr:

All right. Thanks, Kev. I didn't wake up as a 10-year-old boy and say, "I'm going to go sell a billion dollars worth of storage." I don't think that was my dream as a 10-year-old. You know, to see the go-to-market engine evolve and develop has been a really awesome experience, honestly, and to have Kev come in and lead it is a great thing. To go from zero to \$700 million plus in five years is a pretty meaningful achievement. You're watching a market go fast. You're trying to find ways to find leverage, and you're trying to change a model that has been in existence for better than 20 years. So when we think about leverage in our business, I sort of sat down with Coz, our founder, in 2009, and he said, "Hey, we need to rewrite the I/O substack above Flash like we wrote it above disk at Veritas. I said, "That doesn't sound very interesting. Like, I don't want to do that." I'm not going to bore you with the technical specs of that but I am going to tell you about what he said in the follow-on conversation.

His point was that we can actually revolutionize the way customers, the people that are here and what's reflected in that net promoter score, feel about how they interact with this particular piece of infrastructure. That's probably our greatest accomplishment to date and being able to attract that leverage and in the go-to-market engine that we've been able to build is largely dependent upon, or was dependent upon at the time, on how we were able to partner with people like Cisco. So when we think about what the next generation of converged infrastructure really is, you're talking about building things that were replacing something that was in large part the size of a refrigerator with something that's the size of a microwave. That's performance density in a data center.

We have a particular customer up in Minnesota that uses drones to fly over fields to assess crop yield. They're looking at how's the irrigation working in the fields, what type of bugs might be in the field. They're creating a tremendous amount of data that they're then going back and using FlashStack to churn in new application environments and get results on whether they tell that farmer to take the field down early or they pull it in faster. Right? That's the type of thing that we can work with Cisco on, best-of-breed BladeServer platform, big pipes, big networks, big storage on the back end. When we think about what the next generation of these types of analytics are, you move towards the world of massive

parallelism, both at the server layer, and massive parallelism at the storage layer, but that's not enough. There's a network expansion opportunity that has to happen in there as well and what Cisco sees in their relationship with Pure is the ability to move from 40 gig networks to 100 gig networks, which is a huge network opportunity refresh for Cisco. So their investment in FlashStack has largely been focused on about 7 CBDs and 15 or so reference architectures leveraging ACI, UCSD to provide that CI-like experience, and then moving through joint partner programs where we meet in the channels. This is another significant point of leverage.

So, our top channel partners, Dave talked a little bit about trimming the back end of the channel partner. It's amazing that we're doing this inside of a shipping container because the echo up here that you guys, I'm sure when you were talking, it's like you can hear yourself seven times. But the partner ecosystem and the leverage in that partner ecosystem and the drag that we get from being associated with Cisco is really material in how we go to market.

If we think about what those partners are looking for, and these are large, billion-dollar corporations, they're looking for consistency in their channel partner program. So when you talk to the CFOs of these partners, they'll say rather than giving me four or five extra points in an individual year, I want consistency over 20 years where I know what I'm going to have, and we've modeled our program after that model. It's been great for us because those partners that are true north partners and loyal to Cisco have also become true north partners very loyal to Pure which bring us into very big markets.

So from a FlashStack perspective, again, 70% year-over-year customer growth. We have 30 FlashStacks in Cisco's labs globally, so that means that Cisco's sales teams can bring their customers into those labs and actually demonstrate Pure Storage technology in the FlashStack in a Cisco environment. That actually has two benefits. The first benefit is the customer gets to see the technology. The second benefit is we're cross-training a much larger salesforce than we have. So that group of people sits down, that aren't storage people—this is really, really important. They're networking people. Cisco does not have storage depth in their overall employee base, but when their teams sit down and figure out the ease of use of storage, of Pure Storage, they can go to their customers and they become an extension of our salesforce. So it's not just a, "Let's meet for coffee and talk about who our customer lists are," it's technical people getting their hands on cool technology and wanting to talk about that to their customers because it's simple and they can.

So, that 10-year-old boy who didn't get up thinking about selling a billion dollars of Flash storage probably also didn't get up thinking about this. But the 10-year-old boy that's inside of me now sees that this is a bigger opportunity than what the FlashArray opportunity was. That's not to say that the FlashArray opportunity wasn't big. FlashArray opportunity is enormous. It's a \$35 billion TAM combined where 70% of the refresh for storage is still to come, right? There's a lot of space left for latency-driven locked devices, but the world is moving towards throughput. The world is not going to be focused on latency going forward. Throughput is what's going to change our lives.

So, I'm 44 years old. I have a 3-year-old child and I have a 1-year-old child. Neither one of those kids in my mind will have a driver's license. They might have it just because it's a novelty to have, but the reality of it is self-driving cars will be—autonomous cars will be what their future is.

But that's only one example of the types of data sets that are going to be running through this stuff. At the end of the day, how we look at the data stream and how we look at the data path, and being able to process more data in real time is what's going to enable things like a bank in Arkansas to do analytics on loan processing so that they can shorten their loan cycle times by a week.

So, this isn't the most technologically advanced customer that's thinking about AI, deep learning or AI and machine learning. This is a customer that realizes that it's important for them to have it because they will be able to compete better and it's a regional small bank.

The opposite example of that is a major teaching hospital in the Mid-Atlantic where they're working with Welch Allyn and other component providers in their overall hospital and teaching hospital environment to put sensors in every device, every blood pressure cuff, every EKG sticker that goes on your chest. They're going to monitor the drip rate of IV bags and they're going to collect all that data and it's going to stay on-prem and it's going to have a dual purpose. Part of it is going to go to the researchers, to focus on research; the other part is going to go on accuracy of staff. So there's a little bit of an uncomfortable thing here to challenge are our nurses good and are our doctors good, but it's a multi-billion dollar opportunity in the industry to avoid hospital errors and mistakes. Epic has made a big business on the ability to scan bar codes instead of doctors writing prescriptions because doctors writing prescriptions lead to misdiagnoses, lead to misprovided drugs which can lead to bad stuff, and those are liability claims.

When we look at the FlashBlade business—and that's just a commercial on what the potential is for the market opportunity—that market opportunity is nascent. My credit card got stolen about six weeks ago. I called USAA. Anybody a USAA customer? Best ever, right? Like USAA is great. So I called USAA. The gal answers the phone. I said, "I think my credit card was stolen." She looks at my account. She said, "What day—what information is leading you to believe that your credit card was stolen?" I said, "Well, here's this transaction," and she said, "Yup, your credit card was stolen." Minute. "I'm going to give you the money back; it'll be approved on Friday." I said, "Wait a minute. How are you going to give me my money back on Tuesday but it's actually going to get approved on Friday?" She said, "Well, I see that you've made a physical credit card transaction at 12:56 in San Francisco. Apparently, you also made a physical credit card transaction in Miami at 12:58." Her words, not mine. "And unless you've invented a teleporter, you could not have been in two places at once." I said, "Well, how do you know that I wasn't in Miami? What about in San Francisco?" She said, "Ninety-two point six percent of your physical credit card transactions happen in San Francisco."

So, it's pretty awesome that she was able to get that data and empower a customer service representative to make that decision. That was pretty cool. But the next instantiation of AI and analytics involves real-time stream analytics where we have a sliding window over the top of the data path where we don't get to detection, we get to prevention. Those are the types of things, those are the types of applications and those are the types of businesses that will be driving the next generation and that over \$6 billion in business that we talked about earlier.

Back to the go-to-market piece. So we have this new product, right? Everybody has heard of it. What's Pure going to do with their new product? How do you get it into the mainstream? As Kev alluded to, we have 3,400 sellers between us and our partners; how do we go train them and spin them up?

Well, ultimately, we want to go compete in the legacy space that exists today, Oracle data warehousing where you have these monolithic refrigerator-like arrays that suck a lot of power, and by the way, to give you a sense, many of those installations are larger than these three shipping containers. So, we have already replaced the equivalent of two of these shipping containers with four FlashBlades. So, imagine these two rows are completely full, get's replaced with four of them, right here. Visually that means something. So when we start talking about what potential is there for us to slide this into our sales team, so we're going to take our channel, we're going to take our people and we're going to go look at the cross-section of where we can go sell the FlashBlade product and tool up the people that we have to go sell that; 80% rough (inaudible) our business. We're going to take our FlashBlade teams and we're going to focus them on these next gen AI workloads, some of which Dietz referred to earlier as the world's largest supercomputer.

So, you know, when we think about how do we get to that FlashBlade growing at 2X faster than the FlashArray business did, it's actually not that hard to get to that number when you think about I've got a whole bunch of sellers that I can just go turn on instantaneously with a little bit of training, because that's already there.

The FlashArray business grew faster than any infrastructure, storage infrastructure company in history to a billion dollars. We know how to do that. We made mistakes along the way. We learned from those

mistakes. We have adjusted. We have course-corrected. There will be things that we don't know about the FlashBlade business as well, but we'll course correct just like we did here, but we don't have anywhere near the hill to climb that we had when we had the FlashArray business. So having that go-to-market engine in place, leveraging the existing sales team, using our marketing engine, demand gen that we have in place, using our ISO organization, focusing on data science, using our own data science team, which uses FlashBlade to do targeting to understand which of our customers in our install base are going to have the highest propensity to buy is a core focus for us, and of course we'll go enable our channel partners and we'll bring ourselves into new markets, expanding the TAM that we have today through the healthcare and PAX market, which is a very interesting market because there's only about five players and that makes it easy to go focus on those five players and we think that can be a really big space for us.

That's the piece on the FlashBlade business. Thanks for the time. I think we'll move on? Okay. Back to Dave.

David Hatfield:

Just in quick summary, talking about go-to-market executions, A, having a platform that we can go (inaudible) wide and deep customer base, 3,300; adding 300 or so. Sorry. Can you guys hear me? There we go. I'll go back.

So, three key points on the go-to-market execution. First and foremost, a very strong go-to-market platform that we can rinse and repeat all the innovation that's coming out of R&D to be able to go into adjacent markets. Second, wide and deep customer base that buy more stuff. They're delighted with what we've done. They like the model that we've provided and the customer experience is orders of magnitude different than anything they've had. So, go mine that and take advantage of that. Then, sales productivity ramping. Taking a look at how do we bring people on and ramp them faster to drive more profitability to the business while hitting the top line objective.

So, these three things, pretty simple, really. We've proved that we can do it, and I agree with Matt. It's much harder to get from zero to a billion than it is to get from a billion to two billion, if you have the products and you have the go-to-market engine to be able to go capitalize on them and we feel like we've built both of those.

So, without further ado, I'd like to ask Kevin and Matt to come back up and I think we're going to go through the customer interviews.

Matt Burr:

All right, can you hear us okay? No. How about now? All right.

David Hatfield:

You're good. Just put it on.

Matt Burr:

I can do it the old school way. All right. Andy, do you want to start? Could you introduce yourself, your company? Tell us a little bit about the infrastructure that you manage and then what's your team's charter?

Andy Nallappan:

Okay. My name is Andy Nallappan. I'm the CIO for Broadcom. Broadcom is a high tech semiconductor company, and we used to be Avago Technologies. After we bought Broadcom we became Broadcom

Limited and we had number two in the high-tech industry there. I'm responsible for corporate as well as the R&D infrastructure. My team manages about 10 huge data centers and thousands and thousands of servers and petabytes of storage. Our charter is, you know, as you all know that our company has been growing in the last few years. We've been doubling our revenue and we've been focusing on integrating all the acquired companies and containing the cost and the improving our bottom line.

Matt Burr:

The last time I checked, doubling revenue was good. So, nice job. When you—you've been a Pure customer for two, two-and-a-half years now. You just became a FlashBlade customer. Can you tell us a little bit about why you made that choice and some of the results that you've seen through your Pure relationship on the FlashBlade side?

Andy Nallappan:

Within Broadcom and IT I have this mantra. Is it faster, better, cheaper? With that, Pure aligned with all the three. Fast is the faster, the FlashArrays and FlashBlades is 50% faster than our traditional storage there. The high IOPS and all. We have a high-tech NRAM and there we have our chip design and tape outs; it requires a very fast performance. That's the one thing.

The other one is the operating model is a completely outsourced model and I would like to have a solution which is easy to manage, simple to manage. It's better there, and so Pure Storage is much better than what I have there.

With the total cost, it comes to there—because of the high density, and we all (inaudible) and easy to maintain, so the TCO is cheaper there.

With my strategy, the Pure offering is aligned. That's why we've extended from FlashArray, which we went about four years (phon) back and with 2 petabyte there. Now we went to the FlashBlades and that's where (inaudible) 80% of the workloads where we have with our high-tech chip design workloads there with the non-structured data where we are now experimenting for the scratch space where it needs the very high speed performance there. Even minutes and hours improvement, it saves a lot of money and productivity there. We got very good results with the FlashBlades and we're really happy.

I lost this mic. So we are looking forward to expand FlashBlades across the board.

Matt Burr:

Okay. Just to sort of recap that while we're waiting on the mic. It was 2 petabytes of data on the FlashArray plus 80% more to go which will be addressed by the FlashBlade space.

Andy Nallappan:

Right.

Matt Burr:

And I think—if I get this wrong, correct it for me, but one of the advantages of using FlashBlade relative to the scratch space in Broadcom's use case is developer time to productivity. So, these are—imagine if you can—if someone were to take apart your desk at work and you couldn't work for a little while, and then you got back and you still couldn't work, and then you got back and your stuff still wasn't together, and the Internet wasn't connected, whatever it is, you lose four hours of time. Not only is that a bad experience for you—unless you were doing something fun, I don't know—but ultimately it's a bad experience because you're losing four hours of productivity.

What we talk about when we're talking about scratch space is you go from something that might be a six or a four-hour window to rebuild that person's environment in order to go test more code, run more simulations. Shrinking that down to one to two hours is developer productivity.

Andy Nallappan:

And I have 13,000 engineers, so 80% of my employee population is all engineers.

The other one is it's very hard to get downtimes. Even weekends are not good for us because we run long jobs and they're running all hours and days, so the non-disruptive upgrades, maintenance is very critical and that's another factor there; it reduces the downtime.

Matt Burr:

Talk a little bit about security. You hit faster, better, cheaper; we hit those objectives, but I know in your environment security is important. We have some functionality there. What do your users expect out of you from the security side and how have we helped you?

Andy Nallappan:

In our company, we deal with the three different kind of IPs there. One is our own IP we develop, that's our core. Another one is our customers' IP. We do design wins. We do ASIC chips there. We have customer IPs there that are critical, is core of them. And if you work with (inaudible), so we use a lot of their SDK (phon), their IPs to get the tape outs there. So, it's very critical to protect all the three IPs there, our IPs and our customers' and sublayers there.

So, our engineers, users all expect the data to be encrypted at rest. It doesn't matter whether the data is within the firewall or it is outside the firewall, they expect all to be encrypted there. Even sometimes our customers, even auditors then require us to be encrypted. That's the one thing that Pure (inaudible) the dry level encryption that really helps with the (inaudible). Security is critical for us.

Matt Burr:

Andy, when you think about 30,000 engineers ...

Andy Nallappan:

Thirteen thousand.

Matt Burr:

Thirteen thousand, excuse me. Thirteen thousand engineers, and the compute and infrastructure that's required to keep those people productive, how do you view on-prem as a part of your cloud strategy?

Andy Nallappan:

All our workloads, it's very different, very unique. It's very compute intensive and storage intensive. So all of these workloads have to be run on-prem there, but if you imagine there every day we run thousands of jobs and generating a lot of data. It stays in the very high performing storage, and once it's all become cold and glacier, it goes down to our archive storage there, and every year the files keep flowing in the data there, so I keep growing there. So my strategy is keep all the on-prem workload which requires the high IOPS there on-prem there and at the data centers, and all the archived data is moved to the cloud there. Slowly we started that initiative there and all the muted (phon) data, we want to move to the cheaper storage there, and so the mobility is very critical for that. So with Pure helping us to be on-prem and move that data to the cloud when it become muted it's very critical that's all in with our (inaudible).

Matt Burr:

Thanks, Andy. Justin, how about you? One of our longest tenured customers, tell us a little about yourself, your environment, your charter and your business.

Justin Stottlemeyer:

Sure. Hi, I'm Justin Stottlemeyer. I'm currently a distinguished engineer for Intuit. Can you guys all hear me now?

I'm Justin Stottlemeyer. I'm currently a distinguished engineer for Intuit, also a storage and compute infrastructure architect. Formerly I was an early employee at eBay, PayPal, Facebook, ShutterFly. You still can't hear me? All right.

Matt Burr:

There you go.

Justin Stottlemeyer:

Formerly I was a early employee at eBay, PayPal, Facebook and ShutterFly, and I've currently chosen Intuit to try and make a difference. Today, I've completely replaced the disk-based infrastructure with an all Flash one for QuickBooks Online over the last two years; partnered up with Pure. That made such a massive difference to our business, we've had customers that have literally said, "It's finally working the way that it should for the first time ever from an end-to-end perspective." I've actually had engineers on the inside ask me, "Well, now that it just works, what do I do?" So I move them up to Stack, right? We're doing DevOps on some of these places; it makes a big difference.

Sorry, what else did you want me to cover? That really tells a pretty big story, right?

Matt Burr:

Let's talk about ...

Justin Stottlemeyer:

Let's go back and forth a little.

Matt Burr:

Let's talk about QuickBooks specifically, right? QuickBooks is—I think everybody knows what QuickBooks is, but obviously end user performance is materially significant. If you could dive in a little bit to the direction that you guys have taken in order to bring the QuickBooks experience up and what that's meant for the business.

Justin Stottlemeyer:

Sure. In certain cases we actually had applications that had—you guys have used search, you've used Google. We had applications that had a four second search SLA, and that just drove me insane. You know if you've used Google, you've lost attention about a quarter of a second in by the time it hasn't responded.

Matt Burr:

WebCrawler used to actually measure that and put it up there.

Justin Stottlemeyer:

Yeah, that's right. So we actually switched a lot of the search back end as one of the first targets to move and the CTO literally wanted to give me a hug when it responded in 0.2 seconds instead of 4.

Matt Burr:

I bet. That's probably ...

Justin Stottlemeyer:

That type of customer impact is really difficult to quantify as to what it means to your business, but clearly happy customers is what we all want to deliver.

Matt Burr:

What's the future for QuickBooks at Intuit?

Justin Stottlemeyer:

The future for QuickBooks at Intuit strategically today is actually pointing towards the cloud. That's the strategic direction for Intuit today is very AWS based. As the infrastructure and compute guy, one of my jobs is to help them get there. The second of that is to actually provide value to the business, and as I look at that long term, I have an awful lot of very hard questions for people to answer around what the economics and performance of the cloud looks like. Best level SLAs from AWS, you can all look them up. We've got customer that expect sub-millisecond response times, and if you aren't delivering those you don't have customers.

Matt Burr:

Okay, thanks. How about you've dealt with us for a long time. We pride ourselves on being very customer focused, but you've had to have some experience with our customer support organization and what's that been like?

Justin Stottlemeyer:

Just yesterday as a matter of fact.

Kevin Delane:

How'd we do?

Justin Stottlemeyer:

Very well actually, and it ended up not technically being a Pure issue but they did help us troubleshoot it all the way down to the wire.

Kevin Delane:

That's great to hear.

Justin Stottlemeyer:

I've been a Pure customer since very early on. I remember when John Hayes and Kixmoeller came and tried to pitch me on Pure Technology maybe six or seven years ago. I've picked a lot of new technologies over the last several years and Pure is one of the very few that made the cut, and I've been very happy with it today.

I can probably count on three fingers the amount of interactions I've had to have with Pure Customer Support over the last five years, and they've all been very, very good.

Matt Burr:

Great.

Kevin Delane:

Justin, one last question. We believe that performance density is material to the next generation of compute. I would say it's material to the current generation of compute but it becomes even more important over time because you don't want people managing infrastructure; you want people interfacing with your business units and driving application performance.

Justin Stottlemeyer:

Absolutely correctly. Everything old is new again.

Matt Burr:

Yes.

Justin Stottlemeyer:

So we actually went through a period—I've been doing this 20 years. We've been going through a period of storage aggregation, storage disaggregation, and so if you look at things like Hadoop, machine learning, AI, deep analytics, I've got a 10-petabyte Hadoop cluster today and as I look at expanding that or building that out in other places, of course my storage is not linearly scaled with my compute so I now need a lot less storage, a lot more compute. How do I build that in?

I was looking at the numbers that you were talking in the data center and I didn't get 100X, I only got 10X back, but that was with my Array, not the FlashBlade.

Male Speaker:

Sorry.

Justin Stottlemeyer:

It was still pretty useful. We actually got 10X back when we did the FlashArray just two years.

Looking at FlashBlade, I'm looking at something like 20X, 30X. On my deployment model today for Hadoop, when I'm starting to look at storage disaggregation, I'm sitting at roughly 30 terabytes available per 6U, where on a FlashBlade that's ...

Kevin Delane:

A little bit more.

Justin Stottlemeyer:

A little bit more, yeah. And the bandwidth is astronomically higher.

Kevin Delane:

Good. Matt, how about you? I was part of the early sales cycle. Came out and met with you. I'd love to hear from your side why did you choose Pure? You had plenty of options out there. What's been your experience? Besides excellence sales execution.

Matt Harris:

Oh yeah, of course. Two and a half years ago we'd been through a really trying to look very, for us long term; normally it was a few weeks away. As an F1 team we put a brand new vehicle out on the track every week, so for us to turnaround and start looking at something that was going to be two, three years away, it still doesn't feel like it was actually two and a half years ago that we did the initial purchase.

Ours, we started off the discussion around what storage to do. Chris, who's here this week with me, we very quickly learnt that there was no point worrying about performance. There was no point really worrying about necessarily the cost aspect of any of the solutions. What we started to look at was actually the people we were dealing with. It was actually about the support services. It was about how we interacted with the sales people because sometimes they can be a little bit dull, if I'm really honest, and they can be difficult and they can be challenging.

But also the support side; I've got a small team, so these guys are talking about many hundreds of people that are probably in their teams. Realistically, I had one guy who looked after storage and he spent days and days and days looking at storage. What we worked out with the FlashArray product originally was that actually we could get to the point where he didn't have to look at the storage, and things like Pure1, they were actually telling us what was going on. Pure had then turned around and given us a call to say a part and a person is on the way and we're like, "We didn't even realize something was wrong," and it was because of the predictive analysis within the tool, within your software, and this is early door, so this is two years ago, and it's even better to this day.

So, for me and Chris, it became simple. We were looking at a product that we could turn around and stop my guys having to do the mundane. They didn't have to think about the day-to-day. They could now be worrying about the advancement of the team, how it could improve your life. How it could make you faster or better or the results better, and for us that's critical. Everything we do is about data.

If you think about the two cars that go round the track, the several hundred people that are in the factory looking at data all the time. They want the same access, whether they go to it when there's 800 people looking at the data or one person. In the old days, pre-Pure, that was the difference between somebody taking two minutes to open some data or nine minutes. If you imagine, during a race weekend somebody telling me that I had to wait nine minutes for an egg timer for a piece of tape to load, you can imagine what sort of look I was getting from some of my senior management.

But the stranger part of that whole exercise is we moved across the data that we knew was a problem that had contention, but we also moved across other data sets because it was nice and easy. The longest part was actually taking off the old storage. But we took some SQL based—it's SCCM for anybody that's technical, but don't worry about it. It used to take three minutes for a report to run, and we didn't even think about it. We just moved a couple of SQL servers across and we were down to four seconds for the report. We did nothing. No recompiling, no reinstallation, no upgrades. Just the storage change meant that I could go and get a report. Rather than going, "Can I be bothered to wait?" I could make a decision about five seconds later.

I could talk all day, so I'm...

Kevin Delane:

I think my follow-up to that, Matt, would be there's probably no more technologically advanced sport on the planet. So much so than within Formula One, the technological aspect is actually managed. It's actually part of the rules, not just the way that the cars drive from the air frames or anything like that. We're talking about actual when you can take data from the car and when you can't, and when you can make those adjustments and when you can't. The windows to make tweaks to the car from the information that you get are really small. This is analogous to how much faster can people be enabled to make decisions, and in the next regulation in Formula One—well, not the next regulation in Formula One, but the one that's coming—is how long will we allow humans to make those changes versus machines.

We're not going to say that Mercedes is making machine changes today, but, I mean, let's call a spade a spade. You haven't been on the podium that much this year, or at least not at the top of the podium, and this is a perennial powerhouse in Formula One.

Dave Hatfield:

Hold it. Hold it. Hold it. They are number one of the contractors that we talk to.

Kevin Delane:

Yes, they are. Yes, they are. They're one (inaudible). It's been an imperfect season for the Mercedes team who expects perfection. This week was perfection. They went one, two on the podium. Tell us how you got to one, two on the podium?

Matt Harris:

Yes. Coming just the old conversation there about the Championship, we've obviously had three years of just immense success. We still look to our data. We still try to understand what was going on, but it's harder to spot an improvement when you're the number one. When you're looking for differences that you could make, if you're so far ahead, actually staying that far ahead is very hard. The Monaco Race weekend, which unfortunately Hat was there to see...

Kevin Delane:

It wasn't that unfortunate. It was a good weekend, but it was a tough weekend for you.

Matt Harris:

It wasn't the best weekend for us. Without sounding too corny or too cheesy, Pure was fundamental in allowing us to actually do something that in the past we wouldn't have been able to do. During the week after Monaco—so, we normally run some simulations about the next event, and those simulations we, just say we can get 100 for a second for case of a number. We ran in the region of 400 reprocesses of Monaco to try and understand what had happened in Monaco. That was prior to us then thinking about the next event. That was to try and understand what went wrong.

In the past, we couldn't have done. The infrastructure behind the scenes, the storage, wouldn't have been out to handle the extra compute. I think it was the Cisco conversation earlier on the stage, we were talking about the fact that you moved the bottleneck. When we put Pure Storage in, all we did was we started to move our bottleneck to compute or memory or network around anything else. The thing now is that Pure, we don't ever have to think about the storage to get the performance, to get the analysis. We then turned around and did the same analysis, but for Canada, and we actually ended up running the

exact car for the race that actually came out of simulation. We would've been able to do that in the past because we couldn't have gotten the number of CPU cycles we required with our old infrastructure.

Kevin Delane:

Awesome.

Matt Harris:

So, you helped us at the end of the day get one, two.

Kevin Delane:

All right. Matt, so my one other question for you. You talked about ease of use of the overall solution in that. You also talked about how salespeople could be annoying, which is fine, I get it. But one of the things, what we've tried to do is modernize our financial terms and our evergreen approach. So, how did that come into play and what are your thoughts there?

Matt Harris:

Yes. If I'm really honest, when we first did the purchase back in 2015, I think it is now, I never really thought the evergreen thing was going to come in as quickly as it did for me because we originally bought the M70 Series. I think within about six months of actually purchasing it we were upgrading into the 450s, which I didn't really see coming Day 1, if I was really honest. So, we'd already taken account of that.

I actually think I stood on a stage talking about F1 and technology on a day where we were doing an in-place upgrade during the middle of a working day, and I actually think I spoke about it that day on the stage because I was a bit worried that I was going to get some phone calls.

Kevin Delane:

I think you were in the U.K doing it, if I remember.

Matt Harris:

Yes.

Kevin Delane:

Yes, downtown, yes.

Matt Harris:

Yes. But that point of view, there's two things there. There's actually that my team, we don't have to worry about what's going on because you guys were actually telling us the possible pitfalls with the upgrade so that we were able to look at what was going to be going on, what you'd done previously with a similar install base, issues you'd seen for us to make sure we went and looked in our install base and our applications. We were able to pre-check patching and everything else. So, when we did the upgrade during the middle of the day, actually you guys were so confident that you made me actually not worry about it. Otherwise there was no way I would've been on the stage that day.

But still to this day we're still using it. FlashBlade, interesting—I haven't really spoken about that much, so it's probably worth—well, everything we've mentioned so far was FlashArray. We very quickly got into

the FlashBlade with you guys. I think even last year I turned up and I was going to be calling Iridium; stood on the stage and it was only when I got there I was like, "What's this FlashBlade you've got on all these slides?"

But for us it's been a longer journey because we've got some various specialist workloads, but as of last night, we actually finished some of the first proper production runs of FlashBlade with our CFD environment. This comes back to a rule change that's coming up where we're very good as a business when we're restricted on what we can do, so CFD is one of those areas. People in the room, it's a bit like a wind tunnel in a computer, just to try and explain CFD. We got a 25% improvement last night on certain jump runs, purely by changing from our legacy parallel file system just to only a half populated FlashBlade. So it's not even like we've stuck a huge amount at it. It was just a production—it started off actually as a...

Dave Hatfield:

Just like the other half.

Matt Harris:

That's in discussion. Don't worry. But for us, a 25% improvement in a restricted area, money can't buy that. Now, I didn't say that to our salesman. But the thing for us is that you can't buy 25% improvement in F1. That's not something that you can spend money on. To get 1% or 2%, we've spent years and years and years optimizing processes, so for technology stack out-of-the-box, just here we go—and you guys are actually working with us to see whether we can improve that by a few percent as well so.

Kevin Delane:

Well, thank you. How about Justin and Andy, any comments on how it is to work with us from a field side, contractual side?

Justin Stottlemeyer:

Actually, yes. It reminded me, I think I'm actually upgrading Canada tax today, like right now.

Kevin Delane:

You look pretty stressed.

Justin Stottlemeyer:

Yes. Really worried about it. The thing I was thinking about was that my Twitter tagline is actually cheaper, faster, better, which is kind of kindred spirits there. I've really enjoyed working with Pure in general. Very easy going, very easy to get performance. Out of all a lot of applications that we haven't necessarily seen in other places, from our development stack, pre-prod, QA, being able to fail things over seamlessly, being able to upgrade things seamlessly is key for a company like Intuit, as you can imagine, right?

We've got certain timeframes of the year where there's a good 40% chance I've got your tax data somewhere. Whether you use an accountant or not, it's in my infrastructure somewhere. We have very key technology that has to be encrypted, it has to be protected, and Pure is our only Tier 1 technology in play today.

Kevin Delane:

Thank you. How about from your side?

Andy Nallappan:

Sure. On the two sides, one is the support, another one we talk about the innovation side. The support side, now I manage a lot of hardware and our server storage there. I've been a lot of escalations there, but one thing is there was no escalation. For the last four or five years. I've never had to call any one of these guys up. So that's why I prefer. I don't have to work with them, right?.

For me, the keep the lights on is important. A light is always on. I never have to call. I never have to call the help desk. It's very simple to manage and maintain there, even though we're outsourced in India, and the Tansian people there that's able to maintain it. Keep it that way. That's great there.

Now, innovation side, we started with the FlashArray. It was a very small footprint and we know that the big bucks are on the unstructured data inside there, and we needed that. That's where we wanted a lot of innovations there. We worked with them in the last few years there and with the POCs, it's great to work with the team there.

Even before I purchased the first time I went there, they explained the whole architecture and convinced me over the meeting how reliable, how fast, how better it is there. That's how they are transparent and easy to work with. That's why now we made the call to go with them and also partner with them and expanded there. It's a great relationship and great to work with them.

Kevin Delane:

Great. Well, thank you very much. We'll continue to work hard to earn your business and Hat, back to you.

Dave Hatfield:

All right. Thank you. Can you guys hear me? Andy, thank you very much. Matt, thanks. Justin, thank you. Obviously our customers are why we do this and so the fact that these three very busy individuals are here and willing to give us some time, maybe we saved them some time on support, so they give us some time back on sharing their stories but it means the world to us.

As you're walking around and doing your channel checks and talking to customers, talk to all of them. I think this is very representative of what our customer experience is all about. This is not sort of the ones we want to show in front of you. These are folks that obviously are willing to spend some time with you, but spend time talking to everybody out in the show floor and I believe you'll hear something very similar.

From here, a whole bunch of innovation was launched today, so we're going to go into the details with—Kix I think is next, and then we're going to have Q&A at the end for whatever questions you want to get into.

I'll leave one more thing with you. I was thinking about it and I meant to comment on it. A lot of the questions that I get are how long are you guys going to be able to maintain that white glove support, right? Because the belief is we're spending a bunch of dollars on people to deliver support to be able to get these satisfied customers. The reality is—and Tim will walk through the long-term model—as you see the leverage that we're getting from a support gross margin perspective, it's actually outrunning us and in a really good place while we're delivering the kinds of repeat purchase metrics and the satisfied customers that we have. With big data and what we're calling Meta, these guys are going to go through, we can actually leverage AI to deliver automation on a product that is just inherently more reliable.

Packing a bunch more value into that evergreen subscription for less money than what they're spending or comparable money to what they're spending in maintenance will allow us to sort of sustain that over the long term.

Kix, over to you.

Matt Kixmoeller:

All right. Sounds good. Hello. All right. I'm going to try to give a very rapid tour of the 25-plus new software innovations we launched today. I know a lot of you weren't in the keynote so we'll kind of recap that and also give you a sense from an investor perspective how we think this ties to markets we play in today and over time.

The big goal that we've been engineering towards since our inception is to make sure we're continuing to innovate towards the right part of this diagram. Our customers are going through this transition where they have their classic environment they've virtualized, but they're now building new web scale apps. These new web scale apps are totally different from the applications before in terms of how they interact with storage and the developers that develop them also have very different expectations. They don't want to talk to IT people. They want to develop and they want to develop to a well-formed block file and object APIs that just work, that scale, that are effortless, and that, of course, are all programmably controlled into their applications.

Look, our mission at Pure is to help our customers set up that data platform for their developers to be able to manage their existing applications and their future applications, their business apps test of big data, as well as their modern web scale. Our platform provides those software-defined services, scalable block file and object services that are programmably controlled by APIs. We deliver that through our Purity operating environment. The core of our innovation in software and so Purity delivers not only those storage services, but a foundation of key capabilities like data reduction, data protection, security, and assurance. In the most recent releases, we've evolved Purity with Direct Flash to be able to talk directly to the NAND itself, and I'll talk more about that.

Pure1 is our SaaS-based management platform. This is how we deliver simplicity, and we're really upgrading that, taking these same AI technologies and machine learning that you're hearing about it into our data set and taking our role as an IoT company very seriously.

Then finally, our customers want to live in a multi-cloud world. They own their core data centers. They're doing more and more in edge data centers and they want value-added connections to public cloud providers, multiple public cloud providers, and so we announced some integrations today and that's going to be a rich area of our roadmap.

One of the things I think makes Pure really unique is that we innovate at this intersection of three key areas. People often ask are we a hardware company? Are we a software company? Are we a cloud company? The answer is all three. The fundamental DNA at Pure has always been software. Our capabilities are enhanced at software kind of rev cycles, but when we integrate software with cloud-based management, that's how we drive simplicity. That's how we use the SaaS model to reimagine how simple storage can be.

When we integrate the cloud with our hardware, that's where we become an IoT company. So being able to take this rich stream of data and figure out how to make storage self-driving over time, there's just a rich parallel as to what we see happening right now with self-driving cars where you take the modeling of the entire world and get smarter about how you can inject that real-time model in the car. Well, as an IoT company, we see hundreds of thousands of workloads across our platforms and so we can give the information to a storage admin who might see 10 or 12 about how his workload may change over time.

We can help drive the storage rate itself because we know by correlating what we're seeing on that array with the rest of the universe.

Then, finally, we've done a lot over the last couple of years to natively integrate software and hardware, to co-develop software and hardware together. We see a lot of this as the same trend we're seeing right now in the cloud. If you look at what's happening inside the mega-scale data centers, there's a lot of co-innovation of software and hardware. When we fight with folks to hire hardware engineers, it's usually the top three cloud providers we're fighting with and that's because basically being able to integrate software and hardware directly helps take things to the next level and drive a uniqueness of advantage there you just can't get through the open architecture alone, and so Pure is really the synergy of these three areas of investment and so it overlaps where the real power comes.

Today we've organized our talk into kind of two areas and I wanted to just show you how this maps to the three things that are powering our growth. Dietz likely talked earlier about these three key focus areas: going after next gen data use cases, serving the 500-plus leading clouds, and just helping drive the cloudification of classic enterprise architecture.

Today you'll hear us walk through two key areas of innovation. One is this transformation from big data, big intelligence. Rob Lee, who's over here, is going to join in just a second. He's the Chief Architect of FlashBlade. FlashBlade is really at the cornerstone of empowering this and we're really seeing this be picked up in a lot of different areas. The reason it's important is there's a whole bunch of consumers going after these next gen use cases, but the cloud is a key area, and our cloud consumers in particular are taking advantage of big intelligence.

Then the second area is defining a new Tier 1 storage array. This is an area where we're adding a set of features that allow us to go after the highest end use cases that are out there, the highest level of availability with our new active cluster solution and pushing forward with our innovation around NVMe. This also helps us go after cloud providers because it allows us to go after a lot of what used to be the server DaaS workflows that are out there, but also it obviously helps us go out after the classic storage installed base, which is also going to fuel a ton of growth as we going and replace disk with Flash, and we're still in the early innings of that.

Let me start by turning it over to Rob Lee who's going to take you through big intelligence and I'll be back later to talk about the new Tier 1.

Rob Lee:

Thanks, Matt. I'm going to try to speak up to try to outcompete all of the espresso machines and all the other people around us, so apologies in advance of it gets a little bit loud in here.

Before we dive into some of the innovation directions and future announcements in FlashBlade today, I just want to take a moment to kind of set the stage for some of the things we're seeing in the industry that are really driving the direction that we're taking, the innovation investments that we're making in FlashBlade.

You heard Dietz talk a little bit earlier about how data is the new oil, data is the new strategic advantage in business. This is true, but it's only part of the truth, right. The other part of the truth is that that value, that strategic advantage is really in the insights. It's really in the intelligence that's locked away in that data. In many ways, the challenge now becomes finding ways to accelerate and get better at and refine our ability to extract those insights from the massive amounts of data that are being collected in the world.

This movement to focus on and get better and refine our ability to extract those insights and that information, that intelligence from these massive sums of information that we're collecting, this is

something that is now being more commonly referred to as really the fourth industrial revolution. Just like the revolutions that came before it that mobilized us and moved us from rural areas and moved us into cities or changed the way that we build things from manual production into mass manufacturing or digitized the modern office, this technology change also represents a massive change in the way that literally every industry is going to work, everything from advertising to retail to manufacturing to social media to farming to planes, trains, and automobiles.

One of the things that's really driving this, right, one of the things that's really enabling this is really a perfect storm of technology that's coming together. It's the combination of our ability to put more and more data capture devices out in the world. This is everything from video cameras to machine and IoT sensors to log capture devices combined with advances in modern analytics algorithms and artificial intelligence and machine learning, as well as the hardware advances that have come to market to really better support and accelerate those. Taken as a whole, this really represents a massive increase, an exponential increase in computing efficiency that we're seeing today. In many ways, it's kind of a return to the heady days of Moore's Law, something that we haven't really seen as an industry for quite a few years. What it's demanding now is a commensurate increase in storage efficiency and storage performance just to keep pace and just to make that technology be put to work.

If we take a look at some of the places where FlashBlade is being put to work today, at the cutting edge of a lot of these industries and use cases, we're being put to work in everything from training the world's largest AI supercomputers to running real-time sensor analytics to help us. Matt was talking about optimizing the way that CFD calculations and optimizing the way that we build cars and make real-time adjustments between races to launching rockets and doing financial simulations to help guide decisions that are being made by the largest financial institutions in the world.

In the case of one of our largest AI deployments, FlashBlade is being used to feed training data sets generated by hundreds of millions of users into one of the largest deep learning neural networks in the world. This is a deep learning neural network that's being powered by hundreds of NVIDIA DGX-1 GPU accelerated server blades. To put that in perspective, that represents somewhere in the order of hundreds of thousands of cores of computing power. All of that data produced by hundreds of millions of people is highly unstructured. Everything from text data to images to video, all of that data is being fed into these massive compute farms, massive neural network training nets to better train machines to understand and react to human behavior.

Think about that for a second. What we're doing is we're taking unstructured data, we're finding single net and we're training machines to better understand everything from group sentiment to individual buying patterns to how best to advertise to people, all the way down to individual emotional and mental health states. That might sound pretty exotic but in many ways those applications, that working set, those use cases, bear a lot of similarity to the same models that are used to run fluid dynamics simulations and guidance control systems. That similarity is that they're extremely data hungry. They're data starved and data thirsty. Our ability to gather insights and the quality of those insights that we can gather depends greatly on the ability to feed large amounts of data very, very quickly into these modern compute and analytics platforms.

The same groundbreaking performance that's causing FlashBlade to be selected to drive these cutting edge use cases, combined with just overall density, data center reduction, footprint reduction and just overall simplicity is also pulling FlashBlade into a number of other more general purpose enterprise workloads, everything from medical imaging to database backups and restores to software and game development, even video storage for the public sector.

With that in mind, I want to kind of turn our attention to some of the announcements that we're making in Purity for FlashBlade 2 today which are going to help us further our lead in both of these markets.

There's data-driven analytics, there's data thirsty and hungry applications, as well as the next gen and cloud native applications, as well as the general purpose and enterprise workloads.

The first feature we want to turn our attention today is really the next step in the journey of growing FlashBlade to be even larger and even faster than it is today. In the second half of this year we'll be growing FlashBlade to over five times the size and performance that we have today. You'll be able to grow these systems from 7 blades, all the way up to 75, one blade at a time, and of course, non-disruptively, without adding any management or operational complexity. To put that in perspective, what that means is you'll be able to start with systems as small as 30 terabytes and grow them elastically all the way up to 8 petabytes in size, all within a single system, all online, all within about a half rack's worth of footprint. There's no other platform on the market and no other technology in the market that can provide this level of performance and scalability while maintaining simplicity as we can today.

Now, as those data hungry, data-driven analytics workloads continue to improve, they're going to demand to be fed with larger amounts of data even faster, and so by increasing FlashBlade's capacity and performance envelope by 5X today, this is going to allow us to stay well ahead of that curve for years to come.

All right. The second feature we want to talk about today is the introduction of native objects to our interface to FlashBlade. Object storage is really the interface on which next generation and cloud native applications in particular are being written to. Now, historically, object has been generally synonymous with big and slow. What we're doing today is really changing that conversation to big and fast. Object storage really is the primary storage of the cloud era and, as such, it really does demand that same level of performance, scalability, and reliability that you would expect and that we're delivering today for primary block and file storage. In early beta testing with a very large SaaS company that's been using native object store implementation within FlashBlade, they've seen over 10X improvement over public cloud offerings, in particular AWS S3, and over 100X performance improvement over their existing on-premises production system.

Now, what enables us to deliver this level of performance and scalability is the fact that Purity, under the covers—both Purity for FlashArray or Purity for FlashBlade—under the covers internally is actually an object store. It's an object store key value database that we built years and years ago. On top of that core platform, we've delivered block interface for FlashArray, file interfaces for FlashBlade, and now native object interfaces for FlashBlade as well. The reason this is important to highlight is it shows that all of the investments in R&D that we've made over the past five years and the next five years and the five years after that into that core platform are going to benefit all of our customers, regardless of what product or what interface they're using, whether it's block, file, object, or whatever interfaces come after that.

Now, growing FlashBlade and scaling it to 5X and introducing object store, these are features that are going to allow us to further our lead and further our grasp on these next generation workloads, especially these really data intensive data hungry analytics workloads. But at the same time, we remain deeply committed to and rapidly innovating features to help us better address a lot of the more general-purpose enterprise workloads that we're in today. Purity for FlashBlade 2 is also introducing a number of features such as SMB, NLM support, LDAP, snapshots, and HTTP interface, again, a lot of maturing features that will really help us further our reach into those enterprise workloads that we're already in today.

With that I'm going to turn it back over to Matt who's going to talk to us a little bit more about the new Tier 1.

Matt Kixmoeller:

All right. Sounds good. The final thing I'd say about this area is I think as an industry analyst community, folks are really having a hard time understanding the growth potential of some of these next gen use cases. If you look at we're doing around AI and ML, if you look at what we're doing around fought (phon) fast object stores, these are both things that basically just didn't exist before, and so there's been some industry estimates about AI and ML being something that's going to grow enormously fast, be a multibillion-dollar storage market. We don't even know what people are going to do with fast object. We just know it's going to power our next generation of applications, and so we're excited to kind of create these new, fast growth, data-driven opportunities within the overall storage market.

All right. So the second area we wanted to go into is the new Tier 1. This is an area where we're really aggressive about trying to redefine that classic kind of top-of-the-hill storage offering. If you look at the storage portfolio of many of our competitors, you see kind of this compromise, this choice. You have their legacy tier 1 platforms that had the highest reliability, had the sophisticated enterprise replication options, maybe QoS, this is where one suggested to put kind of the most critical apps.

On the flip side, you have modern all FlashAarrays where new innovation is coming that is really designed more for cloud era workloads, things like in-built data reduction, things like simple automation and openness, things like cloud integration and NVMe. Unfortunately, these have been a choice. You either pick reliability or you pick innovation, but it's been hard to get that from the same vendor or the same product line.

At Pure, with this set of releases, we're really ending that compromise and we're trying to define a new Tier 1 that can not only host the highest end business-critical applications, but give modern developers the tools they need to build the next gen of cloud era applications. What we announced today was the Purity 5.0 version of our software, a huge set of upgrades across the board, and I'm going to hit some of the highlights here, but the premier feature within this is ActiveCluster. ActiveCluster is not just synchronous replication. We've been kind of colloquially referring to it as that, but it's really taking sync (phon) up to the next level with a true active, active stretch cluster solution. You can think about this as a cluster where you might stretch an application between two metro distances and two data centers that are up to 250 miles or 5 milliseconds of latency apart. If you look at the traditional implementations from legacy vendors in this space, they're incredibly complex, they're incredibly expensive. This remains one of those features that folks have reserved the right to charge a lot of money for. They also tend to be very services-enhanced where it'll take a lot of time to get them up and running in these competitive solutions. When we looked at this, we basically said, look, we want to democratize this experience. We want to take this signature Pure ease-of-use and bring it to this highest end availability requirement. This functions as just a single stretched volume across two sites. It's active Active. You don't have to configure your application differently. It just believes it's writing to one volume, just happens to be stretched across two data centers.

The other thing that we did that was very novel in our approach was we shipped the cloud mediator as a service. Son in a metro cluster type solution you not only need two data centers, but you need a third data center to figure out in the element or in the situation where a link is lost, which is actually the remaining valid data center. So with competitive solutions you have to run a mediator as either another array or a VM and a third data center somewhere, manage it and deal with the complexity of that three-way integration. We ship that as a SaaS service from Pure1 and so a Pure customer, if they want to take advantage of this feature, just enables it between their two sites, turns on the Pure1 cloud mediator from the SaaS service and they're up and away, and so another way we're driving simplicity into the whole thing.

The final thing I'll say here is, of course, we're doing this in classic Pure style of building it into our base offering, so there's no additional uplift, and we think that's going to be pretty disruptive from a business model point of view for that classic enterprise customer base.

Four steps to configure. We only added one new command. If you look at our competitors, there's often over a thousand pages of manuals to go and configure something like this. We added one new command, about 20 pages to our manual. Because it's active, active it makes it a lot simpler. Competitive solutions that are active passive are a lot more difficult to set up and, again, zero additional cost.

A second key area that we focused on in this release is adding more rich control of QoS. These Tier 1 platforms are about consolidation, hundreds or even thousands of workloads sharing the same array. In our last Purity release we added on always-on QoS, the kind of monitors in the background, for one tenant that might take too much performance. In this release we give more fine-grained control where customers can specify bronze, silver, and gold buckets and put their most critical apps in the gold bucket, or to be able to go after the service provider market in particular, we added fine-grained QoS and the ability to specify limits per volume, and so there's a set of service provider business where a service provider wants to differentiate bill for different levels of service. That enables this business model and enables us to go after that market that we haven't been able to play in before.

Those were the features to really go after the highest end of reliability. At this point we feel like there's exactly one thing that we don't have that prevents us from going after some of these environments, and that's mainframe support. If you're a mainframe environment, which is obviously not a growing market, you might still have to buy a Legacy Array. But if you're not running on mainframe, which hopefully most folks are moving those architectures forward, we think it's time to consider the next gen Tier 1.

Now, let's dive deeper into the set of features that are really designed for empowering cloud data centers. Our core software is the Purity operative environment that we talked about. We really have tried to optimize for these deployment type use cases. In this release we see a number of new features come to bear. The first is deep integration with VMware's VVol Suite. VVols and vREALIZE from VMware is their cloud automation framework. It's kind of an improvement to the storage layer to allow the arrays to siege (phon) VM and do our Purity services on a VM by VM basis, but it's also an integration with storage policy base management, which is VMware's kind of policy automation suite. So there's a large market out there of service providers who run their services based on the VMware infrastructure and we think this is a great fit for those customers.

The second area is going after Docker. In the next gen ecosystem, especially in the DevOps environment, Docker is taking off very aggressively and so adding persistent container support allows us to not only deal with the femoral (phon) containers but persistent containers for higher class workloads.

Then finally on the Microsoft platform we added support for the ODX API which is a dramatic performance boost that accelerates I/O operations in the Microsoft stack. Whichever platform you choose for your cloud, open source, VMware, and Microsoft value-added integrations that help keep moving forward.

The second key area is a set of integrations and a set of future forward-looking integrations here around taking our snapshots to the next level. This is an area where we've announced some new features, but we've also talked about a roadmap over the next year, and so I'll be careful to differentiate between those two areas. But, snapshots are a key engine for data movement and mobility. We've obviously supported for years snapshotting within the current array or a local array and moving with our replication technologies between FlashArrays.

What we've created in this release is a concept called a portable snapshot. That allows us to encapsulate the metadata for recovery right within the snapshot and to kind of frees it from the base array where it was created so it can kind of move around. We're going to take advantage of that portability to open up connections to different places to send snapshots for different types of workflows. The exciting thing about snapshots in general is they could be used to power data protection workloads, things like backup

and recovery, but also DevOps workloads, things like taking a copy of a production database and putting it over in Test Dev to allow the developers to work on it.

We'll be able to now snapshot to FlashBlade. We've got a lot of our customers obviously that are taking advantage of FlashBlade and wanting to use that as a larger scale retention store. We'll be able to snapshot to generic NFS targets. If you have backup plans or you have an old net app filer or something like that, you can take advantage of it and use it as a destination. But, most excitedly, we're opening up connections to the cloud. We've had a lot of customers that want to replicate between primary data centers but have a third copy of data for retention out in the cloud, and local cost cloud storage is really providing that opportunity to replace those former tape copies. We're going to be opening up our first integration with Amazon's S3 to allow snapshots to move to S3 and then to be persisted off to Glacier for lower-cost storage as well, if that's interesting. I'll talk more about that in a second.

Then the final part of the picture is the open ecosystem. If you walk the floors here, you've seen we have a lot of partners. One of the unique APIs we offer is our Delta Snap (phon) API which allows third parties to have that same value of integration to be able to move snapshots in a space efficient way off of our Array so they can only do change block kind of movements.

Let's talk a little bit more about where we're going with the integration with the cloud. As I said earlier, there's really two types of workflows that we're excited about. The first is a simple kind of backup archive and maybe DR workflow where you're moving copies to retention in the cloud. I'd say that's the one that we've probably had the most interest in. The second type of workflow is more of a workload migration for development type use cases. Maybe I want to develop on-prem deploy in the cloud. Maybe I want to develop in the cloud to deploy on-prem. Maybe I want a burst to the cloud. This provides that kind of integration.

We're going to start, as I said, this year by shipping our first native S3 integration. The goal then is to be able to take that data out of S3 and be able to recover in the cloud in the EBS to enable a DR scenario. That'll be something that happens next year. Then this integration with EBS for kind of volume migration will also be something that we're planning on for next year.

Again, this is an area where we aren't shipping this entire picture today. I want to be clear about that. We've had such overwhelming interest from the customer base that we're really going to be leaning into the roadmap here, working with a set of kind of definitional customers and really trying to plow new ground here in integrating natively with multi-cloud.

All right. Another key capability we announced today that I think is going to be very exciting for the industry is Purity Run. This is something that we're doing specifically for developers to be able to open up our platforms as a development ecosystem. A lot of our, especially cloud and SaaS customers, like to think about their storage as a tier that they want to integrate with their applications. Today they talk over storage protocols. But what if they could talk over application-specific protocols? What if they could implement their application logic right on the FlashArray so it can become more of a node in their infrastructure? We're opening up Purity Run, which is an infrastructure that allows anybody to run a third-party VM, container, or application logic right on our platform.

Now, we've stopped short of calling this a hyperconverged platform. We don't view it as a general purpose hyperconverged play. There isn't enough compute in here to be able to run hundreds of VMs and have it be a more broad compute tier. But what it is, is an opportunity to allow our infrastructure to be much more closely tied to next generation applications.

It also has very interesting implications for IoT and for edge compute. If I want to have an edge storage device and take advantage of analytics out of the edge to summarize my data before I send some home,

you could put your code on the platform there. It also becomes a platform we can use to ship enhancements to our Array more agilely as well and come out with faster functionality.

Another thing that we announced today, which I'm not going to talk about in depth, is Windows File Services where we took Windows' server, embedded it on the Array to open up SMB and NSF interfaces for consolidation use cases.

All right. Finally, I wanted to get into a bit more on our strategy around NVMe. We've obviously talked to a lot of you about what we're doing with NVMe. It's an area where we feel like we have a great lead in the market and it's really taking advantage of our ability to deeply integrate software with hardware and own that end-to-end value chain.

This is all about going after that next generation of use cases and we see it both in enterprises and in our cloud customers. Within enterprises it helps us go after higher and higher performance, deliver higher and higher performance and also open up the next decade of higher density on Flash. Within cloud environments, if you were in the keynote this morning, you saw some pictures of ultra-dense infrastructure where you can basically get a rack, have ultra-dense compute and have top-of-rack Flash to be able to feed that to compute with storage. We see that as the building block of the cloud.

Today we're furthering our vision around NVMe to announce our NVMe DirectFlash Shelf. When we first shipped our NVMe product line, essentially NVMe was within the FlashArray X. Our new DirectFlash modules could get slotted into that base module. But many people asked, "Okay. How do I expand beyond that base platform?" What we're announcing today is the Shelf you see here. It connects between the FlashArray X and the Shelf with NVMe/Fabrics 50 gig ethernet, RoCE v2 connections. This is a fully NVMe pipeline that extends beyond and it allows us to continue to scale depth in terms of the size of storage arrays we operate and flexibility over time.

You can also guess that we also want to extend NVMe in the other direction. We're not making a product announcement today, but one of the things you can see here at the show is a joint demonstration of NVMe North up to servers jointly with Cisco. We're showing a Cisco UCS implementation of NVMe through the VIC directly down to FlashArray X. This is something where the NVMe/Fabric standard is still a little early, we're still working through kind of the joint and operability work to make this work on a grand scale, so we're not announcing specific timelines for making it available, but you can see where we're going with this and NVMe, in the front of the Array, out the back, and within every piece of it. This is about removing all the legacy SCSI protocols, opening up the next generation of density and performance.

This is the picture I mentioned earlier, but this gives you a sense for the type of density we can drive. If you talk to large-scale cloud providers, this is how they think. They think in rack scale. I'm going to just line a rack with a lot of compute, a lot of storage, and then print it out and automate across that. This opens up by moving to a top-of-rack Flash solution, the ability to have ultra-dense compute because you no longer need the storage within each of these nodes and so you can have much denser compute down within the rack. You get to un-precedent levels of density here and you get the elastic file, object, and block storage services you need right within that rack, connected in the end (phon) through NVMe.

All right. That is the next generation of Tier 1: redefining availability. Another thing, by the way, that we announced today was achieving our second year of six 9s availability, so that kind of highest mark of resilience in the storage space. We achieved that through non-disruptive upgrades as well, not requiring downtime for upgrades. The brand-new ActiveCluster feature, which really raises the bar in terms of going after those Tier 1 deployments and our plan to not only be very technically disruptive here by driving simplicity into ActiveCluster Solutions, but also to be very economically disruptive here and to make this an included part of the storage solution where many of our competitors still kind of hold this out as a feature that is the pinnacle of their software licensing strategy.

Built-in QoS and VVols to help go after service provider use cases with consolidation, with differential billing to go after service providers and QoS hard limits; entering down the path of cloud integration in our cloud snap strategy to be able to have transparent migration for backup and protection workloads to the cloud, and over time more and more migration and burst capabilities; and then delivery in end-to-end NVMe and doing it all in an evergreen fashion. All these upgrades, totally upgradable from our current FlashArray customer base, particularly the path towards NVMe, which we believe will be very unique.

All right. Then the final area we want to talk about is a fun one. It's a vision around self-driving storage. This is one where as an IoT company we can start to flex our muscles as well. Rob talked about how we're enabling everyone to take advantage of AI at the next level. Well, we can take those same benefits as well. If you think about what's happening today within your cars, your cars are being trained by a very large network and building a self-driving car is very little about understanding how to drive a car; that part's easy. It's about sensing the world around and avoiding problems. It turns out that this is the same problem that storage avenues face, especially at Pure. It's pretty easy to administer a Pure device, but there's a whole ecosystem of stuff connected to the Array and you never know what your database is going to do. You never know what your application is going to do. You never know what the hundred applications that are put on the Array might do today and over time as they grow.

We looked at that and said, "Man, we collect a network of hundreds of thousands of workloads. How do we analyze that data and make it useful to better understand how to build an Array that becomes more autonomous over time and drives up the level of resiliency, drives up the level of automation, drives down overprovisioning?"

The first thing to understand is where this data comes from, right? Since day one we've had call home technologies. In our last year's release we started adding a new capability to fingerprint issues, so we were essentially finding issues and creating fingerprints for every known Sev1, and looking across our install base in real time so we could prevent customers from hitting issues that we might be able to detect before it became a very bad day. This is part of what we've done to drive up resiliency, but these issues to date have only been generated by kind of human cognition. We can understand a problem; we can maybe go upstream and figure out how would we have known this was going to happen; we can write a fingerprint for it; and humans generate that. Well, you can guess where I'm going with that. We can generate a lot more fingerprints if we start to use AI NML technologies to find correlations and issues that humans couldn't discover.

But, the second area, and probably the most interesting one, is really understanding performance. If I want to ask a very simple question, like I've got an Array. I want to put 10 workloads on it. Will those 10 workloads fit? Historically in the storage industry it was relatively easy to figure that out from a capacity point of view—it got a little harder recently with DDUP—but it was nearly impossible to figure that out from a performance point of view. We might have looked and said, "Okay. Well, how many IOPS is it doing or what kind of latency does it need? Let me add those up and try to guess how many IOPS the Array will do." But these were almost childishly simplistic exercises. It was all we could do, but they were often wrong and they often caused customers to basically just buy more storage than they needed because they didn't want to ever run out of performance, or if they under-configured, to have very bad day.

This turns out to be a perfect problem to solve with AI because we can kind of guess of the thousand different performance measures we measure on the Array which ones are the best indicators, but why bother? Just point the computer model at it, have it iterate through, have it look for correlations and have it understand. We did that and the basis of that work is the new AI engine within Pure1 that we call Meta.

Meta is really something that we believe we have the opportunity to plumb into all parts of what we do. We can plumb it into our management tools to take actions faster and be more autonomous; we can plumb it into our analysis to be able to make predictions better; and we can plumb it into our support

process to be more engaging and proactive about support. The first thing that we're unleashing is this new workload planner that runs on a concept called Meta DNA. Our workload DNA capabilities here analyzes those thousand kind of performance measures and figures out which ones of them are more indicative. The answer, of course, is not just one thing, but it uses them all in different proportions and rations, and it allows us to figure out, when I add different workloads together what'll fit on the Array and what won't, and we can start making those predictions very smart.

If I'm an admin that's sitting here just looking at my, let's say 10 instances of Microsoft SQL running on my Array, I might try to guess how SQL is going to grow, but I can probably guess a lot better if I looked at the 10,000 instances that Pure knows about, looks how they've all changed over time, and uses a sophisticated correlation try to give better prediction. What if something changes? What if Microsoft issues a SQL patch and we see that kind of workload DNA start to change for that application? Well, we can detect that at a global scale and then give you better recommendations locally.

The very first kind of instantiation of Meta is our new workload planner tool. This is a great example of something that looks very, very simple on face value but there's a lot of power underneath. I make predictions by just kind of moving the sliders. Underneath it's analyzing all your workloads, looking across the rest of the universe, trying to make better predictions about how they're going to grow. You may or may not be able to see, but at the top there's a great example of something where the capacity is going to be flat, but it turns out you're going to run out of performance. Here's another one where the performance is going to be okay but you're going to run out of capacity. These are two real workloads in our environment right now.

We are just at the beginning of a very exciting journey with Meta and it's a great example of how we're taking our role as an IoT company, leveraging big data, eating our own dog food, you name it, to make storage smarter, and we think there's a very rich road for delivering this vision of self-driving storage.

Just to kind of bring it back to where we started, three key things powering our growth: next generation data, going after the 500 leading clouds, and the cloudification of enterprise storage. The bottom one probably alone could drive our growth over the next few years and make all of us happy, but we're very excited about how the top two are creating kind of a discontinuity and really providing new opportunities that are driven by the next generation of data workloads.

Hopefully you've gotten a sense from the conference today that innovation is part of our DNA. We're proud of originally inventing the AFA recipe, 7.5-odd years ago, evolving it over time. Today we ship next generation cloud era Flash products that FlashArray and FlashBlade and we continue to drive first in the industry. We've obviously done a great job in terms of leadership in the Gartner Magic Quadrant from Day 1, record CSAT (phon), and real depth in our patent portfolio which continues to grow as we continue to innovate.

All right. Just to kind of sum up, for Pure it's all about innovating across these three vectors: being software driven, being hardware accelerated and being cloud connected. We are focused on these markets that matter, particularly driving innovation in data-driven workloads, and we feel like our innovation lead is widening. We continue to see our competitors go after market share where it's kind of easy, going into their installed base, retrofitting their current arrays, putting big SSDs in there, moving on. We're focused on the fundamental innovation that sets us up for the next decade of leading the data industry.

Thank you very much.

Male Speaker:

Thanks, Kix. We're going to take our much-needed break right now, so let's please be back at the room at 3:15. Don't take too long. When you come back, Tim Riitters with the financial update.

Tim Riitters:

My job here is to wrap up for the day. We'll do a Q&A after my discussion today, but really wanted to provide a handful of financial updates about the business. We're going to do them in about three categories today. Number one, we'll start out with a look back. After operating now for about seven quarters as a public company, I wanted to give some perspective and trajectory in the business and also talk a little bit about how we see the rest of the year shaping up. For those of you who've naturally listened to our call, we've talked about it but it's a good opportunity to provide a little bit more color on how we see the fiscal year planning out.

Then, number two, we're going to talk about the road beyond. I think this is what's exciting even more for us. We're going to talk a little bit about the opportunity we see in front of us. You've heard about where the markets are headed; you've heard about the innovation that we're driving; you've heard about the leverage that we're driving in our sales and marketing investments, and so you put those three things together, there's a lot to be excited about as we think about the years ahead, and we'll talk a little bit more about that.

Then, finally, I'll wrap up with a little bit more technical and deep dives. Some of you may or may not find this interesting, but we're going to talk a little bit about the seasonality, nuance dynamics, ASC 0606, a lot of the stuff that the finance guy probably wants to give you a little bit of perspective on. Then we'll turn it over to Q&A.

With that, we'll start on topic number one: from IPO to \$1 billion. I think most of you have seen these numbers, most of you are familiar with these numbers, but, again, it's a good opportunity to reflect back on where Pure Storage has come from. We've been operating now for approximately six years commercially, selling product for six years commercially. If you look at this, we've driven in those roughly six years, \$728 million of revenue in our last fiscal year. That's greater than 150% three-year CAGR in the business over the last three years. It really is best-in-class growth and it's really, from the benchmarking that we do, shows that we are one of the fastest growing enterprise infrastructure companies in history. It really, again, while we're excited about the future, it's good to spend a few seconds and reflecting on the past in terms of the momentum that we've been driving in the business.

At the same time, we look at the two other pieces of information on this slide. I'll focused next on the bottom section. We've been doing and driving this business in a very sustainable and differentiated way, driving 66% non-GAAP gross margins for our last fiscal year. That, again, is industry-leading, best-in-class, and it's something that we've been doing for six quarters now. It really speaks to the discipline that we're driving in the business and, just as importantly, the differentiation and value that our customers see. This is not a commodity product and it's obviously referenced in terms of the ability of the business to consistently be in that long-term target model of mid- to high-60s.

Finally on this slide, we've been doing all of this, while at the same time, driving substantial leverage and benefit in the business from a bottom line perspective. We've talked in the past about our operating principles, about ensuring that every year we're driving measurable improvement on an operating rate basis going forward, and if you look back over the arc in (phon) the last four years, we've definitely done that. We've driven nearly 50%, or approximately 50%—some years a little bit less, some years a little bit more—in our operating profit performance over the last four years.

Last year, posted a negative 13% for the full-year, and as we'll talk about it here in a few slides, we have a negative 7% guide for the full year in the year that we are today, so another roughly 50% improvement in operating margins. Again, it's this combination, it's this recipe of driving industry-leading growth, taking

advantage of a substantial market opportunity in front of us, doing it in a differentiated way, doing it in a sustainable way, and driving leverage in our path to profitability, which I think is unique to us here at Pure Storage.

The last thing I'd add is one more data point on overall operating leverage. We'll talk about free cash flow, we'll talk about profitability here in a moment, but recall—it's not on this slide—but in Q4 we posted a non-GAAP net operating loss of negative 1.9%, so really on the cusp of turning the business both free cash positive on a sustaining basis, as well as profitable. Again, we'll talk a few more minutes about that here in a minute.

Onto the current year, we've had a great start to the year. We posted results for our Q1. That's the period ended April 30 for us. We really had really strong results, top to bottom across the P&L. We guide on three fundamental measures in our P&L—revenue, gross margins, and operating margins—and across all three of them we've delivered notably above the midpoints of our guidance. Revenue, \$183 million, that's 4.3% ahead of the midpoint. It's also 31% year-on-year growth, which, again, at the size and scale we are as a business, really, really good result for the Company. Gross margins, 66%, that's relatively flat quarter-on-quarter. Again, industry-leading, and really done in a period of challenging supply chain environment. For those of you that have been listening to us on the calls, we've talked about this for several quarters now. We've been weathering this very, very well, and we've also been doing it in the context of a product ramp in FlashBlade as well. So, again, really, really happy and proud of the teams, whether be they be our Technology teams, our Support teams, our Logistics teams, and our Sales teams and driving differentiation, driving the best COGS possible and really innovating to drive that differentiation.

Then finally, the operative margin, negative 16.7% for the quarter, that's up from negative 29% the year-ago period. So, substantial leverage, again, year-on-year, notably above the guidance midpoint. We had a really, really good quarter, both on revenue and gross margin, which both sort of combined to deliver really, really nice bottom line results.

Turning to Q2, another sort of strong start to Q2 as well as evidenced from our guidance. So, \$218 million in revenue at the midpoint of our guide; that's 33.6% year-on-year growth, continued that sort of pace of growth at scale, continued consistent gross margin guide at 65%. Then finally operating margins at negative 14% at the midpoint. Again, both improvements quarter-on-quarter, as well as year-on-year, so really, really good start to the year, both in terms of Q1 and Q2.

It's really this first half performance, both the actuals that we've delivered, as well as the confidence that we have in our Q2 guide that has us very confident for the full year as well. You heard Scott talk a little bit about it, both at the keynote here in the main stage, as well as today. One billion dollars of revenue, that's 37% year-on-year growth. We continue to be very, very comfortable with that number. You've heard Dave and his leadership team talk about some of the things that they look at in the measures and the indicators that give them confidence and I think I'll reiterate those and offer a couple of mine as well.

Clearly repeat business is firing on all cylinders in terms of what we've seen consistently; FlashBlade is ramping very, very well. You've heard Matt Burr talk about the FlashBlade momentum as well as the graph earlier in terms of its ramp. Then feature development innovation. It was interesting, on the call a couple of weeks back we talked about feature and innovation that we were driving and we teased people to come here to Accelerate. I hope that you get a sense from the conversations that happened today as well as when you walk the floor, that we were really serious in terms of this year being one of our strongest years of innovation ever. It's those types of feature and functionality that really expands the use cases, expanding the relationships that we can build with our customers, and capture additional use cases in the business. So again, a lot of good tailwinds propelling us forward for revenue for the upcoming year.

Gross margins continue to be right at that target model and we anticipate the next year, throughout the rest of the year, it's no different. It's more of the same there in terms of driving continued efficiencies in our supply chain, selling value with our customers who actually experience the technology, and continuing to be mindful about ramping up our support efficiencies. Dave talked about that a little bit earlier. I think there were some questions on the earnings call about the support margins that we were experiencing. We continue to see some leverage there and it's certainly a testament to driving industry-leading NPS, industry's leading customer satisfaction, and yet doing it by continuing to grow our overall support margins.

Then, finally, operating margins at negative 7%. Again, as I said in the earlier slide, about 50% rate improvement year-on-year in the business. Really, for those of you that have been following the story for the last couple of years, there is inherent seasonality in our business on the operating profit line. We make significant investments in the business at the start of the year. I'll talk a little bit more about that in Section 3 of my talk. But the bottom line is that sort of operating margins do tend to grow over the course of the year, capping at Q4 which tends to be our strongest operating margin results for the company in any given quarter of the year.

I'm going to go on to part two, which, again, I alluded to as being actually what I'm even more interested and excited about for the company. I certainly know the significant number of puritans at the company too are excited about the road ahead. We've never been a company that has been solely focused on any given quarter. Certainly we have an obligation to continue to exceed and do everything that we say we're going to do to the investing community and that ethos hasn't changed, and, yet, really it's about building the long-term capability and momentum in the business. There's fundamentally a significant opportunity in front of us that I want to talk a little bit about today, and translate that into what that means and the opportunity that potentially could mean for the company in the years ahead.

You've seen this framework throughout the conversation this afternoon. It's really how we think about the recipe for success here at Pure. It's ensuring that we're going after good, solid, large, growing, and maturing TAMs, coupling that with strong execution, ensuring that we are driving industry-leading innovation, and putting those building blocks together to deliver the multibillion-dollar profitable company that we've talked about on several of our earnings calls.

In the next couple of slides we're going to talk a little bit about that, but again, I want to sort of, before we go there, talk about each one of these boxes for a few moments. You heard Scott talk about the opportunity. Our TAM has been large. It's we believe one of the largest available in enterprise tech from the day we went IPO. The storage market is vast and large, and yet what's interesting to us over the last, call it, couple of quarters, when we've seen what's happening from a machine learning and an AI perspective, we've seen some of the analyst publications come out about just what that's going to do to the storage market. The truth is, no one knows for sure exactly what it's going to be and what's going to happen, but I will tell you that as we talk to customers that we see the use cases that we're putting some of this new technology into, there are good things that we believe happen for Pure relative to market. So it's an attractive market for us and certainly a substantial market for us as we grow the business.

As I said, coupling that with go-to-market execution, you've heard from Dave, Matt, and Kevin, a number of things to be excited about there. We've driven consistent leverage in the business in the sales and marketing line for quite some time now, but as I think going forward, you'll hear us talk a lot about productivity, you'll hear us talk a lot about channel leverage, those things are playing well and we anticipate them play well in the future.

Then finally, innovation. Not much more to say other than what you heard from Matt and Rob today in terms of 25 software features. Again, indeed, this has been the best year ever from an innovation.

If you put those three things together, we fundamentally believe that we've got a greater than \$2 billion opportunity out in what we call our calendar 2020 or what's also known as fiscal 2021. That means the business is growing at greater than 30% CAGR over the course of those next three years to deliver that market opportunity. Again, it's this recipe of market plus execution plus innovation where we believe we've built the building blocks to make this a reality.

The one thing I think is really interesting though is that there are multiple paths to this success. I'm going to spend some time on the slide that I'm going to show you right here on the right-hand side. As we think about growing the business, as we think about assembling the building blocks, both in our sales and marketing and R&D organizations to be successful, we think about the two different product families and product portfolio is that we're offering today—FlashBlade and FlashArray. We look and ask yourselves are we assembling the pieces, are we making the innovation bets, are we growing the sales and marketing organizations appropriately to grow both of those businesses? What I think is interesting about this graph is if you look, it takes a look at a two-dimensional revenue trajectory for our two primary products—FlashBlade and FlashArray.

On the right-hand side in the rows, we ask the question, “What’s our planning assumption for FlashArray?” This is product growth—product revenue growth over time, and on FlashArray we think that a reasonable planning assumption is anywhere between 10% and 40% in terms of CAGR growth over the course of the next three years in FlashArray. We combine that with the next question that we ask for FlashBlade. “How fast can FlashBlade grow relative to the same three or four years as FlashArray?” You heard earlier today, we’ve talked about this on the earnings calls, right now it’s growing 2X the pace of FlashBlade, both in terms of the use cases that we’re seeing, as well as the go-to-market presence we have now versus what we had as an earlier startup, and so what’s interesting to me about this slide is we set a planning assumption between 1X and 2X, and so the intersections of all those points yield a total Company revenue growth number, so imagine a whole bunch of boxes there yielding different year-on-year growth, and what I think is interesting here is that the vast majority of those intersections drive greater than 30% CAGR over the course of the next three years, and certainly, if we look at the market opportunity, if we look at the building blocks that we’ve assembled from a sales and marketing perspective, if we look at the bias for innovation and the continued focus on innovation that we’re going to drive over the next several years, we definitely see this opportunity is something for us to seize.

Turning to leverage, at the same time that we see this substantial opportunity to continue to grow, we’re also going to continue to stay true to our operating principles of driving leverage in the business going forward on a year-on-year basis, and so fundamentally over the course of the next three years, going from a negative 7% guidance midpoint for the current fiscal, fiscal 2018, we believe that as we continue to grow at a greater than 30% CAGR, we can achieve somewhere between positive 6% and positive 10% operating margins in that third out year, and continue to drive predictable and deliberate momentum in each of the years from here forward. We’re going to do that, obviously, across all areas of our op ex. You can see on the right-hand side here there are a series of bullet points that talk a little bit about how we think about the underlying drivers, the underlying leverage that’s inherent in the business. I’m not going to spend a lot of time talking about every one of the bullet points, but I will spend a little bit of time, both on Sales and Marketing, and on R&D.

As it relates to Sales and Marketing, I think the primary driver of sales and marketing leverage is what you heard about earlier today from Dave and his team in terms of productivity. Productivity is ramping very, very well, but there’s still a lot of room to grow. With 60% of our AEs today still with us for less than 15 months, we know that fundamentally there’s a ramp involved in that. It’s something that you see in enterprise technology. Pure Storage is no different, so just assuming the growth and maturing of that sales force bodes very, very well. You combine that with the platform selling approach, that really has started to take off over the course of the last quarter or so as we start providing a portfolio of products and can solve a variety of customer needs as opposed to a specific use case, and so again, we see that leverage, and then fundamentally, we’re going to keep stepping on the gas in terms of building smart and

wise relationships with the channel, so all of those things help continue to march the sales and marketing number—or line item from a leverage perspective.

Turning our attention to R&D, fundamentally, I think you've heard today, Pure is an innovator at heart. We're fundamentally focused on continued innovation. There is no question that we want to lead the industry in overall innovation, both today and in the future, and so we're going to continue to invest. There's no question about that. We're going to invest in headcount. We're going to invest in capital to continue to develop the product and a whole host of ancillary costs that go along with our overall R&D investments, but we're going to do that while we continue to drive leverage, so in a world where we're growing greater than 30% CAGR, the amount of absolute dollars that we can put into incremental R&D is still notable. We'll still be growing our R&D and engineering teams, and yet we will get natural leverage out of the R&D line item as we continue to grow over the course of the next several years, so again, reflecting back on these last two slides, multiple paths to success in terms of greater than 30% CAGR across our two-product portfolio, and doing it with the same operating principles that you've heard from us in the past in terms of ensuring that we're driving leverage on a yearly basis.

All right, so I'm going to go into part three now, and again, this is—this might get technical for folks, and my apologies. Some people really are interested in this, and we want to make sure that as we offer some thoughts about seasonality, share count, the new rev rec standard, people are aware of it so that we're aligned in terms of how you think about your modeling exercises going forward, so I'll zip through these relatively quickly and we can get back to the more important Q&A and Scott's wrap up.

Fundamentally, I talked about this earlier in the—in my speech about P&L seasonality. I know we've talked about it on the call from time to time, but it's important just to step back and see the underlying seasonality that we see in our business, and there's really two dynamics at play. Number one, revenue. The enterprise storage market has always been seasonal. I was talking to Kevin Delane earlier in the day, and we talked again about this. It's something he's seen for over 20 years; fundamentally, we've seen Pure being no different. These seasonality dynamics have been prevalent in the business since I was here almost three years ago, and so what that means is that we tend to start low in Q1 on a seasonality basis, and build into Q4, and we anticipate those dynamics continuing on a go-forward basis. Now, naturally new product innovations or introductions, different things like that can obviously impact the underlying trends a bit, but the underlying seasonal trend in the business is something that we've seen, and it's consistent with the storage industry.

At the same time, from an incremental op ex perspective, we tend to invest early in the year, and this is a playbook that we've been working on and we've been executing on over the last three or four years, and we don't anticipate that to change. Number one, we hire—we like to hire the vast majority of our sales folks at the beginning of the year. That's when they're coming off good years at our competitors and they want to come and have great years here at Pure Storage, and getting those folks in, getting them indoctrinated into our overall culture and our overall program early in the year while seasonality is building, allows us to get them fully—more fully ramped when the big selling times of the year happen in the latter half of the year, so typically you'll see a lot of our sales and marketing investments happen in that first half.

Similarly, in Q2 we're still picking up some tails of hiring from a sales and marketing perspective, and we have our largest marketing end of the year, which you're all sitting at here with Accelerate, so two—these two quarters are the quarters where you can anticipate the most incremental op ex in the business, and then we tend to taper off a little bit more as we go into Q3 and Q4, which fundamentally means that, really, the leverage in the model happens in the second half. We start out in the lowest period. We combine low seasonality and revenue with high incremental op ex, you're going to have a low operating profit in that first quarter, but it tends to build again over the course of the full year.

I want to talk about a couple of other pieces of seasonality in the business. Free cash flow; we've had a number of people ask this question, and I'll talk a little bit about free cash flow. This is a representative free cash flow walk from last year, so this is not forward-looking. This is last year, and this is as expressed as a percentage of revenue, and one of the things to be aware of is that free cash flow tends to lag operating profit by roughly one quarter, so if you had a strong Q4, the billings and the invoicing that you're going to do tends to manifest itself and you collect that cash into Q1, so relatively speaking we have a Q1 that's high. Our lowest point of free cash generation is in Q2, again, following on low revenue seasonality in Q1, and then it tends to build over the course of the year.

Now I want to stress, this is ex ESPP, so we have an employee stock purchase plan, much like many tech companies, that we make that purchase twice a year, once in Q1 and once in Q3, and so what tends to happen is that the ESPP impact causes some variations in what we think about underlying free cash performance in the business. So for those of you who have looked at our earnings materials, you'll see that we very deliberately put out two bars of free cash flow walk because we want investors to understand what the underlying health of the business is excluding the impacts of ESPP, but we also want to be consistent with how most times people measure free cash generation, so you'll continue to see us going forward, but again, if you think about the underlying performance of the business, think Q1 is a good quarter for us. Q2 is the lowest quarter, and then we start building back from there.

Finally, other seasonality, share grants; we did not get this question at the earnings release, but Q1 tends to be our strongest period for share grants, and that's because we do most of our hiring in the first half. We do our primary employee refresh. We do our Exec refresh, and we do the ESPP purchase that I talked about earlier, and so you tend to have a high, low, medium, low trajectory of overall share grants, so if you're taking the share grant that we did in Q1 and multiplying it by four, you will be high. There's just no question about that in terms of how we ultimately think about bleeding in overall shares over the course of the next four quarters.

Speaking of share count, I wanted to talk a little bit about all the various ways that we can measure share count, and it's particularly poignant to do it right now, because as we start approaching positive operating profit and positive EPS, as we all know, share count does change. We go from basic share count to fully diluted share count, and so right now you've heard me on the call talk about our basic share count being 205 million shares. Again, these are all results as of our last Fiscal quarter. Naturally there will be a creep in these numbers, but as we think and as you model and think about positive operating profit, obviously we need to shift to a different share count, and we did a little bit of a walk here for everybody to make it a little bit easier to understand how we think about this, so right now our fully diluted share count, 279 million shares; that's unvested options, that's vested options, and that's unvested RSUs, 270 million shares in the total cap table, but from an EPS share count, you're going to be somewhere in the range of 230 million to 245 million shares, and that's because we have to employ the treasury stock method, which allows for buyback based on proceeds. Now the thing about the treasury stock method, as you can see, is that it varies based on ranges of outcome, so you can't have a swing of up to, call it circa 15 million shares. Now on an EPS basis, it's not a lot, but we want to make sure that people are aware of how to think about this as they want to start thinking about their models in the out years.

Finally, ASC 606, rev rec standard. If there's a five second headline here, very, very little impact to the company. Unlike full-fledged software companies that have much more profound impact on 606, this is largely a non-issue for us. It's going to be mildly accretive on both a revenue basis and an operating profit basis, and the reason is, is the areas that you can see over on the right-hand side. So, whether it be revenue. We have a modest amount of accretion from some of the support contracts that we do sell. From a contract cost perspective, we will now be amortizing those expenses out over the customer expected lifetime, not over the support contract. The largest individual expense is commissions. Typically we've been amortizing that over roughly three years, which is the normal support contract that we offer. Now we'll probably be doing it over the course of roughly about six quarters—or sorry, six years, apologies. Then, evergreen storage accounting, we take a modest amount of revenue to product

revenue when we start provisioning the controllers, whereas historically we've had to amortize that over the life of the contract as well, again, in accordance with 606.

So that's the spin through the technical stuff. I want to bring it back to a final couple of points before I turn it over to Scott. Again, looking at not only what we've done and what we've delivered and what we look to deliver for the year, but also where we're headed for the long haul, and I think it's really summarized here in these three points, and fundamentally we've had a strong, strong degree of execution since the IPO. We've been operating for seven quarters now and always have at least hit, if not exceeded, our guidance in every one of those quarters, both in the top and the bottom line. We've driven best-in-class revenue and gross margins, and we've driven significant operating leverage in the business. We've had a solid start to the—to fiscal '18, both in Q1 actuals, as well as in Q2 guide, and that really gives us confidence in the full year in terms of a number of the dynamics both on the sales and marketing side and the market side playing in our favor.

Finally, as you heard from me earlier today, there are multiple paths to success in our portfolio of products that deliver us 30% or more three year CAGR, and sees us being a greater-than-\$2 billion revenue company in calendar 2020.

With that, I'll turn it over to Scott to wrap up the day before we start with Q&A. Scott?

Scott Dietzen:

Thank you, Tim, and thank you everyone. I know the acoustics in here are not ideal. I can promise you they're better than the other room where there was a generator out back. So one more slide and then we're going to do a Q&A with the leadership team, so get your questions ready, but I want to put the whole picture back together.

We've got a huge market. A \$35 billion total addressable market to start with, then you add in the disruptors: the growth in cloud, the growth in deep learning and AI centric storage, growing to \$5 billion over the next few years. Remember the partnership with NVIDIA that we discussed this morning for attacking that market. Massively parallel AI needs massively parallel storage. Great go-to-market execution to date. We've built a phenomenal team. We've built a channel with our sky high customer satisfaction. We have the best repeat purchase metrics in the business, and we are in a position to continue to scale, to service ever more of this market, and then most of all, innovation.

I mean this—I am 100% convinced this is the best six months the company has enjoyed in bringing innovation to market, with the launch of FlashBlade, the launch of FlashArray//X, and then the software payloads this week, 25 new software features that are delivered to all of those customers and we'll continue to differentiate, improve repeat purchase rates, improve win rates, and so on over time, and then I think Tim did an excellent job of talking about the financial models and how we are on the right cadence to build a multi-billion dollar profitable business in the years ahead, so since IPO, we made that contract with investors that we were going to maintain best-in-class growth, but make material progress in improving our operating efficiency, and I believe we've delivered on both of those and we're going to continue to do so.

With that, I'd love to call up the rest of the management team and we'll do some Q&A. I think we're going to pass around microphones, and so get your hands up, and management team, come on up.

Male Speaker:

Thanks. Nice to meet you, Scott. Good to see you again, Tim and the crew. A couple of questions; one for you, Scott; one for you, Tim. Scott, on first of all the AI opportunity that you've talked about, help me understand why the economics of doing this on-premise are better than doing it on the cloud. It just feels

to me that it just makes—to put all that stuff on an Amazon, on an AWS, rather than doing this on-premise, so why do you think a lot of that storage is going to be on-premise and not reside in the cloud, and then, Tim, for you, on the operating margins, great CAGR for 2020 calendar, I guess. Why only 6% to 10% operating margin? I mean I'm just trying to extrapolate linearly. Does that mean that \$4 billion is only, at that point, will get to a long-term model? Why cannot that happen faster? One would think that at \$2 billion, you should be a model.

Scott Dietzen:

To start, on the locality, I think what NVIDIA said in the keynote this morning was right. You want to do the AI where the data is, so if the data is in the cloud, if that's where it's being captured and generated, then by all means the AI belongs there, because you're already making that investment. The problem is for video and sensor data that is captured off of Internet of Things devices, there is literally not enough bandwidth in the world—I mean we're 20X short the bandwidth we need to move that data into the public cloud, so either you put it on trucks and you do your learning three weeks in arrears as to what happened, or you have to service it in your local data center, and so one of the demos we'll show tomorrow is deep learning on the edge with a rapidly growing data set, and then we'll pick particular components that we want to do more learning in the cloud and forwarded. That's a use case that co-processing AI on the edge, where the vast majority of the data stays on the edge, and some of it gets moved up to the public cloud, but just to cite the keynote stat again, in 2020, 50 new zettabytes in the world, global Internet capacity for the year is one twentieth of that, 2.5 zettabytes.

Tim Riitters:

On your question on leverage, I guess I would summarize it in one word, opportunity. Fundamentally we are playing for number one or number two in an enormous market, and so when we see some of the statistics that Dave talked about in terms of overall sales force productivity, that causes us to step on the gas and capture an ever larger increasing market opportunity. When you combine that with the things that we're seeing in AI and some of these new use cases, if those markets take off, we're going to want to invest in them, whether they be on the sales and marketing side or on the R&D side, so it definitely gives us flexibility, and certainly I would rather have a lower operating margin rate in the near-in years, with the confidence that we're building something for the larger-out years, than to rush to a target model. There's certainly flexibility in the model and we'll continue to manage through that, but when there's those opportunities out there, particularly given what we're now seeing from some of these early FlashBlade use cases, we want to make sure that we give ourselves the right drive power to invest for success.

Wamsi Mohan:

Hi. Thanks for the presentation. Wamsi Mohan, Bank of America Merrill Lynch. I was wondering, Tim, if for the next incremental billion dollars that you're talking about, how much of that do you think is coming from replacing spinning disk, 10K, 15K rpm drives, versus incremental new opportunities that you have spoken about, like AI machine learning, all the things like that? Secondly, Scott, can you address when you talk about this 50X, the huge delta between the data created, if you think about 2020 or 2025, the total amount of capacity, both from Flash as well as from HDDs, is significantly below that, so what actually is happening to bridge that delta between the true capacity that's there in the market versus the projected data growth?

Tim Riitters:

Wamsi, I'll take your question first in terms of that split, and fundamentally, as I think you saw the graph that we provided in terms of those two dimensions in terms of how you grow, I think, as you heard me say, there's a variety of different paths. I think that while people give a little bit of short shrift to what's happening in, what we call, the cloudification of IT, remember, as Dave said, 70%—we fundamentally

believe that 70% of the refresh is still out there, and so it's a great market for us and we'll continue to capitalize on that, and when we're seeing these new green shoots in terms of some of the CAGRs—I think we said \$5 billion. It's probably even closer to \$6 billion in new AI storage. I think we've got two great markets that we're going after, so I don't think we're going to offer a perspective on how much is each. We're going to go out and attack both of them, no question about it.

Scott Dietzen:

I have not actually looked at the forecast for Flash productivity and hard drives that are going to be produced in 2020. We obviously pay a lot of attention to the flash supply, and we do know there are—there's a phase of retooling happening right now, and that there's a lot of new capacity coming online over the course of the next 12, 18 months. How we project that forward to get to everything we need in 2020, we've not done that analysis yet, but I—there's no question that people are going to try to capture as much of this data as they can because there's so much value in it. This is the consistent learning that we've seen across all these customers that are trying the new data-driven analytics, especially deep learning, is they're finding hidden value inside their data that they didn't know they had, and so there's going to be a desire to capture as much of it as possible.

Katy Huberty:

Is it working? Okay. Katy Huberty, Morgan Stanley. I also have two questions; maybe start with Scott. Who are you competing with for the AI machine learning storage? Then bigger picture question, if you're right about the size and the growth and your competitive positioning, it just makes the company that much more strategic, and I know two to three years ago you weren't that interested in entertaining strategic interest. Is that something that changes as the company scales? Is that a more rational path, and then I'll give you a second to think about that. Tim, following on the first financial question, if you're wrong on growth, if it's 40% or it's 20%, are you planning to flex the op ex model to manage to an operating margin? Are we sort of at a mature growth stage where you can manage that, or will we see you, for instance, push out the profitability timeline if growth is 25% and not 30%?

Scott Dietzen:

I'll start, so then you get the couple of seconds to think. In the cloud use cases and the AI use cases, we rarely see the—our legacy incumbent storage competition, so we still see them heavily in the enterprise, but outside it's generally about—it's us or do it yourself storage, so people are putting SSDs or flash cards, PCIe cards into servers and they're running open source software or other things, but they really don't believe the legacy technology. I mean these are forward-looking companies that tend not to want to look backwards 25 years and get solutions that were designed for a different era, and so we really see those opportunities as much, much more greenfield.

In terms of strategic interest, we think our current strategy is working supremely well. We are driving best-in-class growth in a huge available market. We want to keep that growth going and keep our strategic partners—be able to work with key companies like Cisco, like NVIDIA, to capture as much of this market opportunity as possible, and we think the independent path allows us to serve more customers and provides the best possible return to our shareholders, because they get exposed to this phenomenal growth in an incredibly—a big TAM, maybe the biggest available TAM in at least B2B tech.

Tim Riitters:

Then Katy, on your question, across the spectrum—can everybody hear me? Yes? Across the spectrum of outcomes or scenarios that you saw illustrated there in terms of both FlashArray and FlashBlade ramp, we've got definitely dials in the business to manage to those trajectories, so I don't think you see us pulling—pushing that out. It's really when we see some of the benefits we see in terms of productivity

taking shape, we'll keep going. If we don't like what we see, we have the abilities to still manage those dials.

Male Speaker:

I think the only thing I would add to that strategic sort of conversation is that we're having a hell of a lot of fun. When you see all the kind of trends that you have from customers and you see the impact that we're making, and you see customers and partners that are embracing it to take it to market, it's just fun doing this. I've been doing this 30 years and I've never had more fun in my life than building this company, and so I think the strategic partnerships that we have with NVIDIA and with Cisco and a whole bunch of others make us get the best of both worlds, which is we get the reach and the leverage to get to market, while we get to build the business as we have been.

Scott Dietzen:

Hopefully that's no more fun at work than Pure, but you probably had some more fun. Where are we going next? I think we're here.

Aaron Rakers:

Yes, thanks Scott. Aaron Rakers at Stifel. Two questions as well. First, I might want to build on that last comment. The Cisco partnership seems to be a key headline at today's conference. I think it was noted 1,400+ customers. That'd be 40% of your cumulative customer base, so I'm curious of how deep you think that relationship can expand, and then more importantly, what you're seeing with their other partners. Obviously, they've partnered with Dell EMC, FlexPod from NetApp. I'm curious of what you're seeing in terms of go-to-market with Cisco that might be changing, let's say, over the last six months, and then the second question would be on the sales side, how would you characterize your productivity levels relative to what you had outlined at the point of the IPO, and how far away do you think you are from attaining what you think is a reasonable productivity across your sales force?

Scott Dietzen:

Maybe I'll start. We expect to be able to continue to benefit. Dell EMC is working with Cisco's channel on BCE refreshes, and they have the desire to introduce things like Dell servers, and even Force10 networks into quotes in replacement. Cisco's partners tend not to look on that very favorably, and so Cisco and Cisco's channel needs a storage vendor that is supportive of Cisco networking and Cisco compute, and we are as, at least in our view, Dell EMC's strongest competitor, the natural counterweight to go up into those deals, and so we see that in our global channel traction in particular, and I think there's just a huge amount of more opportunity as we expect them to get more and more competitive over time. I don't know what you guys want to add.

Dave Hatfield:

I think there's a natural alignment with market share, so if you look at the number of the market—the amount of market share that UCS has relative to HPE or Dell, we align pretty well with that, so I think that's just general. I think what happens, based upon the dynamic that Dietz was talking about, is that the sales teams are getting beat. The partners are getting asked to share shift and they don't like it, and so those true north partners or those sales reps that are getting beat are getting more engaged with us, so I think that's number two.

Number three is people are looking for converged simplicity. They're looking for a converged solution that are best-of-breed that are not hamstrung, like a hyperconverged environment for the Tier 1 data center where they have to move, compute, and move storage in big blocks together. They want to be

able to provision each of those independently, so while we think hyperconverge has a play, it's a good technology, it works well down market, it works well in remote office and BDI, it complements what's needed in the data center, and I think the sales teams at Cisco are starting to figure that out, and we're the best-of-breed partner to be able to help make that happen.

Male Speaker:

Yes, and Aaron, on your question on productivity, I guess I would offer two things. Number one, productivity is tracking higher than we thought. We're very pleased with what we've seen in the business, and in terms of the question of when do we cap out, or when do we hit that overall aggregate productivity, I guess what I would offer is that the longest serving cohorts are already there. You saw the graph earlier today, and so really it's not a question of when do we hit a specific magic number. It's just ensuring that that next class that we bring in and the class after that and the class after that continue to follow that nice ramp, if not faster, and as you saw Dave talk about earlier today, the last two years of cohorts are performing. Again, early days, but are performing stronger than their counterparts did, and so it's that machinery, when you see that working and you look at the TAM and the market opportunity in front of you, that's what gets us excited.

Alex Kurtz;

Yes. Alex Kurtz with Pac Crest. Back here, guys, in the back row. Just one question. On this path from \$1 billion to \$2 billion that you outlined, I mean I think the natural assumption would be that maybe a system-wide sales re-org—I mean it would be natural progression for a company getting to \$2 billion, layering in vertical teams, layering in global account teams, building out international teams, so that's very disruptive when that happens, and obviously you guys are executing well, so if you could just kind of give us an update of where that sits today and when you might have to implement something like that domestically and globally.

Dave Hatfield:

Thanks, Alex. We've been implementing it incrementally since the beginning. This is not something that you try and turn on overnight. We obviously saw certain—I mean Nutanix talked about this specifically in the earnings call where you're promoting your individual high achievers into DM roles, and so you flatten out the productivity for a couple of quarters. You've got to be really mindful of how you do that, how you develop high potentials, how you train them, make sure they're management ready, make sure you're not promoting great individual contributors into management jobs that they—actually not going to like. So it's been this sprinkling of incremental segmentation by the areas that we talked about, corporate, sled, commercial, commercial select, enterprise, global, and then trying to get the investments right, and where you're incrementally turning those dials, so I think we haven't been perfect, but we've been incremental in rolling that out over the last five years, and we're very excited about it.

I think as you bring in industry vertical specialists—we've also done that, so for the public sector is the first place that we've done it. Fed is just a very different buying process. You need dedicated relationships that are there with the partners, etc. Same with sled, and then healthcare, but most everything else that we have is very, very horizontal, and so it's more about product marketing and solution marketing and packaging than it is about needing dedicated sales forces by vertical, so I think we've got the incremental verticals laid out, particularly in North America, and as we get critical mass in core international markets, we'd do the same thing when we can, so it's more about incrementality than it is about revolution, because they can be very disruptive. Did I get the second part of the question? Is that all right?

Oh, sorry. The second part that I wanted to make is what—the reason why we had Matt come up, in addition to just being a great puritan and sort of committed to what he's doing—I mean building that

repeatable process for new product introductions, it's one thing to get—have a hit. We know we've got a hit with FlashArray. We've got a hit certainly there with FlashStack. To get a third hit, you need a repeatable process to be able to get it embedded, and so we've done this before at Veritas at scale. We've done it, obviously, with Kevin at EMC at scale, and so one plus one needs to equal more than two, and so creating that repeatable process is what—the infrastructure that we've been building, and so far, so good.

Jason Nolan:

Jason Nolan with Baird. Two questions, and hopefully these are quick. FlashBlade, are you offering incentive compensation for the sales teams, specifically hurdles to accelerators or dedicated quota, and how would that look as we get toward X? Then kind of a big picture question on the 30%+ three year CAGR; what's your underlying assumption for the market growth rate? Is it a growing market, a declining market, and maybe it doesn't matter, given you're clearly taking share?

Dave Hatfield:

I'll give the second one to Tim, and I've got an opinion on it, but he'll be more precise, and on the first one, the answer is—the quick answer is no. We answered that on the earnings call. We don't feel like we need it. We have a dedicated sales force that are our best and brightest folks that were evangelical for FlashArray. They are segmented and aligned to each of our districts and regions, and so they carry the number for that district or for that region, as does the District Manager or the Region Manager, so you've got quota accountability that's there that's discreet, and you've got specialists that are accountable that are only getting paid on FlashBlade, so that combination allows for clear accountability in driving both businesses.

Tim Riitters:

On the market opportunity to grow greater than 30%, I would say two things. Number one, you don't have to assume a lot of market growth in the business to make that happen, because as we talked about, still early innings in terms of that shift from spinning discs. It's remarkable when you think about it, because everyone knows that Flash is the way to go, and yet you talk to customers and there's so much still in the fleet that is spinning disk as they come off in maintenance, and so you don't have to expect a lot, and yet I would also refer you on this, at the same time, to the overall CAGRs that you saw in Scott's part of the talk, growing 22%, 25%, 26%, as well as the AI CAGR, and so you don't have to believe a lot of growth, but fundamentally we believe there's growth out there to be had in the markets.

RK:

Hi. This is RK with JP Morgan. Thanks for taking my question. I was wondering if you could give us more color on your guidance for FlashArray. The guidance range seems quite huge, from 10% to 40%, so could you talk about what kind of scenario in which you'll get the 10% growth, what kind of scenario in which you'll get the 40% growth?

Tim Riitters:

I'm going to probably pass on that. I think what we wanted to do was just illustrate the art of the possible, and I think the key takeaway is not necessarily 10% or 40%, or 15% and 35%, but the fact that the vast majority of those intersections with our two-product portfolios yield us comfortably above negative—or comfortably above 30% year-on-year growth, and so it really was a planning exercise to illustrate that there's a lot of green space there to be had.

Tim Long:

Hi. Tim Long at BMO here. Just a question on the margin side again, on the gross margins, you talk about best-in-breed margins. As you go forward in scale, number one, do you think you're somewhat limited on how high you can go given where the rest of the industry is? Is there an impact there? Then secondly, today you announced 25 new software offerings. Do you think there'll be a time in the future as you get bigger that there might be maybe a better monetization of the software model, either as an extra revenue driver and/or a margin enhancer? Thanks.

Dave Hatfield:

I mean I think overall what's becoming more clear to folks is that what we sell are subscriptions over time, and so we have a cap ex buy upfront, but we sell an ongoing subscription that gives customers a service contract for their gear, as well as a subscription to future software innovation, and a subscription to incremental upgrades to the hardware on an occasional basis every three years, as it is right now, for controller upgrades, and I think we believe that by far the most important thing is to drive more and more adoptions of our subscriptions by getting more terabytes onto the platform, more terabytes under coverage by Pure, and keeping people never looking at going anywhere else off the platform, and so transitioning much more of the commerce mentality to that kind of easy, upfront purchase, ongoing subscription mentality, value add constantly coming in, to make sure we're keeping people on the platform, gaining the incremental terabyte. I mean you have this industry where terabytes are always growing, so as long as you get more terabytes on the platform and monetize them well, you win, and so I think the way we've thought about it is exactly that, and we think there's still a lot of upside in continuing to not only transition legacy environments as they come up for their three year refreshes to go after these next generation use cases.

Tim Riitters:

And to the question in terms of what's the right range, we've been operating in that range now—in that guidance range that we've offered for six quarters, and we really feel that's the sweet spot. I mean it's still 10,15 points above our competition, and what we don't want to do is squeeze too much out of that, because you've got to be really careful in terms of slowing velocity, and I think we feel very, very comfortable being in that sweet spot, continuing to drive leverage as well on the support line, so for those of you that heard our last earnings call, we continue to see a lot of goodness there and we're going to continue to manage that very, very well, deliver investors what they would expect in terms of differentiation, and yet do it in a sweet spot that allows us to keep growing the way we want to grow.

Scott Dietzen:

I just want to add, in the growth phase of the market you want to keep your go-to-market really simple. All of our business is fulfilled through the channel so that it's really simple, that we've got two line items right on our invoice, that product and subscription is really powerful, so keeping it simple, giving all of the software features in order to grow share. We're a top 10 storage vendor in the world with a growth story in the top 10. We've got these greenfield markets that we're playing in that are growing really rapidly. We just want to capture as much market share as we can, especially given, once we capture a customer, it's much more economical for them to stay with us, and they're happy. They're much happier than they were with their prior vendor, and you see the differentiation in the margins that we're driving on the existing systems. We're delivering that much more value to customers that they're getting a great return on their investment, even at the margin rates that we're driving.

Dave Ryzhik:

Thanks. Dave Ryzhik, Susquehanna. Scott, so all of your major competitors now also have a hyperconverged offering, and I understand it's a different product, different use case, but how important is

it? Is this something that you can cultivate internally in the future? For Tim, also some of your competitors have stressed a consumption model. HPE has discussed a flex capacity. Are you working with any—do you have any financing options, or with any partners in advancing this as well? Thanks.

Scott Dietzen:

On hyperconverged, there's no question that the model is popular in a certain end of the market. Where we see it most is our smaller customers; remote and branch office, BDI use cases for smaller businesses where we'll tend to run into it. It's interesting that we don't see it in those cloud deployments, and we don't see it in AI deployments, and I think that is reflective of the architecture not being a great fit for a scaled data center. We talked in the keynote today about fast networks turning hyper-converged inside out, because remote storage in Iraq is now arguably even closer to the CPU than local storage, because it's less overhead to access it. So when you basically make these hyper-dense systems and all of these different cores can access remote storage, we're in a model where we can deliver better performance for lower cost in a converged infrastructure, and to be clear, that's what the cloud is built on, that's what these big AI systems are built on, and that's what enterprises run, is converged infrastructure, so we think that's the sweet spot that we're going to continue to target in the market.

Tim Riitters:

I might turn it over to Matt to talk about some of the product offerings we have from a consumption-based model because if you weren't at our keynote today, we talked a little bit about that, so I'll turn it over to him.

Matt Kixmoeller:

Yes, I mean the other thing I'd say, just on the last question for a second, as we go into more of these data-driven workloads, we almost see the opposite end of the spectrum. A lot of the highest performance data environments, a lot of the next gen web scale apps, they don't even run virtualized. I mean they run physical, and so people are focused on the highest degree of performance, the ability to implement compute and storage independently, etc.

The other thing I'd say is that the hyperconverged movement is also something that's going to be very interesting to see how it happens and rolls out across the channel, where, at the end of the day, one of our competitors, NetApp introduced a semi-conversion infrastructure product, and that's going to put them in competition with some of their partnerships, Cisco most notably, and so how that evolves over time and what that means, we think is interesting. Right now we're very focused on—focusing on advancing our platform, going after the most data-driven use cases, and we think that's going to deliver the highest return and the highest value added storage use cases for us.

Then as far as the consumption models go, I think you might have noticed, and I mentioned earlier, we've evolved our branding a little bit around evergreen where we're really talking about it as a storage consumption service. People want to subscribe to well-structured file and object and storage services, and kind of buy them as a utility, and so evergreen, in many ways, is evolving closer and closer to that full utility offering. Right now what we do is we bundle in the initial cap ex purchase and you subscribe to an ongoing stream of improvements through software and some hardware upgrades. In the future, we reserve the right to go further, and we'll explore that, but for right now we feel like we're in a good position with evergreen, where we're offering that flexibility and something very different from what's out there. If you look at what most of the competitors offer in the market today in more utilized offerings, they're really just kind of commercial lease agreements that are bundled in different ways, and so we think there's a lot of opportunity in the future to rethink those.

Dave Hatfield:

Just two quick add-ons on that. We did announce today that with Cisco Capital, we're working together to develop a utility offering for FlashStack, so full stop, and we have a bunch of customers that are using financing options today through our partners to buy it in that capacity. They are operating leases, at the end of the day that are masked as sort of utility models, but they meet the op ex requirements of customers, so more to come on that.

Kirk Materne:

Scott, Kirk Materne with Evercore ISI. Just a question about your SaaS customers and some of the trends we've seen about them making bigger public announcements about trying to leverage AWS. Could you just talk about how you all see that playing out, because I think it's a question of—people are obviously concerned about your opportunity of getting subsumed by some of the big public cloud vendors, and we have big SaaS companies making these decisions that they want to try AWS at least. Could you just talk about how you fit into that sort of dynamic, or how you think that dynamic will play out?

Scott Dietzen:

Well, there's no question, a public cloud is a phenomena, and it's actually been a great change agent in driving the way people think about IT and the simplicity that they want in the consumption model. Many of our SaaS customers are committed to running their own data centers. You look at the large consumer Internet space. You look at companies like Salesforce Service now, these companies are extremely committed to their own data centers because they will tell you they can offer a higher quality of service, higher performance, and do so for lower cost, and assure their customers of security, ownership of the data all of the time, and so they use that as a selling feature. Again, we're very comfortable in this world where it's going to be multi-cloud. Some of the services are going to go into public cloud, but we know big data, especially edge use cases, can't go into public cloud. So, as companies run hybrid environments, we're going to help. We're going to sell to the cloud, and we're going to sell to the on-prem side, and we're going to help all of that work together.

Simon Leopold:

Thank you. Simon Leopold with Raymond James. I wanted to talk a little bit about how you see the competitive environment evolving. I would think you've benefited somewhat from EMC Dell being in a state of chaos, HPE being distracted, NetApp seems to be turning a corner. So if you agree, one, that the legacy guys are turning a corner, how does that change the landscape for you now that they've sort of awoken to the all-Flash battle?

Matt Kixmoeller:

I'll start, and maybe Hat will have some additions. I'd say if you look across all three of our main competitors you articulated, we see a very similar phenomenon, which is they've all been under a lot of pressure to show all Flash gains, and so they went to kind of the most expedient way to do so, which was retrofitting their legacy platforms. We've seen them all kind of pivot back to that retrofit motion, and in many ways as a result of that, we actually feel like our technology lead has widened against the three of them, and so when I think about how a storage vendor's doing, step one of being a functional storage provider is you keep your customers and you keep their refresh cycle happening, and for many years, some of those vendors weren't doing that. They were even losing existing customers. They've been able to go and start to shift Flash in their systems to try to drive some of those upgrade cycles to Flash and leverage some of their installed base bundling ability to do so, but what we don't see them doing is really going after net new and taking on net new customers, as well as serving some of these net new use cases, and so I think once we see some level of investment in growth in that area, that'll be more interesting, and back to my original point, we have seen them actually back away from innovation in their

next gen platforms, and I think that's in stark contrast to what Pure is doing. You saw me earlier talk about how we have kind of three different pillars of innovation technically, where we're software driven, we have our cloud and IOT stack now powered by AI, and we engineer hardware, and we engineer all three of these things together, and we believe we're building fundamental capabilities to do that, to integrate and innovate in these three ways that we just don't see our competitors doing.

Dave Hatfield:

I think the only thing that I would add to that is that we've taken on sort of the biggest competitors over the last five and half, six years of selling, without any distribution and taking share from them. We talk about our net new customer acquisition rates at four or five new customers a day. All of them have legacy incumbent relationships that go back 5, 10, 15, 20 years, so our focus and our ability to win business is based primarily on our innovation, our business model, and our customer experience being that much better than anybody else has as an alternative, otherwise they wouldn't do it, so I think that that phenomenon, especially with the release of all the software and cloud capabilities that we're delivering in the data platform that you heard about today, just expands our moat, so we're going to stay focused on going and acquiring customers, regardless of what the competitive landscape provides.

We are benefiting from the competitive landscape in disarray as well, and I think that'll continue for the next several years. You're talking about organizations that are having to go through a heart transplant. This is highly cannibalistic to their core business, so in order for them to go create a selling motion that chews away at their top line and their bottom line, that is unnatural. I've run those businesses before. That's hard to do, and that's one of the reasons why I joined Pure, is that—and why we fundamentally believe that an independent company needed to be built to go take advantage of it, and just keep our finger on the innovation curve and pressing the efficiency down to make it harder and harder and more and more cannibalistic to them, and better and better for the market and customers and for us, and so that's—that competitive dynamic—we didn't call Dell buying EMC. We didn't call HPE buying Nimble or SimpliVity, etc., but we did know there was going to be massive disruption, and they're going to have to go through a heart transplant. We'll benefit from that, but they got big because they did something well, and so we have a lot of respect for them, and so we focus very, very much on going and just out-executing them day in and day out, and now we're at \$1 billion in scale and nobody else has done that in 20 years except for NetApp and EMC in this space.

So, now we've got the resources to be able to go take these folks on, and we've got next generation innovations and new products to be able to do it at a platform level, so very optimistic about our ability to continue, but I think we're benefiting from it. I think NetApp, you mentioned, is sort of turning the corner. I think they're the only one financially that is showing some improvements, but they're doing so largely by cutting costs in R&D, cutting costs in support, cutting costs in sales and marketing, which ultimately catches up to you. They were late to the game, as Kix talked about, and sort of moved over flash, and so I think that's—it's kind of good to have a little bit of health, and I think they're benefiting because they've got a good chunk of file customers that, with the disruption at Dell and EMC, they're probably getting looks at the block business and taking advantage of that, but if we get the at-bat against any of these competitors, we feel great, and our win rates are holding, so our biggest challenge is about going and getting at-bats, creating awareness, using a megaphone, things like this, and getting routes to market that can get us more looks. We do that, I like our chances.

Male Speaker:

I just wanted to highlight one thing that we mentioned in the keynote, which is we really look at software as taking center stage, and going back to Kix's three pillars, we really look at our innovation as being software driven and hardware accelerated, and when you look at the actions of our three competitors that you highlighted there, whether they're—however they're trying to turn the corner, whether it's through SSD retrofits or share shifting or so on an so forth, what they're not doing is fundamentally innovating in

the software level, and so in order to really go after those net new workloads, some of those data-driven use cases, those AI use cases that we've talked so much about, they fundamentally are going to need to do that in order to evolve their technology. Without moving that software forward, just throwing SSDs, just throwing more hardware at the problem, really at the end of the day they're going to be left with a half flash array.

Scott Dietzen:

So we're all going to chime in on this question. Last point is one quarter does not a trend line make. If you look across the incumbents, we've seen systematic declines in their business over the last two years. Certainly I expect that trend to continue going forward. We are seeing a whole scale shift away from the 25 year old technology to technology that supports the cloud model, supports data driven, AI, can work in these modern use cases, and it's much easier and denser and ready for modern data center architecture. The legacy vendors are going to continue to hit a wall in these new use cases, and that's all to our benefit.

Michael Pak:

Hey guys. Just FYI, we're going to take one more question, but the panel will stay right here and we're going to naturally turn this into a reception so that you guys can have an extended Q&A with some food and drinks.

Eric Martinuzzi:

It's Eric Martinuzzi with Lake Street Capital Markets. The question's for Tim and it's on the gross margins for the view out to 2020, or I should say FY21. You talk about that range. We're kind of already there, that 63% to 68%, and you just finished a Fiscal Year where you did 66%. A layer deeper there I would have expected some expansion on the support side. Does that imply a contraction on the product gross margin? That's my question.

Tim Riitters:

I don't think that's necessarily the case. I think that we've offered that range of 63% to 68% from a long-term perspective, so obviously there's still two points more in that range. What I would say is that we're going to continue to drive differentiation. We're going to continue to drive discipline from a sales perspective. Dave's teams have been phenomenal about really driving value in each one of the engagements that we have in selling customers, so we're going to continue to stay there, and at the same time we've got that other lever in terms of continuing to drive support efficiencies. Again, I'll reiterate what I said in the prepared remarks. To be able to be at industry-leading margins—sorry, industry-leading customer sat and industry-leading sort of support gross margins is something we're particularly proud about, but I think it goes back to how we fundamentally engineered the technology. You heard a number of people in the panel today talk about, "I don't have to call Pure for a Sev1. The things just don't break," and when you do that, you get leverage on the support side. So, I think both of those dynamics, both of those areas, both product and support, are opportunities to continue to sustain very high.

Dave Hatfield:

Just one added comment. It's easy to sell value when you've got a product this good.

Scott Dietzen:

All right. Thank you so much for spending so much time with us today. We are hanging around. Have a drink with us. Look forward to meeting you, and we'll see you next year. If you have feedback for us on

how we can improve this forum, other than acoustics, we've got the acoustics one, but any other feedback, please give it to our IR team. We'll be hugely grateful.