24-Feb-2021

Seagate Technology Plc (STX)

Analyst Meeting

CORPORATE PARTICIPANTS

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

John Morris
Chief Technology Officer, Seagate Technology Plc

Jeff Fochtman
Senior Vice President, Business and Marketing, Seagate Technology Plc

Ken Claffey
Senior Vice President-Enterprise Data Solutions, Seagate Technology Plc

Ravi Naik
Executive Vice President-Storage Services & Chief Information Officer, Seagate Technology Plc

Sudhakar Chilukuri
Vice President & Chief Information Officer, Lattice Semiconductor

Eric Charles Schwartz
Chief Strategy & Development Officer, Equinix, Inc.

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

MANAGEMENT DISCUSSION SECTION

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

Hello everyone and welcome to Seagate's 2021 Virtual Analyst Event. I'm Shanye Hudson, Head of Investor Relations and Treasury and I'm thrilled that you've chosen to spend time with us here today.

Before we begin, I want to remind everyone that today's presentation will include forward-looking statements. Information concerning our risks, uncertainties and other factors that could cause results to differ from these forward-looking statements are contained in the SEC filings which can be found in the Investor section of our website. We will also be presenting non-GAAP financial measures. You can find reconciliations to the closest GAAP financial metrics in the presentation accompanying today's event which will be published on the website later this afternoon.

We've got a great program planned for you today. I'm joined by several members of our leadership team who will walk you through Seagate's strategy for building on our leadership in mass data storage to capture an expanding set of opportunities and deliver value for all of our stakeholders, customers, suppliers, employees, and shareholders.

Our CEO, Dave Mosley will start things off. He'll be followed by John Morris, our CTO, who will speak to the strength of our technology innovation; next up, Jeff Fochtman, who heads our Global Marketing team, and he's ideal for presenting the breadth of our product portfolio. After Jeff we have Ken Claffey, who leads our Systems Business as well as our CIO, Ravi Naik, who together will talk about Seagate's new Lyve initiative. And finally, CFO, Gianluca Romano will share our financial outlook.

Following today's presentations I'll moderate a question-and-answer session with both Dave and Gianluca. We expect total runtime to be about three hours and include a 10-minute break.

But before Dave kicks off the presentations, we'd like to share a brief video highlighting the vital role Seagate plays in unleashing the potential in maximizing the value of data. Thank you.

[Video Presentation] (00:02:02-00:03:14)

William David Mosley

Chief Executive Officer & Director, Seagate Technology Plc

Welcome everyone to our 2021 Seagate Analyst Day. Hopefully, after today's program, you'll have a refreshed view of our story and our strategy and share our excitement about Seagate's future opportunities in the rapidly expanding data sphere. I'm going to tell you a very simple story this morning and then I'm going to let our outstanding team of technical experts give you all the details on how we're bringing that story to life.

When we're done, this is what we hope you take away: Data growth is exploding and that will only continue; Seagate is the mass capacity leader backed by over 40 years of HDD technology innovation. We have an established supply chain, strong customer relationships with the right products and technologies to address secular demand for mass capacity storage. And as the datasphere scales and the need to securely capture, move, and store data across the distributed enterprise increases, new opportunities are arising for Seagate. We have a plan to capitalize; and in so doing, we believe we're in a great position to drive long-term value creation. That, combined with our demonstrated commitment to capital returns is a compelling story for investors.

Before diving into the future, I want to make a few remarks about the year we just concluded. 2020 was shaped of course by the global pandemic, which had a profound impact on all of our lives. Against this backdrop, Seagate met our customer commitments while managing through many intricate supply chain disruptions affecting ourselves and our partners globally. We executed on our technology roadmap, ramping 16 terabyte drives, launching 18 terabytes and delivering on our commitment to ship the first HAMR 20 terabyte drives by calendar year-end. You'll hear more about our HAMR roadmap today, which will extend our leadership in mass capacity drives. And you'll also hear about our Lyve storage platform which we introduced in 2020. Lyve is a natural extension of Seagate's mass capacity data portfolio, helping customers' unlock data value. We're excited to share more about it today.

While navigating challenges and executing on opportunities, we also achieved financial results for the calendar year in line with our financial model, growing revenue 2%, growing operating profit 8%, and enhancing our capital returns program with an increase in the dividend and a greatly expanded share repurchase program. A difficult, but extraordinary year and one we're excited to build upon in 2021 and beyond.

This chart likely comes as no surprise. The global datasphere that we serve is exploding, doubling every three years. In just four years' time we'll reach 175 zettabytes of data created annually. More importantly, an increasing amount of this data will be put to work. Algorithms and devices will be optimized around processing and learning from that data. This massive growth in data fuels a race amongst enterprises to unlock the potential of their data and gain competitive advantage. And this growth in data continues to drive very attractive mass capacity growth. Seagate's ability to deliver cost efficient, reliable storage technologies is a key driver behind the growing demand for mass capacity storage which continues at a compound rate of about 35%.

We see mass capacity as storing more than 60% of exabytes in 2026. That means the world is relying on Seagate and the HDD industry to preserve a huge amount of its data for years to come. To meet that demand, the HDD industry must ship nearly 4 zettabytes more mass capacity storage in 2026 compared with what was shipped in 2020. Other technologies that support the expanding datasphere will grow as well, but they are

complementary to mass capacity storage and not competitive products. The key point is that all this data needs to be economically stored, sometimes temporarily and sometimes forever.

For the last 5 years to 10 years, the biggest driver of mass capacity storage demand has been the transition to the cloud. In 2020, 40% of all data was being stored in the cloud. Over that same period, Seagate pivoted our product portfolio to address this growing part of the markets, a transition that's been a key driver of Seagate's success and ability to drive value. For cloud datacenter architectures, mass capacity storage delivers the most cost efficient solution for preserving vast amounts of data and complements flash, another emerging memory solutions.

The world's leading cloud data centers rely on mass capacity drives to store approximately 90% of their exabytes. And we expect that trend to continue far into the future. That is due in part to customers being able to rely on Seagate to continue driving innovation and delivering lower total cost of ownership storage solutions.

We have developed a strong portfolio of HDD Technology Solutions, including our common platform drive family that extends from 16 terabytes and 18 terabytes to 20 terabytes using PMR technology and scales beyond that using HAMR and SMR. We delivered on our commitment to ship 20 terabyte HAMR drives in 2020 and we expect to ship 50 terabytes in 2026 with a roadmap to ship 100 terabytes by 2030, achieving roughly 10 terabytes per disk.

Magnetic recording, which has been instrumental in delivering economic value to the datasphere, will continue to thrive well into the 2030s. Our decades of innovation extend beyond HDD hardware, however. We have deep technical relationships with the biggest customers and partners in the datasphere. We design our own ASIC chips which are shipping in many of our high capacity drives today and we have extensive software experience from years of integration into complex application spaces in many markets.

You can imagine the expertise we've developed by integrating into so many chipsets and BIOSs and file systems and object stores, et cetera. We're simplifying some of those challenges now by pioneering a true, open source object store for the future of big data called CORTX. Seagate's technology and product roadmaps are stronger than ever. And when combined with our industry knowhow and customer relationships, we are in an excellent position to deliver value to our customers and capture the growing opportunities in mass data management.

At Seagate we understand device physics very well and we pay close attention to all competing technologies, understanding their respective strengths and weaknesses. With cost being such a key driver behind storage economics, we believe that through Seagate's strong technology roadmap, including HAMR, we are able to deliver on customer big data demand and maintain constant TCO advantage versus competing technologies over the foreseeable future.

Against that technology backdrop, our outlook for HDD demand is strong. We expect the mass capacity TAM to double over the next four years reaching about $24 billion in 2025 and increasing to $26 billion in 2026. The growth is fueled by strong cloud data center demand for nearline drives, but also, with the increase in video and image applications at the Edge. By now we're all familiar with a myriad of different technologies that are driving the creation of data at the Edge. There are Smart City applications looking at traffic patterns, consumer behavior providing safety and security, or the smart factory applications helping businesses drive efficiencies. There are autonomous vehicles collecting tons of information for dynamic learning, or healthcare and genomics applications analyzing enormous datasets to help fight disease.

The list of innovations and uses for data is long and growing and driving the need for mass capacity storage in both the cloud and at the Edge. However, at the heart of it all is the ability to capture, move, and ultimately preserve the right data in a central location to drive actionable insights and extract value.

In the first half of our fiscal year, Seagate shipped about 40 exabytes for video and image applications at the Edge. By our estimates, those drives get overwritten about four to five times each quarter. That would imply 360 exabytes or more of new capacity each year that is trapped in the distributed Edge and not being fully used. The growth in information is also leading to a shift in data gravity; more data being stored closer to the sources of creation.

Over the next five years we project 70% of the data being stored in the cloud and at the Edge will end up in a sprawling, more complex distributed enterprise. The data storage landscape is thus evolving quickly. The Edge is being shaped by latency, security, and economic needs. For example, applications requiring immediate access to data or caching would reside near the end of the network in a Micro Edge, very close to the source of data creation, whereas applications such as AI training that require massive pools of data and heavy workloads would reside in a Metro Edge. These highly interconnected data centers need to support the flow of data from the Edge devices that I spoke about earlier, up to the public clouds.

The distributed environment is creating costs, scale and complexity challenges for CIOs and IT professionals to manage and ultimately put more data to work. As the mass data leader, Seagate has worked with the breadth of all these customers and we understand these challenges.

Starting with our leadership in mass capacity storage and broad capabilities in software host and device silicon design and detailed contributions to many industry standards, we've accumulated over four decades of expertise integrating into many different kinds of systems and architectures. And we have the deep customer relationships from consumers to enterprise and large cloud data centers which collectively make us the right company to address all of these challenges, which is why last year we introduced the Lyve storage platform.

The Lyve platform offers a simple, cost efficient and secure way to manage massive volumes of data across the distributed enterprise. Lyve Mobile takes our years of expertise and feedback from consumers and creates an enterprise quality solution to transfer mass capacity data between the endpoints, edge, and core cloud. Since introducing Lyve Mobile a year ago, we've established 10 important customer engagements, some of which we'll discuss later today.

Lyve Rack, powered by CORTX which is an open-source and object-based software leverages our mass capacity HDD innovations and systems expertise to provide enterprises with the lowest cost per petabyte.

In September we introduced Lyve Rack and in just a few months are running eight customer engagements, including a multitenant co-located data center.

And today we're excited to introduce the next piece of our Lyve platform, Lyve Cloud, an object-based Storage-as-a-Service platform providing mass capacity storage at the Metro Edge. Lyve Cloud is the next logical step in our Lyve platform, and you'll hear much more about it as we move through today's presentations.

At a high level we believe our customers; own their own data, and we've designed a simple, secure, cost effective option to help them capture and activate their data at the Metro Edge and move more data to the cloud. In collaboration with Equinix, the world's digital infrastructure company, we are expanding into four US locations by the end of calendar 2021. With our first customer already on board, there's strong interest from many others. Our

Lyve platform considers the range of cost and complexity issues faced by the distributed enterprise. Rack is designed to provide customers with the lowest TCO option for hybrid or Edge cloud and mass capacity; Mobile, for enabling fast massive data transfer; and our Cloud for capturing and activating data.

With the addition of Lyve, Seagate's mass data solutions are ubiquitous across the new storage landscape. I spoke earlier about cloud demand driving a two times growth in HDD mass capacity to around $24 billion by 2025. We believe the opportunity to bring mass capacity closer to the Edge and combine storage with an innovative data management platform will open up $50 billion market to Seagate in the 2025 timeframe. Leveraging our strength of our mass capacity IP, the Lyve platform offers top-line growth at very attractive unit economics. It's still early days, but the Seagate advantages I've covered today, leadership in mass capacity, our broad capabilities in software, chip design and data management accumulated over four decades ,and our scale and deep industry knowledge make us the right company to embrace and capitalize on this opportunity. We look forward to diving in deeper in today's presentations.

The past performance and strategic vision that I've shared with you today is underscored by Seagate's focus on creating value for all stakeholders. A stakeholder-focused company succeeds only if it adheres to a clear set of core values. These are our values: First, at Seagate, we lead with integrity, striving to balance the planet and its people while generating a profit to build our future. Long-term integrity means sustainability and relationships, and we have delivered that through a commitment to enhancing value for our customers, employees, supply chain partners and investors. For shareholders, we've demonstrated that commitment over time. In the last five years, we have returned $7 billion to investors. This includes 6% growth in the dividend in that period and the repurchase of about 80 million shares. Today, we are building further on our commitment. Our board has approved an increase in our share repurchase authorization of another $2 billion. Above all else, we are committed to allocating all of our capital with the goal of generating the highest, long-term return consistent with our core values. The significant market opportunities ahead, together with our strong business model and cash generation capabilities give us confidence in managing ROI in 2021 and beyond and in our ability to continue to enhance value for our shareholders.

Our second core value is innovation. We cultivate a team with curiosity who works alongside our customers to deliver mass storage economics. And innovation extends beyond the products we build to drive efficiency across manufacturing and supply chains with agility and the right solutions for our customers.

And third, very importantly, embracing inclusion. Continuing to build a more equitable, safe, and inclusive workplace for our employees and expecting the same of our suppliers and partners. We've tried to capture the spirit of these values in the short video I'll share with you now.

[Video Presentation] (00:18:25-00:19:27)

At Seagate, we focus on people, planet and profit. We balance investing in our business and our employees, doing the right thing for our customers, and returning value to shareholders. As you just heard, we have a bright future ahead with opportunities to drive revenue growth, continue strong cash generation and enhance shareholder returns. We would not be a successful company or a productive company without delivering consistently on the commitments we make to all of our stakeholders.

And to wrap up, I want to take a second to recognize the leaders of our organization who embody those commitments and ensure we deliver on them. I am very proud of this leadership team which combines members that have worked together for a long time, as well as new members who further strengthen our bench of talent. Together, we are poised to extend our track record of success.

And now let me hand off to our CTO, John Morris.

John Morris

Chief Technology Officer, Seagate Technology Plc

Hi, this is John Morris, Chief Technology Officer at Seagate. As you just heard from Dave, the datasphere expansion and demand for mass capacity storage continues at an exponential rate and Seagate is in a prime position to move and store more of it on our devices and systems than anyone else. To meet this challenge requires significant innovation and mass capacity storage, and hard disk drive innovation is at the heart of Seagate. With a proven track record of innovation that spans 40 years of industry firsts, from recording technology, to form factors, to interfaces, to performance solutions, Seagate has delivered.

Today I want to share more about our recording technology innovation and the path to significantly higher capacity devices through areal density growth. Areal density growth is fundamental to the long-term value of any storage device. In the next few minutes, I'll show why we have confidence in our ability to grow hard disk drive areal density, and not through a one-time gain, but through strong and consistent growth that can scale with already staged component innovation.

I will finish with a view of infrastructure total cost of ownership that will show why we believe hard disk drives will continue to be the mass capacity storage choice for the foreseeable future.

Let's start with an overview of where we are today in hard disk drive recording technology. First, take a look at the figure on the left where you can see a visual for a track of data. The yellow box represents a bit. This highlights the bit dimensions, the grain structure of the media and how the area of a single bit leads to a definition of areal density, the key enabler for capacity gains. We'll come back to this later.

On the right is a graph highlighting the past 25 years of areal density growth and hard disk drives. You can see that perpendicular recording technology has delivered 15 years of solid capacity gains. But, as with longitudinal recording technology in 2005, there comes a time when it reaches a point of diminishing returns. We are now at about 1.1 terabits per square inch and delivering 18 terabytes on our nine-disk mass capacity drive. And while we will be releasing more perpendicular products in the coming years, it is clear that a transition to a new recording technology will breathe new life into mass capacity storage.

With this in mind, we have staged Heat-Assisted Magnetic Recording or HAMR to enable robust and sustainable areal density growth over the next 15 years. Recall that with HAMR we use a new higher coercivity iron platinum magnetic alloy in the media to allow for grain size reductions. To write data, we require a change in the head to deal with the higher coercivity. The laser light shown on the top of the left image energizes a Near Field Transducer or NFT. This generates an electric field in the NFT which couples into the media. The iron-platinum is very quickly heated and this temporarily reduces the magnetic coercivity in the heated region, allowing it to be written by a conventional perpendicular writer. After cooling, the grains are extremely stable and not switchable by the head field alone. This whole process occurs in under two nanoseconds.

Now, let's take another look at how the two planer dimensions of track density or cross track width and linear density or down track pitch in the media result in a single bit with a one terabit per square inch example. Recall, the data is stored in circular concentric tracks on the disk. And for a given disk drive design, the dimensions of the written data are set by the intrinsic geometry, the design and materials in the heads and media. Today, at 1 terabit per square inch, these tracks of data are just under 50 nanometres in width, and individual bits are just over 10 nanometers in length. This represents about where we are today with our 16 terabyte drive.

Now, we'll look at the last five years of HAMR areal density achievements in our engineering labs. Let's first look at 2016 where we achieved just over 1 terabit per square inch with HAMR – and note that this translates to 18 terabytes on a nine-disk drive. I will now show the progression of engineering lab areal density demonstrations on the left figure, and what was achieved in track density and linear density on the right figure. And note that the graph scale will stay the same.

In 2017, you can now see a cluster of laboratory demonstrations in the 1 terabit to 1.5 terabit per square inch range. And in 2018, a sharp increase in areal density and beginning to get close to 2 terabits per square inch. In 2019, we maintained similar growth to what we achieved in 2018 and demonstrated over 2 terabits per square inch. Occasionally, you'll notice that there appears to be some backward movement in areal density. This is associated with our pipeline of experiments as we explore the design space. And finally in 2020, where we achieved 2.6 terabits per square inch which is capable of supporting 40 terabytes in our nine-disk nearline drive. Here, we collect the past five years and you can now see that the five-year trend of compound areal density growth is over 20% per year. These graphs also highlight that we're at the bottom of a new technology S-curve with HAMR and underscores our belief that this is providing robust and sustainable areal density growth. What we can now do is create a rule of thumb that supports a three-year to five-year projection for the time from an engineering lab demo to a product intercept using our recent 20 terabyte HAMR product shipments as an example.

The gray region captures this three-year to five-year offset from lab demos and highlights the feasibility for achieving a product shipment of 50 terabytes by 2026. By staying on the 22% CAGR line through early next year, we expect to demonstrate the areal density required in our lab. And this demo will establish feasibility for a 50 terabyte drive in 2026 as shown in the figure.

Let's take a look at an outline of four major phases of technology staging in HAMR. Here, we focus on the evolution of our media design, which, in conjunction with head, drive and electronics technologies provides the path to future growth. First, we highlight our current status with technology demonstrations, namely 2.6 terabits per square inch and support for a 40 terabyte drive. We see an opportunity to scale this design space with granular media into the range of 4 to 6 terabits per square inch, at which point we plan to add patterning in one dimension through the use of ordered grain media. This, we expect to be a stepping stone in media to open up the range of 5 to 7 terabits per square inch. Then we will transition to fully patterned media to open up densities to 8 terabits per square inch and even higher. With the areal density CAGR just introduced, we have a path to 10 terabytes per disk by 2030. This then represents our outlook for technology limits over the next 10 to 15 years.

Now, we'll shift our attention to Total Cost of Ownership or TCO. TCO is the fundamental metric that dictates the buying behavior when deploying storage at scale. It encompasses the capital and operating costs of all the components required to deploy the storage including the devices, compute and network, physical enclosures, and cost of power. Here, you can see the TCO breakdown of a hyperscale hard disk drive deployment and compared with a hypothetical deployment of SSDs of equivalent capacity. You can see that the power, compute and networking are only minor components of the TCO in these systems. That isn't to say there isn't significant innovation occurring in these categories, just that they're more impactful to compute-heavy infrastructure than they are in mass capacity storage. Also, while there is healthy innovation in enterprise systems around data reduction, at scale these techniques are applied at a layer above the storage deployment. So fundamentally, when considering mass capacity storage infrastructure, the device level dollar per terabyte dictates the buying behavior and will continue to do so.

Now, if we consider the actual cost as a ratio, we can see that the TCO of SSDs is approximately 6 times greater than hard disk drives of equal raw capacity in a mass capacity deployment. The low to high range displayed here reflects that not all deployments are identical and that applications and workloads can play a role in the overall total cost of ownership.

As we look out for the next decade, we do not see this picture changing significantly at all. We believe that the TCO for hard disk drives and SSDs will stay approximately in equilibrium. Both SSDs and hard disk drives will continue to improve their value proposition, and storage demand for both will continue to grow. They are both critical enabling technologies for the growing datasphere, and their synergistic relationship in the data center infrastructure will persist.

So in summary, we are demonstrating robust and sustainable areal density growth with heat-assisted magnetic recording of more than 20% per year. We can use our recent experience, productizing our 20 terabyte HAMR drive to translate from laboratory demonstrations to products, which puts us on track to deliver 50 terabytes by 2026. We have drive technologies planned which will enable long term HAMR growth and a path to over 100 terabyte devices. And hard disk drives will continue to service the needs of mass capacity storage with the most optimal, total cost of ownership for the next decade and beyond. Thank you.

[Video Presentation] (00:31:46-00:36:30)

Jeff Fochtman

Senior Vice President, Business and Marketing, Seagate Technology Plc

Hello. My name is Jeff Fochtman. I'm Senior Vice President, Business and Marketing. Is a great pleasure to be speaking with you today. First, I would like to walk you through the elements that shape our exabyte storage leadership across all major hard drive device markets. This includes enterprise for cloud and non-cloud mass capacity storage applications, OEM, consumer, branded, and retail product lines. Our current leadership position goes hand-in-hand with executing on and constantly improving on our customer satisfaction metrics. It may sound like a buzz word, but we've made total customer experience a North Star in our go to market approach. And as a result, our customers have rewarded us not only with market share but also by co-planning and sharing their future storage needs with us, helping to shape our path forward and our roadmaps. One step we are taking to is listening to our customers and simplifying our portfolio, focusing our investments in their growth markets and focusing only on big-hitters in declining markets. These decisions benefit Seagate and our investors by improving our return on investment. At the same point, we benefit from supply chain efficiency, as our strongest product brands work harder and harder and become more efficient as they address more of the market.

Finally, our brand. Our brand stands out as a leader across device segments, vertical markets, and emerging solutions. Our 40-year brand journey is at the precipice of another rich decade of helping to store and access the world's data. I will cover our device markets and later will tell you more about how we're extending the brand beyond the drive.

As Dave mentioned in his keynote, we segment our devices business into two markets, legacy and mass capacity, legacy being traditional markets like PC and gaming while mass capacity being storage-dense in growing markets like cloud, and video, image, and applications.

The crossover point in exabytes between these markets happened between 2017 and 2018. And now as you can see, we're riding that wave of growth for mass capacity demand driven inside and outside of cloud hyperscale locales. Through that transition, our portfolio and investment focus has also shifted. Our R&D is firmly aimed at mass capacity opportunities and new market drivers.

While we've minimized investments in the legacy markets, they remain important to Seagate and our customers in a few fundamental applications. We're coming off a wave of success with our latest nearline launches because of faster customer adoption and volume ramp. The 16 terabyte platform was the fastest ramping nearline product in Seagate's history. By leveraging the same platform for both 18 terabyte and 20 terabyte PMR capacity points, we are able to accelerate customer adoption of these higher capacity drives. Customers want plug-n-play solutions, and Seagate is answering the call with this best-in-class product line. Our common platform approach and extensibility into 20 terabytes and beyond on PMR technology offers Seagate the flexibility to meet customer needs and also demonstrates our clear technology leadership. With HAMR technology, we expect to maintain that leadership for years to come. HAMR devices have been in customer hands for several years now, and we just hit a new milestone, 20 terabytes in 2020. We will continue to see this technology across multiple markets and use cases. As we approach the maximum useful capacity of PMR technology, each successive drive increases by 1 terabytes or 2 terabytes at a time. With HAMR technology, it allows us to jump in steps of 4 terabytes, 6 terabytes, or even 10 terabytes at a time. These large steps in capacity gives Seagate a clear line of sight to a decade or more of technology leadership in our core device market. In parallel, we've shipped our multi-actuator technology in volume. We call this MACH.2, and this performance boosting technology will be interspersed across both PMR and HAMR swim lanes.

I'd like to tell you a little bit more about MACH.2. We began shipping in volume in 2019 and we're now expanding our customer base. Well over a dozen major customers have active dual-actuator programs underway. As we increase capacities to meet customer needs, MACH.2 ensures the performance they require by essentially keeping the drive performance inside the storage to your expectations for hyperscale deployments. Although MACH.2 is ramped and being used now, it's also really still in a technology-staging mode. When we reach capacity points above 30 terabytes, it will become a standard feature in many large datacenter environments.

A notable benefit to dual actuator technology standardization is that it drastically cuts down test time, and therefore, hard drive production time is greatly reduced. This is a benefit we're looking forward to recognizing on the cost side of the business. We're extremely proud of our continued partnership with Microsoft on the MACH.2 dual-actuator technology and I appreciate this quote from distinguished engineer, Aaron Ogus. Again, it's Seagate's relationships and deep technical engagements that keep us on the cutting edge, and we have no problems being behind the scenes when we're collaborating with game-changing partners like this.

Now let's talk about where our customers are putting capacity-dense hard drives to solve their most important data challenges. It's in both the Edge and the Core. This graph illustrates why solving the data gravity equation is critical. The future of the datasphere remains complex and sprawling. And while many of the easy applications have already migrated to Core Cloud, we see a myriad of emerging use cases and drivers that will keep data gravity in the Edge as well. Endpoints, including phones, industrial equipment sensors, retail inventory, scanners and other devices, they're all onboard storage-light. Some have no storage at all. Instead, data migrates from these endpoints to aggregation points at the Edge for processing and collection. Aggregating endpoint collection data turns into large volumes of data. And when data gets big, costs matter.

Our record exabyte shipments in Q2 of this year partially reflect this migration into Edge applications, telling us we're on the right track. We're able to address these demands in capacity growth with a massively simplified roadmap of platforms, products and SKUs. This gives us a chance to focus our investments in mass capacity storage delivering a higher return. As I said, our products work harder and more efficiently. With fewer but more powerful product brands, we've had a double benefit of complexity reduction while enabling fast returns for our downstream partners. In the outer years I do expect we'll see mass capacity device platforms evolve further as we hone in on burgeoning IoT and video analytics use cases.

So, we're simplifying while we're winning. How is that possible? I mentioned it previously. We include customer experience in all our decisions. We listen better. As a result, we continue to improve on all of our metrics. We started this five years ago when we realized that while we were strong in OEM markets which had long been a Seagate strength, we had a lot to improve upon in distribution and channel and even with cloud customers. Everyone at Seagate takes part in this, and I want to quickly thank my co-workers in every function for their focus here over the last few years. We're not perfect, but the focus and results speak for themselves and we continue to up-level our internal metrics, challenging ourselves to improve year after year.

Our customers share their feedback with us regularly, helping to shape our path forward. When we started our TCE initiatives a few years back, we set a path for improvement, and this slide shows how we've not only achieved that but we've put ourselves in a great position to be relevant with nearly every leading customer in every market segment in every region.

Seagate was founded on this view of partnership as key intellectual property. We will continue to serve our customers with humility and integrity. I'd like to tell you more about our customers have shaped our perspectives in a few critical key markets. Let me touch quickly on legacy markets. We used to sell inside of client devices like laptop, PCs, or game consoles, but as those devices have gotten thinner and smaller, we now see a robust attachment market connecting to those PCs and consoles.

The first half of the fiscal year our Consumer business is up year-over-year, partially driven by the growth in gaming both in hard drives and in SSDs. As a proof point to our gaming focus and our partnership orientation we launched one of the coolest storage products this last holiday season; a cutting edge, 1 terabyte NVMe Gen 4 SSD product made for the new Xbox consoles. It's a long product name and a cool product. We've also done an excellent job on the placement and promotion in e-tail marketplaces globally. This has driven positive gains, not only in external storage, but in all direct-to-consumer product areas like DIY internal upgrades and NAS markets that typically transact on these major e-tail platforms; really important route-to-market.

While the legacy markets are driven by work-from-home and game-from-home trends, our mass capacity markets are driven in part by the move in video and image applications. We speak a lot about the growth of hyperscale data centers, but there is also a significant exabyte demand generated by the many markets spawning from cameras and sensors. This shift to smart-everything increases capacity demand at the Edge and in the Core. We see AI-lite applications meant for single task processing like security monitoring begin to burgeon into true AI use cases, where multiple forms of value can be created from the same data; efficiency, productivity, safety, you name it. To do this though, you have to save the data.

Video and most other IoT sensor data types are unstructured in nature, and with the right tools, great value can be unlocked with AI processing and machine learning. Unlike traditional structured data, where the value of processing goes down once you've reached a certain threshold, unstructured data continues to yield insights. The autonomous vehicle vertical is a great example where this "store-everything" mindset is already happening. With video and LiDAR sensor data in the test phase, every bit from the test fleets of cars is not only analyzed but stored for long-term modeling and analytics usage.

Customers are storing more data longer creating both the opportunity for new insights and a corresponding data gravity challenge. Our device brands, Exos and SkyHawk AI are meant to handle that challenge. They take center stage in these markets where valuable insights come from data dense modeling. Exos [ph] near line (00:49:07) drives are meant for traditional enterprise workflows and hyperscale density, while SkyHawk video and image app drives are designed for multi-stream HD video recording and data analytics processing.

We will continue to use our deep customer reach and technical expertise in system designs and data workflows to take full advantage of these trends and outpace our customers' needs. I'm very excited about the short-term and long-term future of the Seagate brand as we execute on our device markets as we've done for over 40 years now and also as we extend our great brand and the new offerings and with new customers.

With that, we'll head into a short 10-minute break before our next presenter, Ken Claffey, takes the virtual stage. Ken heads Seagate's systems business and will talk about our Lyve edge-to-cloud mass storage platform. Thank you for your time.

Operator: Ladies and gentlemen, we will now enter a 10-minute interval. The event will recommence in 10 minutes.

[Break] (00:50:21-01:01:30)

Ken Claffey

Senior Vice President-Enterprise Data Solutions, Seagate Technology Plc

Hi. My name is Ken Claffey. I'm the Senior Vice President of our systems business and I'm going to introduce you to our Lyve initiative that really builds on everything you have heard earlier today, adding to our device foundation with our systems, software and services capabilities. I will explain that our motivation is both to help enterprises deal with the explosive data growth in an increasingly edge-centric, distributed, hybrid cloud world, and in doing so, how we can significantly expand Seagate's addressable market opportunity, delivering incremental customer value and further enhance our profitability.

We're going to talk about building upon our HDD device level core technology leadership with further augmenting that through our investment in mass capacity optimized systems, software, open source ecosystems, solutions and our services to solve the very real problems faced by the distributed enterprise. But at a top level, the premise I want you to take away is very simple. We are now the largest provider of bytes in the world and we are executing a strategy that will augment how we best monetize those bytes as we go forward.

So, what is Lyve? Lyve is the culmination of innovation at every layer of the storage stack. It consists of our portfolio of innovative products that are each designed to address specific pain points among the endpoints to edge to core continuum and how that consistency in our technology architecture across all these products enables them to work together to form the mass storage platform that is Lyve; a platform that enables both the secure data movement and storage of massive amounts of data.

The starting point, the foundation of course, is our device level innovation. On top of that, we've been building out our enterprise systems business. This business services a broad range of OEMs, CSPs and enterprise customers. In fact, we just closed a record near 8 exabyte deal. We believe our systems business is on track this year to deliver a significant year-over-year growth. We're using our systems capabilities to maximize the velocity with which we get our core device level innovation to market. A great proof point is that we just shipped our first HAMR [ph] production revenue (01:03:45) drives in our systems to our enterprise channel customers last quarter.

Last quarter, we also announced an important next step in our vertically integrated stack with the launch of CORTX, the world's first and only open source mass capacity optimized object store. It's very exciting for us to showcase our expertise extending beyond drives and our systems, including working with a growing community of collaborative partners and developers, enabling us to further unlock the value of our devices. This software then

combines with the latest in our device and systems technologies to create Lyve Rack, a solution for on-prem or off-prem private cloud storage.

Now, a clear challenge of the distributed enterprise is that data not only needs to be stored but that it also needs to be moved. Sure. Small data, bytes of data can be effectively moved over the wire or even wireless. But terabytes of data? Terabytes of data needs to be physically moved. To solve this problem, we've created a portfolio of data movers, data shuttles in the form of Lyve Mobile, which with further software innovation we call Lyve Pilot, enables centralized orchestration and fleet management. These technologies are the foundation of our new data transportation as a service offering you'll be hearing more about. All these technology innovations, all of them will be used inside out and to enable Lyve Cloud. The devices, systems, software, racks, mobile, all that capability will differentiate the mass capacity cloud, our new storage as a service offering that our CIO, Ravi Naik, will be diving into directly after this session.

Lastly, all these systems and services are all connected. We gather extensive telemetry data which enables AI driven insights into how our devices are used, such as gathering metrics on performance, environmentals, reliability and security. These insights will help fuel our continuous device innovation, thus completing what we call our mass capacity virtual circle.

So, let's put this in context. We would all agree that data is growing to the extent of 175 zettabytes being created in 2025. Now, only some 10% of that data is projected to be actually stored. It should be somewhat confusing to all of us. I mean, we live in a data economy, data is the new liquid gold, AI and ML technologies are helping us to monetize data on the insight it affords us like never before. And across all these application areas, it's clear the larger dataset, the better enabling accuracy in your AI models and the better the outputs. So, again, why not store more of that data?

We asked enterprise customers this very question. And sure they said, not all data was valuable and that data efficiency techniques like compression and dedupe, sure, they play a part. But actually, when it came down to it, they all said, we don't really know what data is not going to be valuable in the future. So, we really would like to save it all, but the costs of storing data so immediate so high that we can't afford it, even though we do recognize that the opportunity cost of that lost data is likely to be significant to our organization.

The takeaway from all these conversations was, Seagate, if you can help make it cheaper to store data, we would store more. We believe this to be a highly priced elastic economic model, in that for any given reduction of the cost of storing data, demand would disproportionally grow. So, we've set about answering that question and have enabled cost effective storage of more data in a way that meets the challenges of the distributed enterprise. Together, let's take a look at those enterprise challenges.

Let's take the example of autonomous vehicles. Now, where these cars are generating 10 to 50 terabytes of data a day, that data needs to be collected across fleets that even in this early stage of autonomous vehicles are hundreds or thousands of cars deep across every global car manufacturer. That sensor data needs to be captured in vehicle, then ultimately moved into an edge processing facility and then on to the main data center or cloud to train their AI driven models.

Today, for customers, there is no easy means to ingest and transport data, that amount of data simply efficiently and securely. Soon enough, these customers will be looking to store exabytes of data. And it's not always competitively desirable or certainly affordable to store all that data in today's public clouds. They need a method of lower cost storage. They need that mass capacity optimized architecture. And they need an architecture that

prevents vendor lock-in. That's what we're offering with Lyve; enable the distributed enterprise to seamlessly move and store data across their hybrid IT environment.

We started with Lyve Mobile and last quarter we added our initial Lyve Rack offering for storing up to a petabyte of data at the macro edge. The physical movement of data is orchestrated and managed by a new software as a service based offering we call Lyve Pilot, which we will be launching later this quarter. Today, we announced Lyve Cloud targeted to support exabyte scale storage as a service and we will extend Lyve Rack to support large, multi-petabyte scale private storage clouds over the coming calendar year. Across all of Lyve, we're working with our hyperscale partners to ensure compatibility and the integration with their platforms.

So, what's the most efficient approach to storing mass data? Luckily, there's a blueprint for that. Let's take a look at it together. We have a great set of customers who have pioneered a blueprint for reducing the cost of storing massive amounts of data. They built a mass capacity centric architecture that uses sophisticated storage tiering 90% HDD, just 10% flash. Central to the architecture and indeed their business models is an imperative to adopt the next generation of storage devices rapidly so you can stay at the forefront of the areal density cost reduction curve. That whole approach is enabled by a software defined storage architecture for which they build their own object stores. They also have a fleet of data movers and shuttles that enables them to efficiently ingest and move all that capacity around. This technical approach and efficiency is not what we typically see in enterprises today and that is why we see that the hyperscalers are roughly two process nodes or generations of HDD capacity ahead of the mainstream enterprise.

Let's take a closer look at how Seagate is helping enterprises to bridge this gap. Our software defined object storage is central to an efficient mass capacity centric architecture. We created the world's first and only open source mass capacity object storage stack. We launched CORTX last quarter. It was co-designed by users, application partners and, of course, our own software engineers working side-by-side with the same device level engineers working on the mass capacity innovations like HAMR. CORTX is truly designed from the ground up to be able to handle the kind of capacities we'll be delivering with HAMR through this decade on beyond, handling tens of thousands of those devices, even hundreds of thousands of them, all working together at scale. It's already the fastest growing open source object storage community.

It was and is important for us to open source this technology. We're big believers in collaboration. Our motivation is to unlock the value of our devices and just make it easier for enterprises to take advantage of their full capabilities and capacity, performance, reliability and security. To be clear, our model is not to charge for the software. It's open source. We're there, we're working to support the community. Our economic motivation is to reduce the cost of storage and enable more storage clouds, all public and private, as we enable the growth of the mass storage market across the enterprise.

Now, let's take a closer look at that mass storage market. Let's look into the big opportunity we see in front of us, the product innovation we're driving to capture that opportunity and the strong economic advantage we will be bringing to our enterprise partners and customers alike. Looking at how the overall enterprise storage systems market is forecasted to evolve up to 2025, we see $5.9 billion opportunity to sell our optimized systems to the public cloud service providers. I would point to our recent 8 exabyte deal as proof of our strong value proposition and our ability to succeed.

We also estimate an additional $5.8 billion opportunity for our systems and solutions in the private storage cloud market, an opportunity that we would address via our large enterprise system portfolio and specifically with our CORTX community powered Lyve Rack solution. Together, these are two of the highest growth segments of the overall enterprise storage systems market and represent an $11.7 billion addressable and incremental opportunity

for Seagate. This does not include Lyve Cloud or Lyve Mobile opportunities that Ravi will be covering in just a moment.

Now, let's take a look at the product innovation we're bringing to bear on that cloud storage opportunity The foundation, of course, is our drives. They represent the vast majority of the physical bill of materials. We start with the latest highest capacity devices that deliver the lowest cost per terabyte. We then integrate them into our mass capacity building blocks, which themselves have unique ASIC driven self-healing technology and provide the most reliable and lowest cost per petabyte. We then marry these systems and building blocks with CORTX software in order to build out our Lyve Rack solution, which is a hyperscale-inspired mass capacity optimized private cloud storage solution available both through traditional CapEx and our future OpEx consumption models.

How attractive will this offer be to the enterprise? It brings a very strong economic value proposition. It will deliver the lowest cost per petabyte, some 33 times lower than all flash based object storage alternatives. It further reinforces how the mass capacity advantage of HDDs especially when coupled with open source software can translate into big savings at the systems and solution level for enterprise customers.

Now, let's take a look at some of the market reactions that validate and actually even enhance what we're doing with Lyve. For Lyve Rack, we've developed a partnership with Raytheon in order to deliver secure private clouds to the US Government and commercial customers. This partnership entails the combination of all that's in our Lyve Rack solution together with Raytheon security suite, including the Electronic Armor and Boot Shield secure root-of-trust technology. We're equally excited to be partnering with Dell on Lyve Mobile.

Again, we believe we are now the largest provider of bytes in the world. We are executing our Lyve strategy that will augment how we best monetize those bytes in the future. We are going to execute on our vision of leveraging our unique vertical integration to address a large and high growth mass capacity market opportunity in the enterprise. We've explained that all these technologies and products are built upon our core device technology leadership and that we are further enhancing our competitive position with incremental innovation across the entire storage stack, from systems to software to solutions and services. Lyve is a unique offering, an open mass storage platform that offers the ability to move terabytes and even exabytes of data and efficiently store it so enterprises can store more of that data, more of that 175 zettabytes, and realize the economic benefits of the data age.

I'd now like to hand it over to our CIO, Ravi Naik, to share more about the newest piece of our platform, our Lyve Cloud storage as a service offering. Ravi, over to you.

Ravi Naik

Executive Vice President-Storage Services & Chief Information Officer, Seagate Technology Plc

Hello, everyone. I'm Ravi Naik, CIO of Seagate Technology. As Ken shared with you just now, Seagate's Lyve Platform is aimed at helping enterprises address the cost and complexity of securely managing mass capacity data in the distributed edge cloud world. I will be sharing with you about Seagate's new storage as a service offering. Data is growing, data is potential, but more data is thrown away due to increasing storage costs and lack of storage solutions near where data is created. It is clear that there is an unmet need in the market.

Lyve Cloud is an S3-compatible storage as a service that addresses these data needs with predictable economics for mass capacity data. Lyve Cloud is about offering choice in a multi-cloud world. Our collaboration with Equinix enables this choice through the power of interconnections and makes it available at the metro edge closer to the data.

The opportunity is big and growing. The addressable market is predicted to grow to $36 billion by the year 2025. We see Lyve Cloud serving the needs of many verticals and use cases. CIOs like myself are facing an acute data problem. There is more data created than ever before and more use cases to unlock value from that data, but no easy answer to store and activate data at scale in a secure and cost effective way.

There are exabytes of data trapped in silos that remain unutilized. Data is being thrown away simply because keeping it is too expensive. And our customers know by doing so they could be losing important insights. CIOs need an easy answer to store more data and harness its power without worrying about the total cost of ownership, security and data mobility.

Today, we're announcing Lyve Cloud, a storage as a service platform, purpose-built for customers to store, manage and activate mass capacity data with predictable cost economics. Lyve Cloud leverages the combined expertise and capabilities of Seagate and Equinix, the world's digital infrastructure company. We are excited to collaborate with Equinix to place Lyve Cloud infrastructure in their colocation data centers.

We see two clear advantages for our customers utilizing Equinix's digital infrastructure solutions. First, with Equinix's global presence, Lyve Cloud customers are now able to consume object storage as a service at the edge, interconnecting sources of data for accessibility and real-time analytics. Secondly, Platform Equinix provides multiple applications for our Lyve Cloud customers to interconnect their cloud storage to compute capability, including Equinix Metal, its automated, interconnected and secure bare metal as a service. You will hear more about this directly from Equinix later in this session.

We're offering a cost effective way for customers to store more data. Data that can help them enable new use cases while also meeting TCO objectives. Further, Lyve Cloud is S3-compatible, making it suitable for many applications. Lyve Cloud is positioned for multiple data-intensive use cases, from backend for backup and archive to big data analytics and several others across industry verticals.

We are on track to have four Equinix locations available across the United States in 2021. Lyve Cloud is not about replacement, but about filling the unmet needs of data storage. Storing and harnessing the potential of mass capacity data. Due to increasing storage cost, customers end up tossing out valuable data. Lyve Cloud fills this gap by delivering a storage-only cloud to work alongside private and public clouds and brings the cloud closer to the data at the metro edge. In the end, we believe Lyve Cloud will become the preferred storage cloud for storing mass capacity data.

Available now, Lyve Cloud delivers on three main value propositions; simple, trusted and efficient. Simple to buy, use and scale. Let me explain what I mean by that. We are focused on creating a simple pricing model for always-on data. This makes Lyve Cloud's economics straightforward with storage costs that is easy to understand and budget for. Customers know exactly what storage they're getting and what it costs them. Data is always accessible and ready to use without delay. Everything is included and there are no extras like read and write access fees or other billing surprises.

Lyve Cloud is frictionless. With S3 compatibility, it's easier to start storing data in no time. With no lock-ins or egress fees, it is also easier to move data out. We hear from our early adopters about the ease of use and the value of predictable pricing with no billing surprises. This is what we seek to build on and deliver to all. Let us hear from one of our early adopters, CIO of Lattice Semiconductor, Sudhakar Chilukuri.

Sudhakar Chilukuri

Vice President & Chief Information Officer, Lattice Semiconductor

We tried Lyve Cloud for our backups. I always say, and most CIOs would agree, I don't have a backup problem, I have a retrieval problem. Lyve Cloud solves the retrieval for us with high data durability. The biggest imperative for us is a cloud solution that enables us to store data without us having to redo the economics, rearchitect and revisit things every few cycles. When we saw Lyve Cloud in action, we recognized how it is built for large volumes of data and enables any application around it. It became easier for us to think of many data types that can benefit from moving to Lyve Cloud, both in terms of cost benefit and from making the data available for analytics.

As a CIO, three things I most seek from any solution are predictability, security and no lock-in. We want clear line of sight on what it would cost us, regardless how much data we store, how many different applications that access the data and how often the data is accessed. The second is security. We want to be assured our data remains our data and protected against ransomware attacks. Lyve Cloud provides both data-in-flight and data-at-rest encryption. And by turning on versioning, we can protect against ransomware attacks.

We do not want to be locked into a long-term contract or tie our data down that prevents us from making future decisions that are in the best interest of our business partners. That's why a cloud storage strategy with transparent pricing and zero ingress or egress fees is the optimal choice for capitalizing on the proliferation of enterprise data. Lyve Cloud's features fit these needs very well.

Ravi Naik

Executive Vice President-Storage Services & Chief Information Officer, Seagate Technology Plc

Trust is a data imperative and Seagate understands that. With Lyve Cloud, data is always encrypted and protected against malicious attacks. And Lyve Cloud's ISO27001 and SOC 2 certifications meet the data security needs of enterprise customers. We also know data access is crucial for customers to not only run their business but unlock its value with analytics. Lyve Cloud offers multiple regions with resilient infrastructure to ensure data is available when needed. Additionally, we believe customer data should remain customer data. Data is always encrypted and Seagate will not look at the data or monetize it.

For enterprise cloud economics, what worked at one petabyte does not work at 10 petabytes or 100 petabytes. Storage costs and other extras add up very quickly and become very expensive. Enterprises pay for storage costs for every second of every day on all the data stored. Lyve Cloud delivers better total cost of ownership, especially when the scale of data grows to multiple petabytes. Lyve Cloud enables customers to be more efficient in managing and using their data. They can store more of it while getting the best value and harness more of its potential.

Perhaps most importantly, we deliver efficiency through the benefit of proximity. Locating with Equinix brings Lyve Cloud closer to data sources and data users at the edge. This allows enterprises to store and process data closer to data sources and enable more edge like use cases with less latency.

Now, let me share with you how we are going to drive growth from this new product by focusing on customer needs and building on Seagate advantages. According to the IDC Rethink Data report, 44% of data generated in enterprises goes uncaptured. And of the data that is captured, 43% goes unused. Over the next two years, enterprise data is projected to increase at a pace of more than 40%, which translates into step function increases in both uncaptured and unused data.

Lyve Cloud seeks to provide a cost effective solution to store and realize more value from the data they're already creating. It is easier to see how this translates into a market opportunity. Based on IDC data for storage as a service market and applying Seagate analysis, the opportunity is predicted to be $36 billion by the year 2025.

Needless to say, this is a big and growing market. We are able to leverage Seagate's unique advantages to deliver value to our customers.

Seagate has been a global leader in providing the building blocks for private and public clouds. And now, with our pipeline of high capacity HDDs, we can deliver the performance and TCO benefits of our innovations directly to enterprise customers through Lyve Cloud. We have crafted a repeatable and scalable 100 petabyte cloud building block that allows us to scale fast and meet customer demands. From one side that is in production today, we will add four more across the United States by year-end with Equinix. We are making Lyve Cloud accessible to more customers and expanding geographically as needed.

Our collaboration with Equinix enables us to offer customers a cost effective storage as a service closer to where their data is created. Lyve Cloud is part of a complete Lyve portfolio of software and services designed to help enterprises find a new way to data. Data transport services powered by Lyve Mobile and fully managed data migration services, key elements of the Lyve portfolio enable offline transport of large volumes of data to and from Lyve Cloud.

Our innovation advantage directly translates to value to our customers. Lyve Cloud delivers customers better cost and more choices. Simple pricing with no extras or egress fees means customers can now store more of their data they are throwing away today without the fear of cloud lock-in.

Lyve Cloud will connect to Equinix Fabric, a software defined interconnection service that allows any business to connect between its own distributed infrastructure to any other company's infrastructure on Platform Equinix. Equinix Fabric is the de facto standard for companies to directly and privately connect to today's leading cloud providers, which allows customers to take advantage of multi-cloud services without deploying additional infrastructure.

At its essence, Lyve Cloud is an object storage as a service that has wide range of applications. The verticals include, but are not limited to, enterprises, media and entertainment, automotive and video and imaging. I would like to specifically call out two unstructured data use cases. The first being big data analytics.

Big data analytics enables customers to harness the power of their business data. With big data, value grows almost linearly with the volume of data stored and analyzed. Today, most big data applications do not run in the cloud because of the cost of storage and additional API charges. Lyve Cloud delivers the TCO customers need to justify storing more data and the reach of Platform Equinix enables storage and analytics anywhere within the platform.

The other is content repository. There is an increasing need to store more data for longer. Media and entertainment is a clear example of a vertical that needs a cost effective cloud storage for its content repository needs. However, customers across all verticals are seeing increasing amounts of images and media, especially given the extensive and increasing use of cameras. We believe Lyve Cloud is an easy choice for these kinds of multi-petabyte scale content repositories.

Lyve Cloud empowers customers to harness the full potential of their data. Enterprises are asking for a simple, trusted and efficient choice for storing mass data, and now it is here, Lyve Cloud from Seagate; the preferred choice for mass capacity data.

And next, let us hear from the Chief Strategy & Development Officer of Equinix Eric Schwartz.

Eric Charles Schwartz

Chief Strategy & Development Officer, Equinix, Inc.

The collaboration between Seagate and Equinix really started when we sat down and they shared their vision for Lyve Cloud storage as a service, where we could all see quickly is that it was a high performance capable platform to address growing storage requirements as all of our customers wrestle with the growing volumes of data that now exist in the world.

For that platform to really have its full value, it needed to be geographically distributed, close to end users. And so, Equinix with our global platform in 26 and soon to be 27 countries really fit together well with the vision that Seagate has for their Lyve Cloud storage platform that's bring the platform capabilities to customers, who are wrestling with these very issues of performance, of latency, of geographic distribution and data sovereignty.

So, bringing together the Lyve Cloud storage as a service platform with Equinix's global footprint of data centers, our broad suite of interconnection capabilities that drive interconnection to clouds to other platforms to other applications and potentially our new Equinix Metal platform which provides bare metal compute services really opens up a new set of possibilities for customers. And it's exciting to see how they'll use that to solve their challenges and bring new capabilities to their end users.

The advantages of storing data closer to data sources really speaks to the requirements of the end users and the applications. End users' expectations for performance are constantly rising. What was tolerable a year or two ago in terms of application latency or delay is now viewed as being unacceptable or worse even now. So, by bringing data closer to the endpoints, closer to the end users, compute platforms can act on that data, whether it's artificial intelligence or other capabilities and respond to the end users even faster providing a better experience, but also providing better value.

Ravi Naik

Executive Vice President-Storage Services & Chief Information Officer, Seagate Technology Plc

I'll now hand it over to our CFO, Gianluca Romano, to walk you through Seagate's financial outlook.

Gianluca Romano

Chief Financial Officer & Executive Vice President, Seagate Technology Plc

Hello, and thank you for joining us today. I will start by discussing our long-term value creation. Today, we operate in an industry that is structurally in a more attractive position with mass capacity storage now representing a majority of revenue, which position us well to deliver future growth. On top of that, we're also excited about our opportunities with our Lyve Storage Platform which are incremental to the hard disk drive growth that we foresee in the coming years.

These trends are complemented by our ongoing focus on operational discipline and ability to convert revenue into strong free cash flow. We have demonstrated our commitment to capital returns, far exceeding our target model and establishing a track record for growing our dividend, while also reducing our share count. I will discuss our future capital return plans in further details later in the presentation.

First, let me review our fiscal year 2020 performance, which was well aligned with our long-term model despite the headwinds of the global pandemic. We grew revenue year-over-year. We delivered strong operating margin and cash flow. We maintained our focus on CapEx discipline. And in terms of capital return, we were once again well above our target, raising our dividend and retiring approximately 10% of share outstanding.

Given some of the positive market dynamics mentioned earlier, including our strength in mass capacity storage and confidence in the future demand of that market segment, we are revising our three to five years financial target framework. Starting with the top line. We are increasing our revenue growth target to a range of 3% to 6%. We are improving our operating margin to 15% to 20% of revenue, well above our prior financial model.

To achieve this profit level, we will continue to demonstrate OpEx discipline and target gross margin in the range of 30% to 33%. Over the past several quarters, COVID-related costs have represented a significant headwind to our gross margin performance. We expect to see incremental improvement in the second half of the calendar year as COVID costs start to abate and we ramp our latest high capacity products into high volume.

These factors along with ongoing mix shift to mass capacity storage support the gross margin performance in our target range toward the end of our fiscal year 2022. We target CapEx as a percentage of revenue to be between 4% and 6%. The lower spending level relatively to the prior model reflects platform simplification and continued manufacturing efficiency improvement.

Finally, we expect to return at least 70% of free cash flow to shareholders, through dividends and share repurchases. We have recently executed well above that threshold, and based on everything I just went through, you can see that we are excited about where we are going as a company and we have felt investing in our shares have been a good use of capital. I will discuss more about our capital strategy later.

Now, for some more detail on how we get to this target. Starting with the top line. We have continued to execute our strategy to capture the significant opportunities associated with secular demand growth for mass capacity storage. Our hard disk drive revenue has now reached an inflection point, with mass capacity representing over 60% of our revenue in the most recent quarter and is expected to reach over 90% by fiscal year 2025, driven by a strong demand from cloud data center customers and video and image application combined with our leading product portfolio. These trends support our expectation for mass capacity revenue to double by fiscal year 2025 compared with fiscal year 2020 with a growth in mass capacity far exceeding the decline in the legacy market.

One element of our business model that is still underappreciated is how the installed base in mass capacity actually create a stable replacement cycle. By 2025, we expect that nearly one-quarter of our mass capacity demand will be from replacement capacity. We look at this as a reliable revenue stream that should compound into the future. Layering on top of our hard disk drive business, we are expecting our known hard disk drive revenue to approximately double by 2025 as we scale our Lyve Storage Platform. Lyve Rack is a natural evolution of our systems business but with significantly more software content and IP to provide a cost effective way to manage and scale petabytes of that.

Furthermore, we are offering services through a subscription model with Lyve Cloud and Lyve Mobile. [ph] These are now as a service offerings better tailored (01:37:53) to address a complexity of moving and storing mass data in the distributed enterprise. We expect the architectural change for edge workloads to be transformational in the coming years. Given hard disk drive our natural backbone of hyperscale cloud storage, we're well positioned to deliver a unique solution to customers consistent with what they have come to expect in terms of performance, but far improved in terms of cost, security, transparency and simplicity. These new businesses offers opportunity to build a subscription based model, which support a more stable revenue stream at attractive profitability levels and generating good return on capital. We expect subscription based revenue to be at least 45% of total non-hard disk revenue by 2025.

On top of focusing on increasing our revenue and making our market position stronger, we also focus on improving our operational efficiency and we will continue to optimize our cost structure by focusing on three key areas. First, we are simplifying our product roadmap, reducing the number of configurations by half by 2025. We are also leveraging our areal density gains to optimize the cost of our drives by reducing the number of heads and disks for different capacity products. We project these cost optimized drives to represent about 95% of exabyte shipments by fiscal year 2025, up from about 20% this year. This enable us to have the greater manufacturing agility to address customer needs, to deliver more value to our customer at attractive cost and better margin.

Second, as we presented through our last analyst event, by doubling capacity every three or so years, we can sustain our cost per terabyte decline that keeps our total cost of ownership well ahead of competing technology for mass capacity storage market, supported by how our HAMR technology. We have proven over the years that we can sustain innovation through market cycle to continue to execute against our cost reduction roadmaps and keep our competitive position.

Third, we have tightly managed our OpEx, lowering spending levels by about $400 million over the last four years. As we look forward, we expect to continue this operating discipline. We will gain R&D efficiency through our common platform and transition more into non-hard disk drive opportunity while keeping total OpEx spend near the current level and align to business progress and results.

Let's now turn the page to capital structure. We continue to maintain a strong balance sheet with ample liquidity, including significant cash as well as undrawn revolving credit facilities. As mentioned on our last call, we raised a total of $1 billion in capital last quarter, each in two tranches of debt as the lowest ever interest rate of any of our prior bonds. Including the new notes, the net debt was $3.3 billion at the end of calendar year 2020.

We've also made a significant progress over the last year in managing our maturity schedule, significantly extending the average tenure of our debt and reducing the annual maturity levels. With last 12 months adjusted EBITDA of $1.8 billion, we are in a strong position in terms of leverage and you can expect us to maintain an appropriate leverage ratio over the long term.

And let's now look at our future capital strategy. Our confidence in the business, both in the underlying market trends as well as our strong track record of execution in delivering consistent cash flow, is a basis for our capital allocation priorities. Our first priority is to maintain this disciplined approach to supporting the growth opportunities we see in our business, in terms of OpEx, CapEx and other strategic opportunities.

Our next priority is to build on our recent track record of reliable annual dividend growth. We have delivered a 3% increase in annual dividend in each of the last two years and we plan to grow our dividend in a programmatic fashion.

Our next priority continues to be our share repurchase program. In the first half of fiscal 2021, we retired approximately 8% of our shares outstanding. As Dave noted earlier, the board approved a $2 billion increase to our repurchase authorization. Our decision to invest in our share in the current environment underscores our confidence in the long-term business outlook and future cash generation abilities. We have recently accelerated our share repurchase activity and expect to return more than 100% of free cash flow to our shareholders for this fiscal year and next. We will do this while maintaining a strong balance sheet and liquidity profile.

Importantly, we are focused on allocating our capital with a goal of generating the highest long-term return. As we have said before, but worth reiterating, the combination of our operational performance and capital allocation discipline has allowed us to be one of the small group of companies in the S&P 500 to be in the top tier in terms of

delivering high free cash flow yield as well as strong return on invested capital. Over the past five years, we have returned more than $7 billion to shareholders via dividend and share repurchases and included a 6% total increase in our annual dividend and reduction in our share count by 27%.

So, in summary, there are three pillars that underpins the value of our financial model today. The first pillar is our opportunity to drive significant revenue growth against an attractive market backdrop for mass capacity storage and entry into new market opportunities via our Lyve Platform. The second pillar is our operational discipline to drive operating profit growth in excess of our revenue growth, supported by technology roadmap execution, cost reduction and prioritizing R&D investment for our new growth markets. And lastly, the third pillar, our strong balance sheet and best-in-class capital return commitment that will augment and enhance the value we are delivering to our shareholders. Thank you.

Operator: Ladies and gentlemen, we will now hand the session back to Shanye Hudson, who will bring you to the next segment of today's event.

Shanye Hudson

Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

Thank you. We'll take questions from analysts here shortly. And again, if you'd like to ask a question and haven't done so already, you can submit your question through the chat function on the web platform. But first, let me hand the virtual stage back to Dave for a few closing remarks.

William David Mosley

Chief Executive Officer & Director, Seagate Technology Plc

Thank you, Shanye. And thanks to all of you for sharing part of your day with us today. I want to circle back to where we began this morning and reiterate the key decisions that I believe Seagate's in a stronger position than ever to capture the significant opportunities ahead of us and deliver value to all the stakeholders, like we discussed. Starting with mass capacity storage, which is vital to the evolution of the datasphere and is now driving overall HDD revenue growth again. Seagate's enhancing its product roadmap to drive technology leadership, address all the mass capacity demand for years to come and pivot.

As the datasphere evolves in complexity, enterprises are increasingly challenged to get the most out of their data. And Seagate is leveraging its broad expertise beyond the drive, even into our IT functions, to help customers securely capture, move and store data across a more distributed enterprise with our Lyve edge to cloud storage platforms. Lyve helps enterprises unlock data potential and it also creates a recurring revenue stream, as Gianluca talked about, opens up great market opportunities for us in the future. So, we're really excited about it.

As we build for that future, we have a strong business foundation to deliver profitable growth, strong free cash flow and continue to return capital to our shareholders, as we have in the past. Finally, I'm confident that we have the right leadership team to execute this strategy, some of whom you heard from today.

So now, let's begin the Q&A session. Thank you.

QUESTION AND ANSWER SECTION

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Thanks, Dave. Okay. Let me start off with a question, you mentioned pivoting and Seagate has successfully pivoted the business towards growing mass capacity storage opportunities and I'd say today we've emerged a leader in mass capacity storage. So, Dave, could you talk about maybe how sustainable technology leadership is in the HDD industry and then how can you leverage that strength in the new opportunities on our Lyve Storage Platform?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Sure. Thanks. It was a very conscious pivot into mass capacity a few years ago when we started talking about the common platform, the 16 terabyte, was many quarters before we ever launched it. And it took quite a while, even though it was the fastest growth of heads and disks that we'd ever done in Seagate, it still took quite a while for adoption.

And I think that very deliberate planning and those long cycle times are going to be the hallmark of the future of mass capacity delivery to the cloud and to the edge. You have to be very thoughtful about your capital, what you're building, because the cycle times are so long through wafer and slider operations and media operations to [ph] kit (01:48:07) the final drives together, have to have very tight discussions with your customers to really understand their demands as well.

So, you're planning has to be very deliberate, you can't flow a bunch of inventory that you don't know if it's going to yield properly and so on. It's enforcing a lot of discipline, a long lead time discipline inside the markets. And because of that, we don't think the market's going to move very quickly. We're very comfortable with our existing CapEx footprint, the way we've pivoted our OpEx and also our customer relationships, as Jeff Fochtman talked about. So, very comfortable with that. We think it's sustainable and we're going to continue to drive this common platform architecture.

The other thing which I just made reference to just a bit ago is the fact that we also run those factories, we understand big data generation in factories, we see how that data has to be used, how much data is created, which is tens of terabytes per factory per day, and how long we have to keep it in order to make decisions. So, our IT department, which is a big department, is actually learning a lot about these problems, these big data problems.

And we can move engineers from our core products back and forth to those solutions, which [ph] becomes Lyve (01:49:16). It's a thing that I think our brand is very well positioned for, our technology is very well positioned for. It's the mass capacity solution that the world really needs. So, we're confident about that.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

That's great. And we have a question here from Katy Huberty of Morgan Stanley. Katy asks, given the significant lead that you could have in areal density, how do you think about balancing market share and margin expansion over the long term? Maybe, Gianluca, you could start with that. It's a two-part question. Her second part is on

another margin related question. But, can you reach your – the low end of the long-term gross margin range in the second half of 2021 depending upon how COVID costs fade? So, maybe, Gianluca, you start and, Dave, you add in.

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

Yeah. Thank you, Katy. Well, I will say market share is not our priority. For us, market share is a result of having the best products at the right time when customers really demand for those products. We are there today with our 16 terabyte. We think we will continue to be there, you have seen our product roadmap, our 18 terabyte, 20 terabyte and, most importantly, our HAMR roadmap. As we commented before, we were able to ship our 20 terabyte HAMR in December and now we are fully focused on developing the second and third generation of those products that will continue to be the most important products for the mass capacity for the cloud part of the business. So, in our model, we don't really look at the market share per se. We look at what we think is the demand and how our product can answer to that demand.

In terms of the gross margin and when we will achieve the range that we have indicated today, I will say depends on many things, as you said, it is not only COVID. The COVID costs that we report every quarter's earnings release are just a direct costs, but there are a lot of indirect costs that are more difficult to explain and to identify, but are really impacting a lot not only on our company but the industry in general. I will say on top of COVID, the mix is for sure moving in the right direction. We expect to continue to go in that direction. That will for sure help our gross margin to continue to improve.

Factory utilization, that is [ph] another (01:51:51) discussion we had about CapEx, how we align supply and demand. Today, supply for us and I think for the industry in general is a little bit higher than demand. This is mainly because of the decline in the legacy part of the business that has free up capacity. I think with the right level of CapEx, that is a fairly high level, it's not [ph] as more (01:52:18) level now, when we say 4% to 6% of our revenue is a fairly big number, but I think with that level, we will be able to align supply and demand.

And as I said before, hopefully, with abating of the COVID costs in the next two or three quarters and with the mix going in the right direction, with the pricing environment, of course, aligning to the supply/demand situation, all those elements will contribute to Seagate's gross margin to go into the range that we discussed. We don't know exactly when it will happen. It can be three, four quarters. We are now almost at the end of fiscal 2021 and next quarter will be our last quarter of fiscal 2021. So, when we say end of fiscal 2022, it's not really far, it's just a few quarters from now.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

Thanks, Gianluca. Anything you want to add, Dave, or should I move on to the next question?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

I think Gianluca answered it well. Just a couple of points I would add. Of course, COVID has impacted so many people so deeply up and down our supply chain. So, we want them to continue put the priorities on safety and security in the factories. And all the freight and logistics issues will still be with us for some time. People are finding ways to be creative around some of those freight and logistics issues, but it's still fairly impactful.

The only other thing I would add to Katy's question is that remember that the mass capacity drives, the long lead times I talked about before, but also we control a significant part of the bill of materials because we make the heads and the disks. So, relative to the BOM, that's why we understand what the big hitters are for cost to scrap and yields and test times and all the other things that people talk about the industry, we understand and control those that become a much more significant part of our bill of materials.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Sounds good. And we have one from Karl Ackerman of Cowen and it's on some of the edge opportunities. He says much of our research and your thought leadership highlights mass capacity storage as a critical enabler of the edge and clearly a unique ability to sell your drives to hyperscale cloud providers, could you discuss at least qualitatively the investment priorities of traditional enterprises and telco providers as they seek to capitalize on the growing service opportunity at the edge?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

It's a very diverse question, depending on where you are in the world and which specific customer you're talking about. But Karl, first I would say that the [ph] Cowen report (01:54:48) is about a year old now, but I still would point people back to it and I just was looking at it a couple of weeks ago, as one of these kind of seminal reports on what's going on at the edge. It's still very relevant too. So, I think that's – it's a great framework to think about.

The way I state the mass capacity problem today is there's too much data being trapped or just frankly deleted at the extreme edges. And so, in the presentation today, we've used the word micro edge for that. That data isn't effectively moved into other instances, whether it's a metro edge or backup into the cloud or the cloud service providers. Getting that data moving and then allowing AI and ML to make decisions on the data, not necessarily snap decisions, but one day, one week, one month later is a big, big opportunity I think in the datasphere. And it's something that really everyone benefits from. So, there's – it's not a zero-sum game.

So, qualitatively, to answer your question, that's the way I think about it is that there's a lot of data today that's not properly utilized. We've talked about that in some of the marketing pitches that Jeff Fochtman has put forward. If we can actually address that in a cost effective, secure, very efficient way, I think it opens up opportunities for a lot of people to add their value in their applications, compute, AI, ML, or even some of the existing cloud service providers they just need more data be fed to them.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Thanks, Dave. And this one I'm actually going to ask both you, Dave and Gianluca, to comment. And it comes in from Wamsi Mohan of Bank of America, and it's a question on HAMR and asks if HAMR will be standard for anything over 20 terabyte or give a little more clarity on the timeline for HAMR to be mainstream? And then, from a cost perspective, can you compare the cost structure for HAMR 20-TB, say, versus a PMR drive that you talked about today?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Yeah. I'll let Gianluca answer the details. He did in his prepared remarks a little bit on HAMR cost. 20 terabytes is at the cusp of a transition. We believe that PMR is actually running out of steam, but we still have solutions that we really like the platform that we're on. So, whether it's 20 or 22 or 18, I think a lot of people get – come to the

wrong conclusions about those things. I think what we need is customer predictabilities in high yields and the deliberate planning that I talked about up and down the supply chain. In order to continue the areal density growth, the capacity growth that the world needs, we're going to have to break and go to HAMR. And so, that's one of the reasons why we're out there in the market today learning from it.

And from my perspective, it doesn't add a lot of costs. There's a lot of maybe misinformation out in the markets about this. It doesn't add a lot of costs. It's not easy to do it, but it wouldn't be easy to do a 25 or 30 terabyte drive to begin with, irrespective of the technology. But we have a lot of confidence in the parts that we have in our lab right now. And we're going to push this out into the market. We're going to learn a lot with our customers. We're going to stage for the future our capitals already for HAMR. So, there's not a whole lot of capital investment that's required. We don't have to change the model that we presented today. So, we're pretty optimistic about HAMR.

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc
A

Yes. We were discussing before about the company's cost decline between PMR and HAMR. We have high expectation in terms of cost decline from HAMR. We show how we expect the cost to decline by more than 40% when we move from a 16 terabyte to a 30 terabyte, 32 terabyte HAMR. We see the similar decline to what PMR was able to generate in that period of time where we were moving from 2 to 4 terabyte or 4 to 8 terabyte, so several years ago. Right now, PMR is not generating that level of cost decline. And so, we expect HAMR to bring us back in that cost decline curve that of course will help us and will help our customers to get the right level of TCO.

In term of R&D, I'd say the majority of the effort has been done in terms of developing the technology. So, we have demonstrated the technology. We are now selling our 20 terabyte HAMR in the market today. Now, we go to a normal or more normal cycle where we develop the product. We don't need to develop that technology and the product. We just focus on the product. And CapEx now, we'll have – of course, we will have some specific tools and equipment that are dedicated to HAMR. We don't think we need to spend more than what we spent in the last few years, as we indicated in our model, 4% to 6% of the revenue. Of course, the revenue is supposed to increase, you need to consider that. But nothing that will drive us very differently from where we are today.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc
A

Thanks, Gianluca. We've got a couple of questions in from Ananda Baruah from Loop Capital. Let me start with Lyve. He says you have a $50 billion TAM, and it's actually served available market, it's our SAM for Lyve. If there was no Lyve, what would happen to that data? And maybe, Dave, I'll ask you to address that. And then, we've got another one on the buyback that I'll kick over to you, Gianluca, when Dave is done.

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc
A

Great. Thanks for the question, Ananda. The way I think about Lyve in the total market is it's a small part of the much, much bigger enterprise IT market. So, this is a mass capacity only focused on the market. And $50 billion sounds like a big number, we talked – Ravi and Ken talked today about the various constituents of that. From my perspective, if Lyve doesn't exist, not as much data is moving around, so the market is probably smaller, but I think people will come up with creative solutions irrespective. I think if we stay with data trapped in ecosystems and getting deleted too much, I think there's a lot of [ph] lousiness (02:01:19) in business models as well, AI and ML business models. So, I think it just doesn't develop as quickly. There will still be mass capacity products. We'll still continue to serve the customers that we serve.

But I think what Lyve actually does is it enables with frictionless data movement and some other opportunities to just have temporary resting spots for mass capacity data. It allows for everybody to garner that. I mean, if you think about edge compute is talked about a lot or AI and ML are talked about a lot, all of those things need to be fed consistently with data. And the more data that we can actually bring at good economics, the more those technologies are going to be able to contribute value back into the world.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc
A

Okay. Thank you. And then, Gianluca, the question is around the $2 billion share repurchase authorization today and that brings the total at least to – at the end of last quarter about $5 billion. Could you talk about the pace of those repurchases over time?

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc
A

Yes. Dave announced earlier today the additional authorization of $2 billion for our share buyback program. This gives us a lot of flexibility on how to execute the share buyback. Share buyback, first of all, is only a part of our shareholder return program. We have a lot of focus on our dividend first of all. We raised our dividend in the last two years. And on top of that, we complement with a share buyback. Dividends are very programmatic in our view. Share buyback are more opportunistic. So, I would not expect a stable buyback in the next quarters or year. We will discuss at the beginning of every quarter with Dave and decide how much we want to allocate in terms of share buyback as part of our cash flow. And see, we will try to maximize the opportunity that we have. It's important to know that we always give priority to that part of capital return that generate the greatest long-term return. So, always looking at the long-term result and not just focusing on the short-term quarterly result.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc
A

Great. Thanks, Gianluca. We've got a question in from Steve Fox of Fox Research. This one again goes back to the network edge story. He says, it seems like there's a significant opportunity there that's already underway. What kind of revenues does the company expect to generate from edge HDDs over the next couple of years? And how well is the company positioned in the specific edge sales channels?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc
A

Thanks, Steve. So, edge is already alive, to your point, there's products that we'd sell today and Jeff Fochtman talked about SkyHawk AI, for example, is a product that is serving markets like smart cities, smart factory, smart healthcare, those kind of applications. And so, we see the growth and we see the diversity of customers around the world and what they're doing to kind of aggregate that data. From my perspective, those customers will still be really valuable to us. We offer up other solutions as well, not just the hard drives that they're using, but other solutions to support their markets, whether it's recovery services or transport services, and now we're talking about all these other things in Lyve. It's a great opportunity for us.

We think about video data a lot, but there's all kinds of other sensor data that's coming ultimately; autonomous vehicles and LiDAR data, healthcare, genomics, and these are large, large datasets. And as those are generated at the extreme edge, we need to move them around or we need to store them in the extreme edge. There's great mass capacity opportunities in for many years. And I think the solutions that come from systems that actually do this economically, quickly and so on, are going to be super impactful to the world.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Great. And maybe here's one from Patrick Ho of Stifel, and it's a little bit of a follow-on, Dave, and a big picture kind of question. So, he asks, Seagate has begun monetizing the opportunity with HAMR and the Lyve Platform and as the Lyve Platform expands Seagate's reach into emerging markets, is it fair to say that Lyve continues to expand and there will be more services based offerings and potentially software enhancements for customers in the future, and overall, would that help to lower overall TCO for customers?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

It is fair to think about it that way, but I think we want our Lyve Platform, especially Lyve Cloud, to be very simple, the interfaces very simple. And probably the best way to think about it, as operations not necessarily just trying to reach a bunch of customers, end customers, and not trying to do applications and things like that, but to just provide the lowest cost per terabyte that we possibly can and ingest from the edge so that you can plug that into clouds. So, we're going to keep the offering very simple and make sure we stay very focused on that value proposition. We think that's where our best opportunity is. It's not necessarily building applications. If that helps you?

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Yeah. Thanks. All right. Here's one from Tom O'Malley of Barclays. And maybe I can start with you, Dave, and then Gianluca, if you want to add some comments as well. So, he says industry data has significant share shift over the next several quarters, or at least that's what many are considering, do you believe you can maintain your dominant share position that you enjoy today? And then, I'll add on, how do you think of share as you've presented in your financial model today looking out over the next several years?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Right. I think Gianluca made reference to this earlier. Market share really is an outcome of great execution, serving the customers really well, having the right product at the right time. And in this really deliberate world, we have to plan that out in time. So, I don't think there're going to be big market share shifts one way or the other. I can't really speak to anyone else's execution. I can just say we're very focused on our own and we're very confident in the investments we're making in our capital, in our programs and things like that that they're aligned well to the market. So – and share will be an outcome of that. I personally think that data is going to continue to grow. If we as an industry do a great job of servicing that, there's going to be a lot of that that we're going to have to answer the call for. And from my perspective, we've done a good job in that pivot to mass capacity and we're positioned really well to take advantage of that.

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

A

Yeah, absolutely. As we said, we don't really drive for market share per se. It is true that we have now the best product roadmap right now and the best technology with HAMR, so we absolutely think that we will take all the benefit from the roadmap and from the technology in the future.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Thank you. We've got one here from Shannon Cross of Cross Research. It's on HAMR. So she says your HAMR roadmap has kind of two major drivers in terms of assumption, one being areal density and the other being time to productize following lab results. Which one do you think is more certain and which one has the opportunity either to outperform or disappoint upside or downside?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

I don't think about it as upside or downside, but I will say that areal density is fairly well-understood right now. And John made reference to this in his presentation. So, you can refer back to that. The time to lag in product if you will, it does – it's not just about the technology, the HAMR technology itself, it's also about how long it takes for us to stage the right product through cycles of learning to get the economics right for us, to get the economics right for our customers to get through the qualification cycles. There are a lot – there's a lot going on there, which is one of the reasons that your investments need to be very deliberate I think in the hard drive business.

I will also say that as we go forward, giving ourselves flexibility in something like Lyve or own systems business we can control some of the variables and get to market faster with that technology because – and we're also much more efficient because we can aggregate some of the challenges that say a hundred different customers are having into 10 different customers' challenges and one solution that kind of meets all of those things. So that's one of the reasons why we're strategically driving our mass capacity and our systems business the same way.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

[ph] Well positioned (02:10:22). I see a couple of questions coming in that are around the Lyve platform. So let me start with one from C.J. Muse of Evercore ISI. And he says, can you help me better understand the new Lyve product. Is the advantage simply no margin stacking or – maybe talk about what you think the advantages are? And importantly, how are your current cloud customers reacting to – are they now considering themselves a direct competitor?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Oh, interesting. So CJ, a couple aspects of this. First, with the shuttles if you will themselves, the Lyve mobile products, they aggregate data at the extreme edge. And there are other products that do this competitively in the market today. But there's a lot of open space, white space I think for a lot of different types of solutions right now. And there's data that's just not being effectively moved and managed the right way. So that starts the discussion about Lyve; getting that data to another point where it actually can be monetized by many people but – be used by many people is the central value proposition of Lyve. The way I think about Lyve Cloud is more of just a temporary resting spot. It could become a permanent resting spot, but – for that mass capacity data, where it goes from there could be up into another cloud, could be hybridized, could be deleted, as people said, well, that's not of any value anymore and delete it there, but once it's aggregated, it can be studied for a little bit of time, a day, a week, a month. And so, it provides flexibility. It's not really competitive threat.

There are obviously some people who say, I want data. Some of those people want data for various reasons. They want to trap it. Other people just want data to be able to put in their ecosystem and show what they – what value they can add. And as the datasphere is just exploding right now, I think everybody is fairly mature about that and they'll say, I'll take more data; please bring it to me, give cost effective ways to bring it to me, and then I can monetize it.

Q

That's great. And maybe this one for you, Gianluca, comes from Sidney Ho of Deutsche Bank. Sidney asks, can you talk about the financial impact of Lyve Cloud in terms of revenue or margin impact? And then maybe a follow on to what you were just talking about Dave, do you expect much cannibalization between your Lyve business and HDDs?

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

A

Yes. So, we gave our estimate for the Lyve Cloud – well, for the Lyve business in general how we expect the revenue to grow. Is maybe a fairly conservative view right now. We are at the beginning of the development of a new business. So, we don't of course want to be too optimistic. The impact is not only in term of revenue as Sidney is suggesting; is also a little bit of a change in our P&L structure. Those kind of businesses are probably higher gross margin kind of businesses, so it is important also to help us and drive the entire company into a different level of gross margin and improve quarter-after-quarter our P&L structure situation. I don't think there is a cannibalization between our hard disk core business and the cloud business. I think now with the time people will get used to see both and is actually a very good use of our production and certain of our drives. And so we think is a perfect complement to our business. It's still based on our core business, that is hard disk drive. So, is a very good evolution of our business.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Thanks, Gianluca. And I don't know, Dave, do you want to add anything on thoughts around potential cannibalization? I think it's sort of a positive for us in both ways as data grows, but – and I know you have a perspective.

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

I think Gianluca answered it well. Look, there – we have an existing business in consumer where we're selling, as an industry, tens of millions of shuttles already. They tend to be small; 5 to 10 terabytes are the maximums, and they tend to be fairly slow with USB interfaces. So you know these shuttles that we're talking about now with Lyve are more enterprise products, much bigger, faster, more secure, things like that. From my perspective, when it gets back up into a data center, is it cannibalization that we have to worry about? Actually I think if we do it economically, there's more opportunity for us for HDDs because that data is then going to fan out and be – and start adding more value. Look, the people who have big built big engines for AI and ML that can process a lot of data with compute, these are awesome capabilities but they need to be fed with data. To the extent that we can do that even more, it probably opens up more opportunity and doesn't cannibalize.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Great. Thank you. We've got one here from Jim Suva. It's a model-related question. And Jim Suva of Citi, he asks, is the stronger sales outlook or expanded revenue growth CAGR mostly driven by the assumption of higher storage shipments, the – demand for data, or is there a pricing assumption that you've applied?

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

A

I would say we were not aggressive with our pricing assumption. Of course, now we see data that is exploding. So, it's mainly coming from the great demand – the great volume demand for our product, but with also an assumption of an improving pricing environment in the next few quarters.

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Yeah. Thanks, Jim. In our mass capacity vernacular that we use, IT 2.0 has – all the disruption that just happened was impacted in the middle of last year. We think the cloud is going to accelerate some of its growth as we come out of the economic situations that we're in. Supply/demand balances may change very, very quickly. And we think that the IT 4.0 world is being accelerated as well. Companies are going to have to make the pivot. Most of them are our customers; we'll help them. We'll be with them through the pivot, but we think there's a lot of opportunity as we do pivot to – especially given the long lead times that I referred to earlier to really change the model.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Okay. We've got a CapEx-related question that's coming from Aaron Rakers of Wells Fargo. He says: I know you mentioned that much of the CapEx investment for HAMR completed by now. But how do you think about your current head and media capacity footprint in terms of Exabytes being produced and where you need to provide incremental investments in these key component areas?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Yeah, we do all the time, Aaron, underneath our existing models. So if we go to 4% to 6%, that's still quite a bit of money. I mean, you can do the math. And as we continue to grow revenue, it'll be even more. So, some of the capacity – some of the OpEx if you – or, sorry, CapEx if you will goes into the factory and provides technology transitions. Some of its just maintenance capital, but there's plenty of opportunity underneath those numbers to do the pivots that we would need for HAMR. It's not a lot of pivot actually. All of the machines we've been buying are with a mind towards where we need to get to with HAMR and that's the way we think about capital utilization. We think our footprint is great.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Great. Okay. Maybe from Mark Miller of Benchmark Capital. He's talking about HAMR; he says, you've emphasized how HAMR drives provide a significant TCO over SSDs. Maybe you could speak a little more about just the TCO advantage that you're seeing and if it's specific to TLC or QLC? How do you think about that advantage over time?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

As John Morris laid out earlier, I think, Mark, you can see that it really fundamentally comes down to the device cost – fundamental device cost. For mass capacity solutions, there's all kinds of other solutions out there in the world that we're not really talking about anymore. And I would always say, Flash is a great technology. It's come of age. It has a lot of utility in many different applications and the data center as well. We understand what mass capacity is. We don't really see that there's much overlap. So, I refer back to John's presentation on details, it is very true; read the fine print. When people come out and say, here's the way I do the comparison, you have to look as a TLC, QLC what's exactly the TCO assumptions that people are making. But I would also say that all of

these customers that are buying mass capacity drives today, they are very, very smart people and they understand how to make all these tradeoffs. It's about how to use the equipment that they have. It's not about some economic tradeoff, one component versus the other. And we think that that kind of market is going to persist for a long, long time given the technology trajectories that are – that both technologies afford. Remember, we have a lot of device physicists in the company. And they watch these things really carefully; they have experience through many, many different industries. And we're very confident in this balance that the data storage is going to have to have for the foreseeable future.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

We have a gross margin question from the line of Toshiya Hari of Goldman Sachs. Is this what drove your decision to share a long-term gross margin target and is this indicative of business leaders at Seagate increasingly evaluating gross margin results?

And maybe Gianluca, you could start and Dave you might add afterwards.

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

A

Well, the decision to share a gross margin target partially come from how we look at the business. Now, we always said that our priority is free cash flow and profitability. And higher gross margin will help both. So we started looking more and more at these metrics. And we see opportunity for improvement based on our hard disk Core business and based on the new Lyve business. So, we thought it was a good idea to share our view with analyst community and investor community. And I think was also important was, I was receiving a lot of question about gross margin, so I think was good to make it clear what is in our view the model for the next few years.

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Yeah I think it's an – actually an interesting question. And if you go back in history, there was a different supply/demand environment entirely. When we were coming down off the back off of client server, our factories were ready to build a certain type of drive that was actually declining, and the cloud just hadn't grown big enough yet – mass capacity hadn't grown big enough yet. As we march down as time progresses, then you get to a different point of the economic equation. And we're now into a point of mass capacity being the dominant revenue source for us. It's going to continue to grow. We are going to have to invest against it. And in the products types and the CapEx that's required, very different to invest; it's largely heads and media now. We talked about that in the bill of material. So, different timing, different maturation, a lot more of confidence in the future. When you don't have so much of a decline going on, it's a big change for us.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

That's great. And while we're on the topic of margins, we have one more that came in from the line of Vijay Rakesh of Mizuho. He says, you know thoughts around the gross margin profile maybe as you ramp HAMR and kind of that mix as you go into the fiscal 2022 and 2023 timeframe and get into a range, could you sort of overlay your outlook for margins and the ramp up of HAMR?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

It's – we are beginning the ramp of HAMR and we've just updated the margins. So, in one sense they are part and parcel of the same discussion. And the way we think about it is, we serve customers. Gianluca said this earlier really well. We focus on that customer predictability; we're staging the products for them; we know it requires investment on our side, not just in equipment but also in the long lead time inventory and things like that. And so, we're going to be a lot more predictable. So I think the refresh of the model if you will is pointing to the expectations that we have for the ramp of our technology. And we'll continue to let this play out over time.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

That's great. And I don't know, Gianluca did you want to add anything or...?

Gianluca Romano
Chief Financial Officer & Executive Vice President, Seagate Technology Plc

A

Not really no. We said before, we expect HAMR cost decline to be important to the company and for sure to contribute to an higher gross margin in the future. And as Dave said, we will ramp the volume of HAMR based on customer demand when it's ready.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Okay. We just have a couple more questions that I can see here and the next couple are back on to Lyve and some thoughts around Lyve. Let me get a question here from Kevin Cassidy of Rosenblatt who asks about the storage-as-a-service market opportunity and how we derive that market opportunity. And then how we're thinking about pricing Lyve Cloud for customers? Is it purely on a per petabyte usage basis, or are there other services or things that we're thinking about in the future?

William David Mosley
Chief Executive Officer & Director, Seagate Technology Plc

A

Yeah, thanks for the question, Kevin. If you think about the source of how we got to an as-a-service model ourselves, these shuttles are big and expensive and we don't want to sell them to the world and then have them frankly shipping around and have them have problems and then ship back to us and then, you know broken and things like that. That happens in the real world. We think an as-a-service model is a much better model than just a straight sales model in that environment. And we think we can actually serve customers well doing that. There's – the data that's on the shuttles needs to be very well protected. So we need to be in some sense monitoring the shuttles all the time. We're developing the software for that as well; it's called Lyve Pilot. Ken made reference to that in his presentation. And then relative to the cloud economics, once it gets there, the Lyve Cloud economics, which is really a Metro Edge cloud if you will, we'll – we will be – we will keep it as simple as possible – and Ravi talked about that. And I think that simplicity, none of ingress and egress fees is really important in allowing people to use it in a very predictable way. Sometimes it's really hard to see the true storage cost in offerings. I think to the extent that we keep that really simple, not only management but the pricing expectations around it, I think people will use it a lot more to then distribute that data out in to the other parts of the ecosystem that they need to. So, thanks for the question.

Shanye Hudson
Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

That's great. And maybe another question on Lyve Cloud from Munjal Shah of UBS, and it's more around, how we're thinking customers will use Lyve Cloud. Dave, you addressed this already, but I think it's important topic.

Does it work for applications there in the public cloud or applications on-prem? How are we thinking about the use cases for Lyve Cloud?

William David Mosley

Chief Executive Officer & Director, Seagate Technology Plc

A

Yeah, to begin with, and I think for some time to come, it will be fairly simple. I mean, it's almost like a data parking lot where the data is put there, stored there temporarily or for a long time – could be a month or a year, probably not permanently, although that's up to other people as well. But there's not a whole lot of applications on top – no data management applications that are complex on top. That's what other people do well. We expect them to be able to do that with this data as well. So we want to keep it really simple, efficient, secure, make sure that we stay focused on those priorities, not building a bunch of new applications and let the other application – people who are really good at applications use the data that way.

Shanye Hudson

Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

A

Okay. And it looks like maybe just one last question on – surprise, it's another gross margin question, but C.J. had one more question on gross margins noting that mass capacity is material – higher today than it was three or four years ago. So, kind of what gives you confidence in being able to hit the higher gross margin levels today versus what we were seeing a few years ago? And maybe, Gianluca, you could you could speak to that. I think you covered a lot of that in your presentation around some of the things we're doing with common platform and certainly our confidence in our cost declines.

Gianluca Romano

Chief Financial Officer & Executive Vice President, Seagate Technology Plc

A

Yes. I think, no fairly similar to what we were discussing before. There are several elements; is not just COVID direct cost that is, that 1% that you are mentioning and you are referring to. We think COVID is actually probably an higher impact to our business. And is not only the COVID situation; is supply/demand alignment, is factory utilization, is the continued move of the mix to the mass capacity. And all those element will contribute to a better pricing situation that overall will help us to achieve that level of gross margin, as we said, in the next few quarters, will not happen tomorrow. We think we will sequentially improve quarter after quarter, and in few quarter we will be in that range.

Shanye Hudson

Senior Vice President-Investor Relations & Treasury, Seagate Technology Plc

Thank you. So it looks like we've kind of exhausted the question list. I want to thank everyone again, and maybe Dave, I'll hand it back to you to close it out.

William David Mosley

Chief Executive Officer & Director, Seagate Technology Plc

I'd like to thank everybody for their support of Seagate. The last year for the world has been very challenging for customers, for employees, for suppliers and their employees as well. And we realize that as a company our job is to make sure that we're sustainable and make sure that we have our facilities and our product lines ready for what's coming in the future. We believe there's a brilliant future for mass capacity storage. We believe we've done a good job of positioning that way. And for our shareholders, we will continue to be the same company we had relative to shareholder return. We're very committed to that. And we look forward to continue to update you on earnings calls and in the next Analyst Day. Thank you for your time.