MOOR INSIGHTS & STRATEGY

EVERSPAN CHANGES ARCHIVING

HUGE DATA SETS & NEW RETENTION POLICIES PRESSURE THE EXISTING ARCHIVE NORMS

EXECUTIVE SUMMARY

The combination of increasing regulatory accountability and significantly larger data sets is driving enterprises and service providers to rethink their storage archiving. Existing methods of archiving—via tape or even hard drives—are starting to show their limitations as the realities of today's needs poke holes in those approaches. Data today is being generated in significantly greater amounts and needs to be retained longer than ever before. The increasing economic cost of a poorly executed archiving strategy is becoming a much more pivotal factor in today's enterprise planning. Optical archiving is now a viable solution to address these storage challenges.

Sony is targeting the enterprise datacenter with Everspan, an optical archival system that complements (and in many cases replaces) expensive, inefficient tape silos and power-consuming hard drive-based systems with a new, more durable, higher performing alternative. By using optical technology, Everspan is able to easily store up to 181PB of data per library (scaling to 724PB per datacenter). Businesses can efficiently maintain archives longer with a lower operational cost and never have to operationally touch or manage that data once it is written. Additionally, for those that need to work with petabyte-sized data sets, Everspan can rapidly load and manipulate huge data repositories over and over without having to ever worry about media reliability.

Previously, archiving decisions were driven by the lowest acquisition cost, but data size and retention demands are putting a new focus on how these decisions are made. Optical delivers against that low acquisition cost and low operational costs with unparalleled longevity / reliability (as environmental conditions change) as well as higher performance to match changing enterprise needs.

TODAY'S ARCHIVING NEEDS ARE CHANGING

Stored data is growing rapidly. Recent estimates put the world's data at 44 zettabytes by 2020, reflecting both a massive increase in conventional data as well as an explosion of Internet of Things (IoT) data, with mobile phones already driving a zettabyte of traffic themselves this year alone. As overall data grows, the least-accessed "cold" and archived data will scale up as fast, if not faster than, primary data. Along with that growth, more data is being retained for future use. Handling cold and archived data is typically driven by economics, favoring the lowest cost per gigabyte. Three key dynamics are now changing how companies view their archived data: more data is coming in, this data now has operational value for decision making (meaning that it will need to be accessed multiple times), and new regulations require that the data needs to be retained longer than in the past.

More inputs and greater depth: Our connected world and the proliferation of IoT (which is expected to soon reach 10% of all data) are generating continuous streams of data, be it customer intimacy through applets or health and telemetry information from embedded sensors. Richer media like video is also becoming more prevalent, and applications like police body cams (at up to 1.6GB per hour per camera) are dramatically scaling those data requirements upward. The terabytes per month of police footage generated per department often needs to be retained for years because of the potential for legal appeals. This high definition footage will soon be replaced by 4K video at even higher GB/hour rates (with 8K video already on the horizon), forcing departments to find both scalable and cost effective methods to address this storage growth. Seismic data from oil and gas exploration now uses 3D technologies that amplify the number of data points as well, resulting in huge repositories of data that must be maintained; each new extraction technology may bring companies back to that older data in search of new opportunities that might not have been originally feasible.

Large data repositories need to be mined: Now that this data is being captured, organizations need the ability to glean insight quickly from these huge data sets. For years, enterprise data mainly consisted of small database transactions and smaller, structured data files. But now larger data sets, rich media, and unstructured object stores are becoming the standard. High performance computing (HPC) workloads are dealing with a larger number of data points as both the frequency of sampling and the depth of sampling increase. Many businesses actually need to work with massive multi-petabyte data sets that could take days to load before the analysts can begin their work, negatively impacting productivity. Past storage limitations drove many organizations to only store data differentials instead of full data sets, restricting both their effectiveness and their analytical depth. As competitive pressures continue to turn up the heat, companies need a better way to deal with these large data sets—and do it faster.

Changes in regulatory environments: The changing legal and regulatory environment is driving greater data retention demands. Incidents like the 2008 financial meltdown are driving new government regulations designed to bring greater accountability and transparency. This increases both the amount of data that needs to be retained and the duration that it needs to be held. Failure to accurately respond to regulatory inquiries within a required time window can carry heavy consequences. Most TCO models are centered on data retention of 5 to 10 years, but what happens when longer data retention is required? No medium can maintain data forever, but some regulations and businesses require perpetual access to data. Further complicating retention is all of this IoT data and real-time telemetry information; it is helping to improve products now, but it needs to be aggregated, maintained, and continually mined over time to discern trends and opportunities. More inputs create larger stores that need to be maintained reliably for longer periods of time: there needs to be another way.

SONY OPTICAL ARCHIVE

While Sony has a strong name in consumer electronics, they have also been instrumental in enterprise storage as the company behind the manufacturing of both tape and optical drives as well as the media for both. Products like the optical-based Sony ODSL30M Optical Disc Archive, which targets the broadcast and media markets, could actually fill in some of the smaller archival needs for certain size data sets in specific markets. Though these individual library units are rack-mountable, they are not truly optimized for some of the more challenging enterprise environments and applications. Because of this need for a more comprehensive enterprise archival solution, in 2015 Sony acquired Optical Archive, Inc. This Silicon Valley startup was led by an executive team with more than a hundred years of combined experience in companies deep into enterprise storage technology needs, including Compaq, Dell, Facebook, IBM, Rackable Systems, and SGI.

Disruption rarely comes from within the industry, as evidenced by companies like Amazon, SpaceX, Tesla, and Über; so changing the dynamics of data archival was probably not going to come from the incumbents like IBM or Oracle. Optical Archive, Inc. set out to disrupt the traditional tape and hard drive archival business with a new optical methodology because they saw unmet needs for handling archived storage among some of the largest enterprises. Their move under the Sony umbrella now allows them to merge their unique approach with Sony’s massive capabilities for robotics, manufacturing and media. As the co-developer of the 5” compact disc technology (and the originator of the first writeable and rewriteable optical drives), Sony has a heritage in optical innovation. The acquisition of Optical Archive, Inc. allows Sony to bring this innovation to the datacenter through OAI's deep enterprise expertise across a wide range of vertical businesses.

SONY BRINGS OPTICAL TO ARCHIVING WITH EVERSPAN

Despite Optical Archive, Inc.'s deep experience with scale-out cloud datacenters, their first product, Everspan, was really designed for a broader set of needs beyond just those massive datacenters. Everspan targets the data that needs to be recorded, will not ever change, and must remain reliably accessible for as long as possible. Everspan's "write once, never touch" methodology helps to ensure that—despite huge amounts of data being archived—operational costs remain low, as IT never needs to physically touch the media after it has been loaded into the system. The proven media, with over 100 years of estimated useful life, enables the data to remain protected for decades after it has been written.

FIGURE 1: FULLY-POPULATED EVERSPAN LIBRARY WITH 181PB CAPACITY (PHOTO SOURCE: SONY)

Everspan is designed around a form factor that can live as easily in an enterprise datacenter as well as a warehouse, or even in the mobile pods that some enterprises have begun to experiment with for datacenter expansion. The core Everspan Library consists of a Base Unit "controller" that houses up to 64 of the Archival Disc (AD) drives, the Robotic Unit that moves media back and forth and an Expansion Unit that houses up to 12.9PB of optical media. The Everspan Library is scalable to allow up to 13 additional Expansion Units (14 total) to be added that will scale to 604,928 total discs, holding 181PB of data. Each library only occupies around 500 square feet of datacenter floor space (fully expanded). The sealed system includes HEPA air filters to ensure that it can run almost anywhere, maintaining both process integrity and location agnosticism, which is a challenge for most tape library systems today. While Everspan does provide a massive storage capacity, it does so with a small impact on the datacenter, minimizing environmentals like weight, floor space, and power.

Inside the Base Unit are up to 64 Archival Disc drives, each with four lasers per disc side that accelerate both reading and writing the data from the 300GB Archival Disc (AD) media. Sony's roadmap for AD media has both 500GB and 1TB disc capacities for the future, increasing the capacity of the Everspan Library over time. Everspan enables scaling at the media level through future higher capacity media and scaling at the system level through multiple Expansion Units. Everspan even supports scaling at the row level, as up to four Everspan Libraries can be deployed in a single datacenter as one addressable system, delivering a 724PB data footprint. As each library is added, the aggregate throughput increases accordingly (for identically configured systems).

The key for Everspan is its ability to fit into existing application environments by using standard APIs and drivers as well as working with common storage tiering methodologies. Everspan offers multiple interfaces providing flexibility to address a broad set of use cases common in the industry today.

- **S3 interface** is based on the popular Amazon Simple Storage Service (S3) API, supports RESTful access (PUTs, GETs), "bucket" style structured hierarchy, and secure access. Additionally, customer data is further protected using software-based data protection (erasure coding) allowing for uninterrupted data access, writes, and reads in the rare event of media or drive outages.
- **LTO drive interface** provides compatibility with well-understood LTO tape library semantics. Everspan will support emulation of multiple generations of LTO drives with support for arrays of up to 16 tape drives, allowing for a non-disruptive entrance of optical technology into IT organizations that are familiar with tape library technology.
- **Optical drive interface** leverages the native Sony MMC-6 SCSI multimedia command set, enabling customers to write directly to optical drives and media, extending their planned or existing software infrastructure to support the Everspan Library System. This option is especially valuable for large scale out cloud datacenters that may have their own custom-written software for archiving and data access.

For its software partners, Everspan is fostering a robust software ecosystem that eases integration for its partners and boosts its value to end customers. Software partners will have native Everspan interfaces that leverage industry standard T10, MMC, and SMC command sets.

EVERSPAN OUTPACES TAPE ARCHIVAL

Tape is a natural comparison point for customers, but Everspan does not always compete head-on with tape because tape is typically viewed as a rewriteable medium, even if few ever rewrite it. Tape is transactional, targeting the “copy-archive-erase” methodology and does not have the true longevity and consistency that businesses need for permanently archived data. Even if set up as WORM, data can still change or be lost over time with tape. To protect against data loss, customers will remaster the tape, writing the data to fresh media, making tape a poor choice for archive data that must be maintained in its original state for long periods of time. Data stored on tapes may be buried deep within the cartridge, requiring the rewinding and repositioning that plagues tape access and reduces lifespan with every read or access. Optical targets the idea of “store everything forever” and is far more flexible for retrieving data multiple times, because optical media does not degrade with use. Archival Discs have a greater than 100-year expected life once written, meaning data written today should be both readable and reliable in the year 2116—a feat that is unheard of in the tape world.

LTO tape readability is only maintained for two drive generations (roughly 5-7 years), meaning that for customers to ensure readability they must migrate the data onto new media, which is an expensive and operationally intensive task. While LTO tape is claimed to have a 30-year reliability, customers are far less trusting because the media must physically interact with the drive head, reducing life expectancy with each use. Ironically, the only way to ensure data integrity and tape reliability is to verify the data from the tape—an operation that also degrades the tape’s lifespan with each access. This is why customers tend to be less trusting of tape as time increases.

To match optical’s 100-year longevity, customers would typically see 14-20 migrations with LTO tapes; each of which is costly, time-consuming, and risks carrying-forward bad data or damaging older data. Unlike optical, magnetic polarity can change inadvertently over time or with use, which impacts data accessibility and reliability. Over time, LTO tape has proven it can typically only sustain about 200 rewinds. In a pure archival situation LTO tape could theoretically last 17 years if written monthly, but if that writing becomes weekly the life expectancy drops closer to only 4 years; every time the tape is accessed, effective lifespan is reduced. As a rule, a single copy of data on tape is not truly considered a copy due to reliability / readability concerns, meaning tape requires a

100% overhead for redundancy. Everspan uses erasure coding for redundancy, reducing that overhead to only 25% (or less), which is just like hard drives. Everspan has even filed patents on its IP that enables erasure coding across multiple libraries.

EVERSPAN VERSUS HARD DRIVE ARCHIVAL

Because of tape's shortcomings, hard drives grew in popularity in recent years as an archival medium. Hard drives do a better job than tape for maintaining a repeatedly accessible record, but they are still a magnetic medium; data written to them, even in the strictest environments, still carries risk as hard drives degrade over time. The hard drive paradigm ties the drive media to the reader, which "traps" the data with the reader in the event of a drive's mechanical or interface failure. By contrast, Everspan redefines the idea of a drive: as a "disaggregated" 19.2TB unit comprised of 64 300GB discs that can sustain failure from either the reader or the media and still continue normal operation.

Because they are constantly powered on and need to spin up / down with access, hard drives will wear out. The typical failure curve resembles a bathtub with 3 types of failures. Early life failures are often tied to manufacturing issues. Over time this flattens out with random failures in the single digit range per year. As drives age beyond 5 years, failures based on usage dramatically increase.

FIGURE 2: HARD DRIVE LIFECYCLE FAILURE CURVE

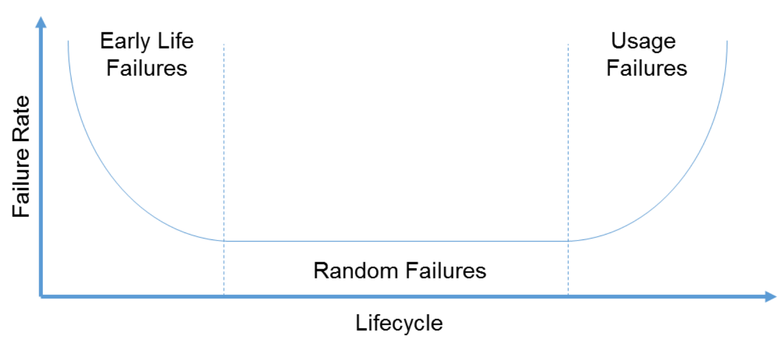

Random hard drive failures happen at a rate of 4-8% per year. One study that looked at drive failures of the latest generation drives (from 2013-2015) found on average the failure rates were 4.81% per year. Google, one of the largest enterprise consumers of hard drives, studied drive failures as well and found that after two years, annualized failure rates were roughly 8%. In practical terms, a typical 60-drive standalone archive unit will probably see up to 3-5 drives per year fail. Even with the more conservative 4.81% failure rate, hard drives are far from being a truly long-term archival medium.

Enterprise hard drives typically carry a 5-year warranty, primarily because after 5 years the failure rates increase quickly (the right hand side of the bathtub curve). Most concerning for customers should be the availability of the proper hard drive models beyond their typical lifecycle; a 5-year-old refurbished replacement drive is typically 2-3x the cost of a brand new drive (while only holding 1/3 of the capacity). Replacing failed drives becomes an increasingly expensive and complex process over time to maintain the commonality and consistency of the solution. Eventually, as drives and interfaces reach end of life, they become impossible to replace.

Whether a hard drive suffers a drive / interface failure or has a data corruption issue, the entire multi-terabyte drive must be replaced and the data rebuilt. To protect from failure, these less reliable hard drive solutions typically use erasure coding for redundancy, which serves two functions: allowing access to data during a drive failure and enabling the rebuild of the failed drive to bring the system back to full redundancy. Rebuilding failed drives can take take as much as 4 days under load. To combat the rebuild overhead, enterprises will typically overprovision networks just to handle the flood of network traffic during the rebuild. Ultimately a hard drive failure increases operational costs and network congestion. For Everspan, failure of a drive or interface can be corrected without impacting data or the network. Everspan also uses erasure coding to allow retrieval of data in the unlikely event that a media location cannot be read. But optical media has such reliability and longevity that media does not need to be rebuilt, it can just continue to be accessed from the erasure coding shards. The disaggregation of the drive from the data and the ability to retrieve data without ever having to rebuild makes Everspan a far better solution for archiving.

Everspan also brings significant environmental benefits to the datacenter over hard drives. In comparison to an EMC Isilon HD, a popular hard drive archival system, Everspan is claiming some significant datacenter advantages over the competition.

TABLE 1: HARD DRIVE ARCHIVING VERSUS OPTICAL

	EMC Isilon HD	Everspan	Optical Advantage
Total Storage	50.9PB	181PB	2.5x more capacity
Storage / Watt	.32PB / kW	16.77PB / kW	51x more storage per watt
Storage / Square Foot	.21PB / ft^2	.36PB / ft^2	71% more storage density
Weight / Square Foot	439lb / ft^2	170lb / ft^2	61% less weight

PERFORMANCE

Through multiple drives, each with multiple lasers, Everspan can sustain transfer rates up to 18GB/s from an archived data store, meaning that 1TB of data could be restored in less than a minute, and 1PB could be restored in less than 15.5 hours—a restoration time unheard of in the industry. If multiple Everspan Libraries are deployed, this transfer rate continues to scale upward.

As data set sizes increase and retrieval becomes more time sensitive, throughput matters. While hard drives are known for being able to retrieve the first bit quickly, the real measure is not the first bit but the sustained throughput, which is where Everspan is roughly 23% faster than typical SATA III hard drives. Everspan is 75% faster than LTO 6 tape and 11% faster than the Oracle T10000D, another prevalent tape technology on the market. Only the most recent (and most expensive) LTO7 drives edge out Everspan slightly in average throughput (although not peak); but even here, LTO7's throughput is only advantaged during a streaming backup or restore. Retrieving random files or pieces or data is slow, in the 90-120 second range for initial access, as the tape drive must rewind through the tape before it can begin retrieving data versus 60-90 seconds of initial access for optical. Once retrieval begins, optical is far easier to position for adjacent reads as the change requires only small movements in the read head versus moving the tape media to reposition for the next set of files.

Everspan has superior sustained throughput relative to both hard drives and tape, making it a better choice when either considering large amounts of inputs to be written or large data sets to be loaded into a working space (as in HPC). This enables better productivity, as highly paid researchers can spend more time working with data instead of waiting on data to load. Additionally, in the financial sector, some regulations require firms to replay transactions or provide millisecond-by-millisecond accounting of fund movements to enforce compliance with complex regulations. Often the response window for regulatory requests is short and time sensitive; slower performance could result in fines or other regulatory actions that can negatively impact firms.

FIGURE 3: DATA THROUGHPUT OF LEADING ARCHIVE FORMATS

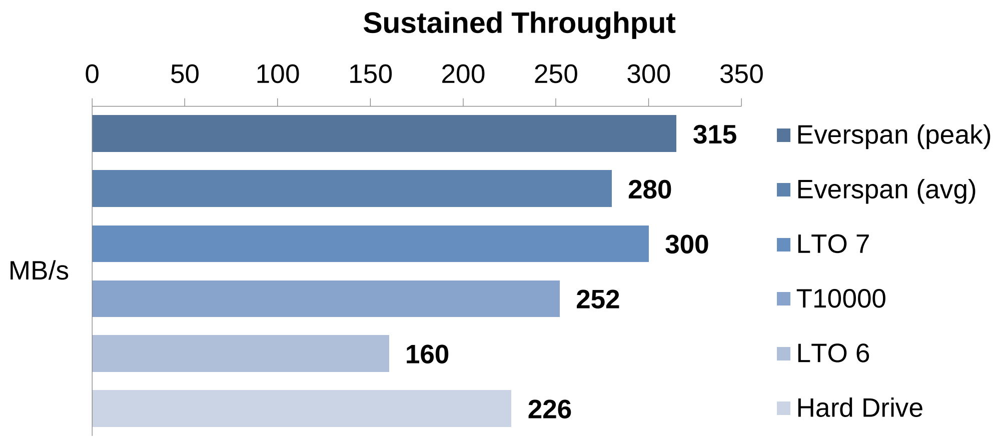

DATA RELIABILITY / LONGEVITY

Retention requirements vary by both region and business, but as the world becomes more data-driven and metric-driven, some retention requirements are extending beyond the previous 7-year standard, forcing tape archives into at least one remastering / migration, if not more. Many governments mandate data retention, like DOL (OHSA) 29 CFR 1910.1020 that requires 30 year+ data retention, far beyond the useful life of tapes or even hard drives. Several states, like Colorado, Connecticut, and Illinois require retention of medical records for 10 years beyond the final treatment or discharge of a patient, while other states like Minnesota require that all hospital records be retained forever. Clearly in these cases the intent is not only to be able to access the data, but the organization must also ensure the integrity of that data. Being able to read data but not being able to trust the results is no better than not being able to read it at all.

Optical media compatibility far exceeds both tape and hard drives. Today's optical drives can still read the first CD ever pressed in 1982, but it would be unthinkable for a business to try to retrieve data from equally old QIC tapes or IDE / SCSI hard drives. The media reliability of these older products would clearly be suspect because it was written magnetically (which degrades), but the larger challenge may be just finding the appropriate application, interface or system to read the device. Because Everspan's AD drives are write once, read many (WORM), the data written to the drives is safe from modification, because it cannot be altered or overwritten.

From a longevity perspective, Sony warranties its libraries for 5 years (and will work with customers on extending that warranty if required); Sony media carries a 100-year warranty. In addressing the reliability of the media, optical is immune to forces that may alter magnetic media. Once data is written it will not change state. The discs themselves are very durable as well. This was proven out when catastrophic disasters resulted in optical media being submerged in seawater for weeks. Despite exposure to the corrosive seawater that would have rendered other media useless, the optical media was still able to be read when it was retrieved weeks later— and will be for years.

Long-Term Cost

Unfortunately, when customers calculate the long-term cost of ownership for archival storage, they rarely consider the changing retention landscape or the costs of migration / remastering to maintain compliance and accessibility. While tape may sometimes appear to cost less than optical, calculations rarely comprehend the cost (both in acquisition and time) for future remastering / migration. Data that needs to be retained beyond the typical 5-7 year life expectancy of tape will require at least one migration, if not more, just to maintain the readability. Additionally, customers may require remastering if there is a concern over data integrity.

In the past, TCO models focused on a very short time horizon (3-5 years), because most IT equipment was considered replaceable / disposable in this timeframe. Recently, with the desire to store more data for longer periods of time, customers have started to model TCO over a 10-year data retention period. This method puts Everspan in a very favorable light, as tape solutions will demand at least one migration and hard drive solutions will demand replacement of virtually every drive as few hard drives live a full 10 years. These extra media costs for tape and hard drive increase acquisition costs, as media typically comprises up to 80% of an archive solution's acquisition cost.

As retention times increase beyond 10 years, Everspan's TCO advantage continues to scale upward. For tape or hard drives to match optical's 100-year longevity, multiple media changes and migrations will be required, astronomically increasing operational costs for media and other expenses required for tape-based or hard drive archives.

Comprehending the long-term cost of archiving when allowing for media degradation / failure nets out a different set of variables depending on retention policies.

TABLE 2: RETENTION IMPACT ON MEDIA REPLACEMENT

Retention Period	Hard Drive	Tape	Optical
5 years	Replace ~20% of the drives (possibly under warranty)	None	None
7 years standard business records	Replace ~35% of the drives	Possibly none	None
10 years most medical records	Replace ~50% of the drives	New media +1 migration	None
30 years many medical records	Replace ~150% of the drives	~5-6 sets of media + 5-6 migration	None
Permanent CPA requirements for many business documents	Media turnover every ~20 years minimum	1 replacement set of media & migration every ~5-7 years	1 replacement set of media every 100 years +1 migration

Maintaining a read-only version of archival data is becoming a critical part of business, and, as more IT organizations are trying to move to a "utility model" for computing, we believe that an additional opportunity exists for archival-as-a-service to better manage the costs of archival. Optical's longevity, capacity, performance, and low maintenance make it a compelling platform for service providers who can productize long-term archival-as-a-service in a metered manner.

CALL TO ACTION

With data requirements continuing to spiral upward and cost pressures driving enterprises to investigate alternatives to traditional archiving methods, it makes sense for businesses to include optical as an archiving methodology because of the clear business benefits it can deliver.

With the ability to bring better performance, reliability, and longevity with a dramatically lower long-term total cost of ownership, optical should be on the radar of every CIO and CFO who is grappling with protecting their company's massive data assets.

Sony Optical Archive has a decided advantage with their Everspan product over traditional methods of tape and hard drive archiving. Sony has the edge when it comes to complex archival of data and overall data longevity as the manufacturer of the drives and, most importantly, the media. Sony is well positioned to capitalize on this market by enabling customers to shift away from traditional archival methods to something that more accurately reflects the changing reality of today's data requirements.

Important Information About This Paper

Author
John Fruehe, Senior Analyst at Moor Insights & Strategy

Technical Advisor
Jimmy Pike, Technologist in Residence at Moor Insights & Strategy

Publisher
Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

Editor / Design
Scott McCutcheon, Director of Research at Moor Insights & Strategy

Inquiries
Contact us if you would like to discuss this report, and Moor Insights & Strategy will promptly respond.

Citations
This note or paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

Licensing

Disclosures
This paper was commissioned by Sony Optical Archive. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

Disclaimer
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.