



2020 Predictions
50 vendors
January 2020

Philippe Nicolas
StorageNewsletter

List of Companies

Summary of vendors' predictions	4
Actifio (Ash Ashutosh, CEO)	5
Atempo (Luc d'Urso, CEO)	6
Bamboo Systems (Tony Craythorne, CEO)	6
Caringo (Adrian Herrera, VP Marketing)	6
Cloudian (Neil Stobart, VP Global System Engineering)	7
Clumio (Chadd Kenney, VP & Chief Technologist)	8
Cohesity (Lynn Lucas, CMO)	9
Commvault	9
CTera Networks (Aron Brand, CTO)	10
Data Dynamics (Piyush Mehta, CEO)	11
DataCore (Gerardo Dada, CMO)	11
Datrium (Sazzala Reddy, CTO)	12
DDN (Alex Bouzari, CEO)	12
DriveScale (Tom Lyon, Chief Scientist, Brian Pawlowski, CTO & Denise Shiffman, CPO)	13
ExaGrid (Bill Andrews, CEO)	13
FujiFilm Recording Media US (Peter Faulhaber, President)	14
Hammerspace (Brendan Wolfe, VP Product Marketing)	14
Hitachi Vantara (Hu Yoshida, CTO)	15
HYCU (Simon Taylor, CEO)	16
Infinidat (Stanley Zaffos, SVP Product Marketing)	17
iXsystems (Mike Lauth, CEO)	18
Kaminario (Eyal David, CTO)	18
Kaseya (Michael Sanders, GM of Unitrends and Spanning)	19
Komprise (Krishna Subramanian, President & COO)	19
StorageNewsletter	2

Lightbits Labs (Eran Kirzner, CEO)	20
MemVerge (Charles Fan, CEO)	21
Minio (Anand Babu Periasamy, CEO)	22
Nasuni (Russ Kennedy, CPO)	22
NGD Systems (Scott Shadley, VP Marketing)	23
Panasas (Curtis Anderson, Senior Software Architect)	23
Pavilion Data Systems (VR Satish, CTO)	24
Portworx (Murli Thirumale, CEO)	24
Pure Storage (Matt Kixmoeller, VP Strategy)	25
Quantum (Eric Bassier, Senior Director of Product Marketing)	26
Qumulo (Molly Presley, Global Product Marketing Director)	27
Reduxio (Jacob Cherian, CMO)	27
Scale Computing (Alan Conboy, office of the CTO)	28
Scality (Paul Speciale, CPO)	28
Spectra Logic (Matt Starr, CTO)	29
StorCentric (Surya Varanasi, CTO)	29
StorMagic (Bruce Kornfeld, CMO & GM-Americas)	30
StorONE (Gal Naor, CEO)	31
StrongBox Data Solutions (Floyd Christofferson)	31
Talon Storage (Jaap van Duijvenbode, Product Director)	32
Virtana (Tim Van Ash, SVP Products)	33
Weebit (Coby Hanoach, CEO)	34
WekaIO (Andy Watson, CTO)	35
Western Digital (Phil Bullinger, SVP and GM)	35
Zadara (Jeff Francis, Sr. Enterprise Solutions Architect)	36
Zerto (Gijsbert Janssen van Doorn, Technology Evangelist)	37

Summary of vendors' predictions

A few days following the 2019 vendors' facts and review publication, we publish our « classic » annual vendors' predictions for 2020. We collect 50 opinions and consolidate them to find some majors trends and directions for the storage industry :

1. Without any surprise, **Cloud** arrives #1 in various flavors (multi, hybrid, on-premises, private, edge)
2. New devices and connectivity around **Flash, NVMe (oF), SCM/PM, QLC**
3. **Kubernetes** of course and containers
4. **Analytics and AI**
5. **Object storage**

Actifio (Ash Ashutosh, CEO)

1. AI will dominate new apps and make data re-use even more important

While the singularity has not happened yet, we're approaching it in software development. As data becomes more heavily governed, controlled and collected, machine learning will be a predominant way for it to be leveraged. Gartner says enterprise adoption of artificial intelligence and machine learning (AI/ML) has tripled in the last year alone, and that 37 percent of organizations have now embraced AI/ML. IDC predicts that by 2025, at least 90% of new enterprise apps will embed AI/ML. Data-powered decisions, including the use of robotic process automation (RPA) will eclipse human-powered decisions in terms of the volume of decisions. Cloud vendors like Google and IBM have dramatically lowered the barriers to adoption and use of AI/ML, leveling the playing field for organizations of any size to become data-driven. We've seen this phenomenon more recently with gene splicing that reduced the cost from millions of dollars to around \$20. MIT's new Schwarzman College of Computing will lay the foundation for computing as a basic skill, just as math is, and further enable exponential adoption of computing in everyday life. Similar to the rapid evolution of test data management (TDM) from an esoteric practice more than 10 years ago to a ubiquitous part of every organization today, I see the rise of analytics data management (ADM) as the new process for rapid adoption of AI/ML and improving the accuracy of data-driven decisions. One thing is certain: all this machine learning will play a major role in making data an even more strategic asset.

2. Hardware lock-in is back . . . pretending to be software

There is a new generation of deduplication appliance vendors looking to displace monolithic, previous-generation hardware platforms. But they bring some of the same dangers that customers came to regret with the legacy storage vendors of the previous era. The wolves are dressed in the sheep's clothing of buzzwords like "scale-out" and "software-based" but make no mistake: users are locked in to expensive "certified" hardware. Nothing scales out more than cloud object storage (which is what Actifio leverages for storing backups while still providing instant mount and recoveries). A truly modern data management platform should run on any compute or storage, on-premises or in the cloud, and work with a wide variety of performance and capacity requirements. In 2020, more IT teams will recognize the trap and embrace multi-cloud copy data management and cloud object storage.

3. Another election year of data insecurity

Hacking, ransomware and data leaks are alleged to have played a central role in the U.S. election process of 2016. Many candidates are back on the campaign trail in 2020, some for the first time and others with legacy election infrastructure. In order to capture and store sensitive donor information, these candidates will need to dig deep into their databases to find the personal identifiable information to reach out. To maintain the best donor experience and the most effective fund-raising and voter turnout operations while improving security and data privacy, campaign managers will turn to software platforms that help them ensure that campaign information stays secure from intrusion, private, up-to-date, and immediately accessible for strategic uses. In November we will see if DARPA's \$10 million contract commitment to secure, open-source election system hardware prototypes will have a positive

effect. Given the volumes of information that will be stored, it will need to be able to securely and rapidly manage the data -- and recover it quickly and efficiently in the case of a breach. Confidence in the integrity of the voting process is the backbone of a functioning democracy.

Atempo (Luc d'Urso, CEO)

1. New Data Management solutions fueled by AI to tackle the challenge of the explosion of unstructured data volume
2. Cyberattacks will continue, tape should be a good choice
3. Local value added MSPs to address data and digital sovereignty challenge

Bamboo Systems (Tony Craythorne, CEO)

The server industry is in constant growth mode, with data centers ever-expanding. These unabated increases have centered around the industry's continued reliance on inefficient legacy server architecture, resulting in the cost of operating data centers to greatly exceed what it costs to build them. Data center power consumption is now a global issue.

The inefficiencies in legacy server architecture design mean that today's data centers currently produce 2% of the world's greenhouse gas emissions and collectively consume over 3.5% of the world's power production. In fact, it is predicted that the server market will overtake the airline industry in greenhouse gas production next year. Something must be done and in the near future. We predict we will soon see more energy efficient data center solutions, with new, ground breaking low power chip designs, as this is now a pressing economic and ecological priority.

Caringo (Adrian Herrera, VP Marketing)

1. Archives will provide increased benefits and value beyond just cheap and deep storage.
2. Object storage will blur the lines between Primary and Secondary storage for use cases that require enhanced throughput.
3. On-demand use cases will increase the need for a 3 tier storage approach incorporating NAS, on premises object storage and cloud archive or on premises tape.

Clouidian (Neil Stobart, VP Global System Engineering)

1. Edge computing will turn the hybrid cloud model on its side

The hybrid cloud storage model is characterised by the ability to seamlessly move data between on-premises and public cloud environments. This data mobility, combined with the flexibility and scalability hybrid storage provides, has made it the solution of choice for many businesses worldwide.

However, the hybrid storage strategy will be turned on its side in 2020 as momentum behind “Edge Computing” continues to build. As a result of the rapid growth of the Internet of Things (IoT) and 5G networks, data collection and the processing of AI applications will be needed at the edge. Analytics will have to be applied near to where the data is being created and where sensors are located – e.g. traffic signals, surveillance cameras and smart cars.

As such, there will be an increasing need for a different form of hybrid storage – one that fulfills the need for robust storage near the edge of the network and enables businesses to move analysed data sets between the edge and on-premises or public cloud systems.

2. Processing AI and ML workloads will require object storage

As data volumes continue to grow, one of the key challenges facing businesses is how to unlock the full strategic value of this data. This is especially true when dealing with artificial intelligence (AI) and machine learning (ML) workloads. That’s why 2020 will see more organisations capitalising on object storage to create structured/tagged data from unstructured information and use metadata to make sense of the flood of data being generated.

With object storage, data is defined with unconstrained types of metadata and located from a single API. This is in contrast to traditional file storage, which defines data with limited metadata tags (such as the file name, date created and date last modified) and organises it into different folders – making it much less searchable and harder to analyse than object storage. For example, a traditional X-ray file would only include basic metadata like creation date, owner, location and size, while an X-ray object could include metadata that identifies the patient’s name, age, injury details and which area of the body was X-rayed. This extra detail makes it much easier to locate via search.

Simply put, object storage architectures make use of metadata in ways transitional file storage doesn’t, making it instrumental in helping to process growing AI and ML workloads.

3. Momentum behind the use of private clouds in hybrid infrastructures will grow

Today, when people talk about the cloud, they usually mean the “public” cloud. This will change in 2020. The term “cloud” will become more nuanced as more businesses deploy private clouds and organisations increasingly pursue a hybrid cloud storage strategy that gives them the best of both worlds.

Organisations in industries such as healthcare, media and scientific research that have large-scale storage requirements already face unique challenges in managing capacity-intensive workloads that can reach tens of petabytes. By providing the scale and flexibility benefits of public cloud platforms along with the performance, access, security and control advantages of on-premises storage, private cloud platforms address these challenges.

Although the public cloud provides businesses with agility and convenience while decreasing infrastructure expenses, it can also come with significant bandwidth and accessibility costs. As such, in 2020 more organisations will take advantage of private clouds in a hybrid cloud infrastructure – storing frequently used data on-prem while continuing to utilize the public cloud for cloud-native applications and disaster recovery.

Ultimately, businesses can't afford to rest on their cloud laurels over the coming year, as 2020 is likely to be another year of significant change. Whether it's recognising the need for edge computing, embracing object storage or adopting hybrid infrastructures, the onus is on businesses to prepare for what's coming their way.

Clumio (Chadd Kenney, VP & Chief Technologist)

1. Continued Growth in SaaS

SaaS will continue to be the easy button for enterprises to remove the complexity of on-premises infrastructure without needing to rearchitect existing applications. Copying the complexity of on-premises infrastructure to the cloud is no longer a viable option. IDC estimates the global market for SaaS applications revenues is expected to be more than \$185 billion by 2022.

2. New Policies Require New Methodologies

Enterprises will be challenged to meet the demands of new data privacy policies such as GDPR and CCPA. There will be a heightened demand for next-generation data protection solutions and methodologies that support these policies and enable fine-grain control over their data assets.

3. Analytics Will Come From Non-Traditional Sources

For many enterprises, backup is just an insurance policy, but in 2020 enterprises will demand valuable insights from backup data in ways that were never possible with traditional on-premises backup solutions of yesterday.

Cohesity (Lynn Lucas, CMO)

1. More and more CIOs will rely on consumer technology to respond to critical data incidents smarter and faster

Consumerization of IT will become even more commonplace in 2020. IT teams in data infrastructure management will increasingly rely on smartphone apps to collaborate and respond to issues quickly and effectively. They will require mobile access to instances of data on-premises and in the cloud enabling them to have 24x7 visibility across their infrastructure.

This infrastructure will enable IT leaders to take action faster, which is critical during time-sensitive occasions such as a potential outage or breach.

2. In 2020, the on-premises vs cloud debate will be much clearer

No doubt that cloud will remain a key part of business IT spend as the public cloud market is expected to globally grow to \$299.4 billion by 2022. However, at the start of the new decade, large enterprise organizations will have a much clearer roadmap of what workloads should be in the cloud and what should be on-premises – and will rely on software that can easily navigate between the two environments. And, for those who are storing data in public cloud environments, more organizations will ensure they are taking responsibility for backing up that data in the cloud to avoid data loss or disruption in service in the event of a cloud outage.

3. Poor data management practices will cause brand reputation nightmares for more large enterprises

Millions of consumers, businesses, and public sector agencies are generating an immense amount of data daily. While the bulk of this content are images and video, which contain massive amounts of information, not far behind is data created from machines, database and application usage. Managing all of this data can be incredibly challenging and oversights can lead to data loss, outages, service disruptions and compliance violations – all of which can tarnish brand reputations in a matter of minutes. In the year ahead, we can expect more enterprises to undergo PR nightmares if they have not taken necessary steps to be exceptional stewards of their data.

Commvault

1. Expect CCPA to Fuel a Data Collection and Processing Backlash (Nigel Tozer, Solutions Director EMEA)

The California Consumer Privacy Act (CCPA) will highlight data collection and monetization in the US, just as the General Data Protection Regulation (GDPR) did in Europe. This will fuel a backlash on data collection and processing in the US, especially around political ad targeting during the 2020 election year. Companies such as Facebook and Google will come under greater pressure to distance themselves from this area, and data analysis companies that are now largely unheard

of will be in the news for the wrong reasons.

2. Data Analytics Moves to the Top of the Companies' Priority List (Matt Tyrer, Technology Evangelist)

Data volumes continue to grow, but we're knowing less and less about that data – which is a huge risk. In 2020, the focus on analytics will be driven by increased regulatory and compliance pressures, risks from data breaches and ransomware, and the need to properly classify data for AI and ML projects. Without “clean” data of value, these AI and ML projects will stumble. Data analytics will support intelligent decision making, feed AI and ML initiatives, and strengthen compliance stances within organizations. Expect more businesses to be hitting this point in their data maturity, where analytics projects take priority.

3. Multi-cloud Adoption Will Increase Demand For More Diverse Data Protection Capabilities (Penny Gralewski, Solutions Lead)

As organizations adopt more clouds for different organizational requirements, the need for fast, flexible data protection - able to protect a diverse set of data workloads - will increase. Organizations are choosing different clouds for different use cases, so today's data protection platforms need to accommodate a wide variety of cloud use cases, including Platform as a Service (PaaS), containers, and massive databases like Microsoft SQL Server, MySQL, PostgreSQL, Splunk, SAP HANA and Oracle.

CTera Networks (Aron Brand, CTO)

1. The decline of centralised cloud computing

The first-generation model of centralized cloud computing and storage has now run its course, and most of the new opportunities for enterprise data management reside at the edge.

2. Data growth outside the datacentre

Consider that 75% of enterprise data is created in branch offices, on mobile devices and by IoT-enabled smart devices. Such data growth outside the datacenter is the new reality, and it's creating a need for enterprises to deploy computing power and storage capabilities at the network edge, aka edge computing.

3. Powerful edge compute will result in more desirable devices

If you take your voice-based home virtual digital assistant from Amazon or Google. More powerful edge compute power would make these devices even more desirable, including:

- Improved response time: your Amazon Echo, Google Home or other smart speaker would perform speech recognition locally and deliver a faster answer
- Offline availability: whether there is a storm, a power cut or a cloud service outage, your device will still be up and running, or 'always-on' even when the network is offline
- Improved security and privacy: by minimising the need to send sensitive information to the cloud, data privacy risks are reduced

Data Dynamics (Piyush Mehta, CEO)

1. Storage management will become part of a broader infrastructure management stack

Storage admins will need to learn more skills tied to compute and networking as you will see management of these stacks to consolidate.

2. Data management will continue to be key focus for storage companies

Legacy and new entrants in the storage landscape have realized that a disc is a disc and parity of features/functionality as it relates to storage have become a 'must' in order to be in the game. The differentiator is going to be those that can help customers with a better data management solution that can help extract business value out of the data that's been created.

3. Multi-cloud will become more prevalent.

Public cloud vendors will have to provide access and allow for an ecosystem of leveraging data to be moved across providers based on usage. The egress charges will need to decrease as data mobility will continue to be a requirement for enterprises for a variety of use cases.

DataCore (Gerardo Dada, CMO)

1. Storage Unification

For decades, the industry has had discrete storage systems for high performance and for lower-cost storage, often from different vendors – a very hardware-centric approach. 2020 will see broader adoption of unified software-defined systems that make intelligent decisions on each piece of data, placing it in the tier it belongs given its performance and cost requirements—optimizing performance, utilization, and cost dynamically.

2. Intelligent Data Placement

Elements including richer metadata management and AI will empower data intelligence, which will increasingly enable software to drive automated decisions on the value of data and where it should be placed – including automatic movement to the cloud.

3. Shift from Hardware-Centric Model Will Accelerate Adoption of Software-Defined Storage

The hardware-centric mindset is not sustainable anymore. IT will increasingly stop thinking about storage systems as discrete, media-centric systems from multiple vendors, as these create islands of storage that become difficult to manage, inefficient in terms of capacity utilization, and nearly impossible to move data efficiently from one system to another. As the industry realizes that it's the data that

matters and not the actual storage system, software-defined storage platforms that control where data is placed will become far more valuable than the type of hardware or media. The industry went through this transition for compute years ago; it's time that storage catches up.

Datrium (Sazzala Reddy, CTO)

1. Ransomware will innovate faster than mechanisms to prevent it

Due to its insane profitability, the proliferation of non-state and state actors, and cybercrimes (including ransomware attacks) will cost the world \$6 trillion annually by 2021. Every business should investigate deploying a quick data recovery infrastructure that can help instantly roll back the IT environment to its pre-ransomware state and recover from an attack unharmed. Ransomware recovery will become a budget line item for the majority of CIOs in 2020.

2. Mainstream enterprises will finally embrace DR to the cloud

In 2020, mainstream business will become open to leveraging the cloud as a DR site and will start shutting down their physical DR sites because new cloud DR technologies will make it possible to leverage on-demand cloud resources during a disaster while keeping cloud costs low during the state of normal business operations.

3. SaaSification 2020

VMware vSphere's move to the cloud will trigger the rise of Cloud 3.0. All products in the enterprise are becoming SaaS products if they're not already. The convenience and on-demand economics of SaaS products in the public cloud make it an unstoppable trend. VMware has 300K+ customers with about 70+ million VMs deployed and has become available as a SaaS platform on most public clouds. With one of the world's biggest enterprise on-prem platforms rapidly transitioning to the SaaS world, it's inevitable that all the other third-party products in the VMware ecosystem will transform as well and "SaaSify."

DDN (Alex Bouzari, CEO)

1. AI-Enabled, Intelligent Infrastructures will get deployed at massive scale and disrupt the High Performance Computing Industry.

2. Multi-Cloud, flexible data center architectures will start to bring true operational simplicity, efficiency and agility to the Enterprise.

3. In 2020 DDN will deliver a full range of new intelligent infrastructure solutions ideally suited for Enterprise and HPC customers' demanding AI, IoT, Multi-Cloud, Big Data and Mixed Workloads at Scale.

DriveScale (Tom Lyon, Chief Scientist, Brian Pawlowski, CTO & Denise Shiffman, CPO)

1. The new private cloud will be created with elastic bare-metal infrastructure

While VM-deployed applications will continue to move to the public cloud, more bare-metal private cloud data centers will be created in 2020 to run performance-sensitive bare-metal and Kubernetes-native applications. This is being driven by the growth and investment in data-intensive apps and the need for scale-out data center architectures.

2. 2020 Will be the year of the Smart NIC

Smart NICs have been used for a few years by cloud providers to offload network switching and encapsulation, but in 2020 we'll see SmartNICs for storage applications and in private data centers. Amazon's Nitro SmartNIC was the first with NVMe storage capabilities, but now both Mellanox and Broadcom have announced similar capabilities in their chips, while high-flying startups Pensando and Fungible are building systems revolving around SmartNIC capabilities. Why SmartNICs now? Aren't these just "front-end processors"? The most compelling feature of SmartNICs is that they provide a separate security domain that cleanly separates user code from provider code. In these days of CPU vulnerabilities like Spectre and Meltdown, it may not be safe to assume the correctness of any security boundaries within a CPU.

3. 2020 will be the Year of Kubernetes

and people realize they have a lot more work to do to deploy Kubernetes for their data intensive applications. The realities of adopting containerization for persistent storage will continue to drive Kubernetes evolution. While the introduction of the CSI API for Kubernetes makes it easier for storage providers to integrate into K8s, requirements for cost efficient scale out application deployment on top of inexpensive commodity hardware remains a challenge in how best to deploy these applications over K8s.

ExaGrid (Bill Andrews, CEO)

- 1.** Public clouds such as Amazon AWS and Microsoft Azure will continue to grow but not at the expense of the corporate data center. Most corporate data centers have concluded that the public clouds cost more than running their own data centers and therefore, the corporate data centers are only doing specific projects with the public cloud providers such as archive data (untouched data), development operations (dev opps) for burstable compute, etc.
- 2.** The small businesses (SMB) will go to cloud solutions but not the public cloud. They are turning to true MSPs (managed service providers) that can provide more hands on value add. SMBs should not be operating data centers as they don't have the IT

expertise. It is better for SMB customers to outsource to MSPs.

3. The move to Kerberos and Containers will slowly begin to show up in production. To date most have been pilot testing. Over time containers will be the wave of the future due to IOT (internet of things) and other distributed data.

FujiFilm Recording Media US (Peter Faulhaber, President)

1. **Software-defined tape for object storage will emerge as a popular solution**, providing the interface to download data from object storage systems to compatible tape systems using standard S3 APIs. Users will be able to write objects directly to tape in native form, in a self- describing, open format. As a result, object storage users can leverage the value proposition of tape including lowest TCO, reliability and long term archivability.
2. **According to IDC, the rapidly growing Data Sphere means demand for persistent storage will be in the multi-zettabyte range in 2020**
The big data growth drivers including IOT, AI/ML, HD video (4k, 8k), surveillance, gaming and other apps will demand cost-effective, long term tape storage.
3. **New hyperscale data center markets will emerge in 2020 led by internet and cloud service providers in markets like China and India**
This will cause a shift in regional storage from USA-centric to global demand and present new opportunities for tape market growth.

Hammerspace (Brendan Wolfe, VP Product Marketing)

1. People will move data out of HDFS and onto NFS environments as they look to update their approach to big data analytics.
2. We will see a lot more companies moving stateful applications into Kubernetes and look for technologies to help them protect and manage the associated persistent data.
3. The three big clouds will ramp up the fight for more footprint in on-premises data centers.

Hitachi Vantara (Hu Yoshida, CTO)

1. The Demand For Large Scalable Enterprise Storage Arrays Will Increase

The increasing demand for core storage is changing. Capacity will be driven by big data and data lakes as well as copies required for exploratory, training, and governance purposes. The trend toward containers will generate swarms of workloads which will require instant storage. IoT will drive real-time applications which will need low latencies. The faster speed of NVMe devices will expose bottlenecks in smaller storage controllers and faster NVMe-oF fabrics will enable increased connectivity to servers. Core storage will also need to connect to edge and cloud and integrate as one pool of managed storage. Large scalable storage arrays with full enterprise capabilities will be the most economical way to address these requirements.

2. Edge to Core to Multi-Cloud Replaces Discussions Around Public/Private/Hybrid Cloud

By now the discussions around public/private/hybrid cloud have become moot as digital transformation has driven users to a combination of different cloud implementations for agility and scalability. Digital transformation has also driven smart applications to the edge for mobility and real-time processing. This has increased the complexity of storage management as users seek to balance their business requirements across this multi-vendor, multi-tier storage landscape in an exploding data environment. Business requirements have also increased in complexity due to looming cyber threats, increasing regulations around privacy and transparency, regional requirements, and democratization of data usage. The only way to manage this environment is to take an open, holistic, edge to core to multi-cloud approach to storage and data management.

3. Edge Computing Will See Greater Adoption with DataOps Packages For The Edge

The introduction 5G and IoT will bring more focus to the edge and this will be the new battleground for cloud service providers who are looking to extend their reach beyond the cloud. The arguments for edge computing are compelling: Low latency for applications like autonomous vehicles and telemedicine, Integration of OT (operational technology) and IT (information technology) for improved business decisions, and a shorter control loop for better QoS. However, the edge is like the wild west. OT (Operational Technology) facilities like factories and manufacturing sites were run in a closed environment and lack the disciplines that have been honed by IT where application and data had to exist in a shared and open environment. The protocols, networks and systems on the edge are very diverse, and low cost, while IT systems are standardized and generally higher cost. An edge gateway must be used to integrate the two types of systems and data to realize the value of combining OT and IT. Tools like avatars, a software representation of the physical machine will help to bridge the gap. While technology is important, edge success depends on the business use case. A software package like Lumada Edge Intelligence can reduce the time to value by providing actionable, real-time insights to help critical operations be more predictable and manageable. It provides the capabilities for local data operations, machine learning, streaming analytics and standalone IoT application solutions. When paired with the powerful services from experienced IT

and OT professionals, customers can develop a comprehensive DataOps strategy that manages assets holistically from the edge to core to multicloud.

HYCU (Simon Taylor, CEO)

1. Multi-Clouds Expand in Use

It comes as no surprise that multi-cloud use is now becoming mainstream in enterprises globally. We are seeing as many as four or more cloud platforms in use with our customers, regardless of size and geographical location. And, this goes for mission critical and numerous enterprise workloads. To support this requires the need for a multi-vendor and multi-disciplinary strategy for cloud usage for IT today. This echoes a prediction from last year where we see more customers that will need solutions that work across multiple clouds, be they on-premises or public. We will continue to see vendors “marketing” their solutions as “multi-cloud” but customers are much savvier and truly understand the value in a cloud-native solution running as a service on any cloud. This will be the year where customers look for ways to help them manage their data the way they want it, in their cloud of choice with a solution that makes it theirs to control.

2. Simplicity Is More Than Just a Buzzword

Customers are smarter than ever before. They understand the full value of eliminating complexity from their IT infrastructure. It’s a major reason why companies like Nutanix are so successful. They have found a way to eliminate unnecessary redundancy and streamline IT infrastructure and operations into 1-click simplicity. That is the driving force behind innovation in the next-generation platforms in the industry today that include AWS, Azure and Google Cloud Platform in the Public Cloud space and Nutanix and VMware in on-premises Enterprise Cloud space. While customers continue to try to consolidate around the platforms that deliver the most value for the enterprise workloads they need managed, protected and recovered, they want to do this with simplicity in mind. No additional hardware, software or infrastructure required. And, that run natively as a service on the cloud platform of choice. There will be fewer standalone companies in 2020 if they don’t address this fundamental customer need for simplicity.

3. Built for Purpose Makes a Difference

When we introduced the idea of being purpose-built for a customer’s cloud platform of choice, first for on-premises and now for public cloud, it was significant. No one data management and protection platform can do it all. We fundamentally believed and still believe that those solutions that complement the underlying cloud platform of choice without compromising the reason the customer selected it are critical. As more and more of the majority of workloads move to the cloud over time, it makes a difference when customers select and deploy data management and protection solutions that are built for purpose. That means they do more than just look and feel like the platform they are helping to protect and support but that run natively, deploy as a service, are lightweight and keep application consistency intact. This in

turn will help companies make the most of their IT investments and allow them to continue to focus on innovation and driving value for their own customers.

Infinidat (Stanley Zaffos, SVP Product Marketing)

1. 5G in 2020: Accelerating the Volume & Velocity of Data Collected

2020 will be the year 5G enters the mainstream. The new wireless standard will start generating real value for companies deploying IoT projects. Gartner predicts 66% of organizations will deploy 5G, and 59% will include IoT communications in the use case for 5G. Companies deploying IoT projects will need to plan for the data deluge that is coming. They'll need to set up infrastructure and processes to filter the data, pre-analyze it, categorize it, store it and dispose of it. Companies that plan well and allow for some flexibility will execute successful projects at reasonable costs. Those that don't will spend only what they can but be limited in the value they unlock in the limited data they capture and store.

2. AI Applications in 2020: Ushering in the Age of 'Smarter Storage'

Artificial intelligence (AI) workloads will continue to generate business value in 2020. But, for organizations to increase their reliance on AI, storage vendors will need to make it easier for AI applications to access more data faster, in turn helping the systems learn faster and unlock the value of the data. As we enter 2020, data sets are getting bigger and demands for instantaneous decision making are becoming more prevalent. This puts stress on the training systems. Expect more demand for smarter storage systems to match the escalating intelligence of the applications themselves. We'll see more investments in tools like software-defined switches to open up more pathways for hardcore analytics; QoS functions to dole out information more strategically; scale-out system architectures; and the ability to deliver data with lower latency.

3. Containers in 2020: Creating a More Competitive Storage Environment

In 2020, containers and multi-cloud implementations will continue to accelerate. More enterprises will push to create flexible computing environments where multiple clouds serve specific strategic purposes. They will embrace the flexibility containers promise, creating set-ups where containers can move freely between public cloud, private cloud and on-premises environments.

Increased use of containers and Kubernetes will help create a more competitive storage environment. Being able to port workloads seamlessly among diverse environments will diminish the strength of vendor lock-in and put pressure on incumbent storage vendors to innovate in areas that improve financial and operational efficiency: lower acquisition and ownership costs; improve staff productivity via more autonomic operation, cloud integration; and new pricing and service offerings.

iXsystems (Mike Lauth, CEO)

1. Software-Defined-Storage (storage software that can operate on third party server hardware) will be the basis for nearly all new storage products

The lines between block, file and object storage are increasingly blurred. Users are benefitting from the agility, expandability and cost structures of Software Defined Unified Storage.

2. Open Source storage will slash the storage costs of many organizations

Data is still growing, but SSD and HDD storage costs are not reducing as quickly. Users will increase their deployments of OpenZFS, HDFS, Gluster, Minio, etc. on industry standard hardware without the vendor lock-in and costs associated with proprietary stacks. Enterprise grade support of Open Storage software is the key enabler of this transition.

3. Hybrid Clouds will be the storage reality of every large organization

The performance of local storage and the long term data retention of geo-distributed clouds are necessary partners. Cloud services have replaced tape as the third copy of data. Data security, storage costs, and migration flexibility are critical.

Kaminario (Eyal David, CTO)

1. Companies will finally enable a true hybrid cloud strategy

Migrating mission-critical applications to the cloud is difficult. By leveraging a common set of shared services that enable companies to decouple the management and movement of data from the infrastructure it runs on, companies will be able to improve business agility and reduce the risk of moving to the cloud.

2. Data Science will be democratized

With the explosion of data and the rise of artificial intelligence and machine learning, IT organizations will have more ways to quickly analyze and make sense of infrastructure data. This will make it easier for stakeholders throughout the organization, who typically don't have access to this information, to leverage it.

3. User Experience will be vastly improved

5G will enable everything/everywhere edge computing that make technologies such as AR/VR and smart, connected devices possible. WIFI6, meanwhile, will offer users more robust outdoor network operations and high performance in dense environments. The combination of the two means lagging/unavailable service will become a thing of the past.

Kaseya (Michael Sanders, GM of Unitrends and Spanning)

both Kaseya companies

1. Pricing and innovation

There will be a major plateau in storage and compute in the cloud as we are approaching a lot of the physical limits of our current technology faster than we thought. The upcoming storage crunch is speeding up innovation; Microsoft and other vendors are experimenting with innovative solutions like glass storage. If they don't come up with a solution quickly, however, cloud storage prices might start to go up.

In addition to the storage limits, there's the CPU side. In 2020 expect more workloads to get pushed back to the edge. This means more devices (endpoints) will need to be protected in the coming years.

2. The rise of multi-cloud toolsets

Multi-cloud is already a reality in 2019, but it's taken a while for the management practices of operating multiple workloads across multiple services, regions, and cloud vendors to mature, while the pressures on business continues to build to do more with less, and yet companies continue to have to "DIY" it themselves.

And as workloads start to dominate the decisions behind "which cloud" to make the best use of each vendor's offerings, expect increased demand for toolsets that can operate natively with different cloud vendors. Sadly, cloud management portals today are generally limited to discovery or cost optimization use cases and fail to tackle the day-to-day operational management pain experienced by today's CloudOps and DevOps teams.

As an example of the rising complexity in the cloud, Amazon Web Services alone has 250 different services (May 2019), each with its own management console and set of APIs.

3. DRaaS is Now Mainstream

Disaster Recovery-as-a-Service (DRaaS) is now mainstream. Large organizations have adopted DRaaS at the highest rates, however, I expect in 2020 to see the adoption of DRaaS by small and mid-sized organizations to drastically increase as organizations discover that not all DRaaS services require their IT departments to become experts in hyper-scale clouds. As a result, SMBs will outsource DRaaS to experts at a fixed price and with little requirement for their time or technical overview.

Komprise (Krishna Subramanian, President & COO)

1. Edge Buildout On the Rise

Edge buildout is already happening, and its pace is accelerating with trends like Internet-of-Things (IoT), self-driving cars, and biometrics. IDC has predicted that by

2023, 50% of all enterprise IT infrastructure will be deployed at the edge – up from 10% today. More apps are generating tons of data at the edge, raising the question of why data should not be better understood and managed directly at the edge. Imagine if you could analyze data, figure out what data is useful and needs to be brought back to the datacenter, and directly process the rest of the data at the edge itself without having to first move it all? This is why edge data management will rise in importance in 2020.

2. **Multi-Cloud Gains Traction**

Most enterprises using the public cloud already have a hybrid/multi-cloud strategy, and enterprises are increasingly choosing to diversify their cloud strategy across two or more clouds. In fact, Forrester Research found this year that 62%+ of public cloud adopters already use two or more cloud platforms. As this trend continues, enterprises in 2020 will need a simple way to understand and manage their data sprawl across clouds and the hybrid enterprise, leading to greater demand for cross-cloud data management solutions that are not tied to any particular cloud or storage vendor.

3. **AI-Driven Intelligence Grows**

For the last couple of years, Artificial Intelligence (AI) and Machine Learning (ML) have been a big theme and this trend is continuing to grow. While these have initially been more of marketing buzzwords, the potential of AI in data management is clear – how can you manage something you don't understand? By using analytics to drive AI/ML, analytics-driven data management solutions can continue to leverage the understanding of data to drive better management. Watch for more exciting developments in this space that leverage deep analytics and data lakes across different storage repositories to help you better manage your data.

Lightbits Labs (Eran Kirzner, CEO)

1. Cloud providers, both large and small, will move to disaggregated storage where storage is handled separately from compute, and away from traditional direct-attached storage (DAS) architecture where storage is located with the compute nodes. From the top-tier cloud providers to the enterprise level, the trend will be moving from traditional all-flash arrays or storage controllers to more of a cloud-like architecture. And more enterprises will adopt of NVMe/TCP as opposed to NVMe over other fabrics.
2. Flash prices will continue to drop as they did in 2019, making it increasingly affordable and pushing more and more companies to adopt flash as their preferred storage media for its affordable price/performance. Newer NVMe/TCP technologies will further enhance the appeal and benefits of flash by helping to reduce latency. Emerging flash solutions, including Quad Level Cell (QLC), QTC, SLC and MLC flash, will also see increasing market traction. QLC was introduced in 2019 and therefore had only minimal market adoption. That will change in 2020, particularly among

companies that have deployed Global Flash Translation Layers (GFTL) to overcome QLC's inherent issues.

3. As Kubernetes continues its successful assault on the data center, more companies will look for storage for containerized environments that move away from traditional Kubernetes deployments with DAS into more flexible and reliable solutions for persistent container storage that NVMe/TCP can provide.

MemVerge (Charles Fan, CEO)

1. Data Center Architecture redefined: Storage Class Memory will make way for a memory-centric data center

With an increasing demand from data center applications, paired with the increased speed of processing, there will be a huge push towards a memory-centric data center. Computing innovations are happening at a rapid pace, with more and more computation tech—from x86 to GPUs to ARM. This will continue to open up new topology between CPU and memory units. While architecture currently tends to be more disaggregated between the computing layer and the storage layer, I believe we are headed towards a memory-centric data center very soon.

A new MCI (memory-converged infrastructure) layer powered by Storage Class Memory (SCM) will boast a larger capacity of memory for applications along with persistence that makes storage tier obsolete. SCM will be the most disruptive new hardware technology within the data center architecture next year. Over time we expect MCI to replace both the existing memory tier and the performance tier of storage and will become a \$10 billion+ dollar market. This transition will take 5-10 years to complete and 2020 will be the first year we start to see any impact.

2. Storage Class Memory technology will evolve and the market will expand

This year marks year 0 of storage class memory (SCM), with Intel shipping Optane DC Persistent Memory in Q2. Over the next 2-3 years, I believe there will be other major semiconductor vendors entering the market, while Intel continues to improve its Optane technology. We expect Intel Optane DC Persistent Memory's capacity to double and its cost will drop to half, every 18 months. In addition to Intel, we will likely see early examples from additional players in this market in the next 2-3 years, effectively validating this new market for SCM. Intel, in particular, plans to release gen 2 of Optane by the end of 2020 – I expect we will see significant improvement in terms of density and speed.

3. Data reasoning will be one of the most critical skills next-generation IT staffers will need to possess

IT is having a coming of age moment. It is slowly improving in terms of deployment of mature data science and machine learning within the enterprise. IT staffers will need to understand all facets of infrastructure and operationalize platforms well. As big data continues to grow, IT staffers will also need to be able to reason about data and pull actionable insights.

Minio (Anand Babu Periasamy, CEO)

1. Appliance vendors hit the wall in 2020 resulting in plummeting valuations. Shifting to a software only model is not window dressing around pricing. It requires a DNA change.
2. Modern databases (the revenue engine for SAN and NAS) go object storage native. This is a real threat to the traditional enterprise storage vendors.
3. The “standards” associated with the cloud are increasingly defined by open source. It will become the primary strategy for the largest players in the space - from VMware, Microsoft, Amazon and Google.

Nasuni (Russ Kennedy, CPO)

1. Prepare for even faster exponential data growth

Data is already doubling every 2-3 years. In 2020, the rate of data growth will significantly accelerate, driven by higher resolution mobile phone cameras, the growth of 4K video, increased video surveillance, genomics research, IoT machine data, 3D medical images... and so many more. Both the number and size of files will increase at an ever faster rate, and IT organizations need to prepare now to store, protect and provide access to it all.

2. Anticipate continued “cloud-washing” by incumbent storage vendors

Incumbent vendors’ storage systems were designed for on-premises environments, but as companies move more of their workloads to the cloud, the market for on-premises gear is shrinking. Unfortunately, vendors’ efforts to connect or port these systems to the cloud doesn't necessarily make things more efficient or easier to use. And in some cases, it actually makes IT harder to manage, especially if applications and data are still partly on-prem, and you’re trying to provide access to users around the world.

3. Consolidation of cloud-based companies will accelerate

The incredible success and growth of AWS and Azure have other companies eager to carve out their own slice of the cloud. As a result, cloud-friendly technologies are highly coveted, especially by large, traditional IT vendors. Expect to see a lot more M&A in 2020 of cloud-based companies, which we recently witnessed through the acquisition of Veeam in 2020.

NGD Systems (Scott Shadley, VP Marketing)

1. Move Less and Analyze More at the Edge

NVMe has provided a measure of relief and proven to remove existing storage protocol bottlenecks for platforms churning out terabytes and petabytes of data on a regular basis. But, is that enough? Even though NVMe is substantially faster, it is not fast enough by itself when petabytes of data are required to be analyzed and processed in real time.

This is where Computational Storage comes in and solves the problem of data management and movement. Computational Storage, especially the way we marry the use of NVMe SSDs and compute power, adds analytical power and speed so that results can be accomplished right away and where the data is generated.

2. Better 5G Connectivity

In 2020, more edge related devices will be needed to process massive data workloads. The advent of 5G is no different. As more cell towers are built to support 5G, there also needs to be more complex infrastructure at each bay station that can manage the data “in and out of the box” so that user data is optimally utilized. Computational Storage with its small form factors and added compute power can pack an analytical uppercut punch in the limited size and power enabled edge-datacenters that live at each of these new cell tower platforms. By providing additional compute to the confined resources that exist is paramount to successful growth of this space. Instead of requiring even more hardware and power to the server, the advent of high capacity Computational Storage provides the needed offload to the system to allow for great deployments.

3. Simplify the Traffic Flow of Content Delivery Networks

Streaming services have continued to dominate headlines this year, with the recent launches of Apple TV+ and Disney Plus combined with Netflix, Hulu and Amazon Prime’s increasing investments. This poses a major hurdle for the content delivery networks (CDNs) – and where Computational Storage can be a major asset. While a typical CDN traffic flow involves lots of data movement and processing spread out over a variety of edge infrastructure, Computational Storage, can simplify this flow.

Panasas (Curtis Anderson, Senior Software Architect)

1. The importance of Total Cost of Ownership (TCO) : HPC storage solutions must deliver value far beyond the initial purchase price

As the requirements for HPC storage systems are becoming more diverse with the addition of new workloads such as Artificial Intelligence (AI), there is an increasing need to start looking at the overall impact on the organisation of the ongoing cost of operations, user productivity and the time to quality outcomes. In addition to evaluating the price/performance ratio, buyers will need to start paying close attention to a range of purchasing considerations that go beyond the initial investment. Those include the cost of unplanned downtime in terms of application

user productivity, the cost of complexity and the headcount required to manage it, and the need for responsive support for mission-critical infrastructure such as your storage.

- 2. As Enterprise's AI projects graduate from "exploratory" to "production" they will leave the public clouds for less costly on-premises solutions, funding a boom in HPC infrastructure build-out, but the requirements for that infrastructure will have changed based upon their cloud experience**

Public clouds are great for learning and experimentation, but not for high-utilisation production operations. Public clouds will, however, have a large influence on the next generation of on-premise infrastructure that is built. The need for the lowest time-to-solution, quickly taking action based upon the insights that AI can give you, drives AI to push the underlying hardware (e.g: GPUs and storage) as hard as it can go. But the simple truth is that the cost of a dedicated resource in a public cloud is higher than the cost of owning that resource. Another simple truth is that the value of AI is the computer deriving information that you can act upon from mountains of data. Add in the fact that AI has an insatiable need for growth of training data, and that public clouds have never-ending charges for data storage, and the costs climb. Put those simple facts together and it's clear that production AI will be less costly if it is performed on-premises. The industry has become used to the extreme flexibility and simplicity of management that public clouds provide, and they will want to retain those characteristics in their on-premise solutions at the lower cost it provides.

Pavilion Data Systems (VR Satish, CTO)

- 1. Convergence of primary and secondary storage ushers NVMe as a cost-effective and performant media for secondary storage.**
- 2. NVMeOF becomes the preferred protocol for all new infrastructure deployments.**
- 3. High-performance edge storage becomes a requirement for hybrid cloud deployment.**

Portworx (Murli Thirumale, CEO)

- 1. Kubernetes will start to be used to manage IT infrastructure, not just containerized applications**

VMWare's Project Pacific is one of the first examples of this, but I expect to see more such offerings in 2020. The move to containerize applications and orchestrate them with Kubernetes is well underway and driving rapid application deployment and portability in enterprises led by DevOps. Project Pacific is a bold move to extend

VSphere with Kubernetes and get traditional IT admins into the mix by having VSphere and Kubernetes now manage IT infrastructure. VMWare is both making Kubernetes a first citizen along with VSphere but also saying Kubernetes via VSphere can now manage not just Apps but also VMs, Storage, Networking and Compute.

2. IPOs for companies with 'cult' CEOs and founders will be discounted by the market as bearing too much risk

The days of CEO worship are over. High profile problems at Theranos, Uber, and more recently WeWork have made investors skittish about relying on charismatic individuals. In 2020, a good product backed by a sustainable business model will be far more attractive to investors than a messianic CEO.

3. We heard a lot about 5G this year but network coverage from the major telecoms providers is spotty

I expect to see this change significantly in 2020. With 5G networks spanning the country, device makers and application developers will start to take advantage of the new high-speed technology. This will mean not just richer smartphone apps but also a range of IoT uses that will reshape computing at the edge.

Pure Storage (Matt Kixmoeller, VP Strategy)

1. Customers Will Demand A Subscription to Innovation with As-a-Service Business Models

As-a-Service models have existed since the beginning of public cloud. For most consumers of storage, hybrid cloud is the reality and the future - and they are looking to get the best out of both worlds; to drive simplicity and automation from their on-premise infrastructure so they can manage it like they manage the cloud, and to get the same enterprise capabilities and control in the cloud, as they have on-premise -- both in a flexible, subscription-based as-a-Service model. In 2020, the demand for as-a-Service in storage will increase and organizations are speaking with their wallets with more investment in OPEX models, but successful models need to balance both the operations and purchasing aspects. From an operations perspective, key attributes include standardization (vs snowflakes), on-demand access, API-driven management, and limitless scale. On the consumption side, key traits include a pay for what you use model, bursting capabilities (flex up/down as needed), and a non-disruptive evergreen experience, services can be grown/evolved over time without disruption. And all this delivered as a 100% pay-per-month OPEX service.

2. Modern Analytics Has Reached Rocketship Status

Fueling the growth for modern analytics is more affordable infrastructure options such as more powerful CPUs, consumption-based infrastructure, available both on-prem and in the public cloud, and lower priced flash memory. There is also a significant growth in stream analytics platforms, both open source (Apache Flink, Apache Beam and Spark Streaming) and commercial (Splunk DSP) replacing more and more batch-based processing platforms. Modern analytics can now reach larger

scale with cloud native analytics architectures comprised of stateless servers and container and high-performance S3 object stores. Additionally, the unbridled growth of data sources including smart devices (smart home, wearables, connected cars, industrial internet, etc.) will drive the adoption of modern analytics in order to drive more insights.

3. AI Operations Will Go From Advisory Roles to Automated Action As Customers Want a Hands-Free Approach

Organizations will be more open to AI making decisions for them. Customers want to set policies and let the vendors implement the policies, which is partially driven by the declarative nature of Kubernetes and container management. The simplicity of containers will enable organizations to define a state, and the container will be the catalyst. The technology should then drive and deliver insights within the whole environment. AI will be applied to efficiently finding where the predictive model performs poorly and augmenting data for that feature space. This is critical for AI applications like anomaly detection and automatic root cause analysis to scale and be applicable in more contexts.

Quantum (Eric Bassier, Senior Director of Product Marketing)

1. Video and Images Represent Biggest Data Generator for Most Enterprises

Between surveillance footage, video for marketing and training purposes across all industries, and the use of high-res image and video content generated by machines in use cases as diverse as movie and TV production, autonomous vehicle design, manufacturing, healthcare – we believe video and high-res image content will represent biggest ‘class’ of data for most enterprises.

2. The Tape Storage Market will Grow, Reversing a Decade-Long Declining Trend

Tape has emerged as a key technology for massive scale cold storage infrastructure – both in the cloud and on-premise. And we believe the architectures used in the cloud will eventually make their way back into the enterprise. So we believe the tape market will grow, and continue to grow over the next 5-10 years, based on a new use case for tape as cold storage for (primarily) video and high res image data.

3. Hybrid- and Multi-Cloud Architectures Become the Norm

Many companies are in some type of hybrid-cloud state, and customers are expecting that vendors provide an even more seamless experience between on-premise hardware infrastructures, and cloud infrastructures. Customers will also expect that vendors can offer a multi-cloud experience, so customers are not locked into a single cloud provider.

Qumulo (Molly Presley, Global Product Marketing Director)

1. NVMe file storage will be adopted broadly for performance starved, low-latency applications in 2020

NVMe is a communications protocol developed specifically for all-flash storage. It enables faster performance and greater density compared to legacy protocols. It's geared for enterprise workloads that require top performance, such as real-time data analytics, online trading platforms and other latency-sensitive workloads.

2. Data-driven businesses will have to shift some workloads to the cloud for data processing, machine learning (ML) and artificial intelligence-(AI) driven workloads

Every major enterprise in the world is going to become a hybrid enterprise. In industries across all major vertical markets including M&E, transportation, bio and pharma, customers are using large volumes of unstructured data in order to accomplish their mission. Despite tremendous downward pressure, IT budgets don't grow at the same rate as the rest of business. The public cloud enables a way to solve that problem with its elastic compute and elastic resources.

3. Scale-out File Storage will become the preferred technology for on-prem unstructured data active archives

Modern file storage solutions deliver performance and economics in a single tier solution managed by intelligent caching. Object storage is not the best fit for on-premises customers seeking simplicity to deliver to performance applications and retain cost effectively, object storage was developed as a precursor to webscale technology and as the storage medium for web technologies. It was meant to be great for datasets that approach exabyte data level and are geographically distributed. In 2020, we believe the on-premises object storage market will evaporate and will become wholly file based.

Reduxio (Jacob Cherian, CMO)

The true power of Kubernetes is to redefine the idea of the cloud for customers by breaking down their infrastructure silos - public cloud resources, and the customers owned and dedicated resources. Customers will increasingly leverage Kubernetes and containers to build infrastructure agnostic IT for their applications - one environment that can be instantiated anywhere rapidly and run all their applications. This is cloud infrastructure as code.

A key requirement for this will be storage that is native and integrated into Kubernetes and provides mobility for applications and data across all the infrastructure pools, stitching the pools together into a single cloud.

Scale Computing (Alan Conboy, office of the CTO)

1. Customers don't want solutions that are complex, difficult to use, require lots of training and cost more because of the additional licenses.
2. Gartner Magic Quadrant's niche players will continue driving innovation with feature rich technologies that simplify IT operations, for example by supporting the growing edge computing industry.
3. Technologies from niche players will make VMware's quasi-vendor agnostic approach look increasingly old-fashioned

Looking ahead to 2020, we will see niche players stay niche. Why? Leaders will always have a place offering well-entrenched legacy infrastructure, but staying niche means being a trendsetter. It means the offering is different, great at what it does, and makes tech better and more accessible to everyone. This is the reason why I believe customers will continue investing in niche players that offer the technology they need in the way that works best for them.

Scality (Paul Speciale, CPO)

1. Object storage at the edge will be on flash

Object storage will move into the edge for applications that capture large data streams from a wide variety of mobile, IoT and other connected devices. This will include event streams and logs, sensor and device data, vehicle drive data, image and video media data and more, with high data rates and high concurrency from thousands or more simultaneous data streams.

These applications will be developed for cloud native deployment and will therefore naturally embrace RESTful object style storage protocols, making object storage on flash media an optimal choice on the edge to support this emerging class of data-centric applications.

2. New ways of identifying patients, customers, and depositors

will be developed in 2020, as the already accelerating pace of hacking and data breaches continues to pick up speed.

There is enormous value in stored data. Until they make these changes, hospitals and medical providers, for example, will remain strong targets due to the value of the data they hold: not only patient health information, but the patient identification that goes along with it (government ID, birth date, address, etc.).

3. Data storage will become massively decentralised

as enterprises leverage a combination of on-premises and public cloud IT resources. This will create a need for a unified namespace and control plane to simplify data visibility and access.

Moreover, corporations will use a variety of public clouds, each one selected to help solve specific business problems, thereby creating a multi-cloud data management problem. In addition, the emergence of edge computing will further drive decentralisation as corporations choose to deploy IT resources “near” the edge devices they manage. These trends all help to create a new and extreme “cloud data silos” scenario, that can only be addressed by solutions that provide global data visibility across these distributed clouds and data centers.

Spectra Logic (Matt Starr, CTO)

1. Digital data ownership is akin to the oil/steel of the early 20th Century

Think about Rockefeller, Carnegie, Ford and the like, industrial giants of the 20th century. In the 21st century it will be the data owner. Like Disney, who recently purchased 21st century fox to become one of the largest Media and Entertainment content owners in the world. Health Care, Oil and Gas exploration, geospatial, it is not just media and entertainment where data is the most valuable asset. Nearly every industry is seeing an acquisition strategy that takes into account the digital asset ownership.

2. Cloud Egress cost will start to drive end-user strategies

Cloud storage is all the rage in 2019, some of the larger users of these Storage as a Service offerings are discovering that not having a local copy is the driving factor in their storage cost, in-turn forcing a closer look at hybrid cloud, especially having a local copy while still utilizing the cloud as the second copy. Allowing a local restore and usage of data without egress costs and a cloud based copy for DR and sharing.

3. Machine learning and AI have moved forward in 2019 and will continue in the next decade

The data sets used to train these new systems will be a leading factor for how far AI can go. It is easy for AI to find copyrighted music on Youtube, and soon AI will be creating more of the musical background we hear. Outside of art, another example is AI used to develop autonomous vehicles will deliver safer vehicles as real world data sets are brought back and worked through AI for better vehicle safety. Look for AI to be in nearly every part of your life soon.

StorCentric (Surya Varanasi, CTO)

parent company of Nexsan, Drobo, Retrospect and Vexata

1. In the enterprise space, on-premises storage will see a comeback

Businesses have been moving to the cloud for primary and archive/DR storage for a long time. In 2020, on-premises storage, whether for active or standby, will see a comeback – especially, as customers are hit with cloud-use bills that are dramatically higher than originally anticipated.

In 2020, to be successful and provide value to the business, enterprise IT will need to be able to straddle the worlds of the cloud and on-premises storage. Software solutions that enable swift mobility between these two domains will become increasingly critical. Only via this bi-model model will IT be able to achieve the highest performance, scalability and capabilities, as well as the safest retention, at the most cost effective price.

2. In the Prosumer and SMB space, multi-layered data management and protection will become priority #1

In the Professional Consumer (i.e., Prosumer) and SMB space, data storage and protection has always been a priority, but cost has been a roadblock for those seeking to employ comprehensive end-to-end data management and protection solutions.

In 2020, Prosumers and SMBs will demand solutions that enable them to seamlessly and affordably layer features and functionality onto their on-site storage, such as integration with off-site cloud and SaaS backup and recovery solutions, with flexible cross-platform support for all major platforms including Windows, Mac, Linux, VMware and Hyper-V.

StorMagic (Bruce Kornfeld, CMO & GM-Americas)

1. More data is being created at the edge than ever before

In 2019, more data was created at the edge than ever before, and that will continue into 2020 and beyond due to a mass movement away from traditional datacenters and public clouds. The overall storage landscape is transforming, and more customers now prefer to process and manage data at the edge of the network because it is more efficient and cost-effective.

2. Key management adoption in edge HCI environments will increase

Security is a growing concern at edge computing sites, where there is typically no IT staff member and no physical security present. As servers continue to get smaller in size the risk of theft dramatically increases. CIOs - particularly in the healthcare, finance and retail markets - will begin integrating data encryption with key management to fully protect data managed at their remote sites and branch offices located at the edge.

3. Virtualization enhancements will be introduced to support physically smaller environments

As IT footprints continue to shrink, virtualization isn't just for your grandfather's datacenter anymore. Smaller-sized servers will continue to hit the market at a 2U, or smaller, footprint. Software companies will respond with lightweight virtual storage solutions that offer a low-cost entry point to suit these smaller hardware products and justify ROI and IT projects as a whole.

StorONE (Gal Naor, CEO)

1. Price per TB

Dramatic reduction due to better software solution with less hardware required. What we have seen in 2019 and we expect will continue through 2020 is inefficient utilization that plagues legacy storage architectures. This will only get worse as customers still aren't getting the rated speed and capacity that they pay for, that the drives are capable of. End users are increasingly frustrated by this, and frustrated having to overbuy. In 2020 companies will look for more efficiency and better resource utilization, for a lower footprint and to achieve the performance and density they're paying for.

2. Built in complete data protection and back up will emerge as part of storage solutions

No longer will organizations have to maintain multiple systems to get complete enterprise data protection and performance. NO need to purchase, integrate and manage standalone systems for data retention/protection, back up, and data integrity. In addition, data protection (RAID) levels based per volume and not the same RAID for the entire storage solution. And no more compromising on data protection to get high performance. Built in data reliability will be a mandatory requirement as part of any storage solution and will be all of these services will be included without extra cost.

3. In 2020 all enterprise storage services will finally be available in the same solution

All the storage protocols (block, file and object) will be included and all drive types (NVMe, SSD, HDD, and Optane) will be supported in the same solution with the ability to simultaneously run different protocols on the same drives. Storage managers will have the full flexibility to manage all their storage requirements and to grow on demand according to their needs. A single environment can manage any use case – performance, capacity, protocol, drive type and data protection based per volume.

StrongBox Data Solutions (Floyd Christofferson)

1. Data management across storage types is the biggest emerging problem

- This is made worse by exponential growth of unstructured data
- Increasingly data managers don't know for sure what they have, and where it should be kept, or if it can be deleted
- Solutions focused on intelligent cross-platform management of data are needed to address this

2. The proliferation of storage types, including multiple cloud and software-defined choices makes the problem worse

- This adds complexity and cost to IT administrators
- Causes them to rely on multiple tools, and manage user experience
- Often leads to over provisioning, and difficulty to contain costs or take full advantage of the savings they seek from more economical storage choices

3. On-prem cloud (object), plus the increasing public cloud offerings cause protocol problems for traditional file-based workflows

- The ability to bridge multiple vendor solutions for on-prem and off-prem without adding complexity or user disruption becomes key
- The ability to understand the data itself, and let data intelligence drive data placement enable true storage optimization and cost savings
- This also reduces the load on IT staff, which has otherwise become a rapidly growing operational expense. Storage and data management OPEX are about 5x the acquisition costs for storage, which therefore neutralizes cost savings achieved by more cost effective storage choices

Talon Storage (Jaap van Duijvenbode, Product Director)

1. In 2020, we will see a continued but even stronger pivot to an enterprise hybrid cloud strategy, one that leverages a combination of on-premises storage and cloud storage. Depending upon the nature of the workload, the data characteristics (i.e. size) or compliance factors (GDPR etc.) can drive a paradigm of regionally located cloud storage footprints, accessed by users/locations in proximity to that storage resource. Backup/DR use cases are the simplest and most forgiving in terms of ongoing operational complexity, and as such are where many enterprises get started in the use of multi-cloud storage resources.
2. Also, we predict that 2020 will see enterprises looking at a varied footprint, with different functions supported from different vendors. From a business perspective most larger enterprises do not want to sacrifice negotiating leverage with regards to their OPEX-based infrastructure. There are advantages in going all in with one vendor. In today's pay-as-you-go world, an enterprise could achieve a volume discount by sending everything in the direction of one cloud provider. It takes work to regularly review if the organization is getting the best deal, and switching costs are high. By maintaining a footprint with different providers, customers stay current with pricing policies/promotions, and have easier flexibility in directing workloads to resources.
3. Compliance will only continue to be an increasingly important factor in the multi-cloud decision. For larger or publicly-traded companies, the requirement to employ best practices of risk mitigation will drive a multi-cloud approach to spreading the risk. For global companies, the need to keep certain data assets in certain geographic domains will factor into the equation.

Virtana (Tim Van Ash, SVP Products)

1. Cloud is not cheaper, and overprovisioning must be brought under control

In 2019 organisations came to the realisation that cloud is not necessarily cheaper. This drove the need for visibility and control of cloud expense from the start to prevent escalating costs. Cloud cost management has become a full-time role as organisations have tried to figure out and get a handle on costs.

2020 will see the continued maturation of the cloud story with a far more considered approach, focused on cost optimisation and cost control. Arguably, cost has always been a driver for cloud, but the highest priorities in 2020 will be determining the real cost of cloud, and developing effective cost governance.

Instead of rushing workloads to the cloud, organisations will seek to understand to how they can optimise their workloads. It will start with right sizing of their workloads, then deciding whether it is a candidate to move to the cloud, or remain on-premises. Enterprises will begin to see the value and savings that can arise as a result of intelligent workload management and data decision-making, leading to more effective outcomes in terms of cost and performance.

Organisations in 2020 will look to reduce or eliminate the over provisioning of public cloud resources that they saw in 2019, which resulted from poor visibility into workloads and their resource consumption. For example, AWS Lambda requires the purchase of a set memory allocation regardless of whether it is needed or used, resulting in a tendency to vastly overprovision memory. This has been impossible to achieve, with cost and performance being managed by separate tools and teams.

In 2020, Hybrid applications and workloads will continue to consume digital infrastructure across both private and public cloud. The need for a single pane of glass that provides continuous visibility into costs, and while assuring performance will be mandatory to assure effective capacity management and optimization across hybrid cloud.

Bearing in mind the considerable transformation taking place in IT environments, the critical issue in 2020 for CIOs, will be seeking to understand how to govern IT spend.

2. Hybrid and multi-cloud will demand consistent levels of visibility to mitigate risk

The start of 2019 was all about multi-cloud. Now, we are hearing people talk hybrid applications running in the hybrid cloud because it crosses multiple clouds, in addition to the data centre. So it is pretty clear now that hybrid is a term used to describe an environment with components of applications and the infrastructure supporting it that spread across private and public cloud. There is debate over whether it includes one or more cloud providers, and one or more private data centres, but in either scenario, anything that resides over multiple clouds still needs to be properly controlled. The lack of common metrics and granularity has contributed to performance and reliability problems across hybrid infrastructures. Everyone is looking for the single pane of glass, and this will be an on-going topic for

at least the next three to five years. The hybrid environment is definitely where everyone will be pivoted towards from 2020 and onwards.

3. The impact of containers on shared infrastructure will start to be understood

There is no doubt that Kubernetes has become the dominant orchestration platform for the cloud. As highlighted by IBM's acquisition of Red Hat, and VMware's acquisition of Pivotal, which positioned VMware to deliver enterprise-grade Kubernetes-based portfolio.

In observing the maturation of the market overall during 2019, we saw that infrastructure teams had learned from the virtualisation experience. With numerous virtual machines running on top of hypervisor, it was very difficult to see which VMs were responsible for the workloads being generated against the back end infrastructure, a scenario known as the 'blender effect'.

Going into 2020, this challenge lies at the heart of infrastructure teams' concerns around container adoption. Although the expectation is that a large percentage of on-premise containers will be long-lived, the ability to dynamically scale up and scale down could dramatically impact infrastructure.

Organisations looking to deploy containers in their data centres must consider and manage the impact and performance of their infrastructure, otherwise the move to containerised applications and deployment will likely fail.

In 2020 IT teams will adopt smart (AI-powered) hybrid infrastructure management platforms that will not only help them understand what applications are generating the workloads, but also how workloads impact every physical element of the infrastructure.

Weebit (Coby Hanoch, CEO)

1. We'll start hearing of more ReRAM commercial deals
2. China will continue its very strong push into memory and storage with big investments and potentially also acquisitions
3. Neuromorphic computing research will move to center stage

WekaIO (Andy Watson, CTO)

1. Smart NIC's (Network Interface Cards)

from companies like Mellanox, with comprehensively extensive capabilities for offloading will bring about revolutionary changes in the infrastructure landscape for the most data-hungry applications. Storage solutions able to leverage this important new networking technology will push the envelope in terms of both performance (i.e., greater throughput and lower latency data access) and improved scalability. Initially the higher cost will likely limit adoption to only the most demanding environments but over time (2021 and beyond) we can expect to see Smart NICs deployed more widely.

2. SCM (Storage Class Memory)

will have a big impact now that Intel has at long last delivered to market Optane DC Persistent Memory. Sub-microsecond latency for persistent storage at meaningful scale (up to 768 TB in a clustered filesystem is already being promised by at least one software provider) will revolutionize the industry. If nothing else, this new "Tier Zero" will enable near-instantaneous checkpointing (and recovery from those checkpoints). SCM can also be used to extend memory complements, and this is already impacting how engineers are thinking about in-memory databases and other memory-intensive applications. Anyone skeptical about this coming tectonic shift need look no farther than MemVerge or Levvyx (or other emerging software companies in this hot new space) who are already offering early access to their software which allows applications to benefit from SCM without any modifications. And as "Tier Zero" is redefined, the adjacent Tier One will also be impacted by intensified performance expectations; at WekaIO our flash-native filesystem is ready, willing and able to assist our customers' exploration of SCM.

- 3.** My third and final prediction here is a familiar one from last year: flash storage pricing will continue to fall faster than HDD storage pricing. This time I'll go further and predict that by Q4 2020 there will be large-capacity SSD's (probably QLC-based) which are irrefutably less expensive per-GB than meaningful HDD alternatives. Considering that SSD's already incur lower OpEx (consuming less power, requiring less A/C, failing less often, etc), 2020 may be the year we later collectively remember as The Crossover.

Western Digital (Phil Bullinger, SVP and GM)

1. In 2020, new data center architectures will emerge to manage the growing volume and variety of data

In the Zettabyte-scale Age, data infrastructure needs to be re-architected to address the growing scale and complexity of workloads, applications and AI/IoT datasets. These constructs will involve multiple tiers of workload-optimized storage as well as new approaches to system software. Zoned Storage, an open-source initiative, will help enable customers to take advantage of zone block management across both SMR HDDs and ZNS SSDs for sequentially-written, read-centric workloads. In 2020,

we'll see a substantial amount of application and storage software investment in Zoned Storage to help drive more efficient storage tiers as data centers are redefined in the Zettabyte-scale Era.

2. In 2020, tiering of data leveraging device, media and fabric innovation, will expand not contract

There will continue to be strong exabyte growth in read-centric applications in the data center, from artificial intelligence, machine learning, and big data analytics to a variety of business intelligence and accessible archive workloads. These at-scale use cases are driving a diverse set of performance, capacity and cost-efficiency demands on storage tiers, as enterprises deliver increasingly differentiated services on their data infrastructure. To meet these demands, data center architecture will continue advancing toward a model where data storage solutions will be consistently provisioned and accessed over fabrics, with the underlying storage platforms and devices delivering to a variety of SLAs, aligned with specific application needs. And while we certainly expect to expand the deployment of TLC and QLC Flash in these at-scale, high-growth workloads for higher performance use cases, the relentless demand for exabytes of cost-effective, scalable storage will continue to drive strong growth in capacity enterprise HDD.

3. In 2020, fabrics and composable will form a symbiotic relationship

Ethernet fabrics are becoming the "Universal Backplane" of the data center, unifying how storage is shared, composed and managed at scale to meet the demands of increasingly varied applications and workloads. In 2020, we'll see increasing adoption of composable, disaggregated storage solutions that efficiently scale over Ethernet fabrics and deliver the full performance potential of NVMe devices to diverse data center applications. Composable storage will significantly increase the agility and flexibility in how enterprises provision and optimize their data infrastructure to meet dynamic application requirements.

Zadara (Jeff Francis, Sr. Enterprise Solutions Architect)

1. XaaS (everything-as-a-service) will accelerate, taking market share from ownership models

Many 3-5 year IT refresh cycles are coming due at a time when multiple options for mature/stable XaaS options exist, plus the rate of technology change is accelerating. It will make less sense – both financially and technologically – to commit to a rigid platform for a span of multiple years.

2. BaaS and DRaaS (Backup and DR-as-a-service)

will grow in adoption, as more customers start to implement their first offsite and truly on-demand Backup and DR capabilities. BaaS and DRaaS bring reliability, convenience and affordability in an on-demand fashion. The combination of backup suite subscriptions and storage-as-a-service makes it possible for enterprises to meet their business continuity and compliance needs in a rapid, reliable and hassle-free fashion. Additionally, the economics for DR are radically improved, as costs for

compute and additional storage only apply for days of “test” and “actual DR” vs. buying and managing that infrastructure all year.

3. Datastores will increase in size

partially due to the increasing use of AI/ML models with large pools of unstructured data, and partially due to the rapidly declining cost for storage. Managed storage offerings (storage-as-a-service) both in the public cloud and on-premises (private clouds) make it feasible and affordable to implement extremely large datastores (from hundreds of petabytes to exabytes) and extract value, insights and monetary value from the larger data sets.

Zerto (Gijsbert Janssen van Doorn, Technology Evangelist)

- 1. Growth in adoption of hybrid and multi-cloud solutions**
- 2. Customers increasingly utilising management systems for disaster recovery and backup as a service**
- 3. Complete visibility over all workloads, data and costs becomes vital as the need for flexibility rises**

One of the things I believe we’ll see in 2020 is the true adoption of things like hybrid and multi-cloud solutions – but the difference in this upcoming year will be that customers and organisations will really begin to look for more management layers on top of their solutions. A lot of companies already have things like backup-in-the-cloud and DRaaS somewhere else, so what they’re now looking for is a uniform management layer on top of that to give visibility on cost, as well as a knowledge of where all data is located. It’s important to know where data lives, whether workloads are protected, and whether they need to move workloads between different clouds if and when requirements change.