A snapshot of market predictions from 37 vendors

January 2019

Companies:

Aptare, Caringo, Cloudian, Cohesity, Coraid, DataCore, Datrium, DDN Storage, DriveScale, Druva, E8 Storage, Hammerspace, Hedvig, HYCU, Igneous, Infinidat, Kaminario, Kaseya, Komprise, Minio, NetApp, Nexenta, Odaseva, Panasas, Portworx, Quantum, Qumulo, RStor, Rubrik, Scale Computing, StorCentric, StorONE, StrongBox Data, Vexata, WekaIO, Western Digital and Zerto

StorageNewsletter.com

Table of contents

Aptare
Rick Clark, CEO

"Obvious" cost savings with cloud or on-prem storage will be scrutinized: IT Administrators and Business Unit Managers will have to provide real data showing the impact of their efforts to move to the cloud and optimize on-prem data, requiring wide-and-deep metrics and more thorough reporting.

RegTech will drive accountability and retention policies: visibility into data protection and backups will need to be increased to meet governance demands. Business Units and Legal Counsel will require more frequent statements of compliance.

More IT organizations will see their storage infrastructure and data protection fracture. As corporations bring in additional products that are best-of-breed technologies, complexity will increase and comprehensive reporting will become even harder.

Caringo
Tony Barbagallo, CEO

As egress costs of enterprise data residing in the cloud are realized, organizations will migrate more of their data back into their on-prem datacenters.

Increase in file sizes, especially digital video driven by 4K and 8K workflows, will continue to outpace the adoption of faster bandwidth connectivity leading organizations to streamline collaboration by bringing data in house.

RESTful APIs (S3) will become more prevalent to all back end storage infrastructures.

Cloudian
Jon Toor, CMO

There's No Place Like Home: Cloud Repatriation Increases
While the growth of the public cloud will remain strong, enterprises will expand their adoption of on-premises private clouds in a hybrid cloud model. This will include repatriating data from the public cloud to avoid the bandwidth, latency and cost issues that can arise when accessing such data.

AI and Object Storage Play Tag
As businesses increase their use of AI to extract greater value from their digital assets, metadata tagging will become an even more critical element of enterprise storage. This will bring more attention to object storage, which is centered on metadata, and the key will be integrating well with AI tools.

What Do You Get When You Mix Blue and Red?
IBM-Red Hat Deal Scrambles the Cloud Landscape – IBM's acquisition of Red Hat will reverberate throughout 2019, giving enterprises more options for designing a multi-cloud strategy and highlighting the importance of data management tools that can work across public cloud, private cloud and traditional on-premises environments.

Cohesity
Rawlinson Rivera, Global Field CTO

New infrastructure models will empower developers to deliver better applications faster, with fewer headaches
In 2019, the application development lifecycle will be revolutionized by data platforms that give developers on-demand access to data and test/dev environments instead of requiring them to ask -- and wait -- for IT leaders to provide a copy of all relevant data sets. By removing these restrictions and delays

in the test/dev process, these data platforms will empower developers to build faster and experiment with broader data sets in order to create applications that provide next-generation customer experiences.

Coraid

Brantley Coile, CEO

Storage companies will continue to either be sold off, or shutdown during 2019, continuing the trend started a couple of years ago.

Users will continue to discover that the smaller, less expensive players are a great deal, because they are more stable as a business and more innovative than the big box tech company hawking acquired second hand technology.

Since everything is software and Moore's Law is dead, there will be a growing focus on software craftsmanship to achieve speed, efficiency, and lower cost.

DataCore

Gerardo A. Dada, CMO

Software-Defined Storage Will Become Increasingly Important in the Modern Data Center
The first step in modernizing the data center is to break the silos, achieve vendor independence, and remove vendor-imposed refresh cycles. This is why software-defined is quickly becoming the foundation for the modern data center. Software-defined storage is very flexible, and enables new storage and technologies to be added non-disruptively. Furthermore, the modern data center will be required to incorporate storage technologies that support synchronous mirroring in local and metro clusters, asynchronous replication for disaster recovery, and continuous data protection, which is like a time machine to undo any damage from ransomware attacks. As IT departments look to reap the benefits of the software-defined datacenter, the advantages of software-defined storage will be quickly realized in terms of performance, uptime and flexibility.

Hyperconverged will Become Hybrid-Converged
Due to challenges including the lack of integration with existing systems and the inability to scale compute and storage independently, many of the traditional hardware vendors in the hyperconverged infrastructure (HCI) market are shifting to a software offering. This has led to an overall evolution from the earlier vision of "hyperconverged," primarily consisting of the convergence of compute, storage and network in one single hardware unit, into a model that is software-driven and software-defined, which is called "hybrid-converged." Hybrid-converged infrastructure provides the same advantages of HCI with additional functionality that allows it to connect to external hosts and to present external storage to the unit. As a result, users no longer have to choose whether to buy into the HCI model or not—they can have the benefits without the limitations.

New Storage Challenges Will Emerge as Container Adoption Continues to Grow
As container technology matures, there are new challenges emerging, primarily in the areas of security and storage. As deployments move to production, IT organizations will require an ability to deliver the same data storage services that are currently provided to monolithic application architectures. The solution has to be capable of providing shared storage to existing virtualized and bare-metal application infrastructures, as well as allow DevOps engineers to consume storage on-demand, ensure stateful application data is persistent. Software-defined storage will enable administrators to present persistent storage to container hosts deployed as VMs on virtual hosts, with the ability to provide persistent storage to container hosts deployed on bare-metal as a next step. As a result, users will be able to manage the provisioning of storage to container deployments, with the same platform as the rest of the application workloads, and provide the same level of enterprise storage services required for all critical production environments.

Datrium
Brian Biles, Chief Product Officer

Enterprise infrastructure for Kubernetes -- not just a stateless toy anymore as enterprises consider mission critical apps on it. Existing storage systems are not built to deal with this. New storage systems that can do primary, backup and DR for Kubernetes will rise in 2019.

More convergence of data services into one platform with respect to combining primary, backup with replication to cloud, and DR runbook automation. Because that is the only way to implement a simple and viable Hybrid Cloud. Any other way is impractical. You can do an expensive product without the efficiencies, but that product will need other products alongside. Datrium has all 3 dimensions, in collaboration with standard VMware and AWS. Others will follow down this path.

DDN Storage
Kurt Kuckein, Sr. Director of Marketing

2019 looks to be an exciting year. With the emergence of large-scale, AI and machine learning environments, the advancement of granular data management capabilities, the move toward cloud-like data management models and autonomous storage as well as the accelerated adoption of at-scale Flash deployments and NVMe the existing market is poised to experience an evolution that will continue to push the boundaries. As the team at DDN looks to 2019 we are focused on the trends that will successfully and efficiently handle the various scenarios presented. Here are our top three predictions for 2019:

The emergence of large-scale AI and Machine Learning deployments
The past couple of years have seen many organizations trialing machine learning algorithms on small data sets, implementing small but growing deployments, and planning at-scale infrastructures. DDN has already seen successful projects emerge for predictive analytics for chronic disease management, workflow enhancement in radiology as well as administrative and financial use cases that bring operational efficiency to these industries. 2019 will be the year that large-scale AI and machine learning environments emerge in mass, with organizations moving from deployments of 4, 8 or 16 GPUs to deployments that range from hundreds to thousands of GPUs.

The advancement of granular data management capabilities for at-scale data systems and private clouds
At-scale data systems and private clouds are increasingly supporting diverse types of data, such as AI and deep learning workflows, that require advanced, granular data management capabilities. These data management solutions will need to deliver simplicity, allow for more sophisticated tagging and searching of the data itself, and provide insight into the types of data that customers have within their systems and within their clouds.

Accelerated adoption of at-scale Flash deployments and the ascendency of NVMe
The adoption of scale-out storage architectures to manage flash on demand and to scale performance as needed will accelerate as flash usage provides an optimal means for balancing performance and cost in long-term storage, and as the price of flash storage will drop significantly in the first six months of 2019. NVMe will be the default media for tier-1 applications (low latency, high IOPs and density – what's not to like?), but NVMeOF will continue to lag as a networking standard as other more established RDMA networks like InfiniBand and RoCE continue to thrive and meet performance demands.

DriveScale
Brian Pawlowski, CTO

HCI use cases converge with a focus on small scale VM cluster deployments
HCI is best suited for pilot and small-scale big data app deployments, while large-scale deployments continue to sprawl on bare metal with platforms like Kubernetes to manage at scale.

Reality around SDS comes home to roost
Customers begin to understand costs and advantages of 'appliances' which are simple to deploy versus 'software' which requires more in-house infrastructure and integration efforts.

Big data keeps getting bigger
Not just in size, but in the complexity of sources and applications that are used to process big data. Pain points start to come into focus. Hadoop deployment growth slowed, and Cloudera and Hortonworks merged. This is not a negative on Big Data use cases, but an indication of increased diversity of apps used to process Big Data.

Druva
Prem Ananthakrishnan, VP of Products

The Rise of Smart Clouds
The adoption of streaming data capture from IoT and sensors, data governance policies, security standards, expanded data curation and compilation and widespread adoption of AI and machine learning have made it impossible to rely completely on on-premises solutions. Technologies such as AI, machine learning, and analytics thrive in environments with expansive amounts of data and compute abilities beyond those available in on-premises solutions. These trends greatly favor cloud-based architectures, and will only increase as vendors offer more advanced solutions.

The Cloud Wars Will Escalate in 2019
Serverless architecture will drive down costs even further, and I would expect hybrid to become more popular with pushes from VMware and AWS. Online marketplaces will shift spending from offline distribution and vendors, and resellers will increasingly adopt adopt digital VAR-like models. Machine learning and AI will continue to rise in adoption and increase the allure of cloud computing. Because of these technologies, public cloud will become the de-facto choice for developers.

2019 is the Year of Government Data Compliance
Data management is no longer simply a consumer vs corporation battle, it has quickly elevated to the country and federal level. In the wake of GDPR, others are using it as blueprint to enact more stringent compliance standards . The California Consumer Privacy Act goes into effect January 2020 and we should expect to see more of the same in the coming years. Such regulations mean company obligations will become more complicated and will need to meet new standards. Having the flexibility and scalability to store data within specific regions will become a key buying consideration and increasingly favor cloud deployments over on-premises solutions.

"Data Protection" is Giving Way to Data Management
The ability to protect data and restore backed up files is no longer sufficient for modern business. Data has become the fuel for company success, driving insights, customer targeting and business planning and even training AI and machine learning models. Any way to extract additional value from it is critical to business success and the shift to data management is key, where data is not only protected, but properly archived, easily searchable, can be leveraged for analytics, and is compliant the entire time.

E8 Storage
Zivan Ori, CEO

2019 is going to be a year of growth and broader adoption of new technologies that have been introduced in recent years.

NVMe drives will accelerate their gain in market share as a SATA replacement in servers, and begin taking SAS share in storage systems. This continues the trends begun in 2017 / 2018, but accelerates into 2019.

NVMe-oF storage arrays will move from testing into production across many customer deployments as more established vendors continue to introduce NVMe based storage systems.

ARM server products will gain traction, with ARM based NICs becoming generally available, and ARM based storage controllers being introduced.

Hammerspace
David Flynn, CEO

With all of the major cloud providers now offering some form of Kubernetes-as-a-Service, will see the industry accelerating the adoption of true multi-cloud environments.

Metadata is the key to hybrid multi-cloud success. More organizations will recognize the need for sophisticated metadata management to cope with issues like data governance and compliance, data accessibility, and data pipeline orchestration.

Finally, in 2019 will see more data operations teams take control of enterprise data, as managing data by copying it between data silos cannot sustain the agility demanded by digital transformation. Kubernetes, metadata management, and data virtualization are all hot trends that will enable this shift in responsibilities.

Hedvig
Gaurav Yadav, Founding Engineer

Now that containers are production ready, think about storage
Containers are finally becoming production ready. As the next big thing in virtualization and resource utilization, they are now being used for mission-critical applications. We'll see a major increase in production deployment containers in 2019, which raise the profile of associated challenges like storage. Containers require persistent storage in order to succeed in mainstream use cases. Most current storage solutions cater to virtualization vendor platforms (VMware, etc.) but aren't a great fit for container technology. In 2019, the problem to solve for is finding one storage solution to support these completely independent virtualization mechanisms so that they can co-exist efficiently. From the application point of view, they are different mechanisms, but that shouldn't matter when it comes to storage. Making progress on this aspect of container deployment will bring much-needed simplification to this infrastructure option.

Get your data governance and security ready for GDPR and beyond
The European Union's sweeping consumer data privacy legislation went into effect in May 2018 and started an important dialogue around the world. Combined with major breaches and sophisticated identity and data manipulation schemes at Facebook, Google, Twitter, and many other major digital players, the global focus on data privacy, security, and integrity is sure to catalyze further regulations. Additional countries and states are already considering tougher privacy mandates and penalties. As European regulators start levying GDPR penalties, companies are finally implementing stricter data protection guidelines. Storage providers will need to have the right answers about data security guarantees and the ability to offer support for best practices. The guarantees include: if the data be accessed, stolen, or leaked, no one will be able to use it or make sense of it; data integrity must be preserved without fail; and data must abide by country-specific data location restrictions. Certain kinds of

data cannot be transported across geographical boundaries (for example, Chinese Internet regulations), so all public cloud providers and cloud-like services will have to guarantee location.

Software-defined storage is the foundation for trust, growth, and innovation
Storage designed with distributed systems in mind makes it possible to quickly provision application specific, policy-based data services. Trust (or lack of it) is becoming a major factor in enterprise digital transformation and brand reputation management. In 2019, infrastructure teams and leadership should carefully consider the significant role storage plays in maintaining public trust and the integrity of data systems used in commercial, industrial, and social settings. Every enterprise has its own IT team, and each team has a unique collection of challenges. These days, making sure distributed applications and business processes run as smoothly as possible is job number one. With modern cloud strategies and software-defined storage solutions in place, these teams can focus on optimizing applications and addressing emerging challenges - there will be plenty of those in the year ahead.

HYCU

Simon Taylor, CEO

Multi-Cloud becomes mainstream
Clouds have been in existence for many years now. From an infrastructure perspective, as many as 90% of customers have some cloud usage going on in their infrastructure today. What we are starting to observe is that many of them are mature and their usage has gotten to the point where they are talking about a multi-vendor strategy for their clouds. This in our opinion is a turning point in adoption of cloud. If not for anything, just to keep things competitive for cloud vendors. Because of this, more customers will need solutions that work across multiple clouds. We will see most of the vendors who have “marketed” their solutions as “multi-cloud” finally will be forced to either deliver or shrink in 2019.

Consolidation around Platforms
It is no surprise that companies are looking to simplify their IT investments and take advantage of advances in reducing complexity. When you look at next-generation platforms in the industry today, the ones that come to mind are AWS, Azure and GCP in the Public Cloud space and Nutanix in the Enterprise Cloud space. While VMware is a great technology, their platform play seems to be around AWS. We would consider a solution a platform when it can handle all of the enterprise workloads. In the enterprise. We are seeing customers trying to consolidate on the platforms they choose and are eliminating siloed infrastructures. One such silo that we see customers trying to eliminate is the separate secondary storage silo. In 2019, we expect a significant number of customers adopting their primary platform also for their secondary workload to significantly reduce their complexity and also increase their operational efficiency.

Adoption of as-a-Service Infrastructure offerings
Infrastructure-as-a-Service has been in use for many years now. For most people in the infrastructure space, they think in terms of instantly provisioning VMs, Storage and Networking as IaaS. While that is good, the bigger value will be realized when customers use higher level offerings like Database as a Service, Backup and Recovery as a Service, Data Warehouse as a Service, and Streaming Big Data Services. When utilized effectively, these would lead to tremendous savings and agility for the business. In 2019, we see customers not just doing a lift and shift to the cloud, but truly using the cloud for burst workloads and also utilizing these higher level services.

Igneous
Kiran Bhageshpur, CEO

Rise of the Data Engineer
Many different roles in various data-centric companies depend so heavily on data and spend so much time in wrangling data from various sources and into the right platform that we will see the rise a a dedicated job function whose job it is to wrangle data.

Unstructured Data continues to grow at geometrical rates
There will be a broader recognition of the number of businesses that are impacted by and dependent on machine generated data. As machines grow in numbers and the fidelity of data from machines improves, the volume of data being generated and stored by various business will weight down on already overburdened IT teams.

Real attempts to regulate "data companies"
Congress in the US will make first "real' attempts around data / privacy issues. Whether it becomes law is an all together different matter.

Expansion of "ML-as-a-service" market place
More companies will arise and offer services with the offer of "Bring me your data and I will provide you insights" along specific verticals.

VR / AR will remain a technology we are waiting for
Magic Leap wherefore art thou?

Infinidat
Brian Carmody, CTO

Raw storage shipments will exceed 700 Exabytes, with over 90% being spinning disk.

Application-level encryption will become pervasive in every industry, driven by aggressive enforcement of GDPR and data breach liabilities, rendering storage array compression and de-duplication ineffective.

The first MAMR (microwave assisted magnetic recording) enterprise storage system is announced.

Kaminario
Josh Epstein, CMO

Software Composable Infrastructure solutions will start to become mainstream options considered by datacenter operators.

Consumption-based pricing will begin to be a standard offering for infrastructure solutions.

The IT channel will continue to transform into service-oriented entities – moving away from traditional VAR models to MSPs.

Kaseya (and Spanning)

Mike Puglia, CPO
Brian Rutledge, Principal Security Manager, Spanning
Matthew McDermott, Principal Technical Marketing Engineer, Spanning

Microsoft Cloud Gains with MSPs
Over the last three years Azure growth among MSPs has skyrocketed, from 47 percent to 67 percent. This will grow even more next year and can be attributed to two major factors: relationships and product familiarity. Some MSPs have been working with Microsoft for 15 years, and have their sales rep on speed dial. Starting over with AWS or Google Cloud Platform is a daunting endeavor. Also Microsoft is making it easy to move on-premises applications to Azure. Microsoft is not just building the infrastructure for MSPs to use, but they are also building the apps and services on top of that, making the transition to the cloud that much easier. Microsoft owns the MSP market and growth will continue in 2019.

Organizations Double-Down on Security Training
In a recent Spanning-sponsored survey of U.S. workers, almost half of all respondents said they would allow a colleague to use their work computer to complete a task. While letting a work friend use your computer might not seem like a risky move, research has found that insider threats account for most security breach incidents and sharing devices — especially for admin holders — is one way that being polite could put enterprise data at risk. In the new year, businesses should look to double-down on security training — not just with phishing tests and lectures — but also focus more security awareness training on accidental lapses in security, because not all security incidents are from malicious actors. IT leaders should emphasize that employees, should never use shared passwords, enable SIEM (Security information and event management) solutions to centrally log actions by administrators and other users with the keys to the kingdom, periodically audit internal controls and implement tools, where possible and financially feasible, for Privileged Access Management (PAM) instead of using administrative accounts directly.

GDPR Makes Personalization of Workflow Easier
Personalization of workflow will become increasingly important in 2019, especially in light of GDPR. This might be perceived as challenging with stipulations such as the right to be forgotten, but GDPR may actually make it easier for organizations to achieve. Because software manufacturers are compelled to identify and implement the components necessary to track personal information, we can take advantage of this data to create richer personalized interfaces. Personalization allows for fewer interruptions of superlative information and increased access to information that will help teams work smarter and faster in 2019.

Komprise

Krishna Subramanian, COO

Improving intelligence
With adaptive automation and machine learning, data management software will start to perform in smarter ways by observing and leveraging patterns. In 2019, AI based data management will start to be able to think outside the box, offering more intelligent ways to manage business needs.

A retraction from rule-based policy management
Rule based policies can be time consuming and inefficient. IT admins have to manually dictate rules, and think about every possible scenario in order to program a rule for it. This has long been the approach for data management, but in 2019 we will start to see a retraction as AI allows us to utilize intelligent software and set goal-based policies.

Improved search and discovery
AI can improve ability to search and discover data pulling from a wider pool. Many organizations today store data both in the cloud and on-premises, but with AI, intelligent software can pull information out of a wider data lake, no matter where it is stored.

Minio
Anand Babu Periasamy, CEO

Object Storage set to replace Hadoop HDFS
In-memory data processing platforms like Spark and Presto continue to shift away from Hadoop HDFS to S3 compatible object storage. Disaggregating storage and compute enables stateless compute nodes to be containerized and managed by Kubernetes. This innovation came from Amazon AWS and public cloud infrastructures and the rest of the world will adopt it.

Enterprises adopt high-performance object storage for AI workloads
Modern businesses need to become data-driven. Traditional data warehouse and big data solutions are inadequate when the volume of data is huge and constantly evolving. This is where machine learning and deep learning algorithms have made significant progress in the last few years to interactively analyze and predict patterns with the help of GPUs, high-performance storage and networking. 2019 will be the year of high-performance object storage and 100GbE+ network.

NVMeoF saga continues into 2019
NVMeoF RDMA is still a long way from getting enterprise adoption. NVMeoF TCP shows promise but its benefits are still unclear when compared to FCP and iSCSI. Independent of the fabric protocols, NVMe block storage systems still need to address SAN features like volume management, snapshots, erasure code or replication. Nevertheless, NVMeoF is a step toward eliminating legacy baggage.

NetApp
Atish Gude, Chief Strategy Officer

Still at an early stage of development, AI technologies will process massive amounts of data, the majority of which will happen in public clouds
A rapidly growing body of AI software and service tools – mostly in the cloud – will make AI development easier and easier. This will enable AI applications to deliver high performance and scalability, both on and off premises, and support multiple data access protocols and varied new data formats. Accordingly, the infrastructure supporting AI workloads will be also have to be fast, resilient, and automated. While AI will certainly become the next battleground for infrastructure vendors, most development will start in the cloud.

The demand for highly simplified IT services will drive continued abstraction of IT resources and the commoditization of data services
Remember when car ads began boasting that your first tune up would be at 100,000 miles? (Well, it eventually became sort of true.) Point is, hardly anyone's spending weekends changing their own oil or spark plugs or adjusting their own timing belts anymore. You turn on the car, it runs. You don't have to think about it until you get a message saying something needs attention. Pretty simple. The same expectations are developing for IT infrastructure, starting with storage and data management: developers don't want to think about it, they just want it to work. "Automagically," please. Especially with containerization and "server-less" technologies, the continuous trend toward abstraction of individual systems and services will drive IT architects to design for data and data processing and to build hybrid, multi-cloud data fabrics rather than data centers. With the application of predictive technologies and diagnostics, decision makers will rely more and more on extremely robust yet "invisible" data services that deliver data when and where it's needed, wherever it lives. These capabilities will also automate the brokerage of infrastructure services as dynamic commodities and the shuttling of containers and workloads to and from the most efficient service provider solutions for the job.

Multi-cloud will be the default IT architecture for most larger organization while others will choose the simplicity and consistency of a single cloud provider
Containers have the potential to disrupt the cloud business model and break vendor lock-in by making development environments highly portable. This will make it easier and easier to move the work to where data is being generated rather than what has traditionally been the other way around. Data is far less portable than compute and application resources and that affects the portability of runtime environments. Even if you solve for data gravity, data consistency, data protection, data security and so on, you can still face the problem of platform lock-in and cloud-specific services that you're writing against, which are not

portable. As a result, smaller organizations will either develop in-house capabilities as an alternative to cloud service providers, or they'll choose the simplicity, optimization and hands-off management that come from buying into a single cloud provider. And you can count on service providers to develop new differentiators and encourage lock-in. Larger organizations will demand the flexibility, neutrality and cost-effectiveness of being able to move applications between clouds. They'll leverage containers and data fabrics to break lock-in, to ensure portability of applications and workloads, and to control their own destiny. Whatever their path, organizations will need to develop policies and practices to get the most out of their choice.

Nexenta

Thomas Cornely, CPO

As service providers worldwide gear up for 5G offerings in 2019, massive software-defined infrastructure opportunities will arise for full stack 5G ready Telco Clouds including the need for full-featured software-defined storage solutions.

Enterprise File service needs will continue to grow at exponential rates, with the ever-increasing use of hyper-converged infrastructures in the data center driving enterprise-class file service offerings that can augment HCI and traditional three-tier architectures.

Adoption of public cloud services will strengthen as an option to supplement on-prem offerings as well as new deployments, which will be increasingly useful with hybrid-cloud and multi-cloud deployments requiring a standard storage solution best deployed in a software model.

Odaseva

Sovan Bin, CEO

Data Rights

More and more individuals in Europe will exercise their data rights to protect their data, such as their right to access personal data or the right to be forgotten.

Data Minimization

Data minimization in accordance with data privacy laws, like GDPR will be become key to successful Enterprises. Abiding by a multitude of different legal retention periods will become a critical challenge.

Data Protection

Enterprises will increasingly look to data encryption and multi-factor authentication as one way to protect their data; and they will invest in better control of the location of their data processing and data storage to facilitate compliance with local regulations.

Panasas

Curtis Anderson, Senior Architect

HPC will be at the heart of AI and machine learning

"HPC" can be thought of as a set of best practices for how to "scale out": how to architect applications, interconnect, and storage that can scale well beyond what a single machine can process. Many AI projects have already surpassed what a single machine can process, and as a result are adopting HPC techniques. Given the growth in data size and increasing complexity of neural networks, that trend of adopting (and extending) HPC's architectural lessons will continue to accelerate. The trajectory of HPC storage architectures will not be greatly affected, however. The amount of AI training data being processed, and the time before NN refinement forces a new training cycle, makes moving the data between different storage systems impractical; the GPUs must be kept fully utilized, so the typical access latencies of Object Stores are unable to keep up; and at these capacities, all-flash solutions are not cost-effective. Performance and capacity will need to go up, but whole new architectures are not required.

The integrity of open source parallel filesystems will remain in the spotlight
There will be continued skepticism concerning the future of open source parallel filesystems due to deployment complexity, overall maturity, and even ownership concerns. The vendor-neutrality of the most mature of the set, Lustre, has been put into question and with it concerns over increases in the cost for licensing and support. As parallel filesystems become more important to Enterprises looking to deploy HPC systems for AI and other advantages, they will be looking for turn-key solutions rather than science projects.

Portworx

Murli Thirumale, CEO

Containerized Applications will pull data management away from storage arrays
Driven by DevOps, containers drive app portability and microservices. Dynamic container placement means data movement and data placement needs to automated and tied to containers not to storage arrays. Just like VMWare virtualized and allowed compute capacity to be managed away from servers, data will be managed by new container centric storage software.

Hybrid cloud deployments will break storage vendor silos
As enterprises deploy apps across on-prem and public clouds or between two public clouds, they cannot depend on meeting data SLAs via just their storage vendor. Automation and data SLAs will be managed at a layer above the storage vendor.

AI and automation will change the economics of IT
Much of DevOps is still driven by people, even if the infrastructure itself is becoming programmable. But data volumes are growing so fast and applications evolving so quickly, the infrastructure must be nimble enough that it doesn't become the bottleneck. In 2019, infrastructure will become increasingly programmable, and AI-based machines will predict storage and compute needs and allocate resources automatically based on network conditions, workloads and historical patterns.

70 percent of customers will opt for the Kubernetes platform from their cloud provider, OpenShift or Tectonic
But consolidation is coming. Most customers will use the Kubernetes distribution and packed services from their cloud providers or a distribution of Kubernetes from RedHat (Openshift) or CoreOS (Tectonic). Smaller customers will probably opt for the fully packaged offering from their cloud provider, even though they will be locked in and find it difficult to implement multi-cloud strategies. Larger enterprises will more often opt for a cloud-agnostic platform not only because such platforms allow for more customization but also because they are less likely to be locked into their cloud providers, something that is helpful when it is time to negotiate price.

Quantum

Eric Bassier, Sr. Director of Product Marketing and Corporate Communications

Video content will grow exponentially, across many industries
Video and video-like data will constitute over 50 percent of all data being generated, from surveillance footage; consumer images, voice, medical imagery, IoT, entertainment and social media. Large and unstructured data is often 50 times larger than the average corporate database. Video is unique, and it is not typically a good fit for traditional backup; it cannot be compressed or deduplicated, it doesn't work well with replication, snaps or clones, and ingest speed is critical. This "unstructured" data – predominantly video and images – is projected to surpass 100 Zetabytes worldwide by 2020. Expect enterprise data services to be increasingly optimized for these data sets, with infrastructure optimized for ingest processing and the full life cycle management of this form of data.

NVMe will revolutionize media workflow infrastructures
NVMe will finally allow our customers to unlock the true potential of flash – dramatically reducing latencies and enabling IP-based infrastructures and workflows, while reducing expensive fiber infrastructure costs. It also sets the stage for truly software defined infrastructures and can free up hardware resources to focus on advanced analytics.

Autonomous vehicle development will drive enormous data creation
As testing continues in 2019, cars are on the cusp of Level 3 autonomy, so we can anticipate new workflows for the test vehicles as they push forward from this point. Watch for other vehicles to become autonomous as well in land, sea, and air. Data storage in the form of flash, disk, tape and the cloud – and the ability to seamlessly move data between these types of storage to balance access and cost factors – will be key in supporting automotive vehicle development.

Qumulo
Bill Richter, CEO

File emerges as the gold standard for unstructured data storage workflows
Customers can take a breath of relief and embrace the things they love about file storage — standard protocols, robust data services, application integration, management and performance.

Software-defined storage goes mainstream
Customers will continue to run data centers for certain workloads, but those data centers need to be software-defined. That means standard hardware supplied by a choice of vendors, coupled with intelligent, API-driven software that can be automated and scaled economically. Software-defined infrastructure makes hardware innovation, such as NVMe, more easy to adopt and consume.

The fog of uncertainty has lifted and the "clouds" have parted
Customers will increasingly demand that their storage infrastructure run in both the public cloud and the data center. "Scale-across" supersedes scale-out as the architecture of choice for the hybrid cloud.

RStor
Giovanni Coglitore, CEO

2019 will be the year that computing comes to you. We predict that while data centers will persist, the trend of multi-cloud deployment spanning on-premise IT to public and private clouds will focus on giving CIOs the control and agnostic platforms they need to manage their hybrid cloud needs.

In that vein, architectures defined strictly as either "core" or "edge" will shift to those that are built on distributed computing, network, storage, and resources, which leads us to our second prediction: CIOs will look for tools and solutions that allow for data to be portable between environments, agnostic to any particular cloud provider and fundamentally more flexible.

This will translate into further momentum around containers and serverless technology. Already, we've seen demand for the integration of Kubernetes and Singularity, a container technology initially developed for high-performance computing, as companies run more service-based compute jobs that involve streaming data and real-time analytics.

Rubrik
Arvind Nithrakashyap, CTO

The future will be hybrid
When it comes to an on-premises or public cloud strategy, enterprises will not choose one or the other. As we've seen through the introduction of new solutions like AWS Outpost and Microsoft's Azure Stack, even the major cloud providers understand the power and promise of hybrid cloud. I expect the major cloud players will continue to introduce and advance hybrid solutions that aim to provide enterprises with a seamless experience across on-premises, public or multi-cloud environment.

Business will capitalize on AI and automation
Next year, you can expect that AI will become more integrated into supply chain business processes. As automation becomes more sophisticated and practical use cases become the norm, CTOs will need to

adjust their thinking and identify opportunities where they can leverage AI to work alongside human counterparts and augment human intelligence.

Cybersecurity will remain top of mind in boardrooms everywhere
In a world in which attackers are seeking out vulnerabilities 24/7, getting hit with a cyberattack is no longer a matter of if -- but when. CTOs will need to focus on implementing stronger security measures to protect emerging technologies, like cloud-native applications, that are beneficial to enterprises but also come with their own set of security risks.

Scale Computing
Alan Conboy, Office of the CTO

Artificial Intelligence and Machine Learning
In 2019 Artificial Intelligence (AI) and Machine Learning (ML) will nearly reach its full potential by connecting and processing data faster over a global distribution of edge computing platforms. AI and ML insights have always been available, but possibly leveraged a bit slower than needed over cloud platforms or traditional data centres. Now we can move the compute and storage capabilities closer to where data is retrieved and processed, enabling companies, organizations and government agencies to make wiser and faster decisions.

Cloud Computing
Next year will be a defining year for edge and hybrid computing strategies as IoT and the global network of sensors pile on more data than the average cloud has had to handle in the past. This transition will officially crown edge computing as the next big thing. According to a study from IDC, 45 percent of all data created by IoT devices will be stored, processed, analysed and acted upon close to or at the edge of a network by 2020. In the process, edge computing will take on workloads that struggle on hosted cloud environments, passing the torch over to HCI platforms.

Edge Computing
According to Statista the global IoT market will explode from $2.9 trillion in 2014 to $8.9 trillion in 2020. That means companies will be collecting data and insights from nearly everything we touch from the moment we wake up and likely even while we sleep. In 2019, edge computing will require a new level of intelligence and automation to make those platforms practical. Where once only a smidge of data was created and processed outside a traditional data center, we will soon be at a stage where nearly every piece of data will be generated far outside the data center. This amount of data will create a permanent home for edge computing.

StorCentric
Gary Watson, CTO

A retreat from the cloud
Cloud was once considered cheap and flexible but many organizations made the move without fully understanding it. As a result, we are starting to see a retraction as businesses look for a more hybrid model in order to reduce costs and regain control. This is something that will continue throughout 2019, businesses need to find a balance by understanding what data is stored and where. The lesson to be learned is to be aware of industry hype around the latest trends.

Consolidations are set to grow
The storage market has been quite noisy with industry leaders touting about the latest trends, but the overcrowded market is starting to make a number of consolidations. One thing to be aware of is the need for meaningful integration. As organizations get acquired, customers need to consider if it's for the right reason, how this will support them, and the type of integration to expect.

High capacity storage will still be a must
We have been talking about data growth for the past couple of years, but it's a major pain point that isn't going away. It is easier than ever before to create vast volumes of data, and to support this, we will increasingly see software and automation develop. In fact, in the past couple of months this has already

started to ramp up. As we store increasing volumes of data moving forward, it's important to compliment storage with data protection, ensuring you have the capacity needed, along with a comprehensive range of data protection features.

StorONE

Gal Naor, CEO

Organizations are tired of bolt-on storage solutions, adding hardware to achieve high performance or high capacity, and having to leverage yesterday's technologies to try to handle today's workloads. 2019 will be a year of innovation. As IT struggles to achieve the results they need in terms of performance, data protection and cost, they will seek out a unified approach using core technologies that can provide immediate ROI. In the upcoming year, we shall see enterprises adopting solutions that help them overcome the following challenges:

Getting more storage with less hardware
Achieving a higher utilization rate from storage hardware allows organizations to use dramatically less of it to achieve the same performance. Whether using AFAs to achieve ultra-high performance with very low latency or high throughput with HDDs, extracting the full value of the drives will result in less overall cost. We will see the emergence of software solutions for hardware challenges that will result in the lowest CAPEX and TCO with less hardware to buy and manage.

Assuring complete data retention and data protection
Enterprise-class functionality will no longer be a nice-to-have feature but something that will be built into next-generation software solutions. Rather than having to buy and integrate external solutions – again, less hardware! – Enterprises will be able to achieve complete data protection without performance degradation within a single solution that provides unlimited snapshots and high-performance erasure coding all included without paying more.

Future-proofing data storage
Getting ahead of the curve is often a pipe-dream in IT. Too much time and effort is spent putting out fires, having to plan for future capacity or performance needs. In 2019, more organizations will be looking for solutions that provide the flexibility they need to simplify their storage environment and allow them to make changes on the fly. A complete and central storage solution that allows for the mix and matching of any drive type (HDD, SSD) from any vendor, and any capacity; any physical hardware or virtual appliance. 2019 will be the first time software will enable block, file and object protocols to run on the same drives included on one platform.

In 2019, a new wave of innovation will emerge, providing organizations with performance and capacity with dramatically less hardware in a single solution while providing all enterprise-class functionality and protocols to satisfy the needs of their mission-critical applications today and into the future.

StrongBox Data

Floyd Christofferson, SVP Products

Cross-platform Storage Automation
Large-scale data users are creating increasing demand for solutions that simplify and automate global data and storage resource management across multiple storage types and vendors, with the objective of reducing the complexity of managing multi-vendor/cloud environments. These new solutions are combining multiple sources of file and storage metadata into powerful policy engines to automate storage workflows and maximize value from the storage and data itself. This is a quantum leap beyond simple HSMs and data movers of the past which had limited policy triggers, such as the age of the file. These new vendor-neutral solutions are driven by rich metadata to meet the demands of AI/ML, improve ROI on storage resources, and to reduce OPEX and complexity for IT, breaking storage vendor lock-in, and increasing storage utilization efficiency and user productivity.

Leveraging AIOps for Storage
AIOps (Artificial Intelligence for IT Operations) is an emerging category that uses Big Data and machine learning technologies based on aggregating multiple layers of inputs to automate global IT actions. This trend is now expanding beyond its traditional use cases for application performance and network monitoring to now include storage and data management tasks. AIOps for storage is made possible by aggregating multiple layers of metadata about files and storage platforms, the combination of which provides the context needed to automate management of both data and storage resources across any storage type, from any vendor. In this way AIOps for Storage provides IT teams with the ability to:

- Automate routine storage IT operations, including file copy management, data migration, data protection, tiering, active archiving, and more.
- Quickly recognize and take action on serious issues driven by real-time feedback from the data and the storage platforms.
- Implement global control across multiple locations and storage platforms.

AIOps for Storage will gain additional traction as IT teams leverage it to reduce storage costs and improve ROI while also reducing operational load on IT staff.

Multi-Cloud & On-Prem Storage
It is no longer enough to simply have a "Cloud" strategy for managing off-premises data. Increasingly storage IT managers realize they need a hybrid approach to data storage that will include a combination of on-premises and off-site data storage, but which does not lock them into a single Cloud storage provider. They need a way to simplify, accelerate, and automate management of data across both on-premises storage and multiple Cloud solutions. Data value changes over time, such that data that might need to be housed in an all flash array or NVMe platform today will need additional protection and lower cost storage options tomorrow. As storage choices proliferate for both on- and off-premises solutions, the need to centrally manage and automate data placement across a hybrid of multi-cloud and multiple on-premises storage types will increase in importance in 2019 and beyond.

Vexata
Zahid Hussain, CEO

Artificial Intelligence, Machine Learning and Cognitive Systems
will drive infrastructure decisions in 2018, requiring that Line of Business and IT management work together to address bottlenecks that exist in current solutions.

In 2018 there will be a continued shift of compute processing and analytics to the edge,
driving infrastructure architecture decisions to de-centralize data processing and positively impact digital transformation, customer engagement and business models.

IoT will continue to drive infrastructure architecture
decisions to deploy solutions that can offer abundant I/O performance at scale to drive down TCO, increase efficiency and deliver performance elasticity to support ever-changing workload demands.

WekaIO
Andy Watson, CTO

Led by QLC, the acquisition cost of flash storage will beat HDD.

Cloudbursting will be more broadly adopted by mainstream customers implementing multicloud strategies which facilitate the flow of data between on-prem datacenters and the cloud, or across multiple geos.

100-GbE and other high-perf networking (e.g., 200 gbit/s InfiniBand) will enjoy broad adoption, driving the last nails into the fading relevance of FC SAN. Applications (especially ones leveraging GPU's, each of which can drive IO at ≥ 6 gigabytes per sec) will no longer be bottlenecked by NFS, which was okay so long as networks were stuck at 10-GbE or less; in 2019 they will look beyond NFS to obtain greater benefits from faster storage over faster networks. (Hint: Call WekaIO!)

Western Digital
Martin Fink, CTO

With workloads constantly changing and processing demands in flux, in 2019, we will see a proliferation of RISC-V based silicon as there will be an increased demand from organizations who are looking to specifically tailor (and adapt) their IoT embedded devices to a specific workload, while reducing costs and security risks associated with silicon that is not open-source.

In 2019 we will realize the first step on a trajectory toward fabric infrastructure, including fabric attached memory, with the wide-spread adoption of fabric attached storage. This may seem like a small step, but the commitment to FAS means we are taking the necessary steps, as an industry, to ensure all components are connected with one another, allowing compute to move closer to where the data is stored rather than data being resigned to several steps away from compute.

Although "composability" is not a new term, in 2019, we will see open composability, versus today's inflexible, proprietary solutions, come of age and start to go mainstream as organizations look to build composable infrastructures on open standards to allow for specialized configurations that are specific to their workloads and address diverse data.

Zerto
Steve Blow, Technology Evangelist

Brexit
Brexit is going to have a significant impact on IT operations next year, and people will start to understand how the UK's decision to leave the EU will impact where they can run their IT infrastructure. Many organisations may find themselves with a new requirement to move things to, from or between different cloud and on-premises environments to ensure their data falls in the right regions for regulatory purposes. And a key to minimising the disruption this could bring is having the right mobility tools in place.

Converging infrastructure
2019 will see a new meaning come to the word "convergence". In 2018 we saw hardware vendors trying to converge the software layer into their product offering. However, all they've really created is a new era of vendor lock-in – a hyper-lock-in in many ways. In 2019 organisations will rethink what converged solutions mean. As IT professionals increasingly look for out-of-the-box ready solutions to simplify operations, we'll see technology vendors work together to bring more vendor-agnostic, comprehensive converged systems to market.

Backup
In 2018, it became clear that current backup solutions are no longer fit for purpose, with nearly half of all businesses experiencing an unrecoverable data event in the last three years. However, preventing data loss requires constant, second by second backup. The big challenge for backup vendors with this is how to minimise the impact such frequent backups can have on production. Some hardware vendors have worked hard to try and minimise the impact, however have created more of an evolution to backup technology rather than revolution. The race is still on to find a way to eliminate the impact entirely.