Market Predictions from 38 vendors

January 2018

Companies:
Avere Systems, C3DNA, Caringo, CloudEndure, Cloudian, Cohesity, Coraid, DataCore, Datos IO, Datrium, DDN Storage, DriveScale, E8 Storage, Elastifile, Hedvig, Igneous, Iguazio, Infinidat, iXsystems, Kaminario, Komprise, Minio, NetApp, NooBaa Storage, Panasas, Panzura, Portworx, Primary Data, ProphetStor, Quantum, Qumulo, Reduxio, Rubrik, Spanning Cloud Apps, StorONE, StrongBox Data Solutions, Vexata and WekaIO.

Avere Systems

Scott Jeschonek, Director of Cloud Solutions
Dan Nydick, Technical Director

Cloud Domination Slashes Cloud Costs and Slows Spending On Data Center Infrastructure
Not so long ago, enterprises around the world were spending over $140 billion on on-premises data center systems. But in the past couple of years, the rapid adoption of public and private cloud services has sparked a major reallocation of IT budgets. Enterprises are moving away from relying solely on on-prem data center systems and are boosting spending on IT infrastructure for deployment in cloud environments. According to IDC, cloud services and infrastructure will reach $266 billion in 2021.

The availability of cloud applications in business environments has helped create a cloud-based workforce, making it is easier than ever for employees or contractors to work remotely. This is just one example of why more companies are operating with no local IT infrastructure. Cloud providers have recognized this reliance on the cloud and we believe they will continue to focus on reducing the barriers to migrating to the cloud in the year ahead. More companies will pull from their data center budgets and apply the funds to building cloud-ready business applications and improving security. Also, based on the competitive nature of today's top cloud vendors, the overall cost structure of the cloud is expected to continue going down.

More Robust Cloud Security Boosts Confidence in Using the Cloud
Worries about data security have made many enterprises hesitant to store sensitive IP in a public cloud environment. But the increasing sophistication of internet attacks have made it almost impossible for non-experts to keep their corporate networks and servers secure. To protect data, we expect more willingness for organizations to move their sensitive IP to the cloud to leverage the security infrastructure and expertise of the major cloud providers. This poses a new challenge for organizations: protecting data from the cloud providers themselves and securing data as it flows to and from the cloud. We expect to see cloud vendors increase efforts to improve organizations' confidence in the security of their data in the cloud.

What will this look like? We expect to see more ability for users to manage data security in the cloud. This will include expanding access to features that protect sensitive data, such as user specified encryption keys, secure key storage appliances, encryption of network traffic, and encryption of data at rest. Additionally, we believe cloud vendors are putting effort into developing radically new (but currently immature) technologies to securely operate on data, such as homomorphic cryptography. Homomorphic cryptography allows an entity (such as the cloud) to do searches and computations on encrypted data, without needing the keys or access to the plaintext. This will allow customers to use the cloud for computation without granting the cloud vendor any access to the actual data.

Cloud Computing = More Serverless Tech and Hosted Services
Organizations often spend significant time and effort managing compute infrastructure, a cost which is not central to their mission. No longer having to manage this infrastructure has always been one of the main benefits of moving applications to the public cloud. The cloud vendors are providing increasingly high level managed services that allow customers to concentrate on their mission without needing to be distracted by management of virtual machines, web servers or databases. We'll see increasing use of hosted, scalable web services (such as Google App Engine and AWS's Beanstalk) and of serverless technologies such as AWS Lambda and Google Cloud Functions as a more cost-effective way to manage and deploy complex enterprise applications.

We expect the cloud vendors to continue the progression to higher level managed services, such as fully managed distributed databases (Google Cloud Spanner), and new abilities for third parties to sell applications hosted in the public cloud (Azure Managed Apps).

Myth: The File System is Dead
The death of the file system and network file system (NFS) has been greatly exaggerated. In 2018, we will continue to see scenarios where organizations want to go all in on the cloud, but they still have file-based applications on-prem. In many companies today, people who have existing workflows using files don't bother making the move to the cloud. In order to reap the benefits of cloud, organizations must go all in on the cloud.

Life Sciences and Explosive Growth in Centralized Research Hubs
We're seeing explosive growth in centralized research hubs, where key institutions are evolving into Service Providers. The ecosystems created by such institutions/organizations offer resources to other organizations that may not be able to afford their own. One key example is the ever-increasing throughput of sequencing technology. Today, a sequencer can handle multiple sequencing jobs simultaneously, dramatically increasing the capability. These sequencers are still expensive, and so a larger organization may have multiple units. Just as the public clouds do, a Life Science Service Provider can distribute the cost of operations by offering sequencing and post-sequencing analytics services to downstream organizations. The downstream organizations benefit by leveraging best-of-breed technologies without having to absorb the entire capital expense. Combined, the ecosystem ends up generating massive amounts of data.

Most significantly, some of these Service Providers have begun looking to using the public cloud as their analytic and distribution layer. The public cloud offers connectivity and on-demand computational capabilities that are simply too expensive for a single institute to scale effectively. This way, a downstream customer can send raw material for sequencing to the Service Provider, who in turn may make the sequenced output available in the cloud. From there, the downstream customer can either choose to acquire the sequenced output and analyze it locally, or could leverage on-demand computation assets to run their alignment and variant pipeline directly in the cloud.

C3DNA
Rao Mikkilineni, Interim CEO

Virtual Private Datacenter and workload mobility will unify Public and Private Clouds

Migration of mission critical core workloads such as Oracle database stack will move to a hybrid with Oracle cloud as a easy home

2018 will be all about applications and their management and not about infrastructure and it's management

Caringo
Adrian Herrera, VP Marketing

The European Union (EU) General Data Protection Regulation (GDPR) and More Major Breaches
In today's digital world, it should come as no surprise that ransomware and major security breaches will continue to occur. Lawmakers around the world will continue to pass legislation like GDPR in response. In 2018 compliance with GDPR and the need to respond to secure data breaches will take IT execs by surprise offering an opportunity for those in the data management and storage space to provide technologies that can quickly identify personal and sensitive data stored.

The Metadata Deluge driven by Artificial Intelligence (AI), Machine Learning (ML) and IoT will Accelerate
As AI, ML and IoT applications continue to become part of the mainstream, the data about interaction, analysis and consumer behavior will increase exponentially. Storage technologies that integrate NoSQL search will help by enabling organizations to capture everything they want, without defining a schema for future, but yet-to-be-defined use.

Major Brands will Continue to Move Off of AWS as Amazon Continues to Dominate Many Industries
While S3 is seen by many as a defacto API for internet accessible storage, there will be more popular brands in the brick and mortar, M&E and web application space moving off of AWS. The benefit that AWS brought by streamlining operations and reducing time to market will be overshadowed by the threat other Amazon business units represent to many organizations' primary business. This will not impact AWS's overall growth but I predict some highly visible migrations to Azure and a move by some to bring the most strategic operations and near-line archives in-house.

CloudEndure
Gil Shai, CRO

We will see a shift of traditional DRaaS providers who are currently using their own data centers and private clouds looking to sell off their data centers and offer their services using public cloud infrastructure in order to reduce capital and operational costs.

The public cloud market will continue to consolidate where the current leaders AWS, GCP and Azure will be joined by public cloud providers from Asia, like Alibaba Cloud. The small providers will continue to disappear

The transition to Cloud Disaster Recovery will continue to strengthen and more and more enterprises will move their disaster site to the public cloud for the benefits of cost, flexibility and visibility (vs. an on premise disaster site). We will also start to see Cloud to Cloud Disaster Recovery

Cloudian
Gary Ogasawara, VP Engineering

2017 saw many changes: exploding data growth in storage driven by AI/ML/DL and data-intensive formats, the rise in cloud storage as a service, and the decline of flash memory prices. In 2018, I believe we're going to see an increase in organizations taking advantage of object storage for structured/tagged data and we'll see object storage creating structured/tagged data from unstructured data; metadata will be used to dig out from the avalanche of data generated by artificial intelligence; and as data volumes explode, organizations will start adopting new economically-friendly storage strategies.

Object storage for structured/tagged data, and object storage *creating* structured/tagged data from unstructured data, will be increasingly important
The recent past saw an explosion of analytics technology as businesses demanded tools to unlock insights from their vast and growing sets of data. Now that those tools are available, the pendulum will swing back to the demand side of the equation and force businesses to pay more attention to collection, management and storage of that increasingly valuable data.

The amount of data that businesses are collecting is spiking – witness Intel's statement late last year that the average car will generate 4000 GB of data per hour of driving. But it's not just the raw data that's causing the spike. Emerging data formats are causing a spike, too, as we collect more unstructured data - think of the video captured by a Tesla and shared with the company. In order to use this data, you have to understand it. That need for rapid understanding is driving a spike in data tagging and an associated spike in the collection of metadata. As data volume increases, the value of structured or tagged data increases disproportionately to the value of unstructured data.

This data is collected by sensors of various types – wear and use sensors built into machines, video and radar recorders and transmitters in cars, and medical equipment that captures patient health information, for example. As that data - and its associated metadata - is leveraged for business advantage, it's likely to spawn even more data.

As businesses analyze these enormous data sets, they will start to identify additional data types that could lead to further correlations and additional business advantages, which will justify the deployment of more sensors to understand more about the real-world experiences of customers. The data from these sensors will generate increased need for greater storage capacity and advanced data storage management, which will lead to more insights and more sensors, and on at on. We're at the beginning of a cycle in which businesses will strive to understand more about their customers digitally, and it will increase the strain on IT to keep growing storage infrastructures while preserving the usefulness of the data.

AI will Trigger an Avalanche of Data – And Help Dig Out from It, with the Help of Metadata
Business is taking advantage of a range of new technologies – artificial intelligence, high-quality video, Internet of Things, analytics and more. The three things all these technologies have in common is that they

are all data-intensive, and they demand ever-greater storage capacity - and they will depend on tagged data to function most effectively.

It does little good to store vast amounts of data if you have no way to access what you need to retrieve, or if you don't have any idea of which data assets exist in the first place. Metadata is the key to extracting value from the data.

Structured/tagged data are a type of metadata, or model of the data. The metadata and models are a higher-level of abstraction beyond the raw object data and are required for analytics. Without metadata, the unstructured data captured by data formats like video becomes an unsearchable liability instead of an asset. With metadata, the data can be navigated, analyzed, understood and put to use.

A good example of this is in the management of video assets; media and entertainment, surveillance and security, and even automotive uses of video are increasing dramatically. But it's not reasonable to expect your employees to watch endless hours of video in search of the single clip you need. Instead, facial recognition software built on AI will be used to sift through tens of thousands of hours of material to tag recognized faces, meaning that when the need arises it's a simple task to locate just the right clip of clips.

As AI/ML generates and uses metadata and models, systems that can efficiently and effectively manage the metadata and models become critical. AI will become an indispensable tool for finding the most valuable data within enormous data sets.

Economics Will Force Businesses to Adopt New Storage Strategies as Data Volumes Explode
Not only is data a tremendously valuable asset, but businesses are creating more of it, faster than ever before. That means that businesses must invest in new strategies for safeguarding and protecting data – but at the same time, they have to store data in economically responsible ways. That's not as easy as it might sound. Configurable data policies that control how the data are stored once inside a storage system will become critical. These data policies can control the durability, cost, availability, and other properties of the data according to dynamic optimization criteria. A simple example is moving data from hot to cold storage. But the optimization criteria are continuously variable, and can reflect tradeoffs based on business priorities.

For example, a user may want to trade data durability for lower cost. One of the paradoxes of data storage is that the more you store in primary storage, the more per unit of storage it costs. That was bearable when a terabyte was considered a lot of storage, but today most large businesses have multiple petabytes of data under management. Is it affordable to keep it all in primary storage? Or is it smarter to look for secure archiving combined with advanced search tools to keep data costs down while making sure you can find the data when you need it?

The concept of tiered storage has existed for 20 years, but it was deemed an unnecessary expense as disk storage prices tumbled in the early 2000s. In 1995, storage costs for a GB of data were around $1000. Within five years, the cost per GB plummeted to less than 10 cents per GB, and today it costs around 3 cents per GB. There was no need to invest in storage management – drive capacity was so inexpensive there was no reason to consider it. Today, however, the volume of data is flipping the economic formula on its head. Storage prices that once seemed negligible are now looking like less of a bargain since the amount of data stored is so enormous.

These economic factors will prompt businesses to revisit these old strategies in the coming year. Harkening back to the past is easier for a couple of reasons: that archived data is much more quickly available today, and search technology has come a very long way, allowing archived data to be examined as easily as data in primary storage.

Cohesity
Mohit Aron, CEO

Data centres will continue to shrink, through the application of hyperconvergence to both the primary and secondary storage parts of the data centre
I believe this will apply across the board to all types of users and use cases. The main reason is that hyperconverged systems are easier to manage, they take less power and space, and are less costly. I often use the smartphone analogy, comparing these devices to hyperconverged systems, and drawing a parallel with what happened in the consumer world with smartphones.

Before these smartphones, we all used to carry separate items, notably, a phone, a music player, a camera, a GPS device, a flashlight, a calculator etc. Then, the smartphone came and consolidated all of that on a single platform, a single device. If you want to take the analogy further, the data centre is the bag you would need to store all of the separate devices, that you don't need any more with the smartphone! It can just go in your pocket. Hyperconvergence is doing the same thing to the various silo'd workflows in a data centre, and over time it will make the need for centralised data centres obsolete.

Not only will data centres shrink in size, but in the next five years, we will begin to see changing data centre models
Thanks to edge computing, even IoT devices will become mini data centres. This can't be done without hyperconvergence. In a traditional data centre, you have IU or 2U or 3U – in racks. In a self-driving car for example, you don't have racks, but still, the IoT device that is this autonomous vehicle, is basically a mobile data centre. More and more, the form factor of data centres will be different, and there won't be a one-sized fits all strategy. These changes will make it possible to have holistic designs and models, perfectly fit for specific needs.

Consolidated backup appliances will become mainstream - the traditional practice of putting together enterprise backup solutions by using multiple silo'd solutions from potentially multiple vendors will be frowned upon
With Cohesity, we started this trend, but now the storage incumbents are responding with similar approaches. I expect other vendors will follow too. The reason this trend will be bigger is mainly because of the simplicity of the approach. Rather than putting together a solution by buying products from different vendors and managing them using different licenses and UIs, customers just find it simpler to deal with one consolidated solution.

Coraid
Brantley Coile, CEO

It's going to be a very interesting year. It is my 40th year in the business and I've seen a lot, but I've never seen the kind of changes that are going on now. There's a lot I see in my fogged up crystal ball. But if I have to pick the top three, it will have to be the following:

Cisco will takes itself private
With shrinking revenues, Robbins will pull enough cash together and get the company out of harms way of renegade shareholders. My old alma mater is making good money but the revenue for enterprise products are on the decline. To avoid a shareholder revolt, and the CEO revolving door disaster that would follow, the company will follow Mikey Dell's lead. As a private company it will happily create wealth for the next forty years.

Pure Storage will be bought by Cisco Systems
I have zero inside insight on this one. But it's a natural for Cisco to take Pure with it as it goes private. Giancarlo knows how to fix operations like Pure, and, being the past heir apparent to the Cisco throne, he knows how to build an organization that will fit smoothy into the big enterprise company.

More and more storage startups will follow Coraid's lead and avoid venture money
No VC. No IPO. Not for sale. With EMC, Dell, and CISCO private, the need to boost ones top line by a couple hundred million isn't needed anymore. Money loosing, VC funded storage startups, public or not, have a

very limited market. Just like most of history, new companies will have to make a profit early and grow on profits. The customers are the big winners here. Companies run by the same people over a long period of time create better technologies, evolve better operations, and just don't stop serving their customers with better and cheaper products. History is full of examples.

DataCore

George Teixeira, CEO

This new Age of Digital Transformation requires Real-Time Response and therefore Software-Defined Technologies become essential to reduce disruption

Digital Transformation can be highly disruptive. In today's 24x7 business environment, real-time response rates will become even more paramount to customer experience. This is a driving factor behind the digital transformation initiatives for many organizations.

Today's enterprises are demanding simplicity and software-defined agility; ultimately, they want to make infrastructure invisible, move to more cloud-like service models and deal with running their businesses and core applications versus getting bogged down within the complexity of the underlying infrastructure, hard-wiring systems or drowning in the details of running IT operations.

However, digital transformation can be both expensive and disruptive. CEOs can't drive operational savings fast enough to fund it and want to be careful about eroding profit margins.
There is a more practical approach to digital transformation: adopting a hybrid approach that maximizes current IT investments with software-defined solutions while successfully evolving and building out digital transformation initiatives.

In more infrastructure-related areas like data storage, where complexity and the degree of IT integration is increased, the power of next generation software-defined and hyperconverged storage approaches can greatly simplify and automate the provisioning, management and orchestration of resources and data access. A smart software-defined approach avoids the "rip and replace" hardware-minded models of the past to better support digital transformation by making the infrastructure more invisible to the applications and users. Moving forward, these new and more powerful software-defined technologies will be supportive of modern cloud-like interfaces and technologies like containers, making them very adaptable to support this digital transformation model.

Application workloads especially those built on critical databases need to be more responsive as business becomes real-time. However, changes and optimizations to databases and especially their associated legacy application workloads can be very disruptive affairs. Therefore, new solutions are out there that can take advantage of innovations like parallel I/O can now be downloaded to allow "plug and play" self-tuning software that requires no programming or hardware changes and yet increase performance and improve response times . DataCore as an example is working diligently to deliver more of these type of 'non-disruptive' solutions to a wider range of applications and use cases in 2018.

It's a Hybrid Cloud World – Software-defined solutions are the bridge

Despite the ever-growing popularity of the cloud, there will always be a need for on-premise technology. Some applications also face larger obstacles in moving to the cloud; latency, intermittent connectivity, and regulation requirements being primary examples. As a result, hybrid cloud technologies will continue to grow in importance in 2018 and beyond. Fundamentally, hybrid cloud applications can help enterprises more effectively achieve mission-critical business objectives such as accelerating response times and providing more efficient disaster recovery as critical data is continuously replicated within this quickly-deployable hybrid cloud configuration.

The infrastructure, whether located on-premise or in the cloud, needs to become invisible. Data can be anywhere, and access and responsiveness to meet customer expectations drive what matters. Software-defined solutions are key to bridging these worlds.

Making Infrastructure Invisible: Software-Defined and Hyperconverged Become Hybrid-Converged
IDC forecasts that the software-defined storage market will grow at a rate of 13.5% from 2017-2021, growing to a $16.2B market by the end of the forecast period.

Software-defined is the bridge to digital transformation that unifies old and new technologies. It makes underlying changes invisible to the applications on which organizations depend. Next generation software-defined storage solutions will continue their momentum as they bridge the gap between complicated legacy infrastructure and modern "invisible" storage infrastructure needs.

However, while the move to greater virtualization has made the data center more agile and easier to deploy, it came with the price of greater complexity. At DataCore, having had a unique industry vantage point where we have seen the evolution of software-defined storage over the past 20 years, what is obvious is that IT teams can no longer deal with the different silos to manage their resources and deal with all of the complexity and details involved. Software-defined architecture is the essential bridge, not only across deployment model alternatives, but also between on-premises and cloud infrastructures. But IT will need help from technologies such as automation and machine learning to take software-defined storage to the next level.

Next generation software-defined storage with its emphasis on analytics, automation and low-latency performance to support data anywhere models is critical to getting complexity under control and breaking the chains to allow greater freedom of mobility from private clouds, to multi-site clouds and the public cloud.

A hot segment of the software-defined storage (SDS) market over the last few years has been hyperconverged storage. DataCore itself has seen well over 60% growth rates in its hyper-converged solutions over the last year alone. In 2018, hyperconvergence and SDS will blur their lines, and hyperconvergence will become a subset of an overall software-defined model where customers can have the flexibility to choose how to deploy, whether on physical hardware, virtual machines, on appliances, or in the cloud. The end result is still optimizing business productivity and agility.

The two methods will continue to blend into more of a "hybrid-converged" model that is part of a larger continuum of infrastructure modernization and convergence, and users will be able to easily move among deployment options – from storage virtualization, through converged/server SAN, to hyperconverged, to cloud, to hybrid-converged – all under the control of a unified management plane spanning existing legacy infrastructure and new hybrid-converged infrastructure, with the software-defined flexibility to absorb future technologies.

Datos IO
Peter Smails, VP Marketing and Business Development

Multi-Cloud Goes Mainstream
To enable their digital transformations, enterprises are increasingly moving away from physical infrastructure to the cloud and in 2018 that migration will increasingly be marked by a multi-cloud approach. This shift is evidenced by public cloud providers all reporting strong growth, led by AWS who reported $18B revenue run-rate and 42% growth! However, rather than relying on one cloud provider, companies are using multiple clouds for different workloads to achieve benefits such as avoiding cloud lock-in, enhancing functionality and improving cost. To support this shift to multi-cloud and hybrid-cloud architectures, legacy solutions will increasingly be replaced by new storage and data management solutions that provide elastic, scalable, cloud-native capabilities and help enterprises accelerate their digital transformations.

Non-Relational Databases Dominate the Hybrid Cloud World
The biggest database event for 2017 was the MongoDB IPO representing strong market validation that non-relational databases are becoming the lingua franca of data in the hybrid cloud world. In the new Year, this new breed of modern databases including DynamoDB, DataStax, and Couchbase will continue to solidify their positions as the standard platform for modern applications.

Data Aware Data Management Becomes Table Stakes
On May 25 in 2018, GDPR, the most sweeping change to data protection in the past 20 years, will go into effect. Under the new set of regulations, both U.S. and European companies will need to demonstrate compliance when it comes to managing, storing and sharing data – no matter how massive the data sets. One of the biggest issues next year will be GDPR Article 17 enabling users' right to be forgotten, which will increase demand for storage and data management solutions that are data-aware. Whether it's application-specific backup and recovery to protect against ransomware, or intelligent query-based data movement to support test/dev, CI/CD, or GDPR initiatives, organizations will require data management solutions that are data aware and enable them to protect, mobilize, and monetize their data across any cloud boundaries.

Datrium

Brian Biles, CEO

Hybrid cloud matures
The hybrid cloud market is quickly maturing as leading HCI vendors invest in disaster recovery (DR) automation that runs across on premise and cloud infrastructures. For example, VMware is increasing investment in Site Recovery Manager for DR automation and orchestration with VMware Cloud. Nutanix is promising DR automation with their Xi DR service. Datrium will release market leading multi-cloud DR automation in 2018 that will grow to complement its new, hyper-efficient Cloud DVX on AWS. The sum of all these advancements will accelerate the hybrid cloud market to even higher growth rates and maturity.

IoT infrastructure challenges push IT infrastructure limits
IoT requires IT infrastructure to ingest multi-petabyte data sets and then immediately make the data available for analysis at a large scale. IoT software from PTC Thingworx, SAP HANA, and Cisco Systems often struggle to find modern infrastructure that can meet their demanding requirements; this data must be captured and processed efficiently. Datrium DVX is designed to satisfy these hyperscale demands with both optimized write throughput to scaleout data nodes and maximum read performance from the compute nodes, so customers can gain the benefits of IoT.

Primary and backup storage converge
Many companies have attempted to combine primary and backup storage into a single cost optimized system. To date, none have succeeded, but leading HCI providers are still trying as we enter 2018. VMware outlined their plans to move toward converging primary and backup at a recent VMworld 2017session. Cohesity suggests customers can run non-critical workloads like trest/dev, archive, and analytics against their backup storage. Simplivity's customer survey found 51% of their buyers retired their backup recovery infrastructure after installing Simplivity appliances. With simpler scalability and always-on erasure coding, compression, and dedupe, Datrium has converged primary and backup storage by leveraging server-based flash, advanced data management software and cost optimized capacity; active data IO processing happens in server flash while all data copies land directly on scaleout backup storage. Indeed 2018 will be a break-out year for the consolidation of primary and backup data.

Blockchain finds its way into mainstream IT infrastructure
Many IT infrastructure providers, both small and large, are leveraging blockchain technology to improve their products and, as an example, secure the identity of online consumers and services. Identity management company Okta is working to help streamline cumbersome and expensive Identity Proofing flows leveraging blockchain. Datrium is also applying blockchain concepts to verify the integrity, location and security of the data across the DVX system.

DDN Storage
Paul Bloch, CEO

Data gets shifted up another rung of the ladder and moves to the top table for enterprises
The business core of global enterprises and their IT departments will be driven closer together by the competitive advantages of Analytics and IoT.

HPC and AI get married
Lots of potential for HPC applications and neural nets to work together and subtly shift our thinking of what EXAScale looks like.

2017 saw the start of NVMe into storage markets
2018 will see NVMe go mainstream and requirements for Specially built appliances to deliver performance for specific applications.

DriveScale
Gene Banman, CEO

Software Composable Infrastructure Becomes the New IT Mandate
In 2018, more organizations will move toward Software Composable Infrastructure (SCI) to manage enterprise-scale data center assets with existing infrastructure – providing organizations with public cloud flexibility as well as the security and cost savings of an on-premise solution. As data projects continue to scale infinitely, cost-effective storage and compute will become more critical to a company's overall efficiency and bottom line.

Companies Don't Become Global, They *Start* Global
The global economy is not going away. That much is clear. What we will start to see in 2018 is the idea of creating a global business at a company's inception. Right now, a large portion of startups focus most, if not all, of their efforts on the U.S. and North America. With such a large concentration of companies in Silicon Valley, organizations will turn their attentions to other areas of the world earlier in their evolution and with more vigor so as not to limit startup momentum and potential growth.

E8 Storage
Zivan Ori, CEO

The NVMe protocol will become the standard for solid state drives over time
Single port 2.5" and PCIe card form factors have already been incorporated into server offerings from all major vendors, with higher performance than SATA for nearly equal $/GB. Dual port 2.5" form factor SSDs will drive the adoption of NVMe in storage systems, and adoption will accelerate as the $/GB approaches that of SAS SSDs.

To go along with that, various standards are emerging to support the NVMe protocol over various transport mediums (RoCE, Infiniband, FC, etc). Early variants of the standards have already been brought to market by emerging storage vendors such as E8 Storage. However, 2018 will see broader deployment of NVMe compliant offerings from a broader range of vendors.

High speed Ethernet continues to outpace other storage protocols, and will increase as data traffic is increasingly moved to Ethernet
Infrastructure hardware such as network adapters and switches capable of 200GbE were introduced in 2017, and field deployments should begin in 2018 enterprise data centers update their networks.

The introduction of very high capacity SSDs in 2016 and 2017, as well as the continued cost reduction curve (2016 supply shortages notwithstanding) are fueling the transition of all flash arrays from niche applications to general purpose computing across the data centers. Many workloads previously run on HDD are either moving to all flash arrays, or alternatively moving to cloud based services.

Elastifile

Amir Aharoni, CEO

The Rise of Hybrid IT Gladiators

Many companies big and small are arming themselves right now for deadly competitive advantage through truly dynamic hybrid cloud workflows. In 2018 we will see the a new wave of battles in the arena, accelerating the business processes across a combination of on-prem and on-cloud resources for maximum agility at the lowest cost. These new Hybrid IT weapons include cloud-bursting for temporary peak capacities, workflows that apply on-cloud compute or specialized GPU resources before moving back on-prem, and integrative analytics that dynamically pull data across sites, clouds, and service providers. We see these gladiators moving aggressively in FinTech, Oil & Gas, Life Sciences, Media & Entertainment, and more. And even though *data* has chained them down in the past, it can now be their sharpest sword with new data fabric solutions... *"Strength and honor!" (Maximus Decimus Meridius, Gladiator)*

Enterprise Eggs in Multi-Cloud Baskets

As customers learn the lessons of their initial forays into cloud deployments, they are moving decisively to multi-sourced cloud IaaS (and avoiding the newest incarnation of vendor lock-in). Most enterprises will have at least dual-source cloud service providers that they will leverage on-demand as pricing, location, features, or regulatory preferences dictate. This will become standard in 2018 with mainstream solutions for multi-cloud app and data portability, standards, DR, regulatory compliance, and even effective hacks to the dreaded egress fees... *"Do not put all your eggs in one basket" (Warren Buffett, The Oracle of Omaha).*

The Prodigal Son Returns

After years of misadventures trying to get critical apps to be refactored and meet higher SLAs for performance-sensitive workloads, object storage will come home, as an integrated part of more holistic file & object services solutions. Enterprises are demanding this on-cloud and on-prem, with new SDS distributed file systems as the standard for "lifting and shifting" existing apps to cloud (without refactoring or compromised SLAs) and object storage as an integrated "cheap and deep" tier for inactive data. Everyone is converging-- object storage vendors are adding file services, file system vendors are adding object tiering, and even analysts like Gartner are combining their Magic Quadrant definitions to include both together..."*The back door beckons to a prodigal son" (Michael Davidow, author and attorney).*

Hedvig

Avinash Lakshman, CEO

Traditional backup vendors feel the heat

As a cloud mentality becomes commonplace, traditional backup software and hardware solutions – a place where enterprises spend significant sums – will lose share faster than in the past 3-5 years. Software-defined storage and distributed systems approaches are new to many organizations, and therefore backup is a comfortable place to start. Backup solution vendors that have been around for a decade or more will find it challenging to hold their ground against agile startups that bring new approaches to the data protection game.

Like Docker for containers, Kubernetes becomes the de facto cloud orchestrator

I could easily have made this prediction last year and it would be just as true. The recent embrace of Kubernetes across the industry – including Microsoft Azure, Docker, and Mesosphere DC/OS – shows that the open-source container orchestration system has proven its effectiveness in providing simpler cloud deployment, better scaling, and more efficient management. It is already baked in to OpenShift and other private label solutions. In 2018, the IT world will become conversant, if not fluent, in Kubernetes, its concepts and terminology and start to expand inside enterprises and clouds.

Serverless computing makes a grand entrance

While organizations have quickly become accustomed to the idea of servers as cloud instances, what if there were no servers, and instead just resources rented to run applications and microservices? Serverless computing promises to remove the overhead of provisioning, scaling, and managing servers, allowing you to focus on the application instead of infrastructure. AWS Lambda is an example of this

approach. In 2017 most of the IT world has become familiar with EC2. In 2018, Lambda will begin to take its place in our our cloud vocabulary and become demystified. This is another step toward cloud computing nirvana, and to me it makes a lot of sense.

Igneous Systems
Steve Pao, CMO

As 2017 draws to a close, we're thinking about what 2018 holds for the data backup and archive industry. Here are our predictions for 2018.

Data management needs to evolve to handle unstructured data
Unstructured file and object data storage will continue growing at a faster rate than unstructured data. A 2017 IDC study states that unstructured data grows at a rate of 29.8%, while structured, or block, data grows at a rate of 19.6%. This trend will continue as more modern applications use unstructured file data.

This means that current data management strategies will need to evolve to better handle large amounts of unstructured data.

The death of software and the rise of services
We predict that the industry will move toward everything offered as-a-service, rather than "software-defined." By 2020, 80% of software vendors will change their business model from traditional license and maintenance to subscription, according to Gartner.

We have seen this shift in both enterprise and personal software; for example, Microsoft Office and Adobe Creative Cloud are delivered as-a-service. Why not have the improved infrastructure that comes with services in your datacenter as well as in other areas of business and life? It will be essential to focus on services in 2018, not just software in the datacenter.

The traditional "headcount per petabyte" ratio cannot scale
Jobs in backup will keep growing as data grows, but not at the traditional "headcount per petabyte" ratio. We asked participants in our first CrowdChat how many people it takes to manage data on-premises. Aaron Cardenas, CEO of P1 Technologies, said for primary storage it takes around 1 person per 500TB and for secondary storage it takes 1 person per 3-4 PB.

However, this ratio of backup admins to petabyte of data doesn't make sense economically when data is growing exponentially. Enterprises will need to employ third party service providers to maintain reasonable economics.

Iguazio
Asaf Somekh, CEO

We will see growth in edge analytics and edge processing
Edge data sources can be anything from a car or wristwatch, to an industrial component which gathers data from several machines in a factory. We will see more and more edge processing in 2018 performing data collection and analysis at the edge of the network, near the source of data. A centralized cloud will be supplemented with distributed "mini clouds" at the edge in order to move high volumes of data across multiple edge devices.

Serverless functions will realize their potential in 2018 and get used across a broader set of applications
Serverless platforms will also extend from the cloud to on premises environments. Further down the road we will see additional servereless patterns, such as serverless networking, serverless databases and serverless data APIs. In general, the growing trend of serverless architecture will highlight our focus on data parallelism in multi-cloud environments, moving towards continuous development and actionable insights.

We'll see a shift from using traditional block or file storage to cloud-native environments which use a data layer that directly attaches storage or flash and leverages the resiliency and data management features built into the data services middleware. Optimized hardware is more cost-effective to use when there's enough data to store, as it has a higher disk or flash density and higher network throughput, as opposed to a generic server/VM.

Infinidat

Brian Carmody, CTO

Digital Transformation (DX)
Digital transformation is the profound and accelerating transformation of business activities, processes, competencies and models to fully leverage the changes and opportunities of digital technologies and their impact across society in a strategic and prioritized way, with present and future shifts in mind. Including in the Digital Transformation there are big trends as Cloud, Big Data, analytics, artificial intelligence mobile/mobility and the Internet of Things. To address these challenges, it is mandatory that a storage infrastructure is reliable, flexible, agile, easy to use, includes high performance and maximum data protection. To support technology innovation, a flexible business model that can scale to large capacity and includes a capacity-on-demand model (e.g. payment *after* the capacity is used), with the possibility to activate new customers or new applications in few minutes.

Companies continuing to invest in DX will also continue to reconcile their need for more storage capacity and greater performance to analyze the data they're collecting. Predictions for the decline of high capacity disk drives are overstated, as the only reasonable economic way to capture, analyze, and retain DX data lies in high-capacity disk-based persistent storage.

Further, the growth of capacity will require more intelligent caching such as machine learning and neural cache to accelerate and improve access to information storage.

Data Security and Protection Take the Top Spot (GDPR accelerates change)
GDPR requires end-to-end encryption, which renders compression and deduplication useless on the server or the storage array; this in turn destroys the value proposition of the flash media focused storage solutions. IT organizations that have declared themselves to be "flash only" will have to rethink their strategy, as the effective capacity of their storage shrinks to 1:1 and the cost per effective all-flash TB soars to 10X the cost of high capacity disk.

Encryption and Compression become commodities
See GDPR above. While this will not happen overnight, it is certainly an attention getting statement, given the state of storage and the "all in" approach the major vendors are taking with flash and NVMe.

Other parts of the data lifecycle will also leverage compression and dedupe (e.g. containers) commoditizing these features in the storage tier. The value proposition will shift [again?] to performance, availability, reliability, and ease of integration with vertical/horizontal solutions.

Cloud Inroads with the Enterprise
Today, enterprises typically view cloud strategy as part of a larger strategy like digital transformation, costs savings and/or new market entry. However, the cloud should be viewed as an enabler of many other business strategies. IDC says that By 2019, cloud IT infrastructure spending will be 46% of total expenditures on enterprise IT infrastructure. CIO Insight says that 81% of IT leaders plan to increase their use of the cloud by at least 25% within the next three years. The first step toward cloud implementation is typically a mix of private and public cloud or to involve customer business applications. Enterprise IT struggles to keep up with the immediate demands of the business. No longer is a 30-day response to a service request acceptable -- customers need it in minutes. To stay in the game, IT must transform, and enterprises rely on their storage solutions to support the rapid transformation pace and increased amounts of data.

While the concept of hybrid cloud with local servers/applications and cloud-based storage has been evolving, requirements like GDPR may flip the design around quickly, resulting in an acceleration in multicloud storage deployments using trusted, co-located storage with high-speed access to temporal

cloud-based apps. This creates a new opportunity for storage service providers in close network proximity to the major (and minor) cloud computing services - Amazon, Google, and Microsoft.

Networking is the New Bottleneck
Standardizing on Ethernet has the potential to dramatically improve stack I/O efficiency.

While NVMe will improve performance, DRAM cache optimized systems (like InfiniBox) will still outperform all-flash systems AND provide the best balance between multi-workload performance and long-term capacity requirements.

iXsystems

Mike Lauth, CEO

More companies "born in the cloud" will begin to recognize the fast ROI and lower TCO of transitioning off of the public cloud into their own data centers and private/hybrid cloud environments.

Proliferation of NVMe will accelerate, and NVMe will become commonplace for hybrid storage arrays. NVMe will become the dominant technology in all flash arrays as well. This is all due to the reduced cost/increased capacity of flash devices and the bulk availability of enterprise-grade U.2 NVMe devices, in tandem with the improved NAND supply in the market.

Open Source-based Software Defined Storage will continue to slowly and methodically pound nails into the coffins of the traditional storage incumbents.

Kaminario

Dani Golan, CEO
Eyal David, CTO

The rise of the public cloud (SaaS, PaaS, and IaaS) will create opportunities for a broad range of xSPs (i.e. Managed Service Providers, Managed Cloud Providers, Hosted Application Providers)
We'll see the growth of custom cloud services offerings that offer the flexibility and agility of AWS with the ability to accommodate concerns around privacy, security, regulation (e.g. HIPAA, GDPR), lock-in, and cost.

The storage industry will continue its trajectory away from traditional scale-up array architectures towards more flexible, software-defined architectures
In particular, we see the composable infrastructure / composable storage paradigm gaining traction as the natural evolution of scale-out storage architectures.

All forward-looking IT organizations – whether they are traditional enterprise, SaaS providers, consumer internet companies, or xSPs – will be highly focused on laying out their next-generation, software-defined-datacenter (SDDC) strategies
SDDC strategies will emphasize the concepts of openness, capturing the benefits of commodity hardware economics, automation, and orchestration – while not compromising on performance and availability. Further, we'll see SDDC incorporating artificial intelligence and automation to simplify datacenter operations and codify devops thinking into the fabric of the cloud-scale datacenter.

Komprise
Krishna Subramanian, COO

Machine-generated Unstructured Data Growth Surpasses User-Generated Data
Trends like Internet-of-Things (IoT) along with more sophisticated instruments in virtually every industry are creating an explosion of unstructured data growth. Whether it is better genomics sequencers or higher fidelity media or self-driving cars or higher density engineering design automation tools, these new machines are generating 2 to 10x more data than before. As a result, businesses are looking for data management solutions that seamlessly handle both machine and user generated unstructured data without any disruption to the applications and users.

Data Management and Governance Becomes a Strategic Line Item, no longer Tactical Spend
Businesses are becoming increasingly data-driven, making data management and governance a key Business priority that drives strategic spend in organizations. Storage is no longer just a tactical line item – instead, it is a cornerstone of key corporate initiatives such as Cloud, Data-Driven Enterprise, and IT Modernization. With this shift, storage infrastructure leaders are increasingly being asked by CFOs, CIOs, and CEOs to provide greater visibility and reporting into data, its usage, data management, governance, and planning.

Another Year of Flat Storage Budgets Makes Analytics-driven Data Management Critical
Storage budgets have remained mostly flat for the past few years, and this trend is expected to continue. Since the cost of managing and protecting data is 4x of storage, as data footprint continues to grow while budgets remain flat, businesses are realizing they can no longer afford to treat all data equally and manage it all with one broad stroke. Instead, since 80% of the data becomes cold within months of creation, businesses are adopting analytics-driven intelligent data management solutions that store and manage data differently based on its current value. Businesses are adopting analytics-driven solutions that adapt the data management dynamically based on the current value of data.

Minio
Anand Babu Periasamy, CEO

Object storage breaks enterprise resistance
Enterprises are beginning to encounter applications that are solely designed with Amazon S3 APIs. The sheer amount of data and security concerns will force enterprises to build private clouds with object storage as its primary storage.

Hyperconvergence and cloud-native are on a collision course
Hyperconvergence is fundamentally incompatible with the cloud-native architecture and DevOps methodologies. Infrastructure architects will be forced to pick sides. 2018 will be the determining year for the private cloud.

Solid state storage enters unstructured data market
Solid state storage will far exceed the density and economics of mechanical hard disk drives and become the only viable solution to build 100 PB datacenters. Given that 80% of the world's data is unstructured, spinning drives will be put to rest in the coming years.

NetApp
Mark Bregman, CTO

As the world is changing fundamentally, data will continue to become even more distributed, dynamic and diverse. To be successful in this environment, enterprises must focus their 2018 strategy on digital transformation – everything from IT infrastructures to application architectures will have to adapt in response to new realities in the hybrid cloud world.

Data will become self-aware
As data becomes even more diverse than it is today, the metadata will make it possible for the data to proactively transport, categorize, analyze and protect itself. The flow between data, applications and storage elements will be mapped in real time as the data delivers the exact information a user needs at the exact time they need it. This also introduces the ability for data to self-govern. The data itself will determine who has the right to access, share and use it, which could have wider implications for external data protection, privacy, governance and sovereignty.

Data will grow faster than the ability to transport it...and that's OK!
Data has become incredibly dynamic and in turn is being generated at an unprecedented rate that will greatly exceed the ability to transport it. However, instead of moving the data, the applications and resources needed to process it will be moved to the data, which has implications for new architectures like edge, core, and cloud. In the future, the amount of data ingested in the core will always be less than the amount generated at the edge, but this won't happen by accident. It must be enabled very deliberately to ensure that the right data is being retained for later decision making.

Data management will need an emergence of decentralized immutable mechanisms
Mechanisms to manage data in a trustworthy, immutable and truly distributed way (meaning no central authority) will emerge and have a profound impact on the data center. Decentralized mechanisms, like blockchain, challenge the traditional sense of data protection and management. Because there is no central point of control, such as a centralized server, it is impossible to change or delete information contained on a blockchain and all transactions are irreversible. Current data centers and applications operate like commercially managed farms, with a central point of control (the farmer) managing the surrounding environment. The decentralized immutable mechanisms for managing data will offer microservices that the data can use to perform necessary functions. The microservices and data will work cooperatively, without overall centrally managed control.

NooBaa Storage
Yuval Dimnik, Founder

Containers footprint will grow with fluid move between on-premises and cloud
As being pushed by the Cisco/Google partnership.

Workloads will stretch and burst environments
Currently hybrid and multi-cloud means some assets here and some there.
With cognitive services offered by different cloud providers, we'll see a primary workload running on site A (or cloud A) and getting some cognitive services on cloud B for the same application.

Data growth will bring back more tape
Most of the data we have is created in the last couple of years, but most of it will never be touched weeks after creation. Tape is still the most economic for archive.

Panasas
Curtis Anderson, Senior Architect

In some segments of HPC, the algorithms and data types are changing too quickly to bother with highly optimizing the codes, eg: genomics and AI/ML, and storage subsystems will need to adapt
Some segments (eg: CFD) are "mature" and the optimal I/O patterns are well understood, while in other segments they are not. New workloads are becoming important that have large numbers of small files, and the less sophisticated users in these new spaces want to concentrate on gaining new insights rather than optimizing their code. HPC storage systems need to widen the sweet spot of their performance profile from the very-high-but-narrow peak of traditional architectures.

HPC storage subsystems will need to rapidly implement several different sets of "data management and compliance" feature sets. Data provenance, where a piece of data came from and who modified it, are going to become critical in all the core growth segments of HPC storage. Specific compliance regimes will be required in several vertical segments as well
In the life sciences and precision medicine space, HIPAA Compliance will be required. In the financial analysis space, at least for personalized analysis, SOX Compliance will be required. In the AI/ML space, a push for Compliance will be driven by the organization's liability lawyers to show that "best practices" were followed when training the neural network.

You may have heard this example before. Assume a self-driving car is moving down a road with a motorcycle on its left and a sidewalk on its right full of pedestrians. Then assume that several pedestrians fall into the lane, the car cannot stop in time, and that a collision of some form is inevitable. The software will hit the brakes, but will then have to decide whether to steer to the right, hitting pedestrians on the sidewalk, steer straight and crush the pedestrians that fell into the road, or steer to the left and hit the motorcycle. Hitting the motorcycle will almost certainly reduce the number of people hurt or killed to just one, but a piece of software actively deciding to target a motorcycle through no fault of that driver's raises hard questions of liability that the AI/ML community has not yet sorted out.

The above two trends are going to push traditional HPC Storage closer to Enterprise Storage in terms of features supported, while retaining the performance required from HPC storage.

Panzura

Patrick Harr, CEO

2018 will be the year of analytics
The year of analytics will be fueled by AI/ML for data. The result will give enterprises greater insight into their businesses and control over their operations.

Multi-cloud will become a formal strategy
According to IDC, 90% of enterprises are building a multi-cloud strategy to prevent vendor lock-in. This will stress data management as the cloud becomes yet another data island. New technologies will emerge to unify, access, and manage data across clouds.

Edge computing will become even more important
We are in the age of digital transformation and the explosion of IoT which puts pressure on real-time decisions and access of data at the edge. Vendors who have hybrid platforms will be benefit.

Portworx

Murli Thirumale, CEO

We believe that 2018 will be the year of Kubernetes, so Murli's predications revolve around that.

The complexity of building and running Kubernetes applications will be addressed by the rise in Kubernetes platforms
The Cloud Native Computing Foundation (CNCF) has realized that implementing Kubernetes is a challenge and has thus created a certification model for platforms, *Kubernetes Certified Service Provider (KCSP)*. Customers looking to adopt Kubernetes will look and use one of these providers.

We will see nearly 50 Kubernetes Certified Service Providers by the end of 2018
Currently, kubernetes.io lists 16 KCSP providers. The largest Kubernetes platform, Red Hat OpenShift, is notably missing from this list, but I see this more as an indication that OpenShift needs less external help from CNCF at this point, since it already established as an authority of running large-scale Kubernetes applications. There will be nearly 50 of these providers by the end of the year.

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.

Primary Data
Lance Smith, CEO

Storage Complexity Simplified with Software
Though a big offender of buzzwords over the past year, we expect to see Software-Defined Storage (SDS) solutions continue to gain traction, especially in the data management segment Gartner calls "management SDS." These solutions will begin to bridge the gap between legacy infrastructure and modern storage needs.

IT Begins to Choose Custom Management over Hyperconverged
Hyperconverged infrastructures (HCI) have plenty of appeal as a fast fix to pay as you grow. Yet in the long run, these systems represent just another larger silo for enterprises to manage. With metadata management software, enterprises will begin to move away from bulk deployments of hyperconverged infrastructure and instead embrace a more strategic data management role that leverages precise storage capabilities on premises and into the cloud.

Adoption of Analytics Fuels Storage and Data Intelligence
Storage has long been blind to actual demands on data beyond what was recently read or written. With metadata intelligence, admins and software can finally see when files have been last opened, how often, by who, when they were changed, modified, and more. This intelligence arms IT with the data needed to manage their storage resources more efficiently, and align the right data to the resource that meets business needs for performance and price.

ProphetStor Data Services
Eric Chen, CEO

Self-driven/automated datacenters through AI-powered software bringing increased resilience and economies of scale in the datacenters

Continued Transition to commodity storage-built true Software Defined Storage Solutions to suit the ever increasing needs in Scale Up and Scale Out storage

Emergence of AI-powered work-load distribution storage software for intelligent fast storage allocation and just in-time provisioning versus proprietary AFA solutions

Quantum
Molly Presley, VP of Global Marketing

Data-Driven Organizations Demand Intelligence in their Storage
Expect greater emphasis on tools to maximize content value, primarily in the form of AI to increase intelligence about content, and data visualization software to improve access and the cost of storing data. As data is increasingly used for strategic decision making for companies – or in cases where data is the product itself - storage strategies will become a more foundational consideration.

Tape Steps in as Ransomware Savior
With the rising tide of ransomware and malware in the news, tape will increasingly be a valued element of many data protection solutions because it offers an offline "air-gapped" backup copy.

More Applications Become Ripe for Containerization
As new applications are being written for the container/cloud-native world, existing applications will increasingly be updated for cloud-native architectures.

Intelligent File Systems Will Push Object Storage in the Corner Where it Belongs
File systems will be increasingly leveraged for data ingest and data management. As a result, object storage will be pushed to the capacity and retention part of storage architectures, where it arguably is best suited.

Cloud-Native Architectures to Gain Traction
As cloud-native architectures become an increasingly common term in the IT vocabulary**,** look for organizations such as the Cloud Native Computing Foundation (CNCF) to become more influential.

Qumulo
Pete Godman, CTO

SSD won't be cheaper than HDD
In 2016, every all-flash vendor in the world claimed that SSD was now cheaper than HDD, based on two nonsensical claims that (1) all data is compressible and duplicated and (2) compression and dedupe don't apply to HDDs. Western Digital's MAMR announcement makes it clear that the ratio of NAND flash capacity cost to HDD capacity cost will remain close to 10x for years to come.

All major cloud vendors will start to build or buy file storage
as they realize that quality, scalable file storage is essential for capturing compute-intensive workloads.

SATA and SAS SSDs will rapidly disappear
as the cost of NVMe and SATA/SAS converges quickly

AMD will make rapid inroads in all-flash storage
due to Epyc's enormous PCIe bandwidth, and ARM will wait on the server sidelines for another year.

Reduxio Systems
Jacob Cherian, VP Products & Product Strategy
Mike Grandinetti, Chief Marketing & Corporate Strategy Officer

Convergence of Primary and Secondary Storage
Data center administrators will be dealing with exponentially increasing scale. Existing data infrastructure architectures that rely on multiple platforms and tools to manage data will become increasingly unwieldy and expensive to setup and operate. Industry analysts from Gartner to Forrester and IDC have all noted in their recently published research the increasing customer adoption for platforms that provide converged management of data through all phases of it life cycle. These platforms will be provided by emerging vendors who are not encumbered by traditional architectural approaches.

Design Moves to the Forefront
For too long, storage vendors have not invested in hiding the complexity of their solutions from their users. The industry has hit an inflection point where users can no longer manage the deluge of data that they need to store, manage and protect using existing solutions. As data grows exponentially and infrastructure budgets increase nominally, users will need to dramatically increase their productivity to keep up. They are demanding solutions with intuitive user experiences and interfaces similar to that of their favorite smart phones, gaming consoles and mobile apps, which require no training or unproductive learning curves. Further, they want to be able to manage their infrastructure from ANY device, including smart phones and tablets. Storage vendors will need to increase their design to developer ratios to keep pace.

Focus on Application vs. Infrastructure
Customers will migrate to solutions that allow them to manage their application without having to know the inner working of their infrastructure. To support this infrastructure will have to become more elastic, self-tuning, and flexible with fewer moving parts and knobs.

Rubrik
Chris Wahl, Chief Technologist

2018 is going to be the year where cloud comes home

Next year, we will see an Armageddon when it comes to delivering cloud services directly to the enterprise
The three major cloud providers have picked their dance partners - Azure with Azure Stack, GCP with Nutanix, and AWS with VMware - to see who can deliver their portfolio of services to disrupt on-prem deployments across private data centers and colos.
Right now, we're only just scratching the surface when it comes to the cloud tools of which operations teams are taking advantage. Soon, those will become ubiquitous both in cloud and on-prem with Kubernetes becoming the clear winner across the industry. As more complexity is removed and governance added to this open source tool, adoption will skyrocket.

Businesses will face a GDPR shake up as they make changes to stay compliant
Businesses will be looking for resources to help them make changes to their data strategies, and will be seeking out solutions to effectively index and make their data searchable - to help them meet this mandate.
Currently, AI and machine learning are basically out of reach for most within the enterprise, except for data scientists. In order to bring AI and machine learning to more business users, organizations will need to implement a sound cloud data management strategy to put all of their data in one location for users to organize, analyze and prepare for AI/machine learning applications.

Spanning Cloud Apps
Mat Hamlin, VP of Product

Adoption of SaaS by the Mid Market and Enterprise Continues to Rapidly Grow
The thing that I'm most excited about from technology in 2018 is the rapid adoption of SaaS by the mid-market and enterprises. In 2016 and 2017, we saw a lot of organizations planning, budgeting and piloting their migrations to SaaS applications.
In 2018 and beyond, the migrations are being executed at an amazing rate.

New Functionality Improving the Way Teams Work: Microsoft Office 365's New Services
I'm also excited about some of the new functionality that Office 365 and Microsoft are bringing to the table. Microsoft Teams is a great collaboration tool that brings together chat, intranet, shared files, project visibility and more. They also launched Planner, an easy to use Kanban tool, and Bookings for scheduling appointments.
As they continue to invest and create new services in their platform, customers are adopting it quickly, and that willingness to adopt new services fundamentally improves the way they work."

More services will leverage Machine Learning and AI and the usage will spread to more and more SaaS applications
Machine learning and artificial intelligence continues to be a differentiating factor in SaaS offerings, especially from the key players like Google, Microsoft and Salesforce.
Simple, everyday examples like suggesting meeting rooms based on availability, location and the number of meeting attendees make our everyday lives easier.
In 2018 we'll continue to see more services leverage ML and AI, and the usage will spread to more and more applications since the availability is readily available from cloud providers like Amazon, Google and Microsoft.

StorONE
Gal Naor, CEO

Software
Software will continue to dominate the storage landscape. The importance of truly efficient software will be the most critical and most interesting area to watch in 2018. A new category of storage software will start to emerge, called TRU (Total Resource Utilization) based software. TRU storage software means you will get the full value of the HW installed on the storage system - a few SSD/NvME will be enough to get any kind of IOPS and the raw capacity of the disks, versus the net capacity reduction games that are now being played.

Hardware
The current drives on the market are now able to achieve very high performance and capacity. Getting speed and capacity is no longer a challenge. With TRU storage software it will become much easier to achieve this speed and capacity with just a few drives and without having a distributed system or by having to fill a system with 24-48 drives to get these results. Most of the HW improvement you will see will be on the density side and will be even more improved on the HDD side.

Pricing
With TRU storage software the price per GB will continue to fall in 2018 and will drop to less than a penny (1 cent), as the amount of HW required will be very low with this new technology. With no need for integration and additional floor space, the ability to achieve high performance with a few SSD/NvME or the ability to achieve high throughput and capacity with HDD will contribute to reduce the price per GB. The end users will understand there is no real reason to continue to pay so much, or to compromise on performance and density since they will get the full capabilities of the drives.

IT budgets will continue to shrink or remain flat year to year, while data will continue to grow. 2018 will be no exception. Efficiency will be the most critical factor as IT departments will be required to do more with less, to make up for massive data growth. TRU storage software, will move into the spotlight. It will be a great year for storage end users, as hardware will become less important, driving down not only the total cost of storage, but the complexity as well. Pricing will move to a simple model, where all storage protocols will be included. 2018 will be the year forward-thinking end users will turn to new, innovative software solutions and new business models.

StrongBox Data Solutions
Floyd Christofferson, SVP Products

At the end of every year the storage industry makes predictions about upcoming trends they expect in the new year, usually citing the exponential growth of unstructured data as a driver for innovative new storage types and/or point solutions that offer enhanced capabilities, cheaper price points, and increasing platform choices to accommodate the ever-expanding use cases for all types of digital assets.

The issue is that even with years of advancements in storage technologies, management of storage infrastructures has only gotten more complex, particularly as IT organizations have to do more with static

or shrinking operational budgets. The problem becomes worse when primary storage tiers get full, or need to be replaced in a tech refresh. This inevitably adds more load on IT Staff, and disruption for users.

Automation is the key trend to address this sort of problem across all industries today, leveraging machine learning and data-driven intelligence. The storage industry is no exception, and the technology has finally gotten to the point to bring cross-platform storage and data automation to this industry as well, driven by new metadata-driven software solutions.

The most pressing need for the storage industry is simplicity and automation. That is, to simplify and automate data management across heterogeneous storage types, while also improving ROI of the entire storage infrastructure. And the key to data management is in its metadata, which enable the following trends:

Cross-platform storage automation
Metadata-derived intelligence combined with advanced software will provide end-users complete control over their data regardless of the storage, providing independence from vendor limitations, proprietary applications, or problems managing multiple protocols. Metadata-driven automation enables seamless data flows across any storage type, and dramatically reduces the overhead for IT staff.

This eliminates data silos with a true global namespace that bridges heterogeneous storage types, blurring the lines between performance and cost, on premises and off premises, public cloud, and much more in to a cross-platform managed fabric.

Auto-data classification
Aggregating all information about the data from multiple metadata sources into an actionable framework that drives storage automation.
Auto-data classification provides the intelligence to:

- Enforce data migration and protection policies across any storage platform with real-time responsiveness.
- Enable continuous data protection, data copy management, and the ability to meet compliance requirements such as the EU's GDPR data privacy regulations.

Real-time Data Lifecycle Management
Global storage management with the intelligence derived from aggregated metadata sources finally enables the promise of cross-platform Data Lifecycle Management.
With the combination of automation and intelligence about the data, IT administrations can apply global data policies across all storage types instantly. Use cases evolve dynamically over time, so metadata-driven automation simplifies management and protection data from creation through to archive across the multiple use cases and storage types that may be required over time.

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

Liran Zvibel, CEO

As 2017 is now behind us, all eyes are firmly on 2018 and what it will bring for the storage industry. With more than two decades working in the storage industry, from my time at IBM to my current role as CEO and co-founder of WekaIO, I understand firsthand of just what lies ahead. Here, are my top three storage predictions for the New Year.

The Changing Face of Workloads

Fast GPUs and other accelerators (like Intel Nervana) will greatly condense compute. Already, servers packing eight of these now provide more compute performance than an entire rack used to not so long ago. As a result these servers have much more intense IO requirements, and will require file systems with previously unheard of capabilities. A single client will need to fill in complete 100gb/sec links (12GB/sec each), and also scaling up to many of these. At the same time, modern machine learning workloads process small files (text samples, voice samples, images), and require high throughput out of many small file accesses.

This new low-latency high-throughput workload requires a new file system design, failing which its expensive GPUs are going to lie idle, or projects won't be able to scale. WekaIO Matrix's parallel FS over NVMe is the only file system that is suitable to these upcoming workloads.

Storage and the Hybrid Cloud

Increasingly we will see local storage transitioning to the hybrid cloud, where the most value and efficient usage can be extracted from both on-premises infrastructure and the public cloud. In addition, when on-premise infrastructure reaches its capacity, the public cloud's elasticity will prove increasingly useful for "cloud bursting", which describes the process of spilling excess work over to the cloud, and bringing the results in. We will also see the cloud being increasingly leveraged for disaster recovery.

The key storage requirement for effective hybrid clouds, is the ability to push snapshots to and from the cloud. In order to achieve efficient compute, both ends must work at local speed, and identify the differences between the copies as work is forked to the cloud.

Stateless Access to Storage

Between serverless compute and IoT projects, there is a strong and growing need for stateless access to storage. Currently, the de-facto standard for stateless access to storage is the HTTP-based S3 interface, where the most popular system is the AWS S3 object storage system. As IoT applications become more demanding, they either require lower latency than is enabled by object storage, or a tool chain that relies on standard applications, which need a file system. For serverless projects — many of these have a legacy component that will have to stay in production for a long while. It is clearly apparent that we will see more and more demand for a unified access file system that supports both file systems access protocols (POSIX, NFS, SMB) and S3 and HDFS. Such systems will allow for the running of different tool chains on the same data, eliminating the need to move these to other storage systems.

Exciting things are happening in storage, the IT world is evolving and we look forward to this new year and what news it will bring.