Storage Industry Facts 2021

January 2022

List of Companies

Summary of vendors' facts for 2021

Like every year, StorageNewsletter asked vendors for a 2021 retrospective. After normalization, we then tried to find some majors topics and patterns identified in these 26 answers :

1. Without any surprise, 2021 was the year of **Ransomware**. Therefore all aspects of **Cybersecurity** are by far #1 even from storage vendors
2. Thus **Cloud** arrives #2 in various flavors (multi, hybrid, private, edge, SaaS and Operating Model)
3. **Data Protection** in all dimensions
4. and then at the same level: **NVMe** (oF), **PCIe/CXL**, **Container/Kubernetes** and **ESG (Environmental, Social and Governance)**

ATTO (Tim Klein, president and CEO)

Fibre Channel is thriving: Mainstream Fibre Channel products have been with us for nearly 30 years and despite some declaring it a dying connectivity technology, activity in the Fibre Channel market throughout 2021 has been extremely healthy. Fibre Channel continually lives up to its well-earned reputation of reliability and stability with predictable high-performance and a flexible protocol stack. In addition to more traditional deployments, Fibre Channel is proving to be very effective in networked NVMe storage. NVMe over Fibre Channel could become a strong alternative to NVMe over Fabrics as system builders look for the best way to utilize NVMe over networks.

Availability of PCIe 4: PCI Express 4.0 is the first major change to the PCIe standard in over a decade, delivering approximately twice as much bandwidth compared with previous generations. Manufacturers of GPUs, NVMe storage, motherboards, servers and adapters are ready to take full advantage of this additional bandwidth to improve data access capabilities and boost IOPs.

Availability of LTO-9: 2021 saw the availability of LTO-9 tape drives which improved upon prior generations with 50% greater storage capacity and faster transfer speeds. Tape storage remains a key element of the modern data centre due to cost and simplicity, and also because of its natural ‘air gap’. Tape backups are physically disconnected from the network and therefore keep an organization’s backup data safe from not only cyber threats but infrastructure failures as well.

***

Catalogic Software (Mike Miracle, Chief Strategy Officer)

Ransomware Now Attacks Backup Datasets - Ransomware attacks have become so sophisticated that they do not stop at exfiltrating your data and then encrypting it on your primary systems like your filers. The cyber-criminals now take their time to search for and compromise all your secondary storage copies and delete your backups so that no data recovery is possible from the backups.

3-2-1 Backup Rule still Rules - The 3-2-1 backup rule is a time-honored strategy for data protection that states that your business should have at least three (3) copies of your data, on two (2) different storage media types, with one (1) of the copies offsite or in the cloud that is locked and immutable. The cyber-attacker should not be able to find or delete that offsite copy, so you will be able to get your data back. However, you really should verify that off-site copy first; therefore, the rule is enhanced to 3-2-1-1, which now includes one (1) copy that is verified to good and work for recovery.

Backup to the Cloud is the Last Line of Defense - Given our inability to stop ransomware attacks, backup and recovery is now the last line of defense when it comes to them. If your backups reside on the same network or the same storage systems as your production data resides, they are also vulnerable to attack. Having backups air-gapped and locked in the cloud, that the ransomware attacker cannot reach, ensures that your data can be recovered.

***

Cloudian (Jon Toor, CMO)

Security experts got ransomware protection wrong: Many security experts continued to focus on updating perimeter defenses and anti-phishing training as the key to ransomware protection.

However, even fully updated perimeter defenses were ineffective in keeping ransomware out, often circumvented through increasingly sophisticated phishing emails. In fact, a survey of ransomware victims earlier this year found that 65% of organizations that were penetrated through ransomware had conducted anti-phishing training beforehand. Fortunately, with conventional strategies failing, organizations began to recognize the need to protect data at the storage layer with an immutable backup copy, enabling quick recovery from an attack without having to pay ransom. Self-managing storage in data centers became mainstream. Automation expanded as a critical component of storage systems to replicate data for disaster recovery, manage copies of data, monitor hardware for potential failures and proactively initiate replacement tickets. Enterprises increasingly leveraged automation to reduce outages and disruptions with predictive maintenance – helping them save costs, enhance security and adapt to evolving workload needs.

Object storage expanded to new use cases: Object storage has traditionally been known as a solution for backing up and archiving data. But that began to change in 2021 as organizations recognized the technology’s effectiveness for supporting other use cases. With its high-performance capabilities, flash-based object storage gained favor for compute-heavy workloads that also have high-capacity requirements, such as AI, ML and data analytics workloads. In addition, enterprises increasingly employed S3-compatible object storage to support modern apps, due to its ability to simplify Kubernetes deployments. This included greater use of on-prem object storage for modern application use cases that demand local data residency and low-latency data access in sectors such as financial services, healthcare, media and entertainment, telecommunications and government.

***

CTERA Networks (Aron Brand, CTO)

Costly Ransomware attacks have exposed basic misconceptions about data protection
Ransomware attacks have reached epidemic proportions, as criminals refine their techniques to target the most valuable data and extract higher payouts. In June 2019, two Florida municipalities paid a total of $1.1 million to cybercriminals to regain access to their IT systems and data following ransomware attacks. In May, the city of Baltimore was hit by a similar attack and decided not to pay the ransom. The result: after 36 days of remediation efforts at a cost of $18 million, the municipality's systems were still not fully restored.

The truth is that many people don't really grasp the concept of backup, and this lack of understanding could end up costing them a bundle. Effective protection must meet two critical requirements:

- It must retain previous versions of your files for a specific retention period (minimum of 30 days), and those files must be in a read-only repository so that they cannot be deleted by a malicious software.
- The archived copy must be physically separated from the main copy of your data.

Hybrid Edge to Cloud
As the 2010s come to an end, virtually all large distributed enterprise companies are in the process of modernizing their networks with hybrid cloud solutions that combine local computing for latency and downtime sensitive applications, backed by infrastructural services hosted in a public or private cloud.

Organizations are deploying cloud storage gateways to solve a wide variety of ROBO IT challenges
According to a recent IDC survey, 91 percent of enterprises have deployed or are planning to deploy a cloud storage gateway at their remote sites to reduce costs, centrally manage users and data, and consolidate infrastructure at the edge.

***

DDN (James Coomer, SVP of Products)

Enterprise organizations accepting technologies that scale over conventional Enterprise systems for AI, analytics and advanced computing (storage (True Parallel File Systems), networking (Infiniband) and compute (GPU)).

The Data Centric Architectural approach gets huge boost due to the promise of AI and impacts organizational approaches to IT – AI forces organizations to consider de-siloization of large scale data assets for manageability, governance, security, performance and cost savings.

Rapid maturation of AI with the entrance of additional AI compute players into the AI ecosystem (Intel, SambaNova, Graphcore, etc) to compete with established leader NVIDIA.

***

Fujifilm Recording Media (Rich Gadomski, Head of Tape Evangelism)

Constrained Storage Budgets
COVID did not go away, nor did repurposed IT budgets focused on WFH connectivity and compute infrastructure. This put more budgetary pressure on storage where low cost tape relieved the pressure. Nowhere was this more apparent than in digital preservation and high performance computing environments with a simple need to offload expensive object storage to cost-effective tape systems using an S3-compatible API.

Cybersecurity Concerns
Ransomware did not go away either and became a top concern for C suite executives in the wake of increasingly high profile victims. The FBI and CISA weighed in with this advice: “Backup your data, system images, and configurations, test your backups, and keep backups offline”. Tape’s inherent ability to provide air gap security in offline, offsite locations proved to be an enduring value proposition.

Global Warming
Climate change certainly did not go away with one climate related natural disaster after another. Curbing CO2 emissions became another C suite imperative and storage did not escape the scrutiny. Reducing CO2 emissions by 95% all while reducing TCO became a compelling reason to move cold, inactive data from energy intensive primary storage to energy efficient automated tape systems.

***

Index Engines (Jim McGann, VP Marketing & Business Development)

Spurred on by 2020's worldwide pandemic-related economic shutdown, cyber criminals emerged from all corners of the globe to wreak havoc on critical businesses, corrupt data and demand record-breaking ransoms.

New Year's did not ring out with the old. The success cyber criminals saw in 2020 drove record-breaking ransomware attacks in 2021. These attacks got bigger, smarter and more expensive for businesses.

Here are 5 key takeaways we learned about ransomware in 2021 and one ominous forecast for 2022:

- **Critical Resources Make Prime Targets**
  Covid 19 put a huge strain on hospitals and other critical infrastructures. Cyber criminals took advantage of this, knowing many of these organizations would have to pay quickly. Attacks on healthcare especially skyrocketed, with more than 1 in 3 health care organizations globally reporting being hit by ransomware in 2020, according to AAMC.org, adding the sector experienced a 45% uptick just since November 2020. Organizations like FIN12 are shutting down systems, eliminating access to patient records, radiology imaging and other functions until a ransom is paid. And beyond hospitals, the Colonial Pipeline and JBS attacks showed cyber criminals can hinder supplies of gas and food with a few clicks.
- **FBI Interest Did Not Dissuade Big Attacks**
  FBI alerts warned of new and impending attacks, clawed back Bitcoin ransoms and sought to disable cyber criminals worldwide. This did little to dissuade cyber criminals from executing bigger, flashier attacks and demanding higher paydays. JBS Meats, Colonial Pipeline, Air India and CWT Global made massive headlines and drew record-breaking ransoms. The Colonial Pipeline drew so much attention, REvil temporarily disappeared but resurfaced, more cunning than ever, in the fall.
- **Cyber Criminals Launch New Attack Vectors**
  As organizations tried to stay in front of cyber criminals, cyber criminals evolved their ransomware attacks, bypassing common detection methods. For years, cyber criminals corrupted data in the same number of ways. As security tools started finding those basic methods, cyber criminals added new approaches. Lockfile ransomware was brought to light this past July, doing something unique in the field of ransomware, "intermittent encryption." This method evades detection of many metadata analytics tools. Other attack vectors also cause significant destruction while avoiding detection, include

Jigsaw (encryption combined with a progressive deletion) and CrypMIC, which corrupts files without changing the extension.

- **Backup has a Bullseye on It**
  Cyber criminals are trying to do as much damage as possible to make organizations as desperate as possible and demand as much money as possible. Disabling, erasing and encrypting backups will make it near impossible to rebuild and recovery without paying ransom, making backup an increasingly common focus for big ransomware attacks. REvil and Conti, a Ransomware-as-a-Service organization led the way, targeting backups including turning off popular backup applications, driving a massive shift towards more intelligent cyber protection solutions.
- **Recovery is Taking Longer**
  Many organizations found that their disaster recovery process did not extend to cyber recovery so attempts to recover grew longer, causing more economic hardship. Average down time is now 23 days, up by two days in 2021. But some organizations take months to get back to normal, especially if they were just relying on disaster recovery backups. Tulsa mayor GT Bynum saw his city attacked mid-April and remarked, “We’re on path to have all our city services restored back to normal by mid-September. That is the goal for us which is six weeks earlier than we were initially projecting. The city has hundreds of different computer systems that our team and contractors and law enforcement personnel that are having to go through and make sure they’re clean before restoring them.”

***

Infinidat (Eric Herzog, CMO)

News of CEO Concerns
In 2021, we saw a much broader realization about the critical nature of data and cyber resilience from the storage estate to be part of the corporate enterprise cybersecurity strategy. The threat of cyberattacks reached such a high pitch that in the Fortune 500 survey of CEOs in mid-2021, 66% of Fortune CEOs said the #1 threat to their businesses is cybersecurity. Similarly, in a KPMG CEO survey in March 2021, CEOs also said cybersecurity was their #1 concern.

2021 Trend to be Cyber Secure
Alarmingly, the average number of days to identify and contain a data breach, according to security analysts, is 287 days. The trend in 2021 has been to ensure a company’s storage estate is cyber secure, aiming to thwart ransomware, malware, internal cyber threats and other potential attacks. It has forced changes in the way companies approach cybersecurity and storage, with implications for 2022.

Adoption of More Advanced Cyber-Resilient Storage Capabilities
In 2021, the increased adoption of advanced capabilities such as logical and physical air gapping, real-time data encryption, and instantaneous recovery, among others, reflected efforts across the enterprise storage market to protect primary and secondary storage. A cyber resilience solution is deemed effective when it provides guaranteed availability and a fully scaled data restoration for business continuity. We anticipate to see an increase in this trend in the new year.

***

Komprise (Steve Pruchniewski, Director of Product Marketing)

File in the cloud won't look like file on-premises
With mature cloud-based file offerings from all the major vendors, it's never been easier to recreate your NAS architecture in the cloud. However: caution abounds. Customers who make simple lift and shift from on-prem to cloud will find higher costs and they will be constrained by the same scale limits with the only benefit being reduced administration.

Unlike on-premises infrastructure where customers could only leverage the technology they had physically deployed, the cloud presents a practically unlimited array of continually evolving options. To take advantage of storage tiers and protocols and services such as ML tools and cloud data lakes, organizations will move granular workloads versus storage volumes as they would have done when moving from an old storage array to a new array in the data center. The Great Cloud Migration will change customers from storage administrators to data architects; from managing infrastructure to generating insight. This data-first approach will provide the mobility to select the right resources at the right time.

Ransomware hits the backup and storage marketplace with vendors scrambling to deliver new products/services to protect and recover data in storage. The pervasive ransomware threat has expanded the discussion outside of security teams to all parts of infrastructure including storage. Having a solid ransomware story will be table stakes for all data and storage vendors in 2022.

Strategies will focus increasingly on using immutable object-lock storage in the cloud as a way to protect data from ransomware attack and enable quick restoration of the data when needed. Organizations which have a detailed understanding of their data – age and usage to start – will benefit from only backing up hot data and moving cold data to a secure, off-site location in the cloud. This will also reduce the backup

footprint, another critical strategy since ransomware attackers are targeting backup systems.

Fast track for cloud data analytics and cloud security companies
Databricks and Lacework were two of the top recipients of VC in the last year (a combined $2.9 billion) as cloud data lakes, cloud data warehouses, and the need for cloud security were top priorities in 2021. IT leaders know they need to modernize their data infrastructure and adopt data lakehouse, data mesh and data fabric strategies to compete. The need to ensure security and access while bringing the right unstructured file and object data into these data platforms will gain greater traction in 2022.

Growing market for smaller cloud storage and infrastructure players: Cloudflare, Wasabi, Backblaze
The pandemic accelerated the adoption of cloud services and it wasn't just the big players that benefited. Smaller cloud storage and infrastructure providers like Cloudflare, Wasabi, DigitalOcean, and Backblaze saw exponential growth. Enterprise and SMBs which began their cloud journey with the major vendors are exploring options with these smaller vendors for lower cost and specialized offerings. Smaller cloud players are introducing new offerings to compete with the so-called hyperscalers and expanding to new regions outside of North America. The year 2022 will see a move from reactionary adoption of the Big 3 clouds (in response to the pandemic) to a more nuanced strategy including smaller cloud vendors for distinct use cases.

***

MinIO (AB Periasamy, Founder and CEO)

Object continues to establish itself as a distinct storage platform.

The concept of a storage product that combines both file and object is antiquated. The object storage space is evolving much faster than file, and any product that tries to continue to serve both masters is doomed to fail on both.

Kubernetes is the way forward and software defined object storage is the natural partner - with RESTful APIs and simple containerization. Trying to patch file's inherent weaknesses when it comes to Kubernetes is a losing proposition and serves as a drag on development.

Further, dedicated object storage vendors are delivering performance that enables them to run all but the most latency-sensitive applications picking up traditional SAN and NAS workloads like databases along the way.

The future is best of breed and it will require specialization. Those who specialize in object storage are positioned to succeed.

***

Model9 (Gil Peleg, Founder and CEO)

Hybrid-cloud models became the standard
This year we saw more and more companies adopting the hybrid cloud model, allowing them to store and manage data in the cloud while continuing to maintain their core systems in the mainframe. This resulted in more movement of data from the mainframe to the cloud, with the result that now more companies have begun to operate on the data in the cloud.

New technologies made it easier to modernize mainframe infrastructure
Companies began to adopt new technologies that are making data movement much easier than in the past—especially when it comes to modernizing mainframe infrastructure. Innovative techniques using zIIP processors and TCP/IP, and leveraging compression and parallelism, have increased throughput between the mainframe and the cloud. This has made data movement—which used to be very challenging in the past—very easy to accomplish today.

Mainframe data silos began to break down
In the past, the inefficiencies in cloud data management caused mainframe data to be siloed. This year, we began to see those silos break down significantly. Companies are now using data storage in the cloud, allowing them to better utilize all of their data in an aggregated manner, unlike in the past when data remained on the platform where it was collected.

This allowed companies to extend their data center to the cloud with a hybrid approach. In turn, we saw a significant increase in the accessibility of data for AI/BI/ML applications, and a corresponding realization of just how valuable this mission-critical data is. In addition, cloud storage allowed companies to use storage tiers to better manage their long-term storage, and also to more efficiently access storage in the cloud for all use cases.

***

Nakivo (Sergei Serdyuk, VP of Product Management)

The rise of ransomware
2021 saw major cyberattacks hitting medium, large businesses and government agencies around the world. Major cyber attacks took place every month since February 2021, with attackers demanding millions of dollars from their victims and ransom demands breaking records. According to Statista, 68% of organizations were affected by ransomware to some extent. So far, 2021 has seen the highest number of attacks compared to previous years.

The on-going adoption of cloud
2021 saw the cloud, mainly related to the enterprise sector, enjoying strong growth. Enterprise cloud storage, which just a decade ago was looked down upon as unreliable and risky, now offers a range of data protection features on par with dedicated solutions. This is not to say that hardware is taking a back seat - the recent ransomware surge may well mean that niche solutions like detachable drives may become a mainstay incorporating data protection strategies.

Global adoption of data-generating technologies
Perhaps anticlimactically, the biggest factor affecting the data storage market has been the on-going adoption of data-generating technologies throughout the world. Many industries have been waking up to the benefits of IoT devices, and there's no reason to expect this trend to reverse anytime soon. At the same time, the past year saw regulatory agencies worldwide recognise the importance of data protection and put forward increasingly strict requirements, driving up the demand for reliable backup management.

***

Nasuni (Russ Kennedy, Chief Product Officer)

Companies will modernize, moving workloads to the cloud
A surprise in 2021 was companies' reluctance to agree on strategies for moving workloads and data to the cloud. Despite teams being forced to work from home, many in the industry had predicted that IT teams would send workloads and data to the cloud more quickly. Cloud storage providers are breaking through 'traditional' thinking on data management and organizations are modernizing with strategies like hybrid approaches - where certain critical assets stay on-premise while other assets/workloads move to the cloud. Perversely, increased supply chain disruption is likely to accelerate the adoption of cloud and managed services – companies facing longer lead times for their on-premise storage hardware orders are having to adapt to cloud service and managed services anyway.

Competition intensifies for global hyper scalers in 2022
The main three hyper-scalers will continue to fight for mindshare and market share – they will gain further market share from on-premise vendors. However, as enterprises are moving more processing and operations to the public cloud, the big three providers know they will fight amongst themselves certainly for new capabilities and offerings – especially over price: organizations realize that a multi-cloud strategy is a good hedging against lock-in with one major cloud provider and vendors will have to become more competitive with new services and package options

IT roles changing as more enterprises move their infrastructure to the cloud
Enterprises will need to be 'cloud savvy' and grasp how to leverage cloud to better deliver IT services to their end users. Organizations will become much more driven by virtual teams and how teams operate collectively. The pandemic has effectively forced organizations to become more virtual and collaborative. The industry will need to grow talent and educate their workforces and especially, young people entering the workforce, on the skills, tools and knowledge needed to be team-oriented and productive in this connected technology driven world. A key skills trend is data engineering, ensuring that data is in the proper location and format to be useful from an AI, machine learning and other analytical viewpoints. So, the data scientist skill set will continue to be prominent and important in organizations, regardless of the industry sector.

***

NGD Systems (Scott Shadley, VP Marketing)

PCIe Gen4 became a standard, but still does not solve the bottleneck in server compute/networking. If you map lanes to throughput and NICs you have bottlenecks of data that are still challenged to do meaningful work with raw data.

Computational Storage has gained more traction, but is still on the edge of growth. With big organizations spending a lot of effort to move the technology forward, including the first PMCSS from SNIA.

Storage Capacity is growing faster than even the massive HDDs can keep up. With the growth of data, and the needs for more capacity, the high-capacity SSDs (both TLC and QLC) are going to become mandatory to just hold the data.

***

Panasas (Todd Ruff, VP Marketing)

A surge of AI/ML workloads in the HPC space due to Enterprise adoption.

Storage class memory was marginalized: Optane ended, but few NVDIMM or CXL choices became available as replacements

COVID affected time and space: component delays, revenue shifts across verticals, and professional adjustments to work-life balance.

***

Point Software and Systems (Thomas Thalmann, CEO)

Tape-based S3 object storage for backup of HDD-based object storage
Many customers have realized that their HDD-based object storage and their cloud storage also require a backup. Object storage systems offer high availability and thanks to erasure coding, object storage can withstand the failure of hard drives, nodes, racks, or even entire data centers. But HDD-based object storage systems and cloud storage are not immune to ransomware or human error. For this reason, many companies have added a tape-based object storage to backup their HDD-based object storage in native format.

Data Management Software to handle data growth and compliance requirements
Data Management Software has evolved as a solution to solve the increasing problems caused by the growth of unstructured data and is used to fulfill archiving requirements. It integrates multiple storage tiers, platforms, and locations (on-prem and off-prem). Additionally, data management software migrates data between storage tiers and technologies without interruption of operation.

Data security as part of data storage
The threat of malicious software such as ransomware has made evident the importance of the issue of data security and its inseparability from data storage. Storage features like immutable objects or object locks have received a lot of attention. Also “air gapping” has increasingly come into focus.

***

Pure Storage (Ajay Singh, Chief Product Officer)

Adoption of the cloud operating model everywhere
Increasingly customers want a 'cloud' experience from their private data centres and on-prem infrastructure. The main attributes of this cloud operating model are storage automation, simplicity of management, infrastructure delivered as-code and an ability to procure and consume storage services primarily based on attributes and SLAs around performance, service classes and tiers, resilience, reliability, and actual usage. What is driving this? Flexibility, agility, and speed to access are now more important than just reliability and performance alone. Organizations are modernizing their infrastructure platforms and operations to support agility and accelerate innovation and they want to avoid heavy upfront costs and expensive overprovisioning -- paying only for what they use. Ultimately they want an infrastructure experience that enables a single and seamless model across private and public clouds.

The rise of modern cloud-native applications
In addition to modernizing their infrastructure, customers also want to modernize their applications. Companies across all industries are finding new ways of collecting, processing and managing their data, and are turning to modern application architectures. The new wave of cloud-native applications, built on containers and managed by Kubernetes, now represents 90% of all new enterprise development efforts. These new architectures enable application portability - such as on prem to cloud, and between clouds. They also provide agile development, with access to open source ecosystems and a more flexible deployment process.

The shift to modernizing today's infrastructure with all-flash
Over the last few years all-flash has revolutionized storage. However, due to price economics organizations have been forced to distinguish between performance-oriented workloads and everything else. Too many companies are trying to modernize fast, but are stuck on antiquated disk-based or hybrid technologies, delaying their ability to transform. QLC-based systems, which have been optimized to leverage gains in efficiency and reliability through sophisticated software, can beat most hybrid disk arrays on performance, density, reliability and price. QLC makes it possible to deliver up to 10 times the performance, density and reliability of the disk systems it's replacing. This makes all-flash accessible for use cases previously relegated to spinning disk or inefficient hybrid solutions, like backup and data protection, test/dev environments, and workload consolidation. And, as flash continues its price-performance improvements relative to magnetic disk, we will witness more and more private and public data centers transitioning to all-flash.

***

Quantum (Tim Sherbak & Diana Salazar, Product Marketing Manager)

Companies struggled to find creative ways to cost-effectively storage, manage, and extract value from its data.

Organizations looked for storage solutions that could be deployed in-house in its own data center, colocation facility, or hosted IT environment.

Enterprises realized the value of tape and no longer looked at it exclusively as an offline storage resource.

The majority of data is unstructured and remains a huge target for ransomware and cybercriminals.

Backup applications that provide immutability are the best method for users to recover from ransomware and get back online quickly.

Companies who understand their data lifecycle and leverage its value are more likely to meet their business goals.

***

Qumulo (*Ben Gitenstein, VP of Products and Solutions)*

“Cold data” is diminishing as increased adoption of active archive technologies keeps data “hot.”
Storage used to be more static, meaning enterprises could throw all of their data in a server and never touch it again. In 2021, we saw that change as customers leveraged more data-intensive workflows (i.e. medical image processing and predictive analytics) that demand huge amounts of available data at all times. Instead of leaving data alone in cold storage, or spending an unreasonable amount of money and computing power to warm it up again, enterprises are increasingly adopting storage technologies that allow you to access all of your data, all of the time. Customer and IT teams are relying on active archive data solutions, so that even “cold” data can still be easily brought back to life.

Multi cloud solutions were at the forefront in 2021, and agility is the priority.
Flexera’s 2021 State of the Cloud Report states that 92% of enterprises reported having a multi-cloud strategy in 2021. A multi cloud solution gives you greater flexibility, more technical control, and better security since you can choose dedicated servers and networks that can restrict access controls. The rise in ransomware

attacks in 2021 taught us to always have a backup plan, and multi cloud enables you to replicate your data in the cloud and failover to it when needed.

Tech companies finally understand that sustainability is a good investment.
We know that this year brought climate change into a stark reality with extreme weather events like hurricanes, droughts and floods. Accenture surveyed over 1,200 CEOs in 2021 and nearly half reported that their business is grappling with supply-chain interruptions because of these extreme weather events. Forrester also reported that empowered customers are increasingly demanding sustainability solutions from their vendors - and the tech industry is finally answering.

***

Robin.io (W. Brooke Frischemeier, Head of Product Management)

Acceleration of stateful workloads on Kubernetes.

Environmental application+storage portability will trend upwards as we see customers exploring multi-cloud solutions as well as undergo cloud repatriation.

As NVMe becomes more mature and affordable, NVMe-based workloads will become more prevalent.

The first trend we see when speaking to our customers and answering inquiry requests is the acceleration of stateful workloads deployed on Kubernetes. While Kubernetes is well known for its scalability, resource utilization and workload orchestration, traditionally it was not seen as the ideal choice for data persistency and stateful workloads. Fortunately, as Kubernetes has matured, there has been considerable innovation contributing to the trend towards stateful applications deployed using over Kubernetes.

In the past, before cloud-native and Kubernetes, there was a simple relationship between applications/virtual machines and their related storage. There were direct and simple host-based relationships. Therefore, when it came to backups, snapshots cloning and disaster recovery (DR), backing up storage alone was good enough and simple.

With cloud-native, when it comes to data protection, the relationship between storage and application becomes more complex. With cloud-native, the application is no longer a homogenous blob as before. All of its roles become exposed and are further broken out into multiple containers, each with a life of its own. On top of this, there are multiple forms of data, application config, Kubernetes config, secrets, metadata,

ConfigMaps and etc., each of which has a different relationship and requirements to/from the application. Thus, just backing up storage devices or persistent volumes alone, does not solve the complexity or RTO associated with application+storage protection and recovery.

As customers scale out their data centers and span multiple clouds, security becomes even more important. Again, customers are pointing to ease of use. Security needs to be as simple as telling the orchestrator, or cloud storage application, to switch it on. Under the covers where there will be volume encryption, security key management and data transfer encryption. But to the user this needs to be completely transparent and dynamically governed by policies.

Furthermore, as multi-tenancy and Roles Based Access (RBAC) becomes more prevalent, it needs to gracefully roll into existing customer security mechanisms such as Lightweight Directory Access Protocol (LDAP).

Further down the security front we see many customers asking for anti-ransomware solutions. Since data is always being backed up this isn't always a huge problem for storage itself. But what if the application as well as the storage is taken over. With an application+storage solution, customers can roll back the entire platform, to any point in time, at the click of a button.

Last, we see Non Volatile Memory express (NVMe) devices becoming more mainstream, where the lower cost is accelerating its availability and prevalence. Hence, there will be more cloud offerings with NVMe devices. With NVMe-based solutions, we frequently see customers looking for a combination of read write to many (RWX) + Graphics Processing Units (GPU) + NVMe to the cloud, where common data is acted upon by multiple worker applications, culminating to storage over NFS. As the inclusion of GPU likely gives away, we see these features tied together in use cases involving Artificial Intelligence / Machine Learning (AI/ML) use cases.

***

Seagate (Ted Deffenbaugh, VP of Cloud Product Line Management and Market Research)

Unprecedented demand for mass-capacity hard disk drive (HDD) storage is leading to **yet another year of record-breaking shipments for hard drive exabytes**. In 2020, the hard drive industry entered the one zettabyte (ZB) milestone by shipping just over 1ZB worth of storage. In 2021 alone, in only the space of the first three calendar

quarters, the industry has already shipped over 1ZB, showing that by the end of 2021 we were on track to smash the record shipped in 2020.

Decentralized, blockchain-based use cases proliferated. Among them was Chia, the storage-based crypto currency considered one of the greener currencies. This has added to the demand for capacious data storage as some HODLed their hard drives.

Decentralized storage emerged to support the new Web 3.0 architectures. This led to new uses for hard drives: Filecoin is a prime example of an exciting, nascent application that became an early candidate for the storage of Non-Fungible Tokens (NFTs) for the Metaverse.

The hard drive industry showed it has a viable and vibrant forward-technology path. Seagate commercialized HAMR and started to ship and sample drives to customers. In addition, the native NVMe HDDs got closer to reality. Seagate unveiled the industry-first demo at OCP of the technology, adding that productization was now closer on the horizon, but still a few years off.

***

SIOS Technology (Cassius Rhue, VP Customer Experience)

We saw increased expansion both into cloud and within the cloud.
Many more companies moved to cloud for the first time while others expanded their cloud ages by bringing mission critical applications that require HA, such as SAP, HANA, and SQL Server, as well as critical data into their cloud environments.

We saw more companies expanding HA, DR and security measures in the cloud.
With a bevy of attacks to availability in 2021, companies moved quickly to secure their mission critical applications against cyber attacks, natural disasters (fires, floods, storms, etc.), and against data center and availability failures.

Simply backing up data storage is no longer sufficient for IT teams given climate change and natural disaster threats which seem to have increased in 2021.
Whether the storage is NFS, SAN, cloud-native shared storage, or replicated local storage, companies need to implement a more sophisticated way to handle DR. We are seeing an increase in protection levels of data storage for large businesses – both on premises and in the cloud – to include high availability and disaster protection.

***

SoftIron (Phil Straw, CEO)

Shift in the “Power Dynamic” between Enterprises and the Public Cloud
Enterprises are becoming increasingly unhappy with the power dynamic in the relationship they have with the big public cloud vendors. Sprawling egress fees and penalties that stop true mobility of data between cloud providers is making them look hard at what data is held “on-prem” and what is shared to the cloud -and at the same timer looking for alternate strategies that can bring these relationships back into more balance.

Long-lasting consequences of the pandemic on remote work, especially for technology workers
The realization by a growing number of employees (especially in the tech sector) that they can be just as productive (or more) without eating the cost of 1 hour daily one-way commute into an office provided fuel for the great resignation. Companies large and small needed to implement new remote work policies to accommodate expected flexibility, and data storage owners had to manage keeping infrastructure running longer with fewer (or less skilled, or less contextually aware) people around to do hands-on work.

Emergence of Optimized Edge Storage
Driven both by high performance workloads such as AI and machine learning, along with shifting work patterns favouring a more distributed IT architecture, 2021 saw the emergence of the first iterations of a class of storage solutions able to handle the unique challenges of delivering the performance required in within the physical (power/ cooling) and manageability (no skilled on-site IT).

Increasing recognition of the value of enhanced security postures in critical infrastructure
In 2021, we saw a combination of high visibility ransomware attacks across critical national infrastructure (oil pipelines, hospitals, etc), multiple open source supply chain hacks resulting in lost or stolen data, the world has reacted by applying more scrutiny on the design, delivery, and operations of IT infrastructure with an eyes on both prevention (securing supply chains) and improving resilience (immutable backups.)

***

StorCentric (Mihir Shah, CEO - Surya Varanasi, CTO & JG Heithcock, GM, Retrospect)

Cybercriminals and ransomware are evolving: from hitting only single organizations and/or individuals to attacking MSPs, where they can target multiple organizations with one fell swoop (e.g., Kaseya ransomware attack perpetrated by the REvil group).

Cyber insurance became increasingly critical: and it wasn't just for large enterprises anymore. Small and medium sized enterprises invested, many for the very first time. Yet, confusion and frustration over what it does and does not cover continues.

Enterprises recognized the need to protect themselves against a ransomware-related class action lawsuit: and began preparations for a worst-case scenario. Enterprises also increased their focus on data protection, particularly PII, as well as their ability to demonstrate that every possible precaution was taken to prevent and recover from an attack.

Unbreakable Backup became an indispensable solution for ransomware attack recovery: thwarting cyber criminals 'attack the backups first' strategy.

Backup copy immutability became non-negotiable: meaning at least one backup copy must be immutable, unable to be deleted, corrupted or changed in any way, even if the ransomware has already infiltrated your organization, and integrated itself into the backup process.

Ransomware as a service (RaaS) is a huge business - with attacks continuing to grow at an alarming pace. Businesses at every size are increasingly exposed to ransomware attacks.

Cyber criminals are attacking backups first: and then once under their control, coming after production data. This means that many enterprises are feeling a false sense of security, until it is already too late.

Recovery capabilities became the #1 ransomware strategy: while prevention and detection remained indispensable, recovery capabilities became the top priority.

***

Tiger Technology (Nikola Apostolov, Head of Business Development)

Hybrid cloud infrastructure gaining momentum

Hybrid cloud is becoming a trendy approach with latency being one of the biggest challenges.

Post-pandemic times increased the need for cloud services to meet objectives that cannot be achieved with on-premises tools. The modernization of several verticals was also forced to follow due to many reasons, from the need to optimize processes and cost to meeting the workflow requirements of the new reality.

Video surveillance/public safety, healthcare, construction and engineering and media and entertainment are a fine example. Modernization in these industries is associated with a massive growth of data that has led to unbearable costs. In addition, early adopters that started this journey one or two years ago are now suffering from latency, lack of functionality and/or inability to connect to an ever-growing number of cloud services (cognitive, analytics, BI etc.) as they are using vendor lock-in solutions. Low-cost tiers in the cloud are used for large data workflows much more, which has increased the need for an automated process that is less disruptive. Not surprisingly, tape companies have started to adopt this and are trying to create similar services.

All this has led to the advancement of hybrid (on-prem/cloud) architectures.

Increasing demand for vendor lock-In free solutions

Avoiding vendor lock-in is a major consideration in building a future-proof architecture.

What's new this year is that many organizations with mission-critical workflows that are early adopters are looking for better ways to fulfill their requirements while those with new cloud installations are staying away from existing vendor-locked solutions. Solutions that involve proprietary data formats and expensive appliances are no longer among the most preferred ones. One of the factors of course is cost but a much more important consideration is finding ways to design future-proof architecture that would allow the use of any service running in the cloud to augment existing workflows.

It is an interesting fact that encryption at rest, deduplication/block-based approaches, and traditional back-ups involve solutions that lock data to a certain vendor. While traditional backup methods are far too commonly used to be a serious consideration, this is a clear disadvantage when encryption and deduplication are in question. To put this into perspective, once considered leading solutions in hybrid technology, StorSimple and TwinStrata have been announced to reach EOL in 2022. The way to leverage the cloud is file-based.

Compliancy and security as a decisive factor in digital transformation

Every cloud transformation depends on existing regulations and security concerns. Some customers are solving this by adopting a private/gov cloud or confidential computing, but some are using public cloud that requires thorough compliancy and security checks. The fewer changes the solution is causing to existing security and compliancy, the better.

Migration

Migration is an important area of interest to cloud vendors, but the important question for organizations is “what should we migrate to?”.

For most non-mission-critical workflows aiming at simple access\file share this is not a serious problem. For mission critical workflows, however, the problem remains as most migration software solutions that exist barely manage to address that.

***

Toshiba Electronics Europe (Rainer W. Kaese, Senior Manager Business Development, Storage Products Division)

The hard disk drive (HDD) market is continuing to grow substantially, against all expectations. It remains the dominant data storage technology, despite solid-state drives (SSDs) taking over as the preferred storage medium in some applications. The reason for this is the immense increases being seen in data storage demand, alongside the need for a low cost per capacity. Only through HDDs is it possible for current demands to actually be met. While SSDs are used in local server storage, boot media and as working storage, active online mass storage relies on HDD as the storage component (due to its ability to be produced cost-effectively in ultra-large volumes). They are being incorporated into nearline enterprise, surveillance, NAS, desktop and external USB connected drives.

Innovations in HDD technology are further underlining its long-term value as a storage medium. Toshiba has pushed HDD capacity beyond the 16TB barrier, by introducing the world’s first 18TB HDD to be based on flux-control microwave-assisted magnetic recording (FC-MAMR) technology.

SAS vs. SATA as the HDD interface technology of choice is still an open-ended topic for debate. While the backplane/expander/JBOD infrastructure is now fully SAS (12GB/s) compatible, the storage component’s interface technology in operation remains two-fold. Nearline SAS drives are utilized for applications that require high availability (HA) and they are also still a dominant component used for passive backplanes in mid-size storage servers. Nearline SATA is often employed for direct (motherboard cable attached) mass storage in small servers and workstations, but

also as leaf components in large scale SAS-based expander/cable/JBOD infrastructure implementations.

***

Veritas Technologies (Deepak Mohan, EVP Products)

Data protection lagged behind cloud transformation, leaving businesses vulnerable to ransomware

The global COVID-19 pandemic was a catalyst for digital transformation throughout the world. However, the need to rapidly introduce new systems to support evolving business practices such as remote work, contactless interaction and providing consumers with online everything meant that IT departments were often forced to prioritize the delivery of functionality over security. This introduced a thunder-and-lightning effect, where we first saw the lightning flash of innovation and then had to wait for the thunderclap of protection to follow. The intervening period is the biggest window of opportunity for failure, where organizations expose themselves to ransomware, compliance failures, downtime and a myriad of other data risks. Veritas research showed that 80% of companies had implemented cloud capabilities beyond their pre-pandemic plans and 56% said that they now had gaps in their protection strategy here – more than any other area.

Everyone upped the ante on ransomware with hackers targeting critical infrastructure and governments biting back

Over the past year, ransomware attacks have increased 185% 1 [1] , with costs expected to surpass $20 billion by the end of this year. 2 [2] The stakes continue to rise, as hackers increasingly targeted critical infrastructure, as evidenced by attacks on Colonial Pipeline in the US or the Health Service Executive in Ireland. With human lives literally on the line, 2021 saw governments step in and up the ante in return. The European Union announced a strategic initiative to fight ransomware including the creation of a dedicated cyber unit. The US government also took a much stronger position, working to deny hackers the profits of their crime by retaking ransom payments. It is also now debating a law that will force companies to disclose their ransomware payments in an effort to force the issue out of the shadows where it can be more holistically addressed.

Businesses overestimated the in-built level of security offered by their SaaS providers

Cloud Providers are outstanding when it comes to ensuring the high availability of their infrastructures. But, when it comes to data loss from incidents like ransomware, many companies discovered in 2021 that they had vastly overestimated the level of responsibility that their cloud providers have for data protection. Veritas found this year that nearly all employees (92%) thought their cloud provider would be able to restore lost data from their SaaS applications. However, this is rarely the case since

most cloud companies make it clear in their terms and conditions that data loss protection is down to their customers. As a result, 52% of people working in office roles in 2021 irretrievably lost SaaS data.

[1] https://www.sonicwall.com/news/sonicwall-record-304-7-million-ransomware-attacks-eclipse-2020-global-total-in-just-6-months/

[2] https://cybersecurityventures.com/global-ransomware-damage-costs-predicted-to-reach-250-billion-usd-by-2031/