Hype Cycle for Storage Technologies, 2024

10 July 2024- ID G00814317- 74 min read

By Julia Palmer

This research explores emerging storage technologies and their impact on businesses, adoption rates and maturity level. By leveraging these insights, I&O leaders can build their strategy for delivering innovative and future-proof storage platforms that align with business needs.

Strategic Planning Assumptions

By 2028, STaaS-based consumption services will replace over 50% of traditional on-premises IT capacity management, budgeting, sourcing and fulfillment activities, up from less than 15% in 2023.

By 2028, over 50% of infrastructure and operations (I&O) leaders will upgrade or replace storage appliances with solutions aligned with the principles of enterprise storage platforms.

By 2028, 100% of storage products will include cyberstorage capabilities focused on active defense, up from 10% in 2023.

By 2028, 70% of file and object data will be deployed on a consolidated unstructured data storage platform, up from 35% in early 2023.

By 2027, 70% of I&O leaders will implement at least one hybrid cloud storage use case, up from 25% in 2023.

Analysis

What You Need to Know

The data storage market faces evolving challenges in enterprise IT. With emerging trends of exponential data growth, public cloud integration, talent acquisition hurdles, emerging workloads, cyberthreats, and the rise of generative AI, modern adaptable storage platforms will be in high demand. Consumption models are changing the way clients source and manage storage infrastructure. Generative AI presents unique challenges as it has a potential to generate large amounts of data, requiring highly performant, scalable and sustainable storage solutions. Security and privacy concerns associated with generative AI content also require robust cyberstorage platforms. Data centers are no longer the sole repositories of data, leading to a need for hybrid cloud storage capabilities.

To address these challenges, Gartner has evaluated 20 hyped storage technologies reshaping the landscape. These technologies offer higher performance, scalability, cloud integration, flexibility, and advanced data management and security to meet evolving business needs, including those posed by generative AI.

The Hype Cycle

I&O leaders responsible for enterprise storage are reacting to the changing requirements of digital business, unpredictable data growth, new workload introduction, and the desire to leverage public and hybrid cloud that include core-to-edge capabilities. This research assesses the business impact, maturity level and adoption speed of innovative storage technologies.

More than half of the technologies reviewed in the 2024 Hype Cycle are poised to mature during the next five years, and 65% of technologies have the potential to deliver major benefits. To provide clearer, more focused research to support your analysis and planning, we have only included a subset of the most innovative technologies and removed those that are well-adopted, understood or no longer relevant for I&O leaders.

In 2024, Gartner observed that certain technologies were gaining significant attention in the market. These technologies addressed important trends like hybrid cloud, container support, data management and as-a-service offerings. Both established and emerging technology providers played a role in shaping the storage market by introducing innovative technologies, services and business models that allowed storage capabilities in various locations, including edge and public cloud. Some of the innovations focused on cybersecurity, sustainability and supporting new workloads such as generative AI.

Notable fast-moving technologies in this space included cyberstorage, storage as a service and data storage management services.

Furthermore, Gartner introduced five new profiles this
year: distributed hybrid infrastructure, cold storage, infrastructure consumption services, c aptive NVMe SSD and enterprise storage platforms.

Figure 1: Hype Cycle for Storage Technologies, 2024

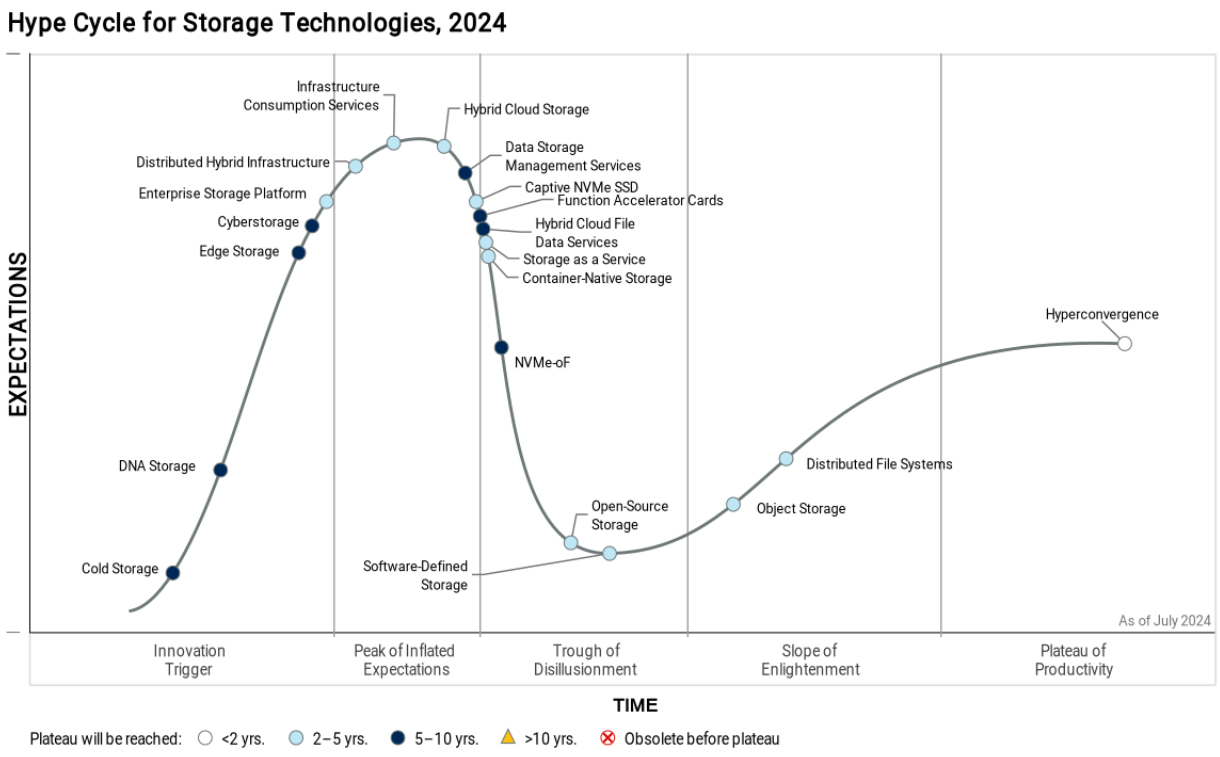

The Priority Matrix

The Priority Matrix is a useful tool for visually representing the maturity timeline of storage technologies or frameworks in a simple format. It helps answer two important questions: How much value can be gained from a specific storage technology or framework, and when will it reach a level of maturity to provide this value effectively? By mapping the maturity timeline on a grid, the Priority Matrix helps organizations assess the potential benefits and determine the best timing for adopting a storage innovation. It provides decision makers with a concise overview, enabling them to make informed choices about leveraging emerging storage technologies for their data management needs. Use the Priority Matrix to prioritize fast-moving, transformational technologies that offer high business benefits and are expected to achieve mainstream adoption quickly. After adopting these transformational technologies, consider high-impact technologies that are likely to reach mainstream adoption soon.

Organizations that have not already done so should evaluate and implement software-defined, hyperconverged and distributed hybrid infrastructure. I&O leaders should consider implementing hybrid cloud storage, edge storage, and hybrid cloud file data services to address the growing needs of data — both on-premises and in the cloud. Due to the increased focus on storage security, I&O leaders

must evaluate cyberstorage products and capabilities that are designed to detect, prevent and recover from cyberattacks, including ransomware and other cyberthreats.

Table 1: Priority Matrix for Storage Technologies Technologies, 2024

Enlarge Table

Benefit	Years to Mainstream Adoption			
	Less Than 2 Years	2 - 5 Years	5 - 10 Years	More Than 10 Years
Transformational		Distributed Hybrid Infrastructure Software-Defined Storage		
High	Hyperconvergence	Container-Native Storage Distributed File Systems Hybrid Cloud Storage Object Storage Storage as a Service	Cyberstorage Data Storage Management Services DNA Storage Edge Storage Function Accelerator Cards NVMe-oF	

Benefit	Years to Mainstream Adoption			
	Less Than 2 Years	2 - 5 Years	5 - 10 Years	More Than 10 Years
Moderate		Captive NVMe SSD Enterprise Storage Platform Infrastructure Consumption Services Open-Source Storage	Cold Storage Hybrid Cloud File Data Services	
Low				

Source: Gartner (July 2024)

Off the Hype Cycle

In 2024, the following technologies have been moved to the new dedicated Hype Cycle for Backup and Data Protection Technologies, 2024:

- Container backup
- Cloud infrastructure recovery assurance software
- Digital communications governance
- Data discovery

- Backup as a service
- Isolated recovery environment
- Immutable data vault
- Data classification

Computational storage has been removed from the Hype Cycle due to its complexity, limited practicality and lack of adoption in the industry.

On the Rise

Cold Storage

Analysis By: Chandra Mukhyala

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Definition:

Cold storage refers to systems that are purpose-built to store infrequently accessed data for long-term preservation. They can be offline systems or online systems with high-latency access. Key attributes of cold storage are durability over very long periods and low cost. With cold storage, data redundancy and data resilience are much more important than data access speed. Cold storage can reside on-premises or in the cloud.

Why This Is Important

Organizations are creating and managing more data now than any time in the past and will continue to do so for the foreseeable future. Organizations retain data for long periods of time to comply with regulatory requirements and internal compliance mandates. These two trends require an approach to storing inactive and archival data separately from production data, on the most cost-effective and durable options for long-term retention.

Business Impact

Organizations with policies to retain inactive digital data for long periods of time — i.e., 10 years or more — will benefit from purpose-built cold storage. Storing data is expensive, requiring businesses to make trade-offs when deciding what data to keep and what to get rid of. Cold storage leverages cost-effective infrastructure with very low management overhead. It preserves inactive data while ensuring immutability.

Drivers

- The digitalization of everything, along with the need to analyze all data for AI use cases, is requiring organizations to manage ever-increasing amounts of data.
- Organizations are facing more regulatory and corporate compliance mandates to store data for long periods of time. Audit, legal, risk mitigation and information preservation requirements are driving these mandates.
- Organizations want to drive new insights and new business models by applying AI to historical data. This practice will require retaining data for longer periods.
- The value of data changes over time, requiring different performance and availability characteristics. The cold data tier warrants the least-expensive storage.
- Offline or logically air-gapped cold storage provides an added level of protection for security purposes.

Obstacles

- Cold storage can become a silo when it is not integrated with end-to-end data life cycle management spanning ingestion, archival and disposition.
- Most organizations do not prioritize cold storage projects unless a storage problem becomes expensive from a cost point of view or an external authority enforces a regulatory requirement.
- The market offers limited options for cold storage, outside of tape-based storage.
- In the market, only a limited number of data storage management services work with cold storage.
- Organizations without large amounts of data and lengthy data retention requirements may see limited ROI from cold storage.
- The underlying technology for cold storage will require multiple refreshes during its lifetime, increasing the operational burden of managing cold storage.

User Recommendations

- Assess storage-as-a-service offerings, as opposed to self-managed products, to eliminate the operational burden of refreshing the underlying storage media supporting cold storage.

- Assess the 10-year total cost of ownership (TCO) of cold storage against the TCO of existing distributed storage. Account for any additional licensing required for long-term retention and immutability.
- Ensure durability of data over long periods by evaluating the technology used to detect and fix data corruption.
- Ensure that the cold storage option minimizes vendor lock-in and provides flexibility to migrate the data to an alternative solution if necessary.
- Invest in an intelligent data storage management policy manager to move data from production storage to cold storage based on high-level policies.
- Select vendors with migration services to move data from existing tape or other cold data sources.
- Use cold storage for better security and to reduce legal liability.

Sample Vendors

Cerabyte; Huawei; IBM, Microsoft; Quantum; Rimage

Gartner Recommended Reading

2024 Strategic Roadmap for Storage

DNA Storage

Analysis By: Matthew Brisse

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Definition:

DNA storage is the process of encoding and decoding binary data to and from synthesized strands of DNA. Any binary sequences can be encoded in DNA sequence, which can then be synthesized and stored. To retrieve the data, the DNA molecule must be sequenced and decoded. DNA is emerging and promises to be an attractive and sustainable medium for the archiving of data, due to its superior density, durability and sustainability.

Why This Is Important

- DNA data storage is important due to its density, durability and sustainability; it addresses long-term data retention and sustainability needs.
- DNA performs error checking and self-repair, making it an ideal data storage medium and computing platform for applications.
- DNA data storage has a long-term, favorable cost profile impact on physical data center space, carbon-dioxide emissions and the avoidance of operating expenditures (opex), by not having to refresh the technology every five to 10 years.

Business Impact

As DNA storage adoption increases, its impact could be transformational for data storage, parallel processing and computing. Use of a complete DNA ecosystem as a consumable enterprise product is likely to occur at approximately eight years for data-intensive industries that are often first movers among new technology. The domains include healthcare banking, finance, insurance, utilities and government. The defense, research and intelligence communities are the most likely to be early adopters.

Drivers

- Future DNA data storage use cases will focus on power and space-sensitive, long-term storage requirements.
- Once DNA storage is written to, synthetic strands of DNA digital data require minimal storage space and almost no power. In theory, they can be accessed for thousands (if not millions) of years by a variety of devices in a future-proof manner, with no necessity of data migration.
- Long-term operational costs are reduced, because data migrations due to technology obsolescence or data degradation will not be issues with DNA-based storage.
- The business drivers for DNA-based data storage are density, stability, durability, sustainability and low long-term operational cost.
- The world creates several hundred petabytes of new data every day, and a single gram of DNA could store all of it. One gram of DNA can store approximately 215 petabytes of data with a minimum life span of hundreds to, theoretically, thousands of years.
- DNA data storage will have a power efficiency profile that could significantly reduce the physical infrastructure space and the carbon-dioxide footprint over the life of the data.

- Data in DNA storage can endure for thousands of years and remain unchanged, free from degradation or drive failure, compared with current technologies.
- In the future, DNA data storage will be used in combination with DNA computing for extremely large, massively parallel, processing use cases.

Obstacles

- DNA technologies face many of the same challenges as any other startups early in their life cycles: speed, time to market, standards and cost.
- DNA data storage patents are likely to cross industry segments, making patents and licensing agreements challenging.
- The creation of synthetic DNA, the medium that will store data as DNA, needs to become efficient and cost-effective.
- Access speeds and throughput rates for DNA data storage must dramatically improve to compete with classical approaches. DNA self-assembly and other similar processes are essentially chemical reactions, which are much slower, compared with today's classical approaches.
- Data security and regulatory challenges will be an issue, because DNA will someday store personal, sensitive and classified materials.
- The industry needs to develop and accelerate standards associated with DNA data storage and retrieval processes.

User Recommendations

- DNA will be one option in the hierarchy of storage use cases and workloads to select from so plan/architect accordingly.
- Prepare for increased hype as technologies mature, realizing that DNA data storage is nascent. Savvy organizations will see through the hype to the practical use-case initiatives that DNA storage may offer.
- Focus on due diligence of startup companies, and align risks with the justification of use-case returns.
- Avoid long-term lock-in with early providers. Startups will emerge and fail until technologies mature, and winners and losers are identified.
- Prioritize DNA storage for early use cases, when available, focusing on write-once, read-never or write-once, read-seldom, if ever large-scale datasets.

- Evaluate DNA data storage viability by gauging when storage prices fall to three to four orders of magnitude the cost of tape archival, and when write speeds reach the megabit/second range.

Sample Vendors

CATALOG; Helixworks Technologies; Illumina; Iridia; Microsoft; Twist Bioscience

Gartner Recommended Reading

Emerging Tech Impact Radar: Compute and Storage

2024 Planning Guide for Cloud, Data Center and Edge Infrastructure

Edge Storage

Analysis By: Julia Palmer

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

Edge storage enables the creation, analysis, processing and delivery of data services at, or close to, the location where the data is generated or consumed rather than in a centralized environment. Edge storage architectures can be deployed at regional or remote data centers or aggregation points in the form of edge servers, edge gateways or software. Edge storage is not delivered by a single technology because it must be tailored to the specific edge computing use cases.

Why This Is Important

Infrastructure and operations (I&O) leaders are adopting strategies to manage data at the edge, acknowledging the persistence of local infrastructure services alongside cloud providers, thus emphasizing the importance of a hybrid cloud approach to data management. Integration of edge storage and data services with cloud resources optimizes flexibility and efficiency, enabling distributed enterprises to reduce latency, enhance privacy and security, enable offline operation, conserve bandwidth, scale easily, and facilitate real-time data processing.

Business Impact

As edge computing becomes a critical component of cloud-connected data center transformations, infrastructure and operations (I&O) leaders are updating their storage strategies. This shift is transforming their role from infrastructure providers to providers of data services everywhere, including at the edge. I&O leaders are revising their storage strategies and vendor selections for deployments outside the public cloud infrastructure, with a focus on specific data services for edge deployments.

Drivers

- Distributed infrastructure: This model addresses the need for infrastructure delivery and management across various locations, including the cloud, the edge, and customer data centers. As cloud migration continues, the infrastructure and operations team is transforming from the provider of data center infrastructure to the provider of the data services everywhere, focusing on select use cases that require storage at the edge. The four most popular use cases at the edge are distributed cloud/data center, data processing at the edge, content collaboration and access, and data ingest and streaming.
- Hybrid infrastructure: Most enterprises require services both on-premises and in the cloud, necessitating hybrid support. Data services at the edge require many factors that preclude deployment in the public cloud: data gravity, cloud and bandwidth costs, application-specific data latency, and the effects on throughput of the speed of light. This is in addition to data security, data autonomy and data governance.
- Edge services: Latency-sensitive and bandwidth-intensive workloads aren't ideally suited for the public cloud or the core data center. Examples of such workloads include real-time data processing, collaboration and the synchronization of massive amounts of data with online storage.
- New workloads: Emerging workloads demand capabilities like global delivery, data portability, Kubernetes support, acceleration, and data management across diverse deployments.
- Evolution of data centers: Traditional data centers are transforming from isolated data hubs to providers of edge data services.

Obstacles

- I&O organizations struggle to determine what actions should be taken now, as opposed to taking a wait-and-see attitude and planning to optimize their IT operating models to mitigate risks and avoid pitfalls that may jeopardize efforts at the edge.

- The diversity of use cases, workloads, volume of data and unique infrastructure requirements at the edge introduces the potential for issues in system management, costs, security and resilience factors.
- Edge storage is not a single technology because it needs to be tailored to the specific locations and use cases.

User Recommendations

- Create edge storage platform initiatives by identifying edge-centric workloads, locations, deployment options, use cases and data service management methods.
- Choose an edge storage topology and platform approach by addressing unique workload requirements that are self-healing, software-defined and power-efficient, and can be elastically scaled up and down cost-effectively.
- Prepare for any new enterprise data center storage deployment to be edge-ready by prioritizing requirements for the edge operating model and public cloud integration.
- Select edge storage products and technologies that focus on addressing key challenges, such as autonomous operations, centralized data management, performance density and data transfer optimization.

Sample Vendors

Akamai; AWS; Broadcom; Dell; HPE; Microsoft (Azure); Nutanix; ObjectBox; SoftIron; StorMagic

Gartner Recommended Reading

Market Guide for Hybrid Cloud Storage

2024 Strategic Roadmap for Storage

Predicts 2023: Edge Computing Delivery and Control Options Extend Functionality

Top Trends in Enterprise Data Storage 2023

Cyberstorage

Analysis By: Julia Palmer

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

Cyberstorage actively defends storage systems and data against cyberattacks through prevention, early detection and attack blocking. It also supports data analytics, forensic attack analyses and storage-specific recovery capabilities. Cyberstorage can be a dedicated solution with comprehensive features, a platform-native service offering with integrated solutions, or a collection of stand-alone products that augment storage vendors with cyberprotection capabilities.

Why This Is Important

Ransomware attacks are more frequent and disruptive, making cybersecurity tools for active defense and recovery essential. While security solutions exist at the network perimeter and increasingly in backup solutions, storage systems often fall short in safeguarding against malicious activities, such as data downloads, exfiltration, manipulation, deletion, destruction or encryption. Cyberstorage solutions offer robust active defense against, and recovery from, cyberattacks targeting data storage systems and their data.

Business Impact

Cyberattacks can severely impact data, causing data loss, business impairment and costly data recovery. These issues disrupt business operations and expose sensitive information, leading to potential legal and financial penalties. Ransomware encrypts data, holding it hostage and demanding payment. These attacks may result in financial loss, reputational damage and regulatory consequences. Cyberstorage aids in prevention and detection, protecting storage systems while minimizing the attack window and reliance on backups.

Drivers

- The prevalence of costly and disruptive cyberstorage threats, combined with increasing regulatory and insurance expenses, has compelled storage professionals to embrace new approaches for active defense.
- The exponential growth of unstructured data capacity across on-premises, edge and public cloud locations increases the need for better protection against cyberthreats such as ransomware.
- Large file repositories are the easiest to attack for both ransomware and data extortion. Hackers leverage elevated credentials to access standard network access protocols, such as Server Message Block (SMB) and Network File System (NFS).

- Protecting against and recovering from a ransomware attack requires a multifaceted strategy that includes multiple market solutions, such as endpoint protection, immutable enterprise backup, and data storage infrastructure detection analytics and recovery, as part of a cyberstorage solution.
- Encrypting data at rest does not address the threat of data extortion, since the data remains unencrypted at the storage access point. Preventing data extortion involves a combination of proactive measures to safeguard data storage and respond to an attempted extortion.
- Gartner forecasts that, by 2028, all storage products will incorporate cyberstorage capabilities centered on active defense, a significant increase from 10% in 2023.

Obstacles

- Most of the innovative, advanced cyberstorage solutions are provided by smaller, newer companies, posing a challenge to established enterprises, which must navigate working with startups.
- The cyberstorage vendor landscape is difficult for enterprises to navigate because they don't understand how specific solutions map to their end-to-end cyberresilience requirements.
- Some cyberstorage vendors integrate their capabilities within a limited range of enterprise storage products and do not offer stand-alone, vendor-agnostic solutions.
- Many enterprises find it impractical to switch data storage vendors solely to enable cyberstorage services, due to the complexity of the process and the depth of their investments in current systems.
- Backup vendors dominate the ransomware protection market, leading end users to focus on recovery solutions for ransomware attacks while overlooking the importance of active defense.

User Recommendations

- Prioritize active defense and security of unstructured and structured data storage systems, because identifying and blocking an attack is just as important as recovering from one.

- Avoid relying solely on data backups and snapshot methods to address cyberresilience concerns. These should serve as a last resort for disaster recovery rather than a proactive defense for your data.
- Select storage vendors based on their ability to support all the pillars of the National Institute of Standards and Technology (NIST) Cybersecurity Framework (i.e., Govern; Identify, Protect, Detect, Respond and Recover).
- Investigate the cyberstorage capabilities of your existing vendors, as many of them are adding cyberstorage services alongside their storage services.
- Do not use cyberresilience products and capabilities as stand-alone alternatives to backup or disaster recovery; they are an additional layer of protection.

Sample Vendors

BullWall; Calamu; Continuity; Deep Instinct; Index Engines; Myota; ProLion; ShardSecure; RackTop Systems; Superna

Gartner Recommended Reading

Top Trends in Enterprise Data Storage 2023

2024 Strategic Roadmap for Storage

Market Guide for Hybrid Cloud Storage

Magic Quadrant for Distributed File Systems and Object Storage

Critical Capabilities for Primary Storage

Enterprise Storage Platform

Analysis By: Jeff Vogel

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

An enterprise storage platform is a modular, software-defined, scalable and programmable solution designed to provide data services and to support diverse workloads. It provides multiprotocol support, intelligent infrastructure management, optimal performance and integrated security measures to safeguard data. It supports hybrid, multidomain IT

operations and aligns seamlessly with cloud operating models to enable SLA-backed outcomes, ensuring integration across diverse IT environments.

Why This Is Important

Enterprises are adopting platform services to modernize their IT operations and proactively manage constantly changing business demands. The traditional approach of acquiring separate storage arrays based on storage specifications is no longer effective due to inefficient processes and siloed infrastructure. Businesses are transitioning from storage-specification-centric buying to infrastructure platform approaches, reflecting the importance of both technological and business requirements.

Business Impact

Adoption of the on-premises cloud operating model is having a profound impact on IT operations and budgets. Conventional requests for information about "speeds and feeds"" no longer deliver optimal results or differentiation between storage products, especially when enterprises are selecting advanced technology infrastructure for IT outcomes based on SLA guarantees such as cyberresilience. Enterprises are using new sourcing initiatives and subject matter expertise to lead the transition to a platform approach.

Drivers

- Enterprises have a pressing need for heightened flexibility in business and delivery models to provide storage that can rapidly support changes in business demands and operational requirements.
- Requirements to extend storage infrastructure to diverse locations, including the cloud, increase the complexity of IT operations and the need for support resources.
- Based on Gartner inquiries with clients, I&O leaders and their teams spend more than two-thirds of their time managing and supporting hardware that will be augmented by vendors' use of intelligent platform telemetry.
- The globalized nature of data is leading to diverse data sources, formats and governance regulations across different regions, requiring IT operations to scale infrastructure and systems to handle the load efficiently.
- The vendor solution landscape is rapidly shifting to accommodate growing demand for integrated data service technologies, requiring agile strategies for vendor platform management.

- Skills shortages have resulted in a lack of storage subject matter experts to address emerging technologies, including cyberresilience, new storage architectures and up-the-stack data services, such as ransomware protection.

Obstacles

- Because of limited skills and misaligned priorities, enterprises must embrace new platform initiatives and practices to accelerate innovation. These efforts require additional investment in modern infrastructure standardization.
- Without investments to support automation tools and technologies, enterprises fail to streamline operations, reduce IT operations overhead and improve efficiency of IT outcomes.
- Concerns about vendor lock-in raise potential risks and dependency on a specific vendor ecosystem, making it challenging or cost-prohibitive to switch providers.
- Enterprises struggle to extract meaningful value from storage platforms due to a lack of IT operating model SLAs.
- IT organizations are unable to restructure and reprogram their methods and processes to support platform-native capabilities, such as AIOps and autonomous storage.
- Vendor capabilities to support platform initiatives vary significantly. Enterprises may require additional proof-of-concept testing and onboarding to vet the new service.

User Recommendations

- Assess current and planned infrastructure demands to identify candidate workloads suitable for migration to storage-platform-native services based on storage-as-a-service requirements and inherent platform-native architecture attributes.
- Evaluate and shortlist storage infrastructure vendors by assessing their ability to address the core requirements and use cases for enterprise platforms against IT operating model requirements.
- Identify and prioritize critical IT operating model SLAs that are crucial to long-term platform requirements that may include two-, three- or five-year refresh cycles.
- Identify data services, such as backup, disaster recovery and ransomware protection, that are core to your platform requirements.
- Invest in training and skills development to enable the transition to platform services.

- Invest in IT operations automation by adopting platform vendor AIOps capabilities. Identify blockers and challenges, and draw up plans to proactively address them.

Sample Vendors

Dell Technologies; Hewlett Packard Enterprise (HPE); Hitachi Vantara; IBM; Lenovo; NetApp; Pure Storage

Gartner Recommended Reading

Stop Buying Storage, Embrace Platforms Instead

Market Guide for Infrastructure Consumption Services

2024 Strategic Roadmap for Storage

Leverage Storage as a Service Platform SLAs and Capabilities to Transform IT Outcomes

At the Peak

Distributed Hybrid Infrastructure

Analysis By: Julia Palmer, Philip Dawson

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Definition:

Gartner defines distributed hybrid infrastructure (DHI) as offerings that deliver cloud-native attributes that can be deployed and operated where the customer prefers — on-premises, at the edge or in the public cloud. DHI establishes the infrastructure framework for deploying applications across distributed environments, while adhering to a cloud-centric approach. This enhances agility and flexibility for workloads beyond the confines of the public cloud infrastructure.

Why This Is Important

The market for DHI has emerged to satisfy the demands of infrastructure and operations (I&O) leaders who are looking for standardized infrastructure platforms capable of accommodating various deployment scenarios. DHI solutions appeal to customers requiring cloud infrastructure as a service (IaaS) solutions, encompassing compute resources (virtual machines [VMs], bare metal or containers); storage; and network services, without the need for platform as a service (PaaS) offerings.

Business Impact

DHI solutions offer a versatile approach to infrastructure provisioning, enabling organizations to leverage the benefits of cloud-like infrastructures, while retaining control over deployment locations and avoiding dependence on specific cloud platforms or PaaS offerings.

Drivers

A key feature of DHI solutions is their flexibility with regard to deployment options. They can be implemented wherever the customer prefers, whether on-premises, in public cloud environments or at the edge. This flexibility enables organizations to tailor their infrastructure deployments to meet specific requirements, whether related to data sovereignty, latency concerns or regulatory compliance.

DHI tackles the constraints of conventional, on-premises infrastructures by offering the advantages of a cloud operating model, and by delivering enhanced consistency and availability across diverse deployment scenarios and use cases. The most prominent use cases for DHI include hybrid cloud infrastructure, multicloud, cloud-native applications, edge and assured workloads.

Obstacles

DHI encounters several obstacles that organizations must navigate to effectively leverage its benefits. These include complexity in addressing skills gaps, managing diverse environments, interoperability challenges, security concerns, cost management issues, performance and latency considerations, governance and compliance requirements, and the risk of vendor lock-in.

Another hurdle in adopting DHI stems from the widespread preference for integrated IaaS and PaaS solutions, where all services are vertically integrated and equipped with APIs.

Addressing these obstacles demands strategic planning, the implementation of appropriate technologies and tools, and ongoing monitoring and optimization to effectively manage DHI.

User Recommendations

- Assess the preferred technological pathway — whether it involves public cloud out of the cloud or on-premises infrastructure into it. Develop high-level use cases for each option, and assess how well they align with your organization’s strategic goals for cloud adoption and on-premises IT infrastructure.

- Compile a shortlist of DHI providers by crafting comprehensive use cases that outline the essential services and compatibility criteria required. Ensure that these requirements adequately reflect your organization's needs and objectives.
- Execute functional pilots or proofs of concept (POCs) by delineating a set of functional and nonfunctional requirements, based on the detailed use cases. Use the pilot or POC to verify API consistency and secure assurances from the vendor regarding compatibility maintenance with existing systems.

Sample Vendors

Alibaba Cloud; Amazon Web Services; Broadcom; Huawei; IBM; Microsoft; Nutanix; Oracle; Tencent

Gartner Recommended Reading

Stop Buying Storage, Embrace Platforms Instead

Infrastructure Consumption Services

Analysis By: Jeff Vogel

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Infrastructure consumption services (ICS) is a consumption-based IT operations model for on-premises mission- and business-critical infrastructures. ICS forms the foundation for a platform-native services delivery model. ICS vendors integrate their platform with core hardware and software product capabilities and data storage services to guarantee IT operating model service-level agreements (SLAs).

Why This Is Important

Platform-native services architectures offer significant advantages for business initiatives, while addressing persistent IT operational issues. ICS vendors enable tangible, measurable, IT performance-based SLA guarantees, using artificial intelligence for IT operations (AIOps) infrastructure telemetry tools and automation to modernize operations. This provides improved productivity, greater infrastructure resilience and efficiency, autonomous operations and continuous cost optimization.

Business Impact

ICS replaces capital expenditure (capex) financing, sourcing and IT operations' activities with consumption-based, as-a-service offerings. Shifting from capex to consumption will eliminate upfront, out-of-pocket expenses and minimize IT operations refresh/upgrade cycles for workloads mapped to vendor technology evolution curves. Vendors and their service partners are contractually obligated to manage life cycles and to maintain the infrastructure environment in accordance with commitments.

Drivers

- Bringing enterprise-class, cloud operating model benefits to on-premises infrastructure provides the impetus to shift the cloud paradigm to centralizing and managing mission-critical infrastructure on-premises.
- Most mission-critical application infrastructure is steadfastly on-premises, with high levels of reluctance and skepticism regarding moving to the public cloud, due to control, costs and complexity issues, necessitating the demand for ICS.
- Vendors are using and investing in AIOps telemetry and machine-based learning models to create IT operating model SLA guarantees that reduce administration and support overhead costs, improve productivity, and support continuous optimization initiatives.
- On-premises infrastructure vendors are investing heavily in platform-native, consumption-based services offerings, including up-the-stack data services, such as backup and ransomware protection.
- IT clients are replacing capex finance and sourcing with consumption-based services as a result of adopting the cloud operating model on-premises.
- ICS vendors now provide the asset management tools and financing support to substantially improve the infrastructure economics that compete with the public cloud.
- This is a path to modernizing critical infrastructure by shifting to a flexible, software-defined architecture that eliminates rigid, proprietary, hardware-centric workload solutions and constrains innovation.

Obstacles

- The maturity of vendor platform services and the breadth of consumption-based, as-a-service offerings.

- The absence of critical platform capabilities (e.g., robust software-defined storage [SDS]) to support hybrid cloud deployment across the on-premises, public cloud and edge infrastructure domains.
- Client finance operations' reluctance to migrate from capex to consumption or operating expenditures (opex), due to balance sheet and income statement leverage.
- Distributed hybrid infrastructure public cloud vendors providing a competitive offering that may be more scalable and diverse in terms of global, on-demand infrastructure.
- The absence of IT operating model SLAs based on AIOps capabilities.
- Business SaaS applications that are operating in the public cloud.
- The C-level "mantra" to find ways to migrate to the public cloud as a long-term IT strategy.

User Recommendations

- Develop a platform-native services strategy that addresses long-term IT operating model requirements and perennial issues, such as subject matter expertise retention that can be augmented or replaced by AIOps.
- Assess ICS vendors based on IT operations priorities and requirements against platform-native capabilities and services offerings.
- Develop an SLA strategy and requirements that substantially upscale and modernize IT operations in accordance with IT priorities, such as productivity improvement, continuous cost optimization, cyber storage resilience (including ransomware protection), and asset management demands to include sustainability.
- Schedule time with finance officers to educate them on the pros and cons of a platform-native ICS strategy and the favorable impact and benefits to the IT operating model.

Sample Vendors

Cisco; Dell Technologies; Hewlett Packard Enterprise (HPE); Hitachi Vantara; IBM; Lenovo; NetApp; Pure Storage

Gartner Recommended Reading

Market Guide for Infrastructure Consumption Services

Stop Buying Storage, Embrace Platforms Instead

2024 Strategic Roadmap for Storage

Magic Quadrant for Distributed Hybrid Infrastructure

Hybrid Cloud Storage

Analysis By: Julia Palmer

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Hybrid cloud storage facilitates seamless data services across different environments, including disparate data centers, edge locations and the public cloud. Hybrid cloud data solutions cover a wide variety of storage protocols. They're offered through such means as distributed hybrid infrastructure, hybrid cloud storage platforms, data transfer appliances, hyperconverged solutions, storage arrays, software-defined storage (SDS) products and data management software.

Why This Is Important

Data and storage services are delivered everywhere. Integration with cloud infrastructure and platform services providers is a critical selection criterion for storage platforms, because infrastructure and operations (I&O) leaders require efficiency and agility of services. Hybrid cloud storage products are enabling multiple use cases to improve data resilience, life cycle management and operation excellence, while leveraging flexibility and scale of public cloud infrastructure resources.

Business Impact

As data services expand to the edge and to the public cloud, on-premises data centers lose centrality. I&O leaders need strategies to manage data across locations. Cloud migration shifts I&O from infrastructure vendors to data service providers, focusing on edge use cases. Edge data and distributed hybrid infrastructure are key to hybrid cloud storage. Although versatile, hybrid cloud storage is mainly used by I&O leaders for disaster recovery (DR), burst capacity and storage standardization.

Drivers

- Providing data DR/business continuity (BC) by snapshotting or copying structured and unstructured datasets from on-premises or edge storage to public cloud storage.
- Bringing data closer to cloud compute and big data infrastructure for processing or analytics.
- Supporting the modernization of unstructured data services, as they are more applicable to hybrid cloud storage workflow.
- Standardizing the underlying storage platform for any deployment scenario — edge, core data center or public cloud infrastructure as a service (IaaS) — to provide consistent, operational use of storage services.
- Addressing the need for increased workload mobility across hybrid cloud locations to flex with business and application demands.

Obstacles

- Integration complexity: Integrating data and applications between on-premises infrastructure and cloud services requires robust connectivity and effective data migration.
- Performance concerns: Maintaining performance levels across hybrid environments, particularly for latency-sensitive applications, may be difficult, due to network constraints and data transfer speeds.
- Cost management: Balancing the costs of additional hybrid cloud storage on top of cloud storage costs can demand careful financial planning.
- Vendor lock-in: Avoiding vendor lock-in and maintaining flexibility to switch among cloud providers or return to on-premises can be challenging, due to dependencies on specific cloud services or proprietary technologies.
- Data mobility and interoperability: Ensuring data mobility and interoperability between on-premises infrastructure and cloud services, while avoiding data silos and compatibility issues.
- Data security: Balancing security protocols between on-premises and cloud environments.

User Recommendations

- Use platform data services by identifying workloads, data types and use cases that benefit from integration with the public cloud.

- Build a business case for hybrid cloud storage that extends beyond terabyte pricing, and emphasizes workflow efficiency and standardization.
- Prioritize hybrid cloud storage solutions supporting cloud-native access to serve public cloud applications.
- Choose a hybrid cloud storage provider based on additional services provided, such as metadata insights, cybersecurity, global access, life cycle management, multicloud support, performance enhancement and data analytics.
- Establish a comprehensive hybrid cloud data services catalog to define and manage global storage services, ensuring standardization and user clarity.

Sample Vendors

Amazon Web Services; CTERA; Hammerspace; LucidLink; Microsoft; Nasuni; NetApp; Panzura; Peer Software; Qumulo

Gartner Recommended Reading

Market Guide for Hybrid Cloud Storage

Modernize Your File Storage and Data Services for the Hybrid Cloud Future

2024 Strategic Roadmap for Storage

Data Storage Management Services

Analysis By: Michael Hoeck, Chandra Mukhyala

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Data storage management services (DSMS) solutions are designed to orchestrate the life cycle of data residing in multicloud, hybrid and SaaS environments. Using data discovery technology, they perform data classification, categorization and tagging to provide insights on structured and unstructured data. These insights enable DSMS outcomes, such as creating storage efficiencies, optimizing costs, enabling analytics workflows, and better managing data compliance and security.

Why This Is Important

DSMS solutions enable enterprises to manage the life cycle of data by prioritizing the value of data, optimizing costs and reducing risk exposure based on its intrinsic value. Organizations are managing an explosive amount of data across on-premises, hybrid, multicloud and SaaS application environments. The perceived easy path of continuing to add more storage is leading organizations down a path of unsustainable growth in capacity, siloed data, complexity, and increased regulatory and security threats.

Business Impact

DSMS solutions offer multiple benefits:

- Improve visibility to the organizations' data through metadata and content-based analytic analysis.
- Categorize and classifies data to align with storage optimization, improved data life cycle, and governance and compliance outcomes.
- Optimize storage utilization of on-premises, cloud and SaaS applications by tiering, relocating, archiving or deleting data.
- Reduce blast radius of breach, ransomware and cyberattacks through proactive management of data.

Drivers

- The constant and uncontrolled growth of data and the resulting investments in storage capacity and cloud entitlements.
- Increasing number of storage resources and the resulting challenges of managing complexity, performance and cost of storage infrastructure and data across on-premises, hybrid, multicloud and SaaS application environments.
- The compounding effect of data growth driving the expansion of backup and recovery infrastructure, services and licensing to protect and recover critical data.
- The likelihood of a threat or security breach and the understanding of "not if, but when" requires organizations to proactively reduce the blast radius for cyberattacks and ransomware exfiltration.
- Enablement of smarter business outcomes through faster analytics workflows using insights from the data content.
- Increasing number of regulatory requirements to manage data retention life cycles, often aligned to privacy legislation.

- Requirement to balance risk versus reward of retaining too much data.
- Aging storage infrastructure and application assets creating technical debt that needs to be eliminated.
- Infrastructure and operations teams' increasing effort to collect data from unmanaged sources at request of legal and compliance to support e-discovery, public record requests and subject rights requests.

Obstacles

- DSMS strategy is easily overlooked in favor of the "simple path" of acquiring more storage infrastructure and cloud entitlements, which only compounds the problem.
- Multistakeholder collaboration is required to establish expected outcomes, ownership and budgeting among IT, security, privacy, data and analytics, legal and compliance teams.
- Challenges exist in gaining executive sponsorship and stakeholder buy-in to a proactive data storage management program.
- Agreement on data classification, categorization and related retention policies is critical for successful deployments, but is often difficult to obtain.
- A new skill set must be established to effectively manage DSMS solutions to attain expected outcomes, such as data discovery, tiering, archive, migration, disposition and other life cycle management processes.
- Limited products in the market provide heterogeneous support of unstructured data stored across storage products of various vendors.

User Recommendations

- Include DSMS solutions as part of your storage purchase decision process to allow proper analysis of existing storage utilization.
- Implement DSMS tools to improve total cost of ownership and organizational data risk position.
- Clearly define scope of data sources and the data life cycle outcomes the organization wants to achieve in managing that data.
- Engage business process owners to align value and use of data to the appropriate retention policy.

- Gain executive team commitment to create or update and, most importantly, enforce data retention policies.
- Compare and contrast vendor offerings for their scope of supported data sources and types and the remediation capabilities to act on the data.

Sample Vendors

Aparavi Software; BigID; Commvault; Congruity360 InfoGov; Data Dynamics; Datadobi; Dell Technologies; Egnyte; Komprise; NetApp

Gartner Recommended Reading

2024 Strategic Roadmap for Storage

Market Guide for Hybrid Cloud Storage

Modernize Your File Storage and Data Services for the Hybrid Cloud Future

Captive NVMe SSD

Analysis By: Jeff Vogel

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Definition:

Captive NVMe SSDs are exclusively developed by select vendors for use in their flash storage systems for consumption. They are optimized for high-performance and low-latency storage, along with enhanced data storage management services. Unlike computational SSDs, they have FPGA or ASICs with extra cores for compute heavy tasks. Captive NVMe SSDs reduce the costs of all flash array systems, opening up the use of flash systems for use cases that have struggled with justifying flash costs.

Why This Is Important

Captive nonvolatile memory express (NVMe) solid-state drives (SSDs) take advantage of vendors' ability to procure NAND flash directly from the fab provider, as opposed to relying on purchasing off the shelf SSDs. Cutting out intermediaries in the supply chain leads to lower costs. This allows more competitive pricing with the storage array, along with driving higher cost efficiencies by leveraging higher capacity densities, reducing cost and power on a flash capacity basis.

Business Impact

Captive NVMe SSDs reduce the costs of all flash array systems, opening up the use of flash systems for use cases that struggle with justifying flash costs. Further, system level direct management of NAND flash memory, versus using an off-the-shelf SSD, maximizes the capabilities of the flash, providing better performance, improved flash cell wear-out management and higher reliability. Onboard computing provides enhanced storage data management operations such as data resiliency.

Drivers

- Flash costs less and has much higher SSD densities required to compete against hybrid hard-disk drive (HDD) arrays for use cases that need the flash performance but can't justify the cost delta.
- Petabyte-scale systems that leverage commodity SSDs can have terabytes of DRAM in commodity SSD drives to maintain their individual flash translation layer (FTL) mappings and metadata.
- The cost per bit of DRAM hasn't improved in the last several years, so high density NVMe SSDs will help offset the efficient use of DRAM as larger capacity becomes critical in larger systems such as AI.
- Controlling the supply chain and reducing dependency on their party SSD drive vendors can mitigate risks associated with supply chain disruptions, price fluctuations and shortages of critical components. Long term cost stability and supply chain resilience will bolster the use of affordable and sustainable flash array systems.
- Each commodity SSD drive contains its own overprovisioned spare space that's necessary for media management by the FTL. Each one of these components is an added cost that will make up a larger portion of the overall media cost as drive sizes increase.
- Captive NVMe SSDs have a much simpler firmware with a much higher reliability versus off the shelf SSD drives. Consequently, overall reliability improves as flash densities scale up.
- Captive NVMe SSDs can take advantage of onboard computing to enhance storage management capabilities, which in turn can improve and simplify IT operations.

Obstacles

- Storage vendors use software and custom logic to overcome quad-level cell (QLC) SSD media limitations to enhance product features, which enable highly valuable SLAs, potentially limiting use cases for NVMe SSDs.
- Captive NVMe SSD requires FTL controller development experience and a minimum three-to-five year development commitment to achieve return on investments and the ability to support and sustain roadmap development efforts.
- Each captive NVMe SSD is unique to a storage vendor and, as such, could limit the flexibility of using other storage vendors' traditional or captive NVMe SSD storage array based on costs and features.
- Captive NVMe SSD storage vendors that don't incorporate computing capabilities alongside NAND flash chipset sourcing and assembly can't provide the comprehensive features of those vendors that do.

User Recommendations

- Reevaluate your application infrastructure environments based on use case and workload requirements to determine the overall price and performance benefits of using captive NVMe QLC-based storage versus hybrid HDD and triple-level cell (TLC) flash based SSD storage.
- Prioritize captive NVMe storage vendors who can provide low-cost QLC flash media and use of software and custom logic to overcome limitations of QLC media, for a wider range of application needs.
- Evaluate and assess the pros, cons and long-term benefits of flash array vendors that both have and don't have a captive NVMe SSD program and roadmap initiatives.
- Establish a sourcing strategy when evaluating enterprise storage flash arrays for mission and business-critical storage workloads. The strategy should capture the benefits of captive NVMe SSD and should be fully rationalized against long-term workload requirements

Sample Vendors

Huawei; IBM; IEIT SYSTEMS; Pure Storage

Gartner Recommended Reading

2024 Strategic Roadmap for Storage

Top Trends in Enterprise Data Storage 2023

Function Accelerator Cards

Analysis By: Anushree Verma

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Function accelerator cards (FACs), aka I/O processing units (IPUs), data processing units (DPUs), distributed services cards (DSCs) or SmartNIC, are a class of devices that have dedicated hardware accelerators with programmable processors to accelerate network, security and storage functions. FACs improve data operations and services, server availability, performance, and security, and enable connectivity to a network. They have onboard memory and peripheral interfaces and can run independently.

Why This Is Important

FACs can improve server performance by up to 50%, via offloading functions such as virtual switching, security and application delivery controller (ADC). They can improve storage performance by offloading storage functions like compression, deduplication and encryption. FACs can host dedicated network appliances, including firewalls. They can also improve security by placing security functions onto a securely booted, locked-down environment. Today, FACs are primarily adopted by hyperscalers and large cloud providers, and we estimate that they will grow at a five-year CAGR of 65% through 2027.

Business Impact

FACs enable cost-efficient and energy-efficient data center environments, while improving performance. By offloading high overhead functions, they allow the server to host more workloads, which reduces the direct cost of additional servers and, in some cases, infrastructure software. In addition, they can facilitate data transmission between remote resources — primarily for high-performance computing (HPC) and artificial intelligence/machine learning (AI/ML) workloads.

Drivers

- Hyperscale cloud providers, such as Amazon Web Services (AWS), Microsoft Azure and Tencent, and other large cloud providers are using FACs today, and growing their implementation to achieve price/performance improvements.

- With the growing need for accelerated applications and services for generative AI workloads at the edge and growing power consumption concerns impacting sustainability, enterprises have started evaluating DPUs as well.
- The rise of AI/ML workloads, solid modeling, seismic analysis and advanced analytics has created unprecedented demand on storage and network, resulting in latency and bandwidth issues that FACs solve.
- FACs can reduce the number of servers and hypervisor licenses by 10% to 30%, and may also decrease the number of application software licenses.
- Moving performance-intensive functions into the hardware reduces the software-based surface area for attack.
- Telecommunication networks are moving toward virtualizing the network edge with 5G adoption, which leads to offloading 5G user plane function (UPF) and 5G network slicing to the FACs to achieve low latency and high throughput.
- FACs are increasingly bundled in high-performance solid-state storage systems to boost input/output per second and minimize latency.
- FACs provide an alternative platform to host network appliances, such as firewalls and ADCs, with price/performance benefits in specific usage scenarios.
- Vendors with a large enterprise installed base, including Hewlett Packard Enterprise (HPE), have invested heavily in the technology and marketed it to organizations with specific usage scenarios.
- Consolidation in the market is increasing, with Advanced Micro Devices (AMD) acquiring Xilinx and Pensando Systems, and Microsoft acquiring Fungible.

Obstacles

- Enterprises perceive FACs as a disruptive and dramatic departure from typical data center networking patterns, which limits adoption due to concerns over risk.
- There is confusion in the market due to vendors using different terminology, and providing different capabilities and architectures.
- Data plane programmability is high-risk, limiting the adoption in enterprises.
- Hyperscale cloud service providers (CSPs) are able to justify the incremental price with the large-scale order and customization benefits they get by adopting FACs. However, enterprises are so far unable to do so, thereby hindering rapid adoption.

- Form factor and power consumption can impact rack, power and cooling budgets, if they occupy a full-size Peripheral Component Interconnect Express (PCIe) slot.

User Recommendations

- Use FACs for specific use cases, such as acceleration of AI/ML and nonvolatile memory express over fabrics (NVMe-oF).
- Engage your existing data center infrastructure vendors on their plans for multivendor interoperability for FACs offloading, prior to your next server refresh.
- Investigate FACs to replace legacy components, such as physical firewalls and reduce the number of application licenses.
- Pilot FAC offerings to improve scale/security needs in a large-scale data center network (1,000 switches), or to support extremely network-sensitive workloads.
- Select FAC-based storage offerings for applications that require microsecond latency performance when processing large datasets.
- Use a cross-functional team that includes networking, compute, storage and security personnel to evaluate FAC offerings.
- Focus on management and orchestration when evaluating FACs, as they are key differentiating factors.

Sample Vendors

Advanced Micro Devices (AMD); Broadcom (VMware); Ethernity Networks; Intel; Kalray; Microsoft; Napatech; Nebulon; NVIDIA; Pliops

Gartner Recommended Reading

Emerging Technologies: Adoption Growth Insights — Function Accelerator Cards (Next-Gen SmartNICs, DPUs, IPUs)

Your Server Is Eating Your Network — Time to Rethink Data Center Network Architectures

Market Trends: Arm in the Data Center: Act Now to Develop Plans to Address This Shifting Market

Sliding into the Trough

Hybrid Cloud File Data Services

Analysis By: Julia Palmer

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Definition:

Hybrid cloud file data services (HCFDS) deliver seamless file data services across disparate data centers, edge locations and public cloud infrastructure. HCFDS can be native to the underlying storage product or a stand-alone offering that is agnostic to the underlying storage.

Why This Is Important

Organizations are increasingly creating and accessing data at multiple locations beyond on-premises, including the public cloud and edge locations. In addition, modern application workflows span multiple locations using data services from multiple locations. HCFDS address this situation by making file data services available wherever there is a need to access or create files.

Business Impact

Businesses can benefit from better data resilience, capacity management, on-demand elasticity, data portability, resource optimization and operational excellence through the six use cases of HCFDS. These six use cases are disaster recovery, burst for capacity, burst for performance, data transport and computation, global data orchestration, and storage standardization across disparate data center locations.

Drivers

- Existing and new unstructured data workloads require hybrid cloud file capabilities to use public cloud for elasticity, operational simplicity, processing and data longevity.
- HCFDS address the challenges that arise as a result of data generated from multiple locations and public cloud migration initiatives.
- HCFDS are expanding from being providers of global data access to providers of platform services, such as data insights, cyber resilience, life cycle management and data mobility across public cloud and on-premises deployments.
- Increasing application analytics tools for business insights is creating more data across edge locations that must be connected with AI services in the public cloud.

- Knowledge workers are distributed across geographical locations but need to collaborate on the same set of data, improving overall productivity among shared projects.

Obstacles

- There is a limited number of mature solutions for global data orchestration leads.
- Solutions based on caching architectures have performance and scalability limitations.
- Existing investments in file or object storage products may limit exposure to vendor-neutral or heterogeneous HCFDS offerings.
- Funding from business is limited for use cases beyond disaster recovery and burst for capacity.
- Hybrid environments create complexities in maintaining consistent security practices, compliance standards and governance policies.
- Robust interoperability solutions are required to ensure data is seamlessly integrated across hybrid cloud elements.

User Recommendations

- Take advantage of public cloud infrastructure and platform services by identifying workloads, data types and use cases that will benefit from integration with public cloud using HCFDS.
- Choose a hybrid cloud provider by its ability to deliver additional services, such as media rendering, data analytics, cyberstorage, data life cycle management, performance acceleration, cloud-native access and other value-added services that enterprises require.
- Build a comprehensive hybrid cloud data services catalog to define and maintain global hybrid cloud storage services and to ensure standardization and customer transparency.
- Build a business case for hybrid cloud data services by establishing the new capabilities and outcomes that come from utilizing the public cloud for disaster recovery, burst for capacity, burst for processing, global data orchestration and storage standardization.

Sample Vendors

Atempo; CTERA Networks; Datadobi; Hammerspace; IBM; Komprise; Nasuni; NetApp; Peer Software; Qumulo

Gartner Recommended Reading

Market Guide for Hybrid Cloud Storage

Storage as a Service

Analysis By: Jason Donham, Philip Dawson

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Storage as a service (STaaS) is a managed service that provides a way for organizations to manage and consume storage without the overhead costs of upfront capital for storage assets and staff time. STaaS provides managers with flexibility, resilience and efficiency. STaaS solutions include both on-premises storage and cloud storage.

Why This Is Important

Infrastructure and operations (I&O) leaders struggle to leverage on-premises cloud operating model benefits while their environment becomes more complex and less agile to business demands. STaaS provides many benefits such as improved asset management through aligning costs to utilization, higher levels of operational efficiency through outsourcing hardware administration and support, and cost savings through a consumption-based, as-a-service model with increased levels of automation.

Business Impact

Budget and spend (capital expenditure [capex]) inefficiencies are driving higher total cost of ownership (TCO) compared to consumption-based (operating expenditure [opex]) spending that is much more in line with storage needs. STaaS solutions enable organizations to:

- Shift from capex to opex to eliminate IT budget inefficiencies.
- Improve workload asset management and reduce capitalization costs.
- Eliminate life cycle management issues and technology refresh cycles.
- Reduce cyber liabilities and threat exposure via data service offerings.

Drivers

- Infrastructure managers need options when reconfiguring and resizing storage environments to meet the rapidly changing demands of applications.
- Threats of ransomware and other cyberattacks require higher levels of data security.
- Traditional storage is inflexible when deployed in a hybrid operating model.
- STaaS can increase labor costs on the back of intensive budgeting cycles.
- Inefficient life cycle management and constant infrastructure turnover conspire to create an inflexible environment that hampers innovation and ability to respond to business demands.
- The lack of subject matter experts (SMEs) or staff attrition issues (offset by moving responsibility and accountability to vendors) around critical elements of the infrastructure lead to a less resilient platform and exposure to untenable events.

Obstacles

- Finance or procurement members believe they are better stewards of capex assets or don't fully understand the indirect cost savings or benefits afforded by STaaS.
- Vendor sales and marketing narratives and selling strategies are reluctant to promote or advocate the benefits of as-a-service models over capex-related product features.
- Vendor business models are immature in terms of back-end operations — for example, metering and billing — and integration with supply chain and logistics.
- Vendor channels are not fully equipped to enable or support the transition to an as-a-service consumption model.

User Recommendations

- Implement consumption-based STaaS to reduce or eliminate capex budgets and IT refresh cycles.
- Add SLAs to drive critical requirements such as ransomware protection to improve security posture.
- Utilize artificial intelligence for IT operations (AIOps), combined with STaaS, to create an intelligent infrastructure platform that proactively and dynamically responds to IT operating model outcomes and business priorities.

Sample Vendors

Amazon Web Services (AWS); Backblaze; Dell Technologies; Hewlett Packard Enterprise (HPE); IBM; Microsoft; NetApp; Pure Storage; Wasabi Technologies; Zadara

Gartner Recommended Reading

Leverage Storage as a Service Platform SLAs and Capabilities to Transform IT Outcomes

Container-Native Storage

Analysis By: Julia Palmer

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Container-native storage (CNS) is designed specifically to support container workloads and focus on addressing unique cloud-native architecture, pod granularity and performance demands, while providing deep integration with the container management systems. CNS is aligned with microservices architecture principles and adheres to the requirements of container-native data services, including being hardware-agnostic, API-driven and based on distributed software architecture.

Why This Is Important

CNS solutions are specifically designed to provide persistent storage to cloud-native applications. The common foundation is typically based on a distributed, software-defined and unified pool of storage and has container-level granularity of data services while providing enterprise data management features. In addition, the entire stack is most often orchestrated with Kubernetes to manage container life cycle integration and enable self-service operations for developers.

Business Impact

Major impacts to the business include:

- The ability of CNS to enable the deployment of stateful cloud-native applications at large scale; this enhances elasticity, availability and multicloud integration.
- It is critical to be able to provide storage solutions that align with container-native infrastructure principles. These solutions enable the seamless support of

stateful applications, efficient sharing of application data, and provision of advanced data services, thereby enhancing operational efficiency and facilitating innovation within the business ecosystem.

- Elimination of the bottlenecks to achieving portable infrastructure agility when building and deploying modern, cloud-native applications.

Drivers

- Organizations are building new containerized applications using cloud-native principles and rearchitecting traditional applications on Kubernetes platforms, both of which are driving significant momentum in the adoption of CNS.
- Due to the increased popularity of deploying and operating container environments by orchestration platform, most IT leaders require a persistent storage solution that can be tightly integrated with container orchestrators, such as Kubernetes.
- Infrastructures and organization (I&O) leaders require new tools and processes for data management to provide storage services accessed by stateful applications running in containers and orchestrated by Kubernetes.
- CNS solutions can be deployed on-premises or in the cloud, making them optimal for hybrid and multicloud deployment infrastructure.
- Because CNS functions are based on software, they can be implemented in containers, enabling them to be managed with the same orchestration functions as containerized applications.

Obstacles

- A CNS solution will not be adopted by every enterprise, because it remains most appropriate for new deployments of cloud-native applications, or for applications that will be revised with significant refactoring.
- Although embracing the CNS paradigm will yield agility benefits, adopting a CNS solution is likely to increase operational complexity in the short term for traditional enterprise environments.
- CNS vendor landscape and technology is constantly evolving with a mix of early- and late-stage startups.
- Given the fragmented nature of the vendor ecosystem, I&O leaders risk creating a technology silo with CNS solutions, which is a common obstacle to large-scale adoption.

- Containers utilize the host kernel, and as such, all share the same underlying hardware system below the operating system layer, making it possible that an exploit (container escape) in one container could affect the shared hardware environment of other containers.
- Storage area network vendors are enhancing the base container storage interface functionality with value-added data management services at container granularity, reducing the need for separate CNS products.

User Recommendations

- Choose storage solutions that align with microservices architecture principles and adhere to the requirements of container-native data services, such as being hardware-agnostic, API-driven, based on distributed architecture, and portable to support edge, core or public cloud deployments.
- Align storage solutions with cloud strategies, while considering CNS applicability in the public cloud and hybrid cloud scenario of your choice.
- Select storage products that are closely aligned with the developer workflow tools and can be directly integrated with the application layer for portability, scaling and data protection.
- Validate your vendor's capability of continuous innovation delivery, quality customer support and a consistent pricing model, given that the container ecosystem is rapidly evolving with unproven vendor business models.
- Ensure that the CNS solution is tested and qualified for specific Kubernetes platforms.

Sample Vendors

Broadcom; DataCore Software; Diamanti; IBM (Red Hat); ionir; Pure Storage; SUSE

Gartner Recommended Reading

Top Trends in Enterprise Data Storage 2023

2024 Strategic Roadmap for Storage

Solution Path for Cloud-Native Infrastructure With Kubernetes

NVMe-oF

Analysis By: Jeff Vogel, Julia Palmer

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Nonvolatile memory express over fabrics (NVMe-oF) is a network protocol that extends the parallel-access and low-latency features of NVMe protocol to Ethernet and Fibre Channel networks enabling the use of storage external to servers. It is designed for high-performance fabric technologies, including remote direct memory access (RDMA) over Fibre Channel, InfiniBand or Ethernet with RoCE v2, iWARP or TCP. NVMe-oF uses the NVMe protocol to extend access to NVMe devices.

Why This Is Important

NVMe-oF, when used in combination with NVMe SSD drives, addresses use cases where low-latency is a critical requirement. NVMe-oF enables the use of high-performance and scalable architectures that can capitalize on the underlying networking capabilities in combination with NVMe flash media.

Business Impact

NVMe-oF substantially reduces data access latency while ensuring more efficient connectivity between storage and servers. NVMe-oF storage targets can be dynamically shared among workloads, thus providing an on-demand or composable storage resource that provides flexibility, agility and greater resource efficiency. NVMe-oF runs on both traditional Fibre Channel (FC) and IP switches. NVMe/TCP offers superior IT infrastructure connectivity and is a good fit for organizations either without legacy FC infrastructure or looking to replace FC switches.

Drivers

- NVMe-oF consolidates multiple applications onto a single NVMe-oF-based storage array due to its superior performance.
- The NVMe-oF protocol takes advantage of high-speed networks and accelerates the adoption of next-generation storage architectures.
- Storage as a service (STaaS) disaggregated storage architectures will pave the wave for an increase in adoption of NVMe-oF as a preferred scale out performance-based hybrid cloud architecture.

- The release of VMware vSphere 7.0 Update 3 for mainstream usage opens up a path for TCP/IP to be a popular data center transport mechanism of NVMe-oF.
- Most storage array vendors now support NVMe-oF-capable products as an alternative protocol for primary storage.

Obstacles

- Depending on the existing infrastructure, the implementation of end-to-end NVMe-oF could require substantial changes to and increased costs of storage platforms, networks and servers.
- Infrastructure and operations (I&O) leaders struggle to justify ROI for end-to-end NVMe-oF deployments.
- The cost and complexity of infrastructure elements, such as host bus adapters (HBAs) and switching devices, impede the adoption of NVMe-oF solutions in mainstream enterprises.
- Some NVMe data storage products on the market deliver only a small fraction of NVMe's potential performance improvements. This is due to end-to-end differences in how NVMe-oF is implemented.
- Software support for NVMe-oF is still relatively nascent.

User Recommendations

- Select workloads where the scalability and performance of NVMe and NVMe-oF justify the costs and complexity of such deployments. Target NVMe-oF for AI and machine learning (ML), high-performance computing (HPC), in-memory databases or transaction processing.
- Define which type of implementation will be used — host-to-controller or controller-to-NVMe media. Consult suppliers to determine which type they support against performance requirements.
- Investigate potential infrastructure bottlenecks, such as applications, servers or networks. Consult suppliers on potential performance and total cost of ownership (TCO) gains to justify the ROI.
- Assess potential storage platform, network interface controller, HBA and network fabric suppliers to verify that interoperability testing has been performed and references are available.

- Verify the availability and support of NVMe-oF networks for hypervisor and operation systems to ensure compatibility and performance improvement.
- Deploy either NVMe-oF with RDMA RoCE v2 or NVMe-oF over TCP/IP-based products to ease the transition and provide investment protection.

Sample Vendors

Dell Technologies; Hewlett Packard Enterprise (HPE); Hitachi Vantara; Huawei; IBM; Lightbits Labs; NetApp; Pure Storage

Gartner Recommended Reading

Top Trends in Enterprise Data Storage 2023

2024 Strategic Roadmap for Storage

Open-Source Storage

Analysis By: Julia Palmer, Vishesh Divya

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Definition:

Open-source storage is a form of software-defined storage (SDS) for which the source code is made available to the public through a distribution license that complies with an open-source definition. It enables many of the same features as proprietary storage, including the support of structured and unstructured data, as well as heterogeneous management.

Why This Is Important

Although open-source storage has been around for more than a decade, it has been adopted mainly by hyperscalers, managed service providers (MSPs) and large organizations. Recent innovations in x86 hardware and flash technology, combined with an innovative open-source ecosystem, are making open-source storage and its licensing models practical for cloud and big data workloads. This makes it an alternative to proprietary storage.

Business Impact

Open-source storage adoption by technology firms, service providers, and research and academic environments underscores the benefit of leveraging the broader open-source

storage developer ecosystem across disciplines. Big data, analytics and private cloud use in enterprises are promising use cases. Open-source storage will enable customers to innovate rapidly in such key storage areas as data management at a lower acquisition cost, with “good-enough” availability, performance and manageability.

Drivers

- Open-source storage is playing an important role in supporting cost-effective, scalable platforms for new cloud and big data workloads.
- Customers are actively evaluating open-source storage products across block, file and object protocols.
- More than 90% of enterprises worldwide use open-source technology in support of their mission-critical IT workloads, resulting in the use of open-source options across the technology stack.
- Cloud computing, microservices application architectures, big data analytics and information archiving are pushing the capacity, pricing and performance frontiers of traditional scale-up storage architectures. This has led to renewed interest in open-source software (OSS) as a means to achieve high scalability in capacity and performance at lower acquisition costs.
- Open-source platforms, such as Kubernetes and TensorFlow, are backed by large, innovative communities of developers and vendors, such as DataDirect Networks (Lustre), IBM and Red Hat (Ceph Storage). They provide enterprises with a broad selection of options to consider for use cases such as cloud storage, big data, stateful microservice workloads and archiving.
- There is also an influx of open-source storage projects for container-based storage, such as Longhorn, MinIO, OpenEBS and Rook.
- Data security capabilities, backed by protection features (e.g., object immutability, encryption, IAM and versioning), have solidified the position of OSS at an enterprise level.

Obstacles

- Onboarding OSS storage requires subject matter experts, because some IT leaders overestimate the benefits and underestimate the costs and risks. Although it requires less upfront investment than proprietary storage, IT leaders should weigh the benefits, risks and costs.

- It's hard to predict cost and ROI of ownership of open-source storage. Pragmatic, long-term OSS investment strategies should balance cost, flexibility and innovation to be successful.
- For mature open-source storage projects, leaders may turn to the community as a knowledge base to augment self-support efforts. However, these communities do not have contracted SLAs, and do not guarantee quick, reliable support.
- Open-source storage licensing can be challenging, due to multiple license types and governing copyright conditions. Complex legal risks may require rigorous due diligence.
- Although vendors often claim to support no vendor lock-in, some solutions might face challenges when integrating with proprietary applications or hardware that requires customization.

User Recommendations

- Allocate resources to invest in open-source storage initiatives that support ecosystem activities.
- Actively deploy pilot projects, identify internal champions, train storage teams, create total cost of ownership (TCO) analyses and prepare the organization for this trend.
- Use commercial distribution and obtain support through a vendor, rather than downloading the source code for free. Open-source storage requires significant effort and expertise to install, maintain and support.
- Evaluate downsides of lock-in against the perceived benefits achieved when deploying "open-core" or "freemium" storage products. The proprietary software version often comes as add-on modules, retained features or management tools that function on top of open-source storage.
- Conduct legal license risk due diligence, and assess usage and compliance. Ensure compatibility with distribution of your product or service.
- Check integration of OSS by running interoperability checks on the OSS solution with your applications and hardware.

Sample Vendors

DataCore Software; DataDirect Networks; iXsystems; MinIO; Openfiler; Red Hat; SoftIron; SUSE

Gartner Recommended Reading

A CTO's Guide to Open-Source Software: Answering the Top 10 FAQs

Toolkit: Gartner's Open-Source Software Assessment Framework — 1.0

Best Practices for Setting Up an Open-Source Program Office

Software-Defined Storage

Analysis By: Chandra Mukhyala

Benefit Rating: Transformational

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Definition:

Software-defined storage (SDS) abstracts storage software from the underlying hardware, providing common provisioning and data services, regardless of locality and hardware technology. It can be deployed as a virtual machine (VM), a container or as storage software on a bare-metal, industry-standard server, providing flexibility to deploy storage wherever the application demands — on-premises, at the edge or in the public cloud.

Why This Is Important

SDS provides flexibility to support applications across the hybrid cloud from on-premises, the edge or the public cloud, with consistent storage service. SDS eliminates dependency on proprietary vendor hardware and runs on commodity hardware, thereby offering a lower total cost of ownership (TCO). SDS can be deployed for a broad range of workloads, such as general-purpose file storage, backup, archiving, analytics, AI, and block storage for VMs and containers.

Business Impact

- SDS offers increased hybrid cloud deployment flexibility and the ability to have common provisioning and data services on-premises, at the edge and in the public cloud.
- SDS breaks the dependence on proprietary storage hardware, supporting lower acquisition costs.
- SDS enables choice of storage compute, flash, memory and networking hardware options to align with application needs.

- Some Gartner customers report as much as a 40% TCO reduction from the use of standard hardware.

Drivers

- Build a storage solution at a low acquisition price point on common platform hardware.
- Decouple storage software from hardware to standardize data center hardware platforms to scale compute and capacity independently.
- Build an agile “infrastructure-as-code” architecture, enabling storage to be a part of software-defined data center automation and orchestration framework that integrates with the public cloud.
- Take advantage of the latest innovations in storage hardware before they are available and supported in traditional external controller-based storage arrays.
- Run the same storage services across on-premises, at the edge and in the public cloud.

Obstacles

- I&O leaders often struggle to navigate SDS vendor solutions as a result of the variety of SDS product offerings, with established and emerging vendors delivering differentiated value propositions and product capabilities.
- Hybrid cloud IT operations is an emerging use case for SDS, because tomorrow’s data center landscape is expanding to include edge and public cloud. Hence, adoption often requires multiple products and complex integration.
- I&O leaders need to invest in personnel with specialized skills to effectively develop and manage SDS in the enterprise. Therefore, the potential cost savings obtained from reduced capital expenditures (capex) requires a cost-benefit analysis for improved operating flexibility.
- Performance and TCO, along with other business value factors, must be considered contemporaneously, creating a more-complex assessment that slows deployment.

User Recommendations

- Recognize that SDS is a growing deployment model that will be focused primarily on web-scale storage architectures, but has applicability at the edge and public cloud deployments.

- Select SDS vendors that provide support for multiple deployment options, and offer validated hardware reference designs that minimize performance and scalability trade-offs.
- Grade SDS products by their ability to be truly hardware-agnostic, API-driven, based on distributed architectures, flexible pricing models and hybrid cloud deployment flexibility.
- Deploy SDS as part of a cohesive software-defined infrastructure design, with an emphasis on delivering uniform storage platforms across on-premises, public cloud and edge environments.
- Before embarking on SDS deployments, recognize that SDS may involve substantial work sizing the underlying hardware and building the total solution on your own, versus a plug-and-play appliance.

Sample Vendors

Broadcom; DataCore Software; IBM; MinIO; NetApp; Nutanix; Qumulo; Scality; StorMagic; WEKA

Gartner Recommended Reading

Magic Quadrant for Distributed File Systems and Object Storage

Market Guide for Hybrid Cloud Storage

Stop Buying Storage, Embrace Platforms Instead

Climbing the Slope

Object Storage

Analysis By: Chandra Mukhyala

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Definition:

Object storage refers to a system that houses data in structures called objects and serves hosts via APIs such as Amazon Simple Storage Service (S3). Objects are composed of unstructured data and metadata, conceptually similar to files. Object storage uses a flat namespace versus the hierarchical structures seen in file systems. Object storage offerings

are available as software-based storage, virtual machines, traditional hardware appliances or managed services.

Why This Is Important

With unstructured data doubling every few years, the need for scalable, resilient and cost-effective storage becomes a critical requirement. A single data lake that captures all of an organization's unstructured data for analysis to provide insights to the business is a key technology enabler. Object storage satisfies those needs through a flat namespace housing key value pairs with rich metadata, all protected with erasure coding.

Business Impact

Object storage helps businesses take control of their data management strategy by attaching rich metadata to objects to enable business insights. Developers' preference for an API-based storage, combined with increasing number of independent software vendors' support for the S3 interface, makes object storage preferable for storing all unstructured data. Its resilient, scalable, software-defined and developer-friendly aspects can deliver cost-effective, flexible and agile infrastructure for businesses.

Drivers

- The primary driver for object storage continues to be the explosive growth in unstructured data, resulting from digital transformation across all industry verticals.
- Organizations are managing larger objects and a larger number of objects, leading to the need for scalable and cost-effective storage. Photos and videos are captured in higher resolution and for more use cases than in the past.
- More application- and machine-generated data is being captured for analytics purposes.
- Application developers want to consume storage through a simple programmatic interface. Modern application developers prefer the programmatic nature of object storage over traditional file-based storage.
- The default storage in the public cloud being object storage leads to more applications preferring it over other forms of storage.
- The rise of ransomware is driving the demand for immutable storage, and object storage addresses that requirement with S3 object-locking for protecting backup data.

- As organizations increase investment in AI and generative AI, object storage is ideally suited to be the underlying data lake for aggregating all the model training data.
- Use cases for object storage are expanding from backup and archiving to primary storage for applications where data processing is done in server memory.

Obstacles

- Legacy applications are still dependent on file semantics (file path names, permissions).
- Modern distributed file systems are highly scalable and available as hardware-agnostic, software-based deployment models that can run on any standard server hardware, on-premises or in the public cloud, making them cost-effective and flexible.
- Most file-based storage offerings also support the S3 protocol for accessing files as objects.
- Vendors are increasingly offering a unified platform for all unstructured data that provide both file and object services, minimizing the need for a stand-alone object storage offering.
- Modern file systems can satisfy the vast majority of requirements for which object storage is considered. Exceptions are when an application depends heavily on the metadata associated with the object or when scaling to several billion objects.

User Recommendations

- Evaluate the suitability of object storage products for all unstructured data, but not when the primary use case requires processing or editing of file-based data.
- Evaluate the product's API support for dominant public cloud providers when building on-premises object storage repositories so that workloads can be extended to public cloud if needed. Amazon S3 has emerged as the dominant API over vendor-specific APIs and OpenStack's Swift (which is in precipitous decline).
- Select object storage vendors offering a wide choice of deployment (software-only versus packaged appliances versus managed hosting) and licensing models (perpetual versus subscription) that can provide flexibility and reduce total cost of ownership.

- Select object storage vendors with advanced API functionality, like Amazon S3 Select, and value-added services for search and discovery.

Sample Vendors

Backblaze; Cloudian; Dell Technologies; Hitachi Vantara; Huawei; IBM; MinIO; NetApp; Scality; Wasabi

Gartner Recommended Reading

Magic Quadrant for Distributed File Systems and Object Storage

Distributed File Systems

Analysis By: Chandra Mukhyala

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Distributed file systems (DFSs) use a single file system to cluster multiple storage nodes together, presenting a single namespace and storage pool to provide aggregated bandwidth for multiple hosts in parallel. Data and metadata are distributed over multiple nodes in the cluster to handle availability and data protection in a self-healing manner. DFSs can be expanded nondisruptively by adding new nodes to the cluster, linearly increasing capacity and performance with each new node addition.

Why This Is Important

Building scalable and cost-effective distributed storage platforms for unstructured data is imperative for infrastructure and operations (I&O) leaders. The accelerated growth of existing file datasets and the introduction of new file-based workloads are bringing distributed scale-out storage architectures to the forefront of IT infrastructure planning. DFSs are designed to address performance and scalability limitations in traditional scale-up network-attached storage (NAS) environments.

Business Impact

DFSs are based on scale-out architecture that scales performance and capacity linearly and nondisruptively, as required by business demands. Many file system products are being deployed as software-only, which provide better flexibility and a lower total cost of ownership compared to integrated storage appliances. The introduction of a DFS

will significantly improve private cloud services for unstructured data, which require highly scalable, resilient and elastic infrastructure.

Drivers

- To address exponential unstructured data growth, I&O leaders are increasingly replacing scale-up NAS appliance systems with distributed scale-out file system products to benefit from the linear scaling of capacity and performance.
- Enterprises are preparing for exponential growth of unstructured data, and seeking data insights and integration with cloud storage for a long-term data life cycle management.
- Analytics, AI and machine learning applications can require file semantics and manipulation of data, thereby increasing demand for DFSs.
- Simulation and data modeling, the traditional use cases for DFSs, continue to be major drivers of DFSs.

Obstacles

- DFS deployments are often complex, requiring careful planning and a lengthy proof-of-concept (POC) process to validate operational requirements.
- Many DFSs are designed to be deployed in a single large data center and cannot efficiently scale down for small deployments at the edge.
- Some DFSs are designed for performance and still require separate tools to do data management, cyber protection and analytics.
- DFSs are not often selected when I&O leaders are looking for hybrid cloud file platforms because they are designed for on-premises data residency requirements.

User Recommendations

- Establish clear workload performance indicators relevant to business and validate all performance claims with POC deployments, with an emphasis on protocol type, file sizes and your choice of underlying hardware.
- Increase agility by integrating DFSs with data insights and life cycle management options. This approach will enable you to get a better handle on data management and provide actionable insights to the business.
- Shortlist vendors with the ability to integrate with the public cloud and enable hybrid cloud storage deployments with tiering, archiving and bidirectional data flow for

data processing. This emerging paradigm is experiencing a positive early traction with enterprises.

- Prioritize vendors that include support for Amazon S3, along with Network File System and Server Message Block, in a multiprotocol manner.

Sample Vendors

Cohesity; Dell Technologies; Huawei; IBM; IEIT SYSTEMS; Nutanix; Pure Storage; Qumulo; VAST; WEKA

Gartner Recommended Reading

Magic Quadrant for Distributed File Systems and Object Storage

Entering the Plateau

Hyperconvergence

Analysis By: Philip Dawson, Jeffrey Hewitt

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Hyperconvergence combines storage, computing and networking into a single system that reduces data center complexity and increases scalability. Multiple servers can be clustered as nodes together to create pools of shared compute and storage resources, designed for convenient consumption. Delivery models include physical and virtual appliances, reference architectures, as a service or public cloud.

Why This Is Important

With changes to VMware licensing effective January 2024, we are seeing a resurgence of hype related to storage and network virtualization in hyperconvergence as clients look at alternative innovations and vendors. Infrastructure and operations (I&O) leaders seeking a cost-effective solution with a single management interface that excludes proprietary, external hardware controller-based storage should consider hyperconvergence as a viable option. Possible use cases include virtual desktop infrastructure (VDI), edge/Internet of Things (IoT), hybrid cloud and cloud-native.

Business Impact

Hyperconvergence enables IT leaders to be responsive to new business requirements in a modular, small-increment fashion, avoiding the large-increment upgrades typically found in three-tier infrastructure architectures. It is of particular value to midsize enterprises that can standardize on hyperconvergence and to the remote sites of large organizations that need cloudlike management efficiency with on-premises edge infrastructure.

Drivers

- Hyperconvergence provides simplified management that decreases the pressure to hire hard-to-find specialists. Adoption is greatest in dynamic organizations with short business planning cycles and long IT planning cycles tied to hybrid cloud delivery. The hyperconverged infrastructure (HCI) market is now trifurcating, focusing on the data-center-led "hybrid cloud" management use case with cloud-native applications, the VDI use case and the "edge/IoT" remote management use case.
- Hyperconvergence leads to lower operating costs, especially as it supports a greater share of the compute and storage requirements of the data center.
- Nutanix, an innovator in hyperconverged integrated system (HCIS) hardware appliances, has shifted to a hyperconverged infrastructure (HCI) software revenue model and continues to increase its number of OEM relationships and partners (including Cisco).
- VMware, a strong player in the HCI software market, made changes to its licensing. It is now bundling its networking (NSX) and/or storage (VSAN) within its two enterprise offerings bundled around vSphere and Cloud.
- Larger clusters are now in use, and midsize organizations are considering hyperconvergence as the preferred alternative for on-premises infrastructure for block storage.
- Hyperconvergence vendors are achieving certification for more demanding workloads, including Oracle and SAP, and end users are beginning to consider hyperconvergence as an alternative to integrated infrastructure systems for some workloads.
- As more vendors support hybrid and public cloud deployments, hyperconvergence is a stepping stone toward public cloud agility as suppliers are expanding hybrid cloud deployment offerings for cloud-native applications.
- A number of niche hyperconvergence suppliers offer scale-down solutions to address the needs of remote office/branch office (ROBO) and edge environments.

Obstacles

- Applications designed for scale-up architectures (as opposed to scale-out ones) are unlikely to meet cost or performance expectations when deployed on hyperconverged infrastructure.
- The acquisition cost of hyperconvergence may be higher, and the resource utilization rate lower than for three-tier architectures.
- While HCI has somewhat matured from a hypervisor compute and storage function, software defined in networking is split between software-defined networking (SDN) and networking around software-defined WAN (SD-WAN), driving edge deployments.
- For large organizations, hyperconverged deployments will remain another silo to manage.
- The refactoring of enterprise agreements to consumption models is not only making HCI bundled in revirtualization initiatives and products, but increases overlap with existing/alternative storage and network virtualization innovations.

User Recommendations

- Implement hyperconvergence for hybrid cloud infrastructure and cloud-native applications for agility, modular growth and management simplicity.
- Hyperconvergence requires alignment of compute, network and storage refresh cycles; consolidation of budgets; operations and capacity planning roles; and retraining for organizations still operating separate silos.
- Test the impact on disaster recovery and networking under a variety of failure scenarios, as solutions vary greatly in performance under failure, their time to return to a fully protected state and the number of failures they can tolerate.
- Ensure that clusters are sufficiently large to meet performance and availability requirements during single and double node failures, which require erasure coding support and/or proofs of concept to reveal any performance anomalies.
- Make sure that HCI as it embeds in consumption models from ELAs does not amplify costs or duplex investments with other similar innovations from alternate vendors.

Sample Vendors

Broadcom (VMware); Dell; Microsoft; Nutanix; Sangfor Technologies; Scale Computing; StorMagic

Gartner Recommended Reading

Market Guide for Full-Stack Hyperconverged Infrastructure Software

Market Guide for Integrated Systems

Manage Your Dependence on VMware vSphere

Quick Answer: How to Handle Broadcom's End of Sale of VMware Perpetual Licenses

Appendixes

See the previous Hype Cycle: Hype Cycle for Storage and Data Protection Technologies, 2023

Hype Cycle Phases, Benefit Ratings and Maturity Levels

Table 2: Hype Cycle Phases

Enlarge Table

Phase	Definition
Innovation Trigger	A breakthrough, public demonstration, product launch or other event generates significant media and industry interest.
Peak of Inflated Expectations	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the innovation is pushed to its limits. The only enterprises making money are conference organizers and content publishers.

Trough of Disillusionment	Because the innovation does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
Slope of Enlightenment	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the innovation's applicability, risks and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
Plateau of Productivity	The real-world benefits of the innovation are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
Years to Mainstream Adoption	The time required for the innovation to reach the Plateau of Productivity.

Source: Gartner (July 2024)

Table 3: Benefit Ratings

Enlarge Table

Benefit Rating	Definition

Transformational	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
High	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
Moderate	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
Low	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner (July 2024)

Table 4: Maturity Levels

Enlarge Table

Maturity Levels	Status	Products/Vendors
Embryonic	In labs	None
Emerging	Commercialization by vendors Pilots and deployments by industry leaders	First generation High price Much customization

Adolescent	Maturing technology capabilities and process understanding Uptake beyond early adopters	Second generation Less customization
Early mainstream	Proven technology Vendors, technology and adoption rapidly evolving	Third generation More out-of-box methodologies
Mature mainstream	Robust technology Not much evolution in vendors or technology	Several dominant vendors
Legacy	Not appropriate for new developments Cost of migration constraints replacement	Maintenance revenue focus
Obsolete	Rarely used	Used/resale market only

Source: Gartner (July 2024)


Gartner