

# DCIG TOP 5 Enterprise Cloud-based NAS Consolidation Solutions

by DCIG Analyst, Todd Dorsey



## TOP 5 Enterprise Cloud-based NAS Consolidation Solutions\*

Panzura CloudFS

CTERA Enterprise File Services Platform

IBM Spectrum Scale

Nasuni

WekaIO WekaFS

\*Licensing vendor is listed first; others listed in alphabetical order

### SOLUTIONS EVALUATED:

- CTERA Enterprise File Services Platform
- Commvault Hedvig Distributed Storage Platform
- DataCore vFIO
- Dell EMC IsilonSD
- IBM Cloud Object Storage
- IBM Spectrum Scale
- Nasuni
- NetApp ONTAP Select and Cloud Volumes ONTAP
- Panzura CloudFS
- Qumulo File Data Platform
- Scality RING8
- Storidge Container IO
- StorONE S1 Enterprise Storage Platform
- SUSE Enterprise Storage
- Tiger Technology Tiger Bridge
- Tintri NexentaStor
- Veritas InfoScale Enterprise
- VMware vSAN
- WekaIO WekaFS
- Zadara VPSA

### FEATURES EVALUATED:

- *Deployment Capabilities*
- *Data Protection*
- *Product and Performance Management*
- *Documentation*
- *Technical Support*
- *Licensing and Pricing*

## The Expanding Volumes of Unstructured Data

Many enterprises struggle with managing expanding volumes of unstructured data throughout the organization. Storing, protecting, and securing this data creates challenges around cost, complexity, and scalability. Addressing these issues take cycles away from other important IT projects bringing innovation and future profits. A new generation of software-defined storage (SDS) offers solutions to these challenges.

## The Difficulties Created by Unstructured Data Growth

Unstructured data growth is filling the enterprise data center and its branch offices. This growth brings many difficulties.

**Filled data centers.** Data growth is creating demand for more storage capacity in the data center. However, floor space or power constraints may hinder expansion. Higher density storage systems may be available to these organizations, but the budget may not allow for the acquisition.

**Terabytes have become petabytes.** Many legacy storage systems were designed when 100 TB was a lot of data, but now many organizations need to manage multiple petabytes of data. At petabyte scale, storing, protecting, backing it up, and recovering it all is problematic using legacy solutions. Many IT departments struggle with finding or training staff with the experience to manage these environments. If staff leave, departments may be unable to replace these individuals, placing operations at risk.

**Remote office file data.** Managing unstructured data at remote offices brings its own issues. Technology at these locations often lags behind the data center, resulting in different technology throughout the business. Branch office locations typically lack IT staff; thus, remote offices must involve others to recover lost files. File data may use up storage, and backup solutions may rob bandwidth during copies. Consider that a one terabyte file directory can use two additional terabytes of storage for backup and offsite recovery. Further, inactive data grows as no one wants to delete files when unknowns exist. For team collaboration of files hosted at remote offices, latency can become unworkable.

***“Storing, protecting, and securing unstructured data creates challenges around cost, complexity, and scalability.”***

**Storage proliferation.** Over time, a company grows and needs expanding storage. However, the current system has not reached end-of-life. So IT adds to it. Now there are two systems instead of one to manage. Two systems become three, four, five, and ten. This growth may take place because of capacity requirements or because of protocol requirements. The effect is the same: An organization ends up with a proliferation of data silos without global visibility into the file estate of its multiple, underlying systems.

**Total cost of ownership.** These difficulties increase the total cost of ownership for unstructured data management. Enterprises must factor in the day-to-day expenses of ongoing maintenance and upkeep. In addition, there are the not-so-obvious costs associated with acquiring new storage infrastructure, including cycles for budgeting, researching, evaluating, procuring, installing, and migrating data among multiple storage systems.

The explosive growth of unstructured data creates costly, difficult-to-manage problems using legacy approaches to file data management.

## SDS-based NAS Consolidation—Two Key Approaches

Cloud-based NAS consolidation, based on enterprise-class SDS, offers a solution for providing fast, flexible, usable access to all of an organization’s file data for all of its end users.

There are two key approaches to NAS consolidation:

**Replace.** One approach replaces multiple filers with a single solution designed for multi-petabyte management. In some cases, that single solution uses object storage as the actual backing data store, whether on-premises or in the cloud.

**Integrate.** The other approach integrates the enterprise’s existing filers under the domain of a

global file system. This enables visibility and unifies management across the data landscape. In other words, aggregating, enhancing, and overlaying the management and visibility of the file data without necessarily replacing the existing storage system.

Many SDS NAS solutions address both approaches to NAS consolidation. In addition, these products bring benefits through the following capabilities.

### SDS-based NAS Consolidation Benefits

**Elastic, scalable storage capabilities.** Especially for solutions using public cloud providers as the back-end object store, the public cloud represents unlimited storage that can scale on demand. If new capacity is needed, an administrator can quickly turn this up.

**Offloading storage infrastructure management.** For public cloud storage, much of the burden of storage infrastructure management moves from the customer organization to the cloud provider. Cost management (space, power, labor), complexity (performance management, networking, administration), and capacity planning (monitoring, acquiring, implementing) shifts out of the enterprise. By offloading these activities to the cloud provider, enterprises reduce IT costs for on-premise, file data management.

**Fast, remote office file sharing.** Without additional aids, end-users may experience latency issues when accessing and modifying files stored in the cloud. To nullify this, many of these TOP 5 solutions use local edge appliances to cache frequently accessed files, saving changes to the cloud where the master copy is stored. This gives enterprises LAN speed access to active files wherever the office may be located, while preserving cloud benefits.

**Global file system view.** Rather than different data silos where IT administrators have disparate views of file storage, these solutions provide a single view of the entire file storage landscape. A single data repository replaces data silos. This global visibility also unlocks opportunities to derive new value from the data.

*“NAS Consolidation solutions give organizations a single view of all their unstructured data.”*

**Global file management capabilities.** Global views, including permissions management, capacity utilization, and analytics, enable new opportunities for ensuring optimal performance and cost of managing an organization’s unstructured data. Recovering corrupted or deleted files becomes simple. If a location needs access to a particular file directory, it is as easy as configuring this group’s permissions. While the dynamics causing file growth still occur, these solutions give enterprises the tools to discover and manage files globally.

### Distinguishing Features of DCIG TOP 5 Enterprise Cloud-based NAS Consolidation Solutions


DCIG evaluated twenty SDS solutions for a cloud-based NAS consolidation use-case. Using feature-based analysis and comparisons of defensible data derived from publicly available sources, vendors, and DCIG’s own experience, DCIG’s TOP 5 Enterprise Cloud-based NAS Consolidation Solutions evidence these characteristics in contrast with the other evaluated solutions.

**Robust support.** DCIG TOP 5 providers display robust support capabilities. All TOP 5 vendors provide 24x7x365 technical support and one-hour support response times compared to 75% of the other evaluated providers. Many of these vendors offer community support forums and knowledge-bases for self-service support. In short, they evidence a greater breadth of technical support options in comparison with the other solutions.


**Virtually unlimited capacity.** Another characteristic of TOP 5 NAS consolidation solutions is virtually unlimited capacity. While there are a few exceptions, these solutions provide near unlimited capacity for the maximum number or size of files, directories, and volumes. Unlimited capacity means organizations can dynamically adjust to changing business requirements.

**Public cloud support.** Cloud-based NAS consolidation provides the opportunity of centrally storing an organization’s file data in the cloud for the benefits this architecture brings. DCIG TOP 5 solutions support multiple public cloud providers. Such broad support offers flexibility in matching a cloud provider’s capabilities with the needs of the business.

#### Into what public cloud environments can the product be deployed?



#### Which object or public cloud storage providers are supported as targets by the product?



Source: DCIG 2021-22 SDS Research

**Global namespace.** A key feature of a global namespace is a single presentation of an organization’s file system data. All DCIG TOP 5 solutions provide this feature. The resulting global visibility greatly simplifies the management of unstructured data.

**File locking capabilities.** Distributed workforces value file locking for avoiding editing conflicts while working with centrally stored files. Robust file locking capabilities prevent more than one person from modifying and updating a file, or a specific portion of a file, at the same time or provide mechanisms for reconciling such changes. Each of the DCIG TOP 5 solutions support file locking capabilities.

## Similarities among DCIG TOP 5 Enterprise Cloud-based NAS Consolidation Solutions

In addition to the major characteristics that all DCIG TOP 5 Enterprise Cloud-based NAS Consolidation solutions generally share, the solutions have these traits in common.

**Amazon and IBM data targets.** All DCIG TOP 5 solutions show wide public cloud support as a centralized host for file data, including Amazon and IBM cloud storage. Such wide support offers many possibilities for hybrid-cloud and multi-cloud deployments.

**Encryption.** Data encryption is understandably important to enterprise customers. The financial, legal, and brand reputation cost of data leaks and breaches is well-known. All DCIG TOP 5 solutions provide support for array-based, data in-flight, and data-at-rest encryption. Not all the remaining evaluated solutions provide these features. For example, only 63% of the other solutions support data-in-flight encryption.

**Automation.** Enterprises look for solutions that integrate well with their existing infrastructure. REST APIs facilitate communication of data storage with external applications. Each of the TOP 5 solutions provides REST APIs for integrating with infrastructure automation frameworks and applications.

**Non-disruptive migrations.** Enterprise end-users and customers expect 24x7 availability of the applications they use and a seamless end-user experience when migrating file data from one storage system to another. All DCIG TOP 5 solutions support non-disruptive data migrations.

**SMB and NFS support.** All of the DCIG TOP 5 solutions support both SMB 3.0 and NFS protocols. These are the protocols used by most enterprise applications to address file-based storage.

## Differences between DCIG TOP 5 Enterprise Cloud-based NAS Consolidation Solutions

The TOP 5 solutions differ from one another in the following ways:

**File systems characteristics.** Each of the DCIG TOP 5 vendors characterize their NAS solution as distributed file systems. In addition, several of the providers characterize their solutions as parallel file systems. Parallel file systems, in particular, are associated with high-performance computing (HPC) use cases.

**Customer-owned encryption keys.** Several solutions allow the customer to own the encryption keys. Some solutions provide customer-owned encryption keys on-premise but not in the cloud. Some provide support for both.

**VMware integration.** Organizations have widely adopted VMware for its virtualization software. Many of the TOP 5 solutions provide support for VMware deployments.

**S3-compatible object storage.** Many enterprises have embraced a cloud-first, cloud-native approach to new application development. These applications are designed for web access and object storage. Most DCIG TOP 5 solutions provide S3-compatible object storage. The combination of file and object protocol support enables storage systems to support legacy enterprise applications and new cloud-native applications.

## DCIG TOP 5 Enterprise Cloud-based NAS Consolidation Solution Profiles

Each of the DCIG TOP 5 Enterprise Cloud-based NAS Consolidation Solution Profiles highlights three notable solution features that make the product attractive to enterprise organizations.

### Panzura CloudFS

Panzura built Panzura Cloud File System (CloudFS™) to securely consolidate an enterprise's file storage as a single data source in the cloud. Organizations benefit from local file performance as CloudFS caches frequently used files on physical or virtual edge appliances. Panzura deploys on a physical Panzura Filer™, as a virtual machine in an ESXi environment, or as an in-cloud instance running as an Amazon AMI, Google VM or Azure VHD. Panzura also supports any public or private cloud object store, including AWS, Azure, Google Cloud, IBM iCOS, Cloudian and more.

Notable features that earn Panzura a DCIG TOP 5 award include:

- **Ransomware mitigation.** Panzura CloudFS never modifies cloud-stored data. Instead, CloudFS stores new or changed data as additional objects, further protected by read-only snapshots taken at configurable intervals. If ransomware should ever get past an organization's first line of defense, CloudFS can use snapshots to revert files to before the ransomware attack.
- **Storage analytics.** Panzura Data Services allow administrators to search, monitor, audit, and analyze their entire file network from one dashboard. Administrators can search for any file or file event and display results in moments. Metrics can be sliced and diced as needed to support Key Performance Indicators. The Panzura analytics dashboard provides historical, current, and forward-looking analysis for optimizing performance and forecasting capacity.
- **Efficient global deduplication and compression.** Panzura optimizes capacity by running inline block-level deduplication and compression on any stored data on the network to remove blocks common across different files. Panzura's deduplicates at the block level (128KB) instead of the chunk level (4MB) for greater efficiencies. Leveraging the cloud as a distribution point, Panzura embeds the deduplication reference table in metadata, which then shares this table among all Panzura appliances. These methods remove data redundancy globally across all controllers, rather than just a single controller. Panzura's deduplication features enable firms to reduce their total storage footprint and file transfer times.

### CTERA Enterprise File Services Platform

CTERA's Enterprise File Service Platform enables organizations to simplify the cost, scale, and complexity challenges of traditional NAS in the new era of distributed and remote work IT. The CTERA platform connects core, branch, and work-from-home users to a single namespace with fast and secure data access from any edge location or device, enabling enterprises to transition from legacy on-premises file storage to a cloud-based file solution without sacrificing performance or data privacy. CTERA supports public, private, hybrid, and multi-cloud deployments and also offers a managed service solution for organizations seeking a hands-off approach to file storage management.

Notable features that earned CTERA a DCIG TOP 5 award include:

- **Unified file services.** The CTERA platform extends cloud file services to any office or endpoint while preserving folder structures, permissions, and access control. CTERA Edge Filers (virtual, physical, or hyperconverged) enable secure LAN-speed access for headquarters and branch offices to the organization's cloud-based file system. The CTERA Drive client delivers the same user experience to end-user devices, including mobile. CTERA also supports cloud VDI deployments. In addition, CTERA offers macOS, Adobe, and Office 365 integration.
- **Military-grade security.** Founded by cybersecurity experts, CTERA offers a private, in-firewall deployment model with end-to-end

encryption (FIPS 140-2 certified). The platform is built on a zero-trust architecture and is validated as 'military-grade' through inclusion on the U.S. Department of Defense Information Network Approved Products List (DoDIN APL). Core, branch, and home offices are connected to a single namespace with consistent access control. Additional security features include antivirus and DLP tools, multi-factor authentication, auditing capabilities, and geo-segmentation for data sovereignty compliance.

- **Fast data synchronization.** CTERA's proprietary patent-pending edge-to-cloud file transfer protocol, CTERA Direct, supports file transfer rates of up to 30 TB per site per day.<sup>1</sup> CTERA Direct deduplicates and compresses data at the source and provides continuous, bi-directional data synchronization across distributed locations, even with limited bandwidth. These capabilities enable data-intensive workloads such as video, medical imaging, and IoT/ML data at the edge.

## IBM Spectrum Scale

IBM represents a well-known name in computing and storage products and services. The IBM Spectrum Storage suite of products supports a wide range of SDS capabilities. IBM Spectrum Scale™ is the center of IBM storage for data and AI information architecture. It provides a scale-out high-performance global parallel file system (GPFS) with access via file and object storage protocols. It is a big data solution that manages data across multiple tiers of storage, including tape and cloud. Spectrum Scale deploys on-premises or in the cloud and is widely adopted for technical computing, analytics, and content management workloads.

Three features that earned IBM Spectrum Scale recognition as a DCIG TOP 5 solution include:

- **Accelerate AI adoption.** IBM ranks #1 in AI market share and supports some of the world's largest and leading-edge AI/ML applications today.<sup>2</sup> Enterprises can accelerate AI adoption through IBM's Cloud Pak® portfolio of pre-built applications, tools, and runtimes with IBM Watson. Spectrum Scale offers one-click integration to the IBM Cloud Pack® for Data and IBM Watson, which supports AI data use cases through a cloud-native data and AI platform.
- **IBM Cloud®.** Spectrum Scale brings the benefit of deploying on IBM's owned and managed public cloud platform. IBM Cloud contains over 170 products covering data, containers, AI, IoT, and blockchain applications. Within IBM Cloud, IBM Spectrum Scale runs on bare metal servers bringing hybrid cloud benefits such as data mirroring, tiering, recovery, and consistent storage function between on-premises and the cloud.
- **Mature platform for innovation.** With over 20 years of production deployments, IBM Spectrum Scale is a proven file system for legacy workloads that also provides a scalable foundation for modern applications. For example, IBM Spectrum Scale now offers Container Native Storage access, native object storage, and NVMe access to object storage data. The IBM Spectrum Scale Hadoop connector reduces the need to move data for Hadoop workloads, without requiring changes to applications.

## Nasuni

Nasuni stores and synchronizes files across any number of locations at any scale. The heart of the Nasuni file services platform is its patented cloud-native global file system, UniFS®. It unifies enterprise NAS, backup, and disaster recovery infrastructure while consolidating all of an organization's files in cloud object storage. The Nasuni solution deploys at the edge as a physical or virtual appliance on-premises or in the cloud. The Nasuni edge appliance caches frequently used files, giving office users a local file-sharing

experience. Nasuni works with all major object store vendors, so organizations can choose which cloud backend to overlay with the Nasuni platform.

Three features that earned Nasuni recognition as a DCIG TOP 5 solution include:

- **Efficiency.** With Nasuni, the 'golden copy' of the file resides in cloud object storage. Nasuni's edge appliances cache frequently accessed files locally. New or updated blocks of data within any file are compressed, globally deduplicated, and encrypted before being sent to cloud object storage. This minimizes bandwidth requirements and egress fees while maximizing effective storage capacity.
- **Data encryption and protection.** Nasuni extends protection of cloud-stored data by encrypting data both in-flight and at-rest. The customer controls the encryption keys. Organizations can take advantage of the data protection inherent within cloud architecture for safeguarding data against loss. Cloud data storage can materially save enterprises the time, money, and effort associated with backups and disaster recovery plans for their unstructured data.
- **Fast ransomware recovery.** Nasuni provides continuous file versioning to capture changes on every edge appliance as they occur. UniFS then stores file system deltas as a snapshot in object storage. Recovery points can be up to every few minutes, and recovery times just a few minutes more to restore a single file, a whole directory, or the entire file system. If an enterprise experiences a ransomware attack, it can quickly revert the affected files to a point-in-time just before the attack.

## WekaIO WekaFS

WekaIO's software defined data management platform provides extreme scale and speed while reducing management OPEX. The Weka name derives from the wekabyte, a million petabytes, conveying the idea of unlimited capacity. Enterprises may deploy WekaFS on bare metal, VMs, containers, and in the cloud. WekaFS supports NFS and SMB file and S3 object protocols. WekaIO's partner ecosystem allows enterprises to acquire WekaFS as an integrated solution from popular OEM server providers.

Notable features that earned WekaFS a DCIG TOP 5 award include:

- **GPUDirect storage.** WekaIO targets technical compute use cases requiring high performance and low latency. WekaFS is the only solution in this report that supports NVIDIA® GPUDirect® Storage. This enables direct memory access from storage to GPU, relieving the CPU I/O bottleneck and enabling increased I/O bandwidth and capacity. Customers experience notable performance gains over workloads that otherwise may take hours or days using x86 processors. This translates into real business value for Weka customers through better infrastructure efficiency and reduced time to insight.
- **High-performance cloud bursting.** When demands grow greater than on-premise servers can supply, or circumstances require even faster results, enterprises can tap into cloud services from AWS for high-performance cloud bursting. Organizations may tier, snapshot, or stage data locally for movement to the cloud. A user may spin up i3 and p3 instances for work on the S3 data. When the cloud-based processes are complete, enterprises may move the resulting data back on-premises.
- **Parallel file system performance.** WekaFS further focuses on high-performance use through its distributed, parallel file system. Parallel file systems enable high performance by storing data across multiple networked servers. Distributed file data allows for parallel tasks through concurrent data access. The more compute available, the faster work completes. Enterprises may scale-out their WekaFS parallel file system up to 8 exabytes<sup>3</sup> for even greater performance and data resiliency.

## Inclusion Criteria for Enterprise Cloud-based NAS Consolidation Solutions

In this report, DCIG specifically focused on enterprise cloud-based SDS NAS consolidation solutions possessing the following characteristics. DCIG identified twenty different solutions meeting these inclusion criteria:

- Commercially available as of February 1st, 2021.
- Sufficient, publicly available information available for DCIG to make an informed decision.
- The solution must support NFSv3.
- The product may be available as a pre-integrated software and hardware appliance from the solution provider.
- The product must support deployment in at least one public cloud.

DCIG evaluated each of these solutions in the following areas:

- 1. Deployment capabilities.** Evaluate the capabilities concerning on-premise deployment options, cloud provider deployment options, cloud provider targets supported, storage protocols supported, virtual environments supported, and certifications with equipment, operating systems, and applications.
- 2. Data protection capabilities.** Evaluate solution capabilities supporting availability, encryption, replication, and snapshot features.
- 3. Product and performance management features.** Evaluate options to manage the underlying hardware and optimize it for performance. Examples include dashboard views, predictive analytics, storage optimization, quality of service features, auto-tiering capabilities, and directory service integration.
- 4. Documentation.** Evaluate the breadth and depth of documentation the provider makes available to customers. Examples include whitepapers, knowledgebases, online manuals, and community forums.
- 5. Technical support.** Evaluate the availability and technical support options of the solution provider. Examples include support availability, response time commitments, options to open cases, escalation support, and proactive problem resolution.
- 6. Licensing and pricing.** Evaluate the relative ease of doing business through flexibility and simplicity in contract lengths, pricing elements, and bundled pricing options.

## DCIG Disclosures

Vendors of some of the solutions covered in this DCIG TOP 5 report are or have been DCIG clients. This is not to imply that their solution was given preferential treatment in this report. In that vein, there are some important facts to keep in mind when considering the information contained in this TOP 5 report.

- No vendor paid DCIG any fee to research this topic or arrive at predetermined conclusions.
- DCIG did not guarantee any vendor that its solution would be included in this TOP 5 report.
- DCIG did not imply or guarantee that a specific solution would receive a TOP 5 designation.
- All research is based upon publicly available information, information provided by the vendor, and/or the expertise of those evaluating the information.
- DCIG conducted no hands-on testing to validate how or if the features worked as described.
- No negative inferences should be drawn against any vendor or solution not covered in this TOP 5 report.
- It is a misuse of this TOP 5 report to compare solutions included in this report against solutions not included in it.

No vendor was privy to how DCIG weighted individual features. In every case the vendor only found out the rankings of its solution after the analysis was complete. To arrive at the TOP 5 solutions included in this report, DCIG went through a seven-step process to come to the most objective conclusions possible.

1. DCIG established which features would be evaluated.
2. The features were grouped into six general categories.
3. DCIG weighted each feature to establish a scoring rubric.
4. DCIG identified solutions that met DCIG's definition for an Enterprise Cloud-based NAS Consolidation solution.
5. A DCIG analyst internally examined the feature data for each solution and completed a survey for it based upon the analyst's own knowledge of the solution and publicly available information. Providers were also given the opportunity to complete surveys about their products.
6. DCIG evaluated each solution based on information gathered in its survey.
7. Solutions were ranked using standard scoring techniques. ■

## Notes

1. <https://www.ctera.com/technology/ctera-direct/>
2. "How to scale the AI ladder: Watch these enterprises." IBM Journey to AI Blog, May 21, 2019. <https://www.ibm.com/blogs/journey-to-ai/2019/05/how-to-scale-the-ai-ladder-watch-these-enterprises/>
3. [https://www.weka.io/wp-content/uploads/files/2017/12/Architectural\\_WhitePaper-W02R6WP201812-1.pdf](https://www.weka.io/wp-content/uploads/files/2017/12/Architectural_WhitePaper-W02R6WP201812-1.pdf)

## About DCIG

The Data Center Intelligence Group (DCIG) empowers the IT industry with actionable analysis. DCIG analysts provide informed third-party analysis of various cloud, data protection, and data storage technologies. DCIG independently develops licensed content in the form of TOP 5 Reports and Solution Profiles. More information is available at [www.dcig.com](http://www.dcig.com).