

DCIG TOP 5

Enterprise Multi-site File Collaboration Solutions

by DCIG Analyst, Todd Dorsey



TOP 5 Enterprise Multi-site File Collaboration Solutions*

Nasuni

CTERA Enterprise File Services Platform

NetApp Global File Cache and NetApp Cloud Volumes

Panzura CloudFS

Qumulo File Data Platform

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

SOLUTIONS EVALUATED:

- CTERA Enterprise File Services Platform
- Commvault Hedvig Distributed Storage Platform
- Dell EMC IsilonSD
- IBM Cloud Object Storage
- Nasuni
- NetApp Global File Cache and NetApp Cloud Volumes
- Panzura CloudFS
- Pure Storage Pure FlashArray//X, FlashArray//C and Pure Cloud Block Store
- Qumulo File Data Platform
- Scalify RING8
- Storidge Container IO
- StorONE S1 Enterprise Storage Platform
- SUSE Enterprise Storage
- Tintri NexentaStor
- WekaIO WekaFS
- Zadara VPSA

FEATURES EVALUATED:

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

File Collaboration Challenges in the Multi-site Enterprise

In many organizations, core business processes rely on effective file-based collaboration. This was fine when most employees worked at a headquarters facility with fast network links to the enterprise data center and its filers. However, with an organization's users now spanning the globe, legacy systems fail to meet the collaboration needs of this distributed workforce. This frustrates end users, wastes time and money, and increases multiple risks to the organization.

The challenges around effective file collaboration include:

Version Control

When a distributed organization lacks an effective file collaboration solution, troubles occur. An employee discovers they are working on the wrong version of a file. Or worse, they discover this after they have sent an incorrect version to a client. Team members lose time when they compare versions to understand differences between two possible documents. Then users must spend cycles to merge and resolve different versions into a correct one.

Unmanageable File Data Growth

IT experiences an unending growth of unstructured data. End-users and IT are reluctant to remove files for concern of deleting something necessary. Employees copy files temporarily with the intention of reviewing later, only to never do so. Organizations watch their backup and archive storage grow along with their active file storage. These dynamics contribute to file clutter and increase storage volumes dramatically.

Sharing Files and Folders

Sharing files and folders for collaborative work brings its own concerns. Sending files through email presents security risks, delivery failures, and out-of-date files floating around. If a team uses email to send documents, they must spend time with messaging, making changes, then emailing files back. Enterprises can create VPN's or other shares for outside partners; however, this is frequently with manual activity and possible mistakes.

Data Security and Control

Traditional file sharing approaches often lack security and compliance features. Employee negligence, poor security, or compromised storage media can result in data breaches. A notable reason is the lack of visibility and management over the file data by the IT department. A single breach or attack can devastate a business and its reputation.

Handling Large Files

Whereas a traditional local NAS infrastructure handles large file sets with ease, this becomes problematic when sharing large files across the wide area network (WAN). File-sharing can become slow or even impossible when distributed teams are involved. And end-users must completely rule out emailing large files for collaborative work.

Latency

Cloud file storage opens possibilities for multi-site file collaboration but also obstacles. WAN transfer speeds, mobile access, and competition with other applications over the WAN link can result in latency problems that make collaborative work tedious. Cloud storage by itself does not resolve other issues noted above. For files hosted on-premises, a slow user experience is still typical for anyone except for those local to the hosted files.

Frustrations, Costs, and Risks

Teams lose time and money because of these problems. Additionally, legacy approaches do not provide opportunities for enterprise-wide automation for efficiency benefits. These issues create frustration, lower productivity, increase costs, elevate risks to data security, revenue, and brand reputation.

“DCIG TOP 5 solutions enable fast file collaboration, improved end-user experiences, new file management capabilities, improved security, and reduced storage needs.”

SDS-based File Collaboration Benefits

Along with the change to a distributed workforce, enterprises are adopting Software-defined Storage (SDS) for its flexibility, agility, and increasing capabilities. Many SDS-based file-storage solutions offer features that enhance multi-site file collaboration and bring a number of benefits.

Effective Version Control

As a primary feature, these solutions provide global file locking. Global file locking ensures only one person can access and edit a file. Others who try are notified the file is locked. This feature prevents file collisions and unintended overwrites. These solutions track and retain previous versions for restoral when there is a need.

Modern File Collaboration

These solutions enable file and folder sharing with internal and external stakeholders outside the department. The degree of access can be customized based on need. Changes to files by collaborators are updated automatically with private or public cloud storage used as the authoritative source. To speed synchronization, only the portions of a file that have changed are transmitted across the network.

Reduced Storage Capacity

By centralizing file storage and managing effective version control, companies realize savings on file storage. Many offerings utilize deduplication and compression for efficient storage and reduced data transmission. Thus organizations save file storage costs and reduce WAN bandwidth needs.

Data Protection, Security and Control

Cloud storage and a global file system can consolidate islands of file data into a unified platform that is easier and less expensive to deploy, operate, and scale. For many organizations, cloud storage architecture offers more than sufficient data protection and additional back up can be added. With these solutions, IT administrators have access to tools to monitor their entire file data landscape. IT departments experience control through their abilities to assign file permission attributes at a granular level. API features enable integration with the customer's IT infrastructure.

Fast File Access

Many of these SDS solutions use edge appliances or solutions to provide fast access to end-users for active data files. While the authoritative file is stored in the private or public cloud, active data is cached locally to speed up performance when users or applications access data. Changes to files are updated on the back end and invisible to the end-user. Stakeholders have a near-immediate view into the file updates.

In summary, SDS-based file storage solutions enable fast file collaboration, improved end-user experiences, new file management capabilities, improved security, and reduced storage needs for enterprise organizations.

Distinguishing Features of TOP 5 Enterprise Multi-site File Collaboration Solution Providers


DCIG evaluated sixteen SDS-based solutions for a multi-site file collaboration use-case. Using feature-based analysis and comparisons of defensible data derived from publicly available sources, vendors, and DCIG's own experience, the TOP 5 solution 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. All of these vendors offer community support forums and knowledgebases 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 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 file storage 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.

Cloud deployment options— Enterprise Multi-site File Collaboration Solutions



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 the TOP 5 Enterprise Multi-site File Collaboration Solutions

In addition to the major characteristics that all DCIG TOP 5 Enterprise Multi-site Collaboration solutions generally share, the solutions have these traits in common.

These shared features include:

- **Amazon and Google Cloud Platform targets.** All DCIG TOP 5 solutions show wide public cloud support for both deployment and as a centralized host for file data. Uniquely all five TOP 5 solutions

support Amazon and Google Cloud Platform as a cloud storage partner. Such wide support offers many possibilities for hybrid-cloud and multi-cloud solutions.


- **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 TOP 5 solutions provide support for array-based encryption, data in-flight, and data-at-rest encryption. Not all the remaining evaluated solutions provide these features. For example, only 45% of the other solutions support encryption for data in-flight.
- **Data automation support.** 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 applications with storage managed by the solution.
- **Non-disruptive migrations.** Enterprises desire a seamless end-user experience for moving their file data from one system to another. All DCIG TOP 5 solutions support non-disruptive data migrations for a positive conversion for their workforce.
- **SMB and NFS support.** All of the DCIG TOP 5 solutions support file sharing protocols up to SMB 3.0 and NFSv3. File protocol support beyond this enables enhanced performance through greater parallelism and security that does not require custom drivers.
- **VM integration.** Organizations have widely adopted VMware for its virtualization software. All of the TOP 5 solutions provide support for VM deployments.

Differences between the TOP 5 Enterprise Multi-site File Collaboration Solution Providers

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

- **File systems characteristics.** Each of the DCIG TOP 5 vendors characterize their solution as distributed file systems. In addition, several of the solutions are characterized as parallel file systems. Parallel file systems, in particular, are architected for high-performance computing use cases.
- **Customer-owned encryption keys.** All solutions support data in-flight and at-rest encryption. Several solutions allow the customer to own the encryption keys. Some solutions provide customer-owned encryption keys for the premise but not the cloud. Some provide support for both.
- **S3 Compatible object storage.** Many modern applications are designed for web access and object storage. Most DCIG TOP 5 solutions provide S3-compatible object storage. Both file and object protocol compatibility make it possible for enterprise systems to support legacy and new cloud-native applications.
- **Multitenancy support.** Some organizations require multitenancy for isolating storage for different business units. Solution providers depend on this feature to keep customer's data segregated. Many, but not all TOP 5 solutions, support multitenancy.
- **Storage capacity optimization.** Enterprises of all types continue to generate and accumulate data needed for storage. Organizations look for data optimization techniques such as compression and deduplication to minimize capacity needs and save storage costs. TOP 5 solutions vary for their support of data optimization techniques such as compression and deduplication.

Object or public cloud storage providers supported as targets by the product



Source: DCIG 2021-22 SDS Research

- **Pre-integrated appliances.** Some organizations prefer integrated, turn-key solutions that include the hardware and software needed for an enterprise-wide file storage solution. Many of the TOP 5 solution providers offer turn-key solutions that cover the hardware and software needed for on-premise deployments. Such solutions provide single ownership for implementation and ongoing technical support.

TOP 5 Enterprise Multi-site File Collaboration Solution Profiles

Each of the DCIG TOP 5 Enterprise Multi-site File Collaboration Solution Profiles highlights three notable solution features that make the product attractive to organizations.

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:

- **Nasuni Global File Lock™.** Nasuni differentiates itself with a patented global lock feature that enables large scale file collaboration without data conflict worries. Capitalizing on its cloud-centric approach, Nasuni's solution offers redundant lock services and enhances availability through intelligent failover across regions. Nasuni's Global File Lock also includes intelligence that determines when a "local lock" at the edge of the network is sufficient or when a full global lock is required. Global locking can be configured at a folder level as needed. Locks can be viewed and, if needed, cleared through the Nasuni management console and filer interfaces.
- **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. Data remains secure through AES-256 encryption.

- **Nasuni Continuous File Versioning™.** 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.

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. In the process, CTERA enables 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.** CTERA's global file system technology enables secure and real-time multi-site file collaboration. CTERA Edge Filers deployed at headquarters and branch offices serve as access points to facilitate local file access to the cloud-based file system. Similarly, CTERA Drive Share, an enterprise-grade file sync and share tool, provides home and roaming end-users the ability to access, edit, and share files from any laptop or mobile device for secure, internal, and external collaboration. CTERA's proprietary WAN-optimization protocols ensure fast file transfer across globally distributed sites.
- **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 connect to a single namespace with consistent access controls. Additional security features include antivirus and DLP tools, AD/LDAP integration, two-factor authentication, auditing capabilities, and geo-segmentation for data sovereignty compliance.
- **Scalability.** CTERA scales file storage to infinite capacity by utilizing the cloud, and also scales out to support very large numbers of users and locations. CTERA provides enterprise scalability, with global enterprise deployments supporting thousands of sites and tens of thousands of users. This scalability is key for providing storage access to large and growing organizations, as well as organizations with widely distributed employees.

NetApp Global File Cache and NetApp Cloud Volumes

The NetApp® Global File Cache (GFC) with Cloud Volumes ONTAP® extends an organization's NetApp data fabric to edge locations by deploying a Windows virtual machine to those locations. As the name

implies, the solution caches files locally. This approach allows a scalable, flexible unstructured data management strategy that addresses the needs of many distributed enterprises. NetApp Global File Cache connects to NetApp Cloud Volumes, a fully managed cloud storage solution for AWS, Azure, and Google Cloud. Cloud Volumes can grow seamlessly to 100TB.

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

- **Extending NetApp data fabric to the edge.** For organizations using NetApp on-premise and cloud solutions, organizations can extend file management capabilities to the edge when they deploy NetApp Global File Cache. Global File Cache provides branch offices visibility and access to an enterprise's consolidated file storage in the cloud. To give end-users a local LAN experience, active datasets are cached locally while GFC sends changes to the cloud.
- **Authoritative file locking.** Rather than a 'replication/synchronization' architecture for file storage, GFC is based on a single authoritative instance of the data. Thus for global file locking, there is no need for distributed locks to be kept in sync via metadata synching. When any user accesses the central authoritative copy of the file directly or through a GFC edge, a file lock is immediately placed. Since GFC doesn't replicate locking information across all sites, it doesn't rely on maintaining lock synchronization databases, which may introduce the opportunity for data inconsistency and loss.
- **Windows-focused solutions.** NetApp's Global File Cache integrates with Microsoft's suite of enterprise products allowing enterprises to leverage other embedded services such as Microsoft Active Directory, DNS/DHCP, DNS, Microsoft Distributed File System (DFS) Namespaces, and Software Distribution Service in their branch office IT image. Global File Cache integrates fully with security principles such as those used in Active Directory, access control lists (ACLs), NTFS permissions, and DFS Namespaces.

Panzura CloudFS

Panzura CloudFS™ helps companies with distributed workforces accomplish work faster. Inherent within Panzura's collaboration solution is the Panzura Cloud File System (CloudFS™), a distributed file system specifically designed for the cloud. For end-users, CloudFS looks and behaves like local file storage. Intelligent caching keeps the most relevant data close to users at the edge. Panzura backs this edge storage in a centralized cloud data repository on public or private cloud storage. Enterprises may deploy the Panzura CloudFS three ways: on a physical Panzura Filer™, a virtual machine, or as a virtual instance in the cloud. CloudFS supports the latest NFS and SMB protocols.

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

- **File and byte-range locking.** In addition to locking at the file level, Panzura Distributed File Locking supports locking down files to the byte-range level—for applications that support it—to enable users from different offices to work in the same file simultaneously. Byte-range locking lets extended work teams edit different parts of large, complex files without data collisions, data corruption, or file versioning. Panzura updates changes in real-time, so every user always has access to the current version.
- **Real-time global data consistency.** Panzura's hub, spoke and mesh architecture provides immediate global data consistency to all locations in a file network through peer-to-peer services, even for data not yet sent to the cloud. Panzura exchanges lock ownership, changed data and file pointers between filers. It sends only unique

blocks of data over the network, reducing both network traffic and cloud storage. Panzura also uses deduplication and compression to accelerate and optimize data transfers.

- **Panzura Mobile.** Panzura Mobile™ natively integrates with Panzura CloudFS to add enterprise file sync and share to an organization's file services. It provides real-time data access for end-users and third-party collaborators on Android, iOS, Windows, macOS, and web-based clients. Policy-based controls restrict file and folder sharing to specific domains. IT departments retain complete control of shared content through public or password-protected links to Active Directory and non-Active Directory users. These features enable an organization's workforce and partners to collaborate effectively from any device anywhere.

Qumulo File Data Platform

Qumulo's file data platform enables enterprises to consolidate unstructured data with a single solution. Organizations can run Qumulo on industry hardware provided by Qumulo, by partners such as HPE and Fujitsu, and in the public cloud with AWS and GCP. Qumulo supports Windows, Mac, and Linux clients. Enterprises can access their data over standard protocols such as SMB, NFS, and FTP to support a variety of enterprise applications and user environments. To make costs transparent, Qumulo's offers its SDS solution as a subscription that covers everything, including upgrades, new features, and technical support.

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

- **Integrated real-time visibility.** The Qumulo file data platform allows organizations to monitor petabytes of data usage and performance across on-premises and the public cloud with real-time operational analytics (ITOps.) Customers may see, near instantly, capacity growth (or shrinkage), applications consuming resources, and the active parts of the file data platform. This enables organizations to troubleshoot, manage capacity, and plan using real-time data.
- **API-first data platform.** Qumulo provides programmable API endpoints for system creation, data management, performance, analytics, authentication, and accessibility. Qumulo allows enterprises to download and explore its self-documenting API. Customers can try each endpoint and see sample JSON outputs. Through Qumulo's API features, IT departments can build integrated solutions that automate administrative tasks, workflows, configurations, and data movements to enhance agility and save organizations valuable time.
- **Data-intensive file collaboration.** Qumulo's SDS solution provides a modern infrastructure for data-intensive media organizations for delivering animation, visual effects, gaming, and video content. Qumulo CloudStudio allows creative agencies, post-production studios, broadcasters, and visual effects studios to work on editorial projects such as creative content development, rendering, and production by moving applications and data to the cloud. Qumulo partners with a number of media and entertainment technology companies for integrated, collaborative solutions wherever creative teams may be located.

Inclusion and Evaluation Criteria for Enterprise Multi-site File Collaboration Solutions

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

- Commercially available as of May 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.
- If available as a pre-integrated appliance, the product must also be available on servers from multiple OEMs, as software that can be installed on hardware from other providers, or run in the cloud.
- At least one public cloud provider for support for deployment or a target.

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 and its merit.

- 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.

DCIG wants to emphasize that 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. 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.
4. DCIG identified solutions that met DCIG's definition for Enterprise Multi-site File Collaboration.
5. DCIG weighted each feature to establish a scoring rubric.
6. DCIG evaluated each solution based on information gathered in its survey.
7. Solutions were ranked using standard scoring techniques. ■

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.d cig.com.