

DCIG TOP 5 Storage for Healthcare Solutions

by Mike Matchett



TOP 5 Storage for Healthcare*

Nasuni

Dell PowerScale

IBM Spectrum Scale

Pavilion HyperParallel Data Platform

Quantum StorNext

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

SOLUTIONS EVALUATED:

- DDN EXAScaler
- Dell PowerScale F900
- Hitachi Vantara HCP
- HPE ClusterStor E1000
- Huawei OceanStor 9000
- IBM Spectrum Scale
- Nasuni
- Netapp HPSS (E-series)
- Panasas ActiveStor Ultra
- Pavilion HyperParallel Data Platform
- Pure Storage FlashBlade
- Quantum StorNext
- Qumulo File Data Platform
- Quobyte
- VAST Data Universal Storage
- Weka WekaFS

SOLUTION FEATURES EVALUATED:

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

Description of the Use Case

Healthcare IT must tackle significantly more challenging data handling requirements than most enterprises. In this environment, data storage is perhaps the most difficult architectural component to structure and operate both effectively for the mission and correctly for security. Not only are ultra-high availability, remote collaboration and fast data access critical to treating patients and achieving good patient outcomes, but also data protection needs to be paramount. And both compliance and data privacy regulations ensure that Healthcare IT has a lot of demanding work cut out for them in managing data access, integrity/retention and disaster recovery.

Because of these challenges, Healthcare is lagging behind many industries in IT modernization efforts. But there is incredible opportunity. Many healthcare organizations have now adopted key transformation initiatives to provide connected/remote care, conduct clinical and population health analysis, coordinate patient care across locations, improve patient engagement and foster more personalized/precision medicine.

Specific Healthcare workloads requiring aligned modern data storage can include PACs/VNA (medical imaging and image archives), telemedicine/telehealth distance sharing and collaboration (VDI etc.), EMR/EHR (patient records) and big data powered research analytics. These applications require local performant access to large media files (both images and video streams), secure sharing and collaboration across distances, and large on-line archives.

Imagine a common radiology scenario in which a successful patient outcome might depend on accessing and sharing a critical patient medical image in real-time between a hospital and an off-site pathologist. Large image data needs to be available on-demand, the image application needs fast streaming quality access, and online sharing needs to be supported and secured—all can be challenges for healthcare IT stuck with weak data storage solutions. In another common scenario, patient health records may be hard to access even internally (much less transferred to the point of patient care) due to overburdened or legacy storage systems that are still in use because of the fear of failing to meet legal compliance or security concerns (i.e. protecting data in a cloud or cloud-like distributed system).

Benefits of an Effective Storage For Healthcare Solution

The modern storage solutions evaluated in this report all have features that address IT modernization challenges and concerns. From a healthcare business perspective, they support and enable many of the desired Healthcare IT initiatives mentioned above. For example, these storage solutions can be used to modernize and consolidate unstructured file data storage that is often found locked inside aging NAS arrays local to each hospital data center. By aggregating local data into a global, distributed “file system”, data management policies can be applied immediately and universally while ensuring wide secure access for collaboration and sharing. Many of the solutions address performance with “edge” devices that can provide flash-speed streaming media access to local applications.

“It’s hard to imagine truly proving compliance and really protecting critical data without the support of a modern healthcare storage solution.”

It’s hard to imagine truly proving compliance and really protecting critical data without the support of a modern healthcare storage solution that can natively capture immutable copies, audit access, enforce retention/deletion, and recover on-demand. Storage with native compliance features and policy-based management directly addresses the data concerns found in HIPAA, HITECH, GDPR, PIPEDA and other regulations. In addition, data with provable integrity provides a better foundation to ensure healthcare insurance and billing are correct, complete and defensible.

Cost is a major concern in the healthcare industry; data storage can become a significant expense as useful medical data is growing at an exponential rate. The best storage solutions can help address growing storage cost/capacity concerns with features like deduplication, copy reduction, inherent snapshots or built-in backup/archive. Some will tier from flash all the way down to tape, while most enable integration to cloud storage to take advantage of utility scale and

pricing. And top storage solutions will make storage management easier on staff with many consolidating end-to-end storage management into a single pane-of-glass.

Finally, data protection requires security from both hackers and insider threats. Encryption features are critical. And ransomware attacks can be a life or death issue in healthcare. Top storage solutions will offer native features to help protect and/or recover quickly from ransomware attacks like immutable snapshots, WORM options, and quick restore to fine-grained recovery points.

Distinguishing Features of Storage For Healthcare Solutions

In addition to the broad feature sets mentioned above, all of the storage solutions evaluated in this report have a few things in common that distinguish them for the Healthcare use case from the wider market of IT data storage solutions:

- **Scalable capacity.** They scale capacity (to many PB) cost-effectively, enabling global NAS consolidation and the creation of a single file system architecture.
- **Performance at scale.** They have the ability to deliver low latency and high throughput streams to support the most demanding data-intensive applications even as they scale.
- **Interoperability and access.** They provide multiprotocol support and a wide range of connectivity to meet a variety of workload requirements while delivering IO to a large set of simultaneous, distributed users.
- **Architectural flexibility.** They generally can adopt or include a variety of underlying storage media, storage server node types and cloud extensions as storage tiering needs grow and expand with a proven track record of refresh and updating to new technologies.
- **Enterprise management.** They offer solid, mature management approaches and policy-based management for data protection, security, replication/synchronization, access and audit, monitoring and lifecycle data governance.

Similarities among the TOP 5 Storage For Healthcare Solutions

In addition to the distinguishing features above that all evaluated storage solutions share, the select DCIG TOP 5 Storage for Healthcare solutions all have the following traits in common.

- **Application-focused performance.** These storage solutions focus on delivering file shares to maximize consuming application (e.g. medical imaging, records retrieval) performance at scale and across many simultaneous distributed users. There is support for big data analytics and ML/AI performance too, but all the solutions baseline on delivering low latency user application requirements.
- **Designed for data-intensive verticals/Industries.** These solutions have been adopted across several different data-intensive verticals because of the benefits of having a globally accessible namespace, creating an enterprise-wide single data and storage management solution, and generally accelerating aggressive IT modernization initiatives.
- **Cost/capacity focused.** These solutions are all based on cost-effective and performant scale-out architectures. Additionally they all

offer cloud integrations (to AWS S3 and Google Cloud) for additional capacity and functionality.

The TOP 5 Storage solutions also deliver the following product features:

- Low latency and streaming throughput for unstructured data/media
- Hybrid storage options
- Ransomware mitigation features
- Pre-integrated appliance options
- NFS and SMB multi-protocol support
- Auto-tiering (including flash)
- Secure encryption for data at rest

Differences Between the TOP 5 Storage for Healthcare Solutions

However, not all storage solutions are the same. Many of these solutions have evolved from differing original design implementations. The DCIG TOP 5 Storage for Healthcare solutions can differ from one another in the following ways.

- **Deployment directly into cloud environments.** Most storage these days has been recast as software-defined even if mainly purchased via fixed appliances, but some of our top solutions are also certified to run—entirely or in part—in cloud environments with varying levels of functionality.
- **Ability to pool third-party storage.** Some of the solutions are capable of layering their file systems over underlying third-party storage of varying types, with some differences in tiering and data migration between those systems (including cloud storage services).
- **Provide S3 compatible storage API.** Some of the solutions now provide direct S3 API compatible object storage in addition to file storage. Solutions may or may not provide simultaneous access to the same data through both file and object protocols.
- **Space efficiency features.** Our top solutions all offer capacity optimizing features, but not all offer the same efficiencies or approaches. The details of deduplication, compression, incremental snapshots, versioning, and other features as well as implementation design and policy-manageable processes can vary greatly.
- **Encryption in-flight or in-cache.** Data protection is critical for Healthcare Storage, and it's important to note how, when and where each solution supports data encryption and key management to make sure that the storage architecture will support complete end-to-end privacy and compliance constraints.
- **Levels and types of tiering.** It's normal for various designs of storage solutions to support different levels and types of tiering. When evaluating a large capacity solution, one may need to dive into the details of the various quantities and types of media available or required, the ability to expand or upgrade tiers over time, and the relative intelligence of any auto-tiering or policy-based tiering provided.

Top 5 Storage for Healthcare Solution Profiles

Each of the DCIG TOP 5 Storage for Healthcare Solution Profiles highlights three notable solution features that make the product attractive for this market.

Nasuni

Nasuni is enterprise file storage delivered through SaaS. Nasuni cleverly combines the best of distributed file and cloud object storage together to bring their customers both the unlimited on-demand capacity scaling and utility economics of cloud. Nasuni's integrated storage simplicity comes from having a single management console for global file system provisioning and management. The single SaaS Nasuni Management Console covers all storage, edge appliances, volumes, snapshots, protocols, and shares in the system, across all cloud and on-premise deployments.

Three of the key features that earned Nasuni a spot among DCIG TOP 5 Storage for Healthcare solutions include:

- **Distributed access with fast local performance.** One of the main architectural advantages to Nasuni is the ability to collect regional NAS consolidations into a global storage domain based on back-end cloud storage. Nasuni then fosters sharing and collaboration of that cohesive data set through global distributed access using both multi-file synchronization and global file locking. For performance, the Nasuni Edge Appliances cache active data to deliver fast local data access.
- **Natively integrates storage, backup, DR and file sync.** At the edge, Nasuni's cloud-based global file system, UniFS, de-dupes, encrypts and stores all incremental changes with immutable copies in the cloud. This obviates the need for separate backup (and DR) processes or systems. On the backend, Nasuni stores all files, metadata, and nodes as cloud storage objects, with the benefits of inherent data protection and fast recovery of past versions. Data sets can be directly restored or cloned in the cloud for on-demand analytics. UniFS ensures compliance with end-end encryption of data both at rest and in-transit to the customer's cloud object storage, whether on-premises or in the cloud.
- **Cloud-based scaling, capacity and economics.** Nasuni can utilize cloud-side storage across all three major cloud platforms (AWS, MS Azure, and GCP) or any cloud object storage platform of choice. This flexibility eases migrations and avoids vendor lock-in. Nasuni Edge Appliances can be run as VMs on virtual or hyperconverged infrastructure, either on-premises or fully in the cloud.

Dell PowerScale

Dell EMC PowerScale is a scale-out NAS array based on the latest version of OneFS (developed originally as part of the Isilon architecture which is now a subset of the larger PowerScale storage line). OneFS effectively creates a unified pool of storage from across scale-out storage server nodes and can be implemented as high performance storage using NVMe PowerScale nodes, tiered hybrid storage, or capacity-oriented storage using archival nodes running inline compression and data deduplication. A mature enterprise platform, OneFS offers a large portfolio of software options ranging from CloudIQ performance analytics to SyncIQ secure data replication and CloudPools for tiering off colder data to cloud storage.

Three of the key features that earned Dell PowerScale a spot among DCIG TOP 5 Storage for Healthcare solutions include:

- **Scalable data management.** PowerScale OneFS now scales up to 252 storage nodes using any mix of PowerScale and Isilon all-flash, hybrid, or archive-focused nodes, offering wide choice to meet specific workload requirements and extending the longevity of existing Isilon investment.

- **Availability, redundancy, data protection and replication.** With the large software portfolio available, PowerScale users can deploy and configure a wide range of proven enterprise storage management features as needed. Software packages include dedupe, analytics, heterogeneous data management, performance monitoring, on-demand snapshots, multi-site replication, compliance locking, load balancing/failover, cloud tiering and smart storage tiering.
- **Multi-protocol simultaneous access.** An advantage in supporting mixed workload workflows and enabling workload NAS consolidation, OneFS natively supports simultaneous data access through NFS, SMB, S3, HTTP, FTP and HDFS (Hadoop) protocols. This multi-protocol support coupled with scale-out performance characteristics is one reason why historically Isilon has been a top pick for key mission-critical unstructured data storage.

IBM Spectrum Scale

IBM Spectrum Scale is an enterprise parallel file system solution based on IBM General Parallel File System (GPFS) designed to enhance resiliency, scalability and control. Spectrum Scale can unify flash, disk, cloud and tape into a high performance, high capacity, and cost-efficient single storage solution. A single namespace, single management GUI and policy-based management make Spectrum Scale enterprise-ready. Parallel file system clients achieve high performance using the appropriate parallel file system drivers, accessing metadata from designated metadata servers but streaming file data at high speeds directly from parallel storage nodes.

Three of the key features that earned IBM Spectrum Scale a spot among DCIG TOP 5 Storage for Healthcare solutions include:

- **High performance parallel file system access.** GPFS is a widely recognized HPC-class clustered file system, capable of delivering world-class performance from either shared-disk or shared-nothing distributed parallel storage nodes. Scale-out clusters can grow into the thousands, and although GPFS really shines at serving large files (e.g. video/media) at high speeds to many users at once, there have been recent improvements in serving larger numbers of smaller files as well (variable block sizes, metadata performance, et al.).
- **Global file access.** IBM Spectrum Scale can provide the performance and scale cornerstone of an "enhanced global data fabric" leveraging other Spectrum solutions. In addition to file and container native support, Spectrum Scale now incorporates high performance S3 access API for fast multiprotocol access across workflows (e.g. ingest as files from data sources, serve as objects to analytics packages).
- **Simplified data management and security.** IBM Spectrum Scale offers a single GUI and storage wizards to ease management across potentially yottabytes of storage capacity across multiple environments and data centers, tiering even into exiting NAS, object storage, tape, and clouds. IBM Spectrum Scale can now be thought of as a hybrid cloud file system, and the Spectrum family offers mature architectural options to meet any enterprise requirement.

Pavilion HyperParallel Data Platform

The Pavilion HyperParallel Data Platform provides massive IO parallelism through its unique architectural approach that basically hyper-converges a high-speed network switch with a unified flash storage array. The base 4RU system can hold over 2PB and up to 20 controllers, while serving

high-performance block, file and object I/O as every controller can talk directly to any disk. Pavilion can cluster an unlimited number of these base arrays with linear (or better than linear) performance scaling.

While the Pavilion HyperOS and its inherent Pavilion HyperParallel File System provide a combined NFS/S3 global namespace across an unlimited number of array units and assure both performance and availability at scale, Pavilion's flash arrays can also host other high-performance file systems (e.g. IBM Spectrum Scale) if desired.

Three of the key features that earned Pavilion HyperParallel Data Platform a spot among DCIG TOP 5 Storage for Healthcare solutions include:

- **Unlimited linear performance (and capacity) scaling.** A remarkable storage density enables Pavilion HyperParallel platform users to start fairly small but grow both scaling up (to 2PB per 4RU node) and scaling out (unlimited) cost-efficiently as necessary. Because of the integrated network switch design, performance across clusters can be linear with growth.
- **Built for serious workflows.** The Pavilion HyperOS and HyperParallel File System provide a native multi-protocol access to enable a single, unified storage repository to power both very high-speed data ingest (90GBps writes) and parallel client consumption of fast data in disparate forms. Protocols include NVMe-oF (RDMA, IB or TCP), iSCSI, NFS and S3 with native support on any combination of controllers across any number of arrays. Pavilion can be easily connected to any cloud in just a few clicks.
- **Integrated enterprise data services and application plugins.** Tiering, replication, snapshots, clones, security, encryption and application plugins (popular S3-based apps, ML, Big Data, Spark, TensorFlow, Kafka, Splunk, Teradata and more) are all included in the core platform. Hardware upgrades and cluster expansion can be made non-disruptively, and the platform design eliminates single points of failure common to other arrays.

Quantum StorNext

Quantum StorNext is a shared-storage platform that layers a high-speed file system over many different types of underlying storage. StorNext can serve a single namespace with optimized metadata handling, and natively integrates key data lifecycle management. StorNext is well known for its support for large repositories of unstructured data including media files, research data and digital imagery that can be mission-critical to many industries (like healthcare). StorNext offers a wide range of client connectivity options, ensures desired data protection across the data lifecycle, and has the performance to serve large numbers of client video streams.

Three of the key features that earned Quantum StorNext a spot among DCIG TOP 5 Storage for Healthcare solutions include:

- **Broad client connectivity options.** StorNext provides a direct client (NFS/SMB) for parallel client access across MacOS, Windows and Linux clients and supports a wide range of network connectivity options including iSER/RDMA, Fibre Channel, iSCSI, and Infiniband. StorNext also supports S3 for multiprotocol access to help foster data sharing and collaboration. Up to 2,500 parallel video streams can be served to StorNext Direct Clients with sub-millisecond latency.
- **Unstructured data consolidation.** StorNext can be implemented across Quantum's own storage appliances, VM's, and in the cloud. It can also readily integrate archival solutions like tape or long-term cloud

storage. StorNext features include Stripe Groups that enable array aggregation and intentional data alignment to performance needs along with file system pooling to enable transparent primary storage tiering.

- **Enterprise data lifecycle and services management.** The StorNext File System (SNFS) is POSIX compliant and enables data movement across primary storage tiers based on policies for aligning security, integrity and accessibility with price/performance. StorNext also offers FlexTier for secondary storage tiering, FlexSync for replication, Quotas, and QoS solutions.

Inclusion and Evaluation Criteria for Storage for Healthcare Solutions

In this report, DCIG specifically focused on Storage for Healthcare solutions possessing the following characteristics. DCIG identified fifteen different solutions meeting these inclusion criteria:

- Commercially available on December 1, 2021.
- Sufficient, publicly available information available for DCIG to make an informed decision.
- Clear vendor intention to support Life Sciences and Healthcare use cases as evidenced by publicly available solution-focused content.
- Demonstrated business in these industries.
- Capable of supporting global initiatives, meeting demanding application file performance requirements, and scalable to PB's.

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, QoS features, auto-tiering capabilities, and directory service integration.
4. **Documentation support.** Evaluate the breadth and depth of documentation the provider makes available to customers. Examples include whitepapers, knowledge bases, online manuals, as well as 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.

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 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

DCIG wants to emphasize that no vendor was privy to how DCIG weighted individual features. In every case the vendor only found out the ranking 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. Providers were also given the opportunity to complete surveys about their products.
4. DCIG identified solutions that met DCIG's definition for a Storage for Healthcare solution.
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.dcig.com.