


White Paper

Nasuni UniFS[®] – a True Global File System

File systems are the time-proven way to store, share, and protect unstructured data. But traditional device-based file systems no longer meet the needs of the modern digital enterprise. This white paper provides an in-depth look at new global file systems designed for cloud storage, and how Nasuni UniFS sets the standard.





What is a Global File System, and Why is One Needed?

The traditional technologies being used to store and protect file data, combined with the need to share files across multiple locations, have created enormous management challenges for IT organizations already strained by relentless data growth. Enterprises facing the combined pressures of increasing file sizes, workforce mobility, and an expanding global footprint require a next-generation file system that can scale beyond the bounds of an individual disk array or even a Network Attached Storage (NAS) cluster.

The solution is a global file system. This next-generation file system resides in the only storage scalable enough to meet the requirements of the modern digital enterprise – cloud object storage. By completely separating file data and metadata from any hardware dependencies and living natively in the cloud while extending on-premises using local appliances, a global file system can elevate IT conversations from the nuts and bolts of file storage (provisioning, file sharing, protocols, backup, DR, etc.) to a strategic discussion about unstructured data management.

IT organizations can now align with line of business managers to answer questions like:

- Who needs access to different file data workloads?
- Where do they need it?
- What level of performance is required?
- What level of collaboration is required?
- What are the recovery point and recovery time objectives?
- Can users recover their own files through self-service, or is IT assistance a requirement?

This white paper examines the requirements of a global file system. It then provides a deep look at the first global file system designed for the modern cloud era, Nasuni UniFS®, to see how it measures up.

Global File System Requirements

Public and Private Cloud Object Storage Integration



For a file system to be truly global, it needs an affordable, limitless reservoir of storage to contain all the possible files, directories, and metadata generated by a global enterprise. Traditional NAS devices and array clusters cannot meet these requirements, especially when files need to be accessed from multiple locations and replication or remote access is required (adding significantly more cost).

Object storage has emerged as the affordable and infinitely scalable alternative to on-premises storage. Azure Blob, Amazon S3, IBM Cloud Object Storage (COS), Dell EMC Elastic Cloud Storage (ECS), and other public and private cloud options now offer bottomless pools of globally accessible object storage. All that's needed is a global file system that integrates with them to make these platforms truly useful.

A global file system integrates with all leading object storage platforms to give enterprises the option of using one or more of them to meet location-specific performance, data residency, and other regulatory requirements.

Unlimited Scalability and Global Access

A global file system provides a hypervisor-like layer that separates files from storage resources, managing one master copy of data in public or private cloud object storage while distributing data access where it's needed. The global file system manages all metadata such as versioning, access control, audit records, and locking and provides access to files via standard protocols such as CIFS/SMB and NFS.

A global file system is different from a distributed file system or a global namespace. In these models, files are tied to a specific piece of hardware, and the distributed file system or global namespace finds the file and directs access to it. Sometimes replication is added on top to bring file access closer to the user.

With a global file system, files are untethered from hardware and stored in public or private cloud object storage. Active files are then cached locally so users continue to enjoy high performance access through existing drive mappings and share points. A caching algorithm keeps the active, in-use data in cache so it is always locally accessible.

All files, including files in use across multiple local caches, have their master copies stored in cloud object storage, so they are globally accessible from any access point. With this model, the global file system provides the unlimited scalability of the cloud while maintaining the local performance of a traditional NAS device.

A global file system is also different from a simple cloud storage gateway or cloud-integrated storage. In these models, the cloud simply acts as a secondary



tier of storage tied to the single gateway device. By contrast, a global file system treats the cloud as the primary tier of storage, and treats the caching devices at the edge as stateless, secondary tiers whose primary function is to provide fast file access in any location.

The global file system provides unlimited file storage scalability, global access and mobility, and shared use of the same data across all offices. Yet, end users in every location feel as though they are simply using a local file server.

File Locking for Conflict-Free Collaboration

Since a global file system enables any number of users in any number of locations to access the same file, it must provide a mechanism to prevent file contention and version conflict that could result in data loss. That mechanism is a global file lock.

With global file locking, if a user opens a file in one location, any subsequent opens by other users will result in an error message that the file is already in use. To provide a truly global file lock that spans any number of locations, the global file system typically implements file locking as a scalable cloud service.

If a file is left open by a user or a local network or server event, a global file system allows IT to break a lock. Even in the case of a manually broken lock, any file contention is handled by the file system with conflict resolution procedures so no data is ever lost.

Right-Sized Performance with Caching at the Edge

The global file system stores the master copies of all files in the cloud to provide true global scale. However, users accustomed to the speed of local file servers and NAS devices shouldn't have to deal with the latency of fetching files from the cloud every time they need file access.

With a true global file system, they don't have to. The global file system extends on-premises to any location to provide cached, high performance access to active files. This is typically done with small local appliances installed in data centers, remote and branch offices, or even temporary job sites. These local appliances, which can be virtual or physical, look like a local NAS device or file server to users, and to IT administrators as well.

Files are served up via standard CIFS/SMB and NFS protocols, or even FTP. Windows users access files through the usual "letter" drive (e.g. G:, S:) while Mac users access files through the usual "Connect to Server" choice on their Finder. Mobile device and desktop users can also use Web browser and local apps for file access.

A global file system stores only the small amount of metadata needed to access the active files in each local cache. This ensures optimum data propagation performance and reduces the amount of cache storage needed in each location. Gateway solutions that don't offer a true global file system store all the metadata



for the entire file system on every device, increasing both latency and the size (and cost) of every local cache.

Since each local appliance makes the global file system look and act like a traditional file server, applications using the files in a global file system function as they always have. Appliances can be scaled to fit the application workloads in each office to provide both the I/O performance and capacity required, regardless of whether users access files from an application or directly through Windows Explorer or Mac Finder.

Multiple caching appliances can be deployed for high availability to ensure continuous data access, or to meet specific performance requirements and user service levels.

Infinite File Versioning in the Cloud

Another function provided by a global file system is versioning. As users change file data on an appliance, just the changes are sent continuously to the cloud. With the economics of cloud storage and the ability of the global file system to only send the tiny “shards” of data that have changed, versions of every file can be kept forever (unless requirements call for purging data after a certain time).

By using the cloud’s inherent reliability, durability, and geo-redundancy (physical copies of data stored in multiple locations far apart), infinite file versioning eliminates the need for costly backup and archive systems, saving massive amounts of time, money and frustration. File versioning also offers more recovery points (every change) and improved recovery times (previous file versions can be quickly pulled back from the cloud through user self-service or IT assisted service).

Centralized Management - Software Defined File Services

The term “software defined” is now applied to almost everything in IT. The reason for this is the power of using software-based policy and configuration to control things that previously required hardware reconfigurations.

Although overused, the term describes an important benefit of virtualization. The increased flexibility of software-defined architectures enables more efficient use of resources, and enables management to be centralized so all resources are coordinated in a more efficient and secure manner.

A global file system applies software-defined principles to enterprise file services. Policies can be defined and implemented through a single management console for many petabytes of data consisting of millions, or even billions of discreet files, instead of having to reconfigure silos of NAS devices and backup products in many locations.

File accessibility, access control, protocol access, locking, quotas, audit policy, share/directory configuration, volume provisioning, caching appliances, and more



can all be centrally managed, and the state of the file infrastructure in any location can be centrally monitored.

With a global file system, IT can efficiently manage and control the global file infrastructure and ensure business needs are being met with much less cost and must less burden on staff.

Improved Security

A global file system can actually increase data security compared to traditional file servers and NAS devices, even when a public cloud is used to provide the object storage.

The reason is encryption. A global file system encrypts all data using encryption keys owned and stored by each customer. When data is changed, or created on a local caching appliance, it is immediately encrypted before being sent to the cloud. No data is ever in the clear while it is being sent or retrieved (in motion), or while it is stored (at rest). In addition, no keys are stored with the encrypted data.

For local access and authentication, the global file system integrates with and joins on-premises Active Directory and LDAP infrastructure. This enables the existing enterprise permission system to be applied to all data stored on the local caching appliances and prevent unauthorized access, just as with local file servers or NAS devices.

Nasuni UniFS – the First Global File System

Now that the requirements of a global file system are well understood, it's time to look deeply at Nasuni UniFS, the first global file system designed for the cloud.

Public and Private Cloud Object Storage Integration

Nasuni UniFS integrates with the leading object storage platforms, with hundreds of enterprise deployments on Azure Blob, Amazon S3, IBM Cloud Object Storage, and Dell EMC ECS. Customers are not locked into any one provider. They can switch at any time, or use more than one at the same time by mapping different UniFS volumes to different object stores.



Figure 1: The Nasuni hybrid cloud platform is powered by Nasuni UniFS, the first cloud-resident global file system.

One Master Copy in the Cloud – Accessible Everywhere

More than just integrating with the leading object storage providers, UniFS uses a “cloud-first” architecture that stores the master copies of all file data and metadata in the cloud. By using the cloud as the primary tier of storage rather than a secondary backup tier, file infrastructure becomes infinitely scalable, well-protected, and globally accessible, since data is not dependent on any given device.

UniFS is organized into volumes that contain the structure of directories and files, as well as metadata such as access control lists, modified date, creation date, lock state, and more. When a volume is instantiated on a local caching appliance (Nasuni Edge Appliance), the volume structure is built in the cloud. As files and directories are created, data is chunked, de-duplicated, compressed, and encrypted, then stored as objects in the cloud.



Mappings between the volume structure of files and directories, metadata, and objects representing the file data are all stored in the cloud. Any Edge Appliance given access to a UniFS cloud volume can map to the volume and read and write files and directories using the same process. In this way, access to the master data in the cloud is local to any user, in any location.

UniFS operates using WORM (write-once, read-many) principles. Once an object is written to the cloud, it is permanent. This makes UniFS stable and resilient. As data is changed on Nasuni Edge Appliances, the changes are sent to the cloud as new versions, enabling the current or any past version of any file to be quickly retrieved.

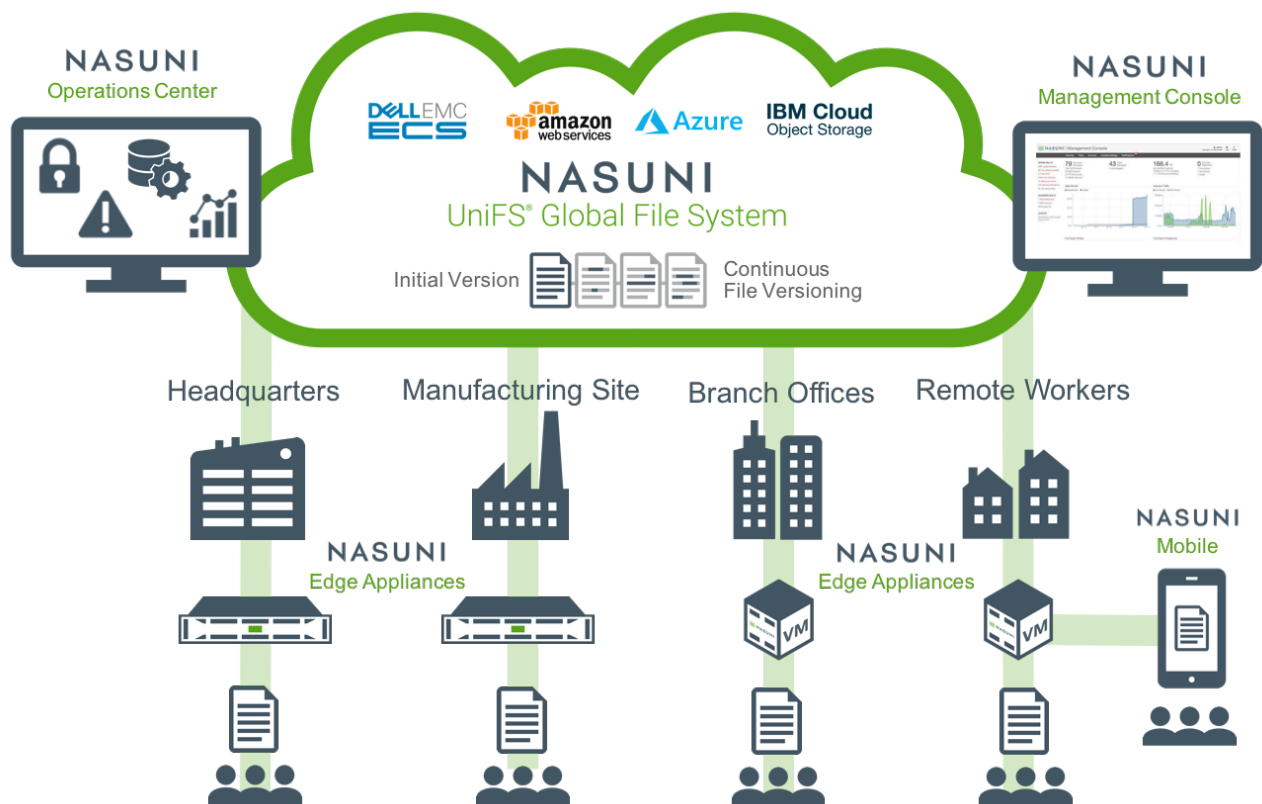


Figure 2: Nasuni UniFS extends from the cloud to on-premises locations using small Edge Appliances, which cache only the active files and the fractions of metadata needed to represent them.

Infinite Versions and Instant Restores

Leveraging the scalability and low cost of cloud storage and efficient “sharding” technology that only captures the tiny fragments of files that have changed, Nasuni UniFS can version files as frequently as once a minute, and keep them forever.

Third party backup and archive infrastructure, along with its high cost and complexity, is no longer needed, since a complete version history of every file is

always available at any time. UniFS enables files to be restored to any point in time at the file, folder, or volume level instantly, just by restoring a small amount of metadata first.

End users can navigate to previous versions and perform self-service file restores themselves based on access control permissions. Or, IT can choose to require all recoveries be helpdesk-assisted.

Local and Global Locking

UniFS supports file locking on any CIFS/SMB or NFS share. When an application opens a file with a lock, it will appear to the application exactly as it does on any other NAS device or Windows file server. The application will behave the same as it always does, and there is no change to the user's experience whether collaborating with someone in the next office or across the globe. This also ensures that file locking with UniFS is compatible with every application and requires no special integration or management.

When users who write data are located in many different sites, UniFS integrates with Nasuni's cloud-based Global Locking Service to extend the lock globally. This enables distributed users to collaborate on files wherever they are without risk of losing or corrupting any data, while preserving all data changes. The Global Locking Service runs in the cloud so locking is always available and scalable to the global enterprise without dependencies on any device to maintain lock states.

Directory Integration

As a global file system that exists fully in the cloud but extends on-premises, Nasuni UniFS augments the security features of its cloud partners with strong on-premises security. Nasuni Edge Appliances join Active Directory and LDAP domains to leverage each customer's existing authentication and access control procedures.

Data is accessed just as it would through a Windows File Server or any traditional NAS device using existing credentials and identities. All identity and user controls already in place still apply. Existing ACLs can be migrated in with the data, or the migration to Nasuni can be used as an opportunity to clean up polluted ACL structures.

High Performance File Access

When files are cached locally by UniFS on Nasuni Edge Appliances, users and applications have the exact same experience they would have with traditional, full-sized NAS controllers and file servers.

The big difference, of course, is the Nasuni Edge Appliances are typically only 10-20% of the size of a full-sized NAS device. Even though customers store hundreds of terabytes on Nasuni, UniFS only caches the most important and most frequently accessed files on each appliance.

By using sophisticated algorithms that drive a 98% cache hit rate across all Nasuni customers, users gain all the benefits of local file server performance even though only a small percentage of the data is in the local cache.

This reduction in the amount of hardware needed for file sharing is just one of the ways Nasuni reduces costs up to 60% compared to traditional full-sized NAS infrastructure.

Advanced Security

Nasuni's data security model begins with a solid foundation of strong encryption. The Nasuni Edge Appliances that send file data to the cloud while caching the active data locally for "NAS-like" access use encryption keys generated by customers. Encryption with customer-controlled keys ensures data can never be viewed or used by anyone outside the organization. Neither Nasuni nor the cloud provider (e.g. AWS, Azure, IBM, Dell EMC, etc.) can "see" the data.

Each Nasuni Edge Appliance performs the encryption on-premises before sending the data off-premises, so the data is always encrypted both in transit and at rest.

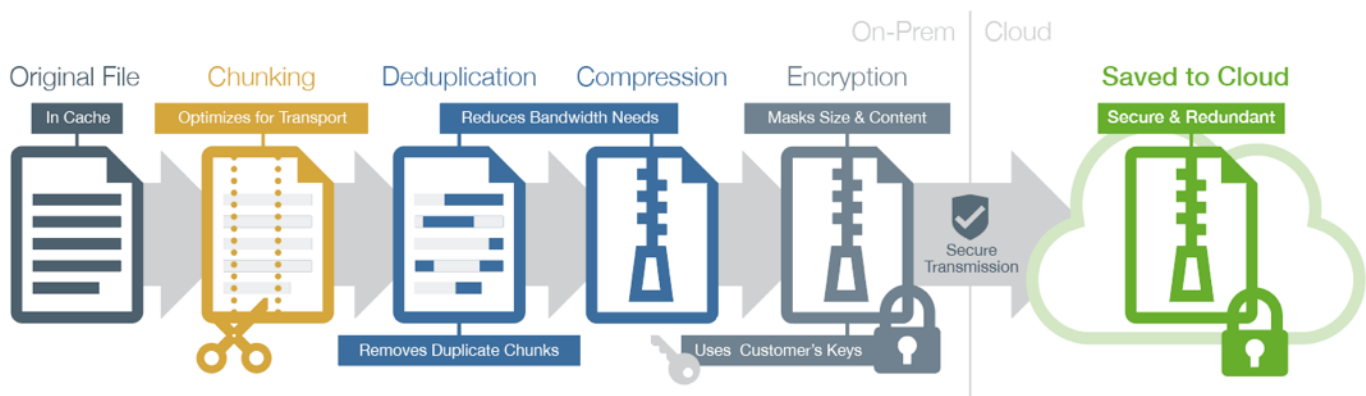


Figure 3: Flow of data as it originates on-premises on Nasuni Edge Appliances until it reaches the cloud where it is securely stored.

Nasuni employs the non-proprietary OpenPGP protocol for public-key-based encryption and decryption. OpenPGP establishes a framework for how to combine widely available security algorithms into a secure system. OpenPGP's open standard and source code are continually enhanced through an extensive and thorough review process.

OpenPGP combines symmetric and asymmetric encryption technologies that not only protect the data, but do so without compromising performance. Using fast symmetric encryption to encrypt the data and slower asymmetric encryption to encrypt the keys enables data to be encrypted efficiently and at a high level of granularity. OpenPGP also specifies several important details, including proper salting (inputting random bits to a one-way cryptographic hash function) and cipher modes. OpenPGP's cipher feedback (CFB) mode also avoids the drawbacks of less secure techniques, such as Electronic Codebook (ECB).



Along with OpenPGP, Nasuni employs the AES-256 standard for symmetric encryption. AES is the first publicly accessible and open encryption standard approved by the US National Security Agency (NSA) for top-secret information.

In addition to encrypting the data itself, Nasuni Edge Appliances also encrypt metadata, both in transit and at rest. This means no identifiable information – not even file names or timestamps – is decipherable once it leaves customer premises. Encrypted file metadata includes the file name, file size, timestamps, access control information and location within the directory tree.

Bandwidth Efficiency

A perpetually cloud-connected storage system would saturate an enterprise network, especially if the local storage controllers need to frequently talk to each other.

This is not the case with UniFS. Integrated WAN optimization enables UniFS to be deployed all over the world – even in offices and locations with limited and inconsistent connectivity.

UniFS running on Nasuni Edge Appliances communicates directly with the cloud – never directly to other appliances. This leverages the bandwidth of the cloud instead of the limited connectivity between offices.

In addition, before any data is transmitted, it is chunked, compressed and de-duplicated against the global file system so that nothing is sent twice.

Global deduplication means UniFS is comparing file changes not just locally, but to every file stored in the cloud – dramatically reducing the amount of data that needs to be sent. Sub-file chunking ensures UniFS transmits only the de-duplicated parts of a file that have changed instead of the whole new version. Compression of every chunk before transmission removes extraneous data, further reducing bandwidth needs.

This reduction in the amount of network bandwidth needed for data propagation by using the cloud is another way Nasuni and UniFS reduce costs up to 60% compared to traditional multi-site NAS infrastructure, which requires costly MPLS bandwidth or WAN acceleration to bring users closer to their files.



Conclusion

In the same way server virtualization enables compute needs – not server hardware – to be at the center of application deployment and management, a global file system enables file sharing and data protection needs – not storage hardware – to be at the center of unstructured data management.

Hardware becomes a stateless resource that can be adjusted to deliver the required level of performance and secure access. Central management moves to a core global service that can monitor and manage every component regardless of its location.

To architect a global file infrastructure, one must think beyond the confines of any single physical appliance or data center. This is what a global file system enables: a single source of file data that is accessible everywhere, is well-protected, and is tightly secured.

Nasuni UniFS is the first global file system designed to live in the most scalable storage available - cloud object storage – and extend on-premises to meet the security and performance requirements of traditional NAS and file servers. The benefits of UniFS covered in this white paper can be summarized as follows:

- Unlimited, elastic capacity for file shares, project directories, and user home drives
- Performance of local NAS with “right-sized” edge appliances in every location
- Global file access and global file locking for multi-site collaboration
- Continuous File Versioning with advanced recovery points and recovery times, eliminating the need for traditional file backup
- Simplified, central management of all volumes and shares
- Disaster recovery in 15 minutes in any location that can connect to the cloud
- Single authoritative source of all files with complete audit trails
- Dramatic cost reductions by consolidating NAS, backup, archive, DR, WAN acceleration, and replication into a single cloud-scale technology.



About Nasuni

Nasuni (“NAS Unified”) is transforming the way enterprises store, share, protect, and manage fast-growing file data. Powered by Nasuni UniFS®, the first global file system as scalable as the cloud itself, Nasuni’s hybrid cloud file services platform combines the limitless capacity, geo-redundancy, and low cost of object storage with the security, performance, and flexibility of local file servers. By using Nasuni and their preferred cloud provider for Network Attached Storage (NAS) consolidation, multi-site file collaboration, and active archiving, Nasuni customers are meeting global growth, workforce productivity, and “cloud-first” objectives, while also realizing massive IT cost savings. Nasuni is based in Boston, Mass. For more information, visit www.nasuni.com.