



# How to Replace Windows 2008 File Services with the Cloud



***Organizations using Windows 2008 for their file services have a decision to make. On January 14th, 2020 support for Windows 2008 and 2008 R2 ends. Since Windows 2008 is a favorite file server solution, organizations need to find a new place to store and share files. Cloud Storage is an option that many IT planners are considering, but most cloud services only provide the most basic capabilities. Organizations need to look for a solution that fully takes advantage of the cloud to solve the historical challenges with NAS systems. Organizations also need to look for a provider that understands Windows.***





# CHAPTER 1:

## The Requirements of Next Generation File Servers

File servers used to be a way for users to share information within the organization, typically within the same building. Organizations of all sizes use Windows File Servers to meet this requirement. Times have changed and the need to collaborate continues to increase encompassing more users and locations. Traditional file servers can no longer keep pace with these demands.

Users, and in fact organizations, are no longer self-contained within a single building. Organizations have multiple data centers and their users work from a variety of locations around the world. At the same time the workload is often shared between these offices with users working on the same files at different times. Finally, the 9 to 5 workday is a relic of the past, with users connecting to each other at all times of the day, and global clients requiring round the clock service.

### The State of the File Server

Organizations need to reconsider their file server strategy for two reasons. First, traditional file servers can't keep pace with the new expectations of them. Just serving files is table stakes, the focus is now on enhancing user productivity and securing data.

Second, and potentially a more pragmatic reason to reconsider the organization's file server strategy is that Microsoft announced the end of support for Windows Server 2008 which means the thousands of organizations world-wide that use Windows File Servers 2008 need to look for something new. The default upgrade is to use the next version of Windows Server. But as even Microsoft points out, this often requires a hardware upgrade from the systems that ran 2008. IT planners instead should take advantage of the opportunity to explore new options that leverage the cloud to provide a truly global distribution of data and improve data protection.

### The Requirements of the Next Generation File Server

Modern File Servers must provide file sharing and collaboration across multiple sites. The challenge with providing multi-site collaboration is how to properly share data across those sites. Forcing users through a virtual private network (VPN) to a centralized server can significantly slows down performance and do not always deliver maximum security. Instead, modern file serving solutions should leverage the cloud as the primary storage area and as a mechanism to share data securely and efficiently.



Leveraging the cloud also addresses the next requirement of next generation file servers, scaling. Unstructured data continues grow at an unprecedented rate with no end in sight. Traditional file servers quickly run out of capacity, which leads to the implementation of more and more file servers and storage. The cost and complexity of managing multiple NAS solutions is prohibitive. By contrast, the cloud provides a single method of storage that can scale almost infinitely.

The cloud, however, causes a challenge in meeting the third requirement of next generation file servers, high performance. Performance is increasingly a critical element in file server consideration, both to meet user expectations and the very real demands of applications that interact with network shares.

To overcome the performance challenge, a cloud centered solution should provide an edge capability so that for each location's most active files the users experience fast, local performance. These edge devices can be flash based without breaking the budget. All files go immediately to the cloud, but active files are available quickly. Much of the IP of the next generation file server is in managing the distribution of data to the edge so that its capacity is used efficiently, and the users rarely experience cloud latency.


Another key requirement is reduced infrastructure and administration costs. Since the cloud can be the centralized repository for file data, capacity management should be greatly simplified. Administrators don't have to play a guessing game regarding how much capacity they may or may not need at each location. The typical 3-year planning cycle can be abandoned for a much simpler pay as you go model. Perhaps best of all, there is no wasted capacity – every TB purchased can be based on need. Provisioning doesn't require physically installing disk racks and can be done with the touch of a button.

The ability to maintain security and access control is probably one of the most important capabilities for next generation file servers. Moving from one solution to another, for example from Windows File Server 2008 to a cloud file service based in Azure, should preserve all security protocols while enabling users to continue working uninterrupted.

Our next chapter will discuss the next requirement; data protection which as the data set grows is more challenging than ever and as new data privacy laws and regulations come into place is more important than ever.







## CHAPTER 2: Protecting Distributed File Data

Where previously unstructured data comprised a minority of the business's data and was not strategic to the organization, that scale has tipped dramatically. Unstructured data may account for two-thirds or more of the data that businesses are collecting, storing, and through analytics and machine learning, uncovering insights that are moving the needle when it comes to competitive advantage. The desire for the insights is also increasing unstructured data retention times.

The robust growth of unstructured data is contributing to complexity when it comes to serving file data. File types vary in size and type, they are distributed globally across the organization's locations, and must be accessed by a user base that is equally as dispersed. Previously, Storage Switzerland discussed the collaboration and performance challenges that are inherent with trying to serve this environment with legacy network-attached storage (NAS) arrays. In this installment, we will discuss how to protect this data – which is no easy feat, especially in the face of strict data privacy laws.

### Protection is No Longer a Nice-to-Have

In the past, many organizations left their unstructured file data unprotected due to costs and complexities. The scale of a required data protection implementation often outpaces budgets and, from a management perspective, IT staff resources.

In today's world, however, this is no longer an option. In addition to increasing criticality, data privacy regulations including the European Union's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) requires an organization to gain comprehensive visibility into the data they are storing, and how that data is being stored as well as accessed. Furthermore, these regulations are causing consumers to be more mindful of how their data is being stored and used by the business. Business requirements including analytics and DevOps are furthermore creating the need for more data to be stored in a way that will preserve its integrity.

## Traditional File System Backup

Backing up unstructured data is time consuming. Also, using a typical image-based backup technology sacrifices file granularity, which makes adhering to regulations like GDPR more difficult. Administering and recovering from backup data such as snapshots is expensive and cumbersome. The snapshot technology that typical backup solutions rely on consumes expensive primary storage capacity. Furthermore, most snapshots rely on metadata for tagging, which greatly taxes the storage system's Input/Output Operations Per Second (IOPS). In fact, many legacy file systems cap the number of snapshots that can be retained in the interest of preserving system performance. Snapshot vaulting and mirroring may alleviate some of these headaches, but may still impose retention limitations, and require the need to invest in additional capabilities such as replication software and wide-area network (WAN) acceleration – not to mention storage infrastructure resources and software licenses to host snapshots.

## A Newer Approach: Continuous File Versioning

Continuous file versioning is emerging as a less expensive and more agile approach to ensuring data availability and recovery. Like a more traditional snapshot technology, continuous file versioning starts off by capturing an image of an entire volume. From there, however, subsequent snapshots capture only changes to the storage volume – thus facilitating more efficient use of storage capacity and network bandwidth (the latter being especially true when images are stored in a cloud service). Storage managers should look for a continuous file versioning technology that employs a data sharding process, which enables a smaller amount of changes to be captured and processed in parallel, to avoid an impact to performance and to increase scalability.

## Rapid Disaster Recovery

In today's around-the-clock business world, server downtime and loss of data access can wreak havoc on the company's operations and credibility. The impact can be even more devastating when multiple departments and locations are relying on a particular location for key information. When it comes to ensuring speedy, secure and high-performance access to file data, infrastructure planners and storage managers should consider investing in a solution that preserves multi-site access to files, and that allows the "down" site to be restored as quickly as possible. Users throughout the enterprise should be able to access files related to their workflow regardless of the outage. Likewise, the standard for restoring onsite file access should be minutes, not hours or days. Solutions that are designed to rapidly restore edge resources and rehydrate these resources with frequently accessed files offer strong advantages here.

Our next chapter will explore the value of the cloud for facilitating distributed file data availability and protection, as well as scalable and global secure access to this data.

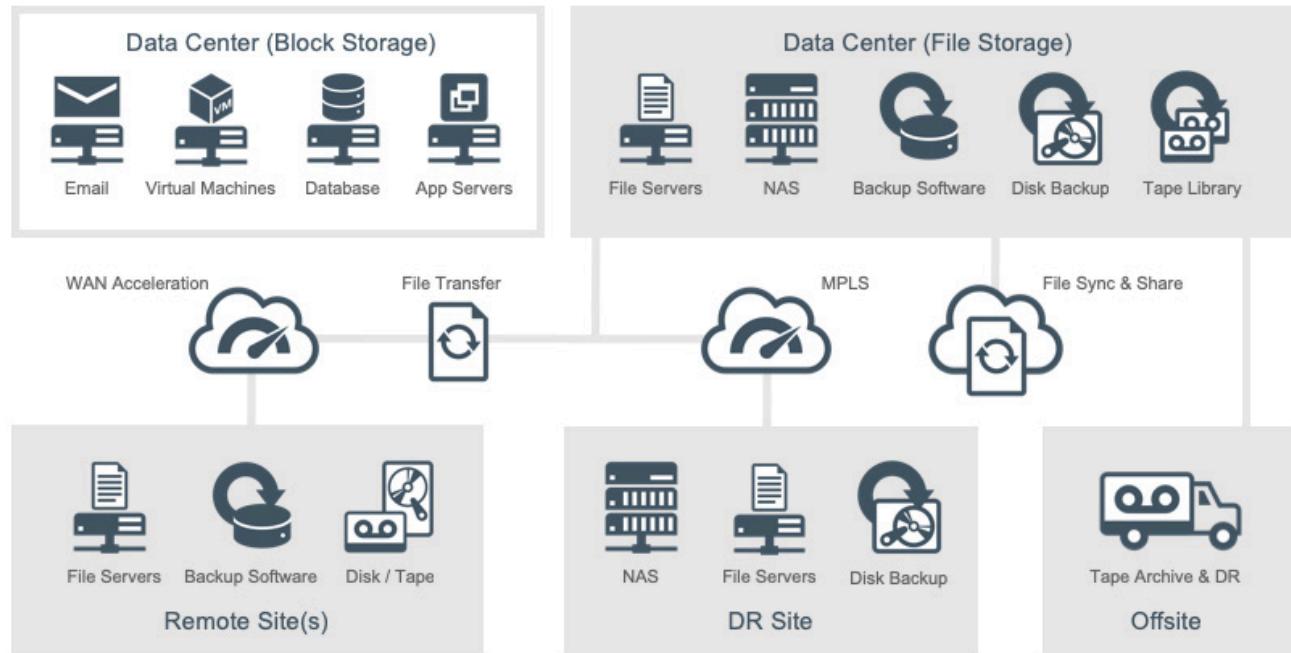
“ In today's around-the-clock business world, server downtime and loss of data access can wreak havoc on the company's operations and credibility. ”



---

## Traditional file infrastructure creates data silos and requires complicated data protection strategies

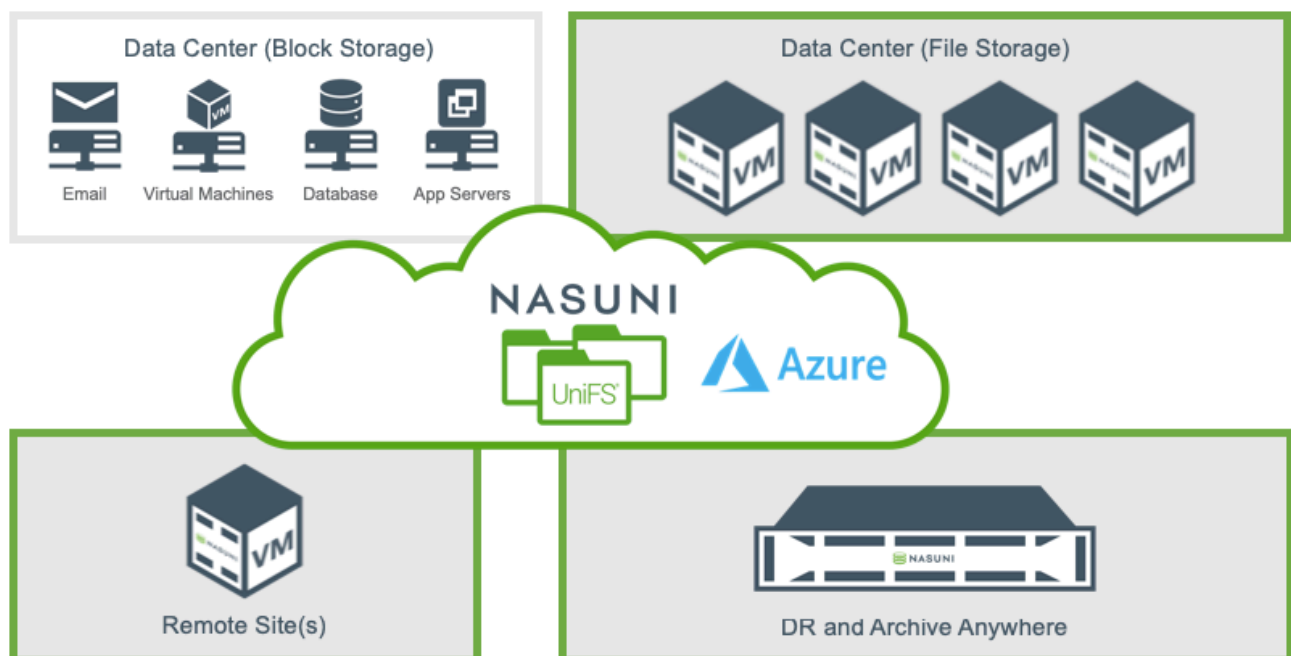
---



---

**Modern file infrastructure lives in the cloud which provides a single source of “truth” for all data and makes it quickly accessible for analytics.**

---







## CHAPTER 3:

# Understanding the Value of Cloud Based Services

File serving is evolving to file collaboration and the cloud is an ideal destination for this evolved use case. The file system, the software behind the file server, needs to change to keep up with this evolved use case. Simply adding a sync and share service on top of a legacy file system and file service won't solve the problem.

### A Cloud File System

The edge appliances act as a cache keeping a copy each location's most active data on-premises. When a user creates or modifies a file with a cloud native file server, the new or modified data is stored in a local cache on an edge appliance. The edge appliance can be a physical or virtual machine. The edge appliance then replicates that data to the centralized cloud repository. Once in the repository the new or modified data is available to the organization's other users. A copy of the data stays on the edge appliance while it is active and until the appliance reaches its watermark. Users though see a global view of the organization's data beyond what is available just on the appliance. They can, if permissions allow, access data that another location creates or modifies, and they can access their location's older data, which is not on the edge appliance.

The cloud file system also respects security services like, Active Directory, to ensure that users can only access the data to which they have the necessary access rights. The result is users have access to all the data they need with a performance experience that is at least as good as their legacy file server and certainly better than accessing another location's data via a VPN or other WAN connection.

### The Value of the Cloud

The primary value of using a cloud-based file system is the reduction in data center resource requirements and a reduction in IT personnel resources. To deliver the same or better level of performance the organization needs only a small edge appliance in each location. Storage Switzerland finds that most organizations only use 10% or less of their data on a day to day basis. That active data is available on the edge appliance, and again it is specific to each location which increases the chances of the right data being on the appliance. Organizations can easily increase the size of the edge appliance's cache for specific use cases but in almost all cases, a large percentage of data is moved to the cloud freeing up a significant amount of floor space and budget. The result of using the cloud as a repository with edge



devices also means that the organization saves a significant amount of operational time. IT administrators no longer need to manage dozens of independent NAS shares. The cloud file system is the equivalent of managing a single system. The days of multiple NAS systems from a single vendor or worse, multiple systems from different vendors based on use case, are gone.

Using the cloud as the basis for the modern file server also means nearly limitless storage capacity and locations. Solving the capacity problem means the elimination of costly and time consuming upgrades. The organization no longer needs to worry about migrating data or re-routing users to a new file server. Solving the locations issue means the organization can open as many branch offices as it needs without impacting user productivity or adding to IT user management burden.

Another value in using a cloud file system is the burden of data protection, at least for unstructured data, is handled in large part by the cloud file system. The cloud file service, depending on the cloud provider's redundancy setting, automatically replicates data to "N" number of locations as the edge device sends data to the cloud repository, creating a multi-level disaster recovery strategy. If a location's site goes down IT simply deploys another edge device and once metadata is replicated to the edge appliance, a process that takes a few minutes, users can immediately start accessing data. If the cloud provider has an outage, one of the edge appliances is pointed to one of the replicas.

Finally, the cloud file system is secure. In addition to respecting Active Directory settings it also makes sure that all data is encrypted both in transit and at rest. It also limits access to the encryption keys so that only the organization controls who can access what data and for how long. It can also provide protection from rogue administrators by requiring two factor confirmations of deletes, or retention policies that can't be overridden by administrators.

---

## Conclusion

Organizations are dealing with two major changes when it comes to unstructured data. The first is unstructured data's growth. Growth of the unstructured data set has been an issue for years. The new challenge is the rate at which that growth is occurring and how that growth rate is accelerating. The rate of growth puts pressure on IT and current file server technology simply can't keep pace with it. The second challenge is caused by the organization. Most organizations now are multi-office and multi-country but each of these locations require access to data and often the same data as the other offices.

A Cloud File System that intelligently leverages edge appliances addresses both of these issues without impacting performance. Leveraging the cloud to create a hub and spoke file system enables the organization to reduce on-premises storage investment, ease the burden of data protection and provide all locations unfettered access to all data they are authorized to access at all times.



# CHAPTER 4:

## Azure + Nasuni - The Future of NAS?

File storage capacity requirements are growing exponentially, and at the same time, these files must be protected and accessible by a range of users across the globe who must also collaborate on these files, in real time. Legacy network-attached storage (NAS) arrays are not equipped to meet these requirements. The cloud stands to help, but only if used strategically. In this chapter, we will explore the integration of Nasuni and Microsoft Azure as a solution through combining the advantages of the cloud with the advantages of a local file system.

### Microsoft Azure Storage Service

Cloud object storage services such as Microsoft's Azure Blob services are increasingly being used in place of disk or tape secondary storage infrastructures. They enable the storage manager to limit (or entirely bypass) capacity planning, investing in infrastructure, racking and stacking equipment, and migrating data as capacity requirements increase. These cloud object storage services also improve data consistency and integrity. Also, because Microsoft stores multiple copies of customers' files in data centers that are spread out around the world, it improves disaster recovery and availability (and again, without the need to purchase and manage additional infrastructure).


### UniFS® Global File System™

Nasuni provides a software-defined file storage platform that adds value on top of Microsoft Azure in the form of global file access and synchronization for collaboration, and centralized data management. It also facilitates fast backup, recovery and data protection. Nasuni does more than lift and shift a legacy file system architecture to the cloud. It moves the index node (inode) data structure, which stores metadata and other attributes that describe the file, to live natively in cloud object storage while removing limitations resulting from hardware, such as directory sizes and the number of locations that can be integrated. It also facilitates data protection, compliance and continuous versioning; changes to data made across any location are captured in an immutable Write once read many (WORM) format. An infinite, immutable version history is created and users can quickly roll back or recover to any point in time.

Customers can deploy virtual or physical appliances at edge locations to serve as a cache that provides fast access to data that lives in Azure. Customers may deploy appliances on existing virtualized infrastructure (such as a VMware or Nutanix hyperconverged appliances), or they can purchase an appliance from







Nasuni, that runs on Dell EMC server hardware. Files live in Azure, but metadata is cached in the Nasuni appliance for fast access. Edge appliances can be centrally managed through the Nasuni management console.

According to Nasuni, its approach of blending edge and cloud resources enables its customers to achieve faster performance on an infrastructure footprint that is only 1/5th of the legacy implementation. Meanwhile, it reduces the volume of data that must be pulled back from Azure, thus helping the organization to save on egress fees. In the event that a disaster is declared, the appliance can rehydrate data from metadata within 15 minutes, according to Nasuni. Users can also spin up the virtual machine in Azure, and have Azure be the disaster recovery site. Nasuni algorithms determine what metadata should be cached, and can also automatically archive data to less expensive classes of Azure storage as it ages – thus further saving on cost. Data retrieval remains rapid due to the on-premises cache. Data is encrypted at rest and in transit with customers holding the key,

---

## Conclusion

Cloud object storage solutions provide capacity scalability, a centralized namespace, granular searchability, a low-cost footprint and data resiliency that are required to meet modern file storage demands. However, they require some help in the form of performance acceleration and enhanced disaster recovery capabilities. Storage managers should consider the integration of Nasuni and Microsoft Azure to modernize their file storage approach.



## The Firm

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.



## About Our Partner

Nasuni delivers a single software platform to store, protect, synchronize, and collaborate on unstructured file data at scale. Nasuni Cloud File Services™, powered by the patented UniFS® global file system, leverages cloud storage to modernize primary NAS and file server storage; file archiving; backup; and disaster recovery, while offering transformational new capabilities for multi-site file collaboration. By combining the low cost, unlimited capacity, and durability of private or public cloud object storage with the high performance, security, and broad application compatibility of traditional disk-based file storage, the Nasuni subscription service improves workforce productivity, simplifies IT operations, and reduces IT costs. Nasuni operates globally from its worldwide headquarters in Boston, Mass., USA.