




# Solving Your Linux File Services Challenges in Azure





## Contents

<b>Executive Summary .....</b>	<b>4</b>
Limitless Scalability .....	5
The Azure Difference with Azure NetApp Files.....	6
<b>The 7 Challenges of Using Linux File Services in the Cloud.....</b>	<b>7</b>
1. Extreme File Service Performance.....	8
2. Host Client Data Compatibility .....	8
3. Guaranteed Business Continuity .....	8
4. Data Protection.....	8
5. Migration.....	8
6. Keeping Data in Sync.....	8
7. Data Security .....	8
<b>Three Major File Services for Linux in the Azure Cloud .....</b>	<b>9</b>
<b>Azure Files/Azure Premium Files.....</b>	<b>10</b>
File Share Services for SMB .....	10
Easy Set-Up.....	10
Azure File Sync .....	10
Azure Files/Azure Premium Files at a Glance.....	10



<b>“Roll-Your-Own” NAS Solution.....</b>	<b>11</b>
Reliable Access.....	11
Manual Administration.....	11
Complexity Over Time.....	11
Performance Constraints .....	11
A Use Case - GlusterFS.....	11
“Roll-Your-Own” at a Glance.....	11
<b>Azure NetApp Files .....</b>	<b>12</b>
Unprecedented HA Performance – Benchmark-Tested in Labs.....	12
NetApp Snapshot™ Technology.....	13
Encryption of Data .....	13
Azure NetApp Files at a Glance .....	13
<b>Conclusion .....</b>	<b>14</b>



# Executive Summary

*"Linux usage on our cloud has surpassed Windows."* Sasha Levin, Microsoft Linux kernel developer<sup>1</sup>

It's no surprise that Linux is dominating in Azure. According to research from IDC, Linux owns 68% of the OS market share globally<sup>2</sup>. That is because enterprises love to build applications on Linux. The system is robust, flexible, proven and, as a "free" open source technology, does not require permissions or command licensing fees. Any innovation built on Linux remains your own. That is a big advantage for companies that are relying

more and more on software applications to remain competitive in the digital age.

As companies make the decision to move more of their applications<sup>3</sup> to the cloud, they are often confronted by barriers.


The vast majority of applications have some amount of required file space. Even your core, business-critical applications have a requirement for files. As you try to resolve the complexities in moving these Linux-based applications to the cloud, you have to look through all the options that are available to you.

"Every business will become a software business, build applications, use advanced analytics, and provide SaaS services<sup>4</sup>."

**Satya Nadella**  
CEO, Microsoft

## WORLDWIDE INFORMATION TECHNOLOGY (IT) INFRASTRUCTURE SPENDING BREAKDOWN (BY VALUE) FROM 2014 TO 2023, BY DEVELOPMENT TYPE\*\*

SHARE OF IT INFRASTRUCTURE SPENDING WORLDWIDE 2014-2013, BY DEVELOPMENT TYPE<sup>5</sup>



**Note:** Worldwide; 2014 to 2018; IT infrastructure includes servers, enterprise storage, and Ethernet switches  
**Source(s):** IDC, Statista estimates; ID 486586

### Limitless Scalability

File shares provide data access and management capabilities for files to any client, in any geography and data format, at any scale. Migrating file-share workloads to the cloud offers the benefits of limitless scalability.

With Linux-based (NFS) file shares, file systems can give concurrent access to hundreds or even thousands of user machines. When these file services are cloud-based, it brings in added scale with support for media processing workloads, big data analytics, backups to offsite repositories, and more. With such a wide range of use cases, a large number of factors must be considered when deciding on a cloud-based file service:


- Will it provide the performance levels that enterprise workloads require?
- Would opting for a fully managed service be better than relying on a file share that needs resources and labor to maintain?
- Is speed-to-service a critical factor in rolling out new applications or workloads?
- Will I have peak demands that need a scale-up or scale-out high-performance infrastructure?

Answering these questions is an important first step in choosing file services. This white paper looks at shared storage for file services in the Azure cloud.

This paper identifies the seven key challenges of running shared storage for file services in the cloud and then introduces the three major offerings available in Azure today, including an “at a glance” summary of the benefits and considerations for each choice.

## FORECAST COMPOUND ANNUAL GROWTH RATE OF THE CLOUD APPLICATIONS MARKET BY SEGMENT FROM 2017 TO 2022

CLOUD APPLICATIONS MARKET CAGR 2017-2022, BY SEGMENT



**Note:** Worldwide; 2018

**Source(s):** Apps Run The World; ID 475768

The paper concludes by exploring the benefits of Azure NetApp Files in more detail and how this Linux file service provides the features and performance to address even your most critical business applications in Azure.

### The Azure Difference with Azure NetApp Files

Azure NetApp Files offers something brand new: a fully managed Azure service built on NetApp technology running in the cloud. It supports the most prominent Linux file

services protocols including NFSv3 and (as of this writing, coming soon) NFSv4. It also allows full data mobility between any end points for hybrid and multicloud flexibility.

In addition to Linux environments, Azure NetApp Files also supports SMB (Windows) file access such that you can access data from different user platforms. The service combines NetApp's data management capabilities and Azure's unlimited availability and scalability, without any of the usual administrative tasks and overhead.

This white paper contains research from third-party sources as well as the opinion of the NetApp product team. I hope you enjoy this white paper. Please reach out to me directly if you have any questions.

### Jeff Whitaker

Senior Manager,  
Cloud Data Services, NetApp



### Jeff Whitaker

As the Senior Manager, Cloud Data Services at NetApp, Jeff is a perennial disruptor with a passion for emerging technologies that let start-ups grow up fast, and traditional companies re-discover their entrepreneurial essence. Cool ideas that Jeff considered (but didn't invent) include machine learning, virtual reality and a drone that fetches tapas on demand. You can reach Jeff directly at [jeff.whitaker@netapp.com](mailto:jeff.whitaker@netapp.com).



# The 7 Challenges of Using Linux File Services in the Cloud

When it comes to shared storage in the cloud, enterprises – even those with cloud-first mandates – are struggling to move mission-critical workloads built on Linux. In addition to the obvious challenge of running a major workload off-premises, there is a need to scale-up unlimited amounts of storage while still providing availability guarantees as well as meeting the compatibility needs with data formats and operating systems.

In the following, I explore the top 7 challenges facing enterprises moving Linux applications to the cloud.



### 1. Extreme File Service Performance

File services are extremely dependent on the high level of performance of the storage environment in use. Some workloads may require a high-level SLA to maintain peak operability and IOPS levels.

### 2. Host Client Data Compatibility

An enterprise's file services need to be compatible with all their host clients' data formats and operating systems. Linux environments are the de facto standard when it comes to enterprise applications and having a robust NFS file service solution is a requirement to moving any of these applications to the Azure cloud.

### 3. Guaranteed Business Continuity

When it comes to running an enterprise file service, any disruption in normal operation can negatively impact the business. Whether an outage is caused by a disaster or through an update process, it is essential to ensure the availability of the file share, with zero downtime and no data loss.

### 4. Data Protection

File shares require companies to comply with industry-specific data security and data protection regulations, and because file shares are often the most important part of a business's operation, the ability to automatically create copies is mandatory to meet the stringent RPO, RTO, and backup requirements for most industries.

### 5. Migration

When using a cloud-based file share, it's important to find a solution for migrating data to the cloud without having to rewrite the applications that depend on the data. If applications need to be rearchitected or refactored, the enterprise has to commit significant amounts of time and money to design and test the "new" applications to ensure that the end user isn't negatively affected. Migration needs to be as seamless as possible.

### 6. Keeping Data in Sync

Following up on the previous point, any data that moves to the cloud in the file share must be kept synchronized among the disparate environments in use, whether it is data housed on-premises or in one or more cloud environments.

### 7. Data Security

With shared storage, data security becomes a major concern. It's important to make sure that access to the file system is in the user's control at all times. Ensuring data security at all levels is crucial to preventing losses before they happen. Role-based access and encrypted data can keep data more secure.

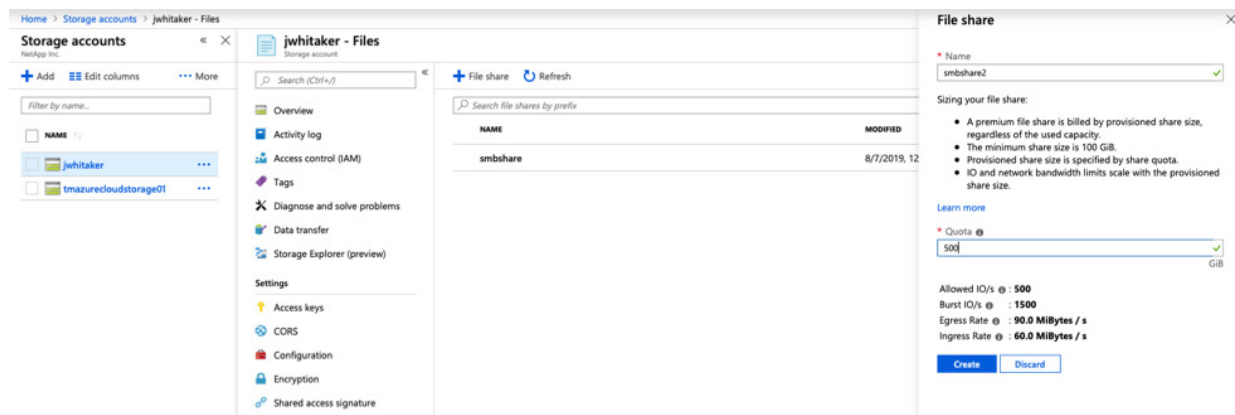


# Three Major File Services for Linux in the Azure Cloud

This section profiles the cloud-based file services available in Azure, including Azure Files, Roll-your-own file service, and Azure NetApp Files.



# Azure Files/Azure Premium Files



## File Share Services for SMB

Microsoft has solved files in the cloud with Azure Files, which provides file share services for SMB v3.0. Given that SMB is primarily for use with Microsoft systems, it's notable that Azure Files can be mountable for read/write operations with MacOS and Linux operating systems via a SAMBA server.

## Easy Set-Up

Support for the newer version of SMB enables features such as encryption in-transit, meeting the file service demands for high levels of data security. However, the same level of protection is possible using the RESTful calls via HTTPS. Although they are easily creatable through the user interface, Azure Files file shares can also be created by using the Azure CLI or PowerShell.

## Azure File Sync

In addition to Azure Files, Microsoft also offers Azure File Sync. This service offers more opportunities for Azure Files to work with on-premises storage architectures. Through an agent installed in the on-premises system, Azure Files can cache data locally, increasing access speed and allowing writes to sync transparently with Azure Files.

On top of this, Azure Premium Files offers enhanced performance for SMB-based Windows applications. Configuring additional servers provides a way to unite storage environments located in multiple geographic areas.

Azure Files snapshots gives users the ability to create read-only Snapshots of the Azure Files file share that can be used for backup.

Charges for Azure Files are based on two factors: storage consumed and file access. Storage costs also accrue when utilizing snapshots of data in Azure Files (at a reduced rate). Additional costs can accrue for the use of the additional features, including Azure File Sync.

## AZURE FILES/AZURE PREMIUM FILES AT A GLANCE

### Benefits:

- Managed by Azure
- Easy setup
- New services adding backup and faster speed

### Considerations:

- This is not a Linux file-service (Windows SMB service)
- Snapshot capability (added storage consumption costs)
- Max volume size 5TB
- Native backup system remains manual
- Active Directory support through the use of Azure Active Directory

# “Roll-Your-Own” NAS Solution

An alternative to using a managed, cloud-based file service—and until recently, the only solution for Linux files service support in Azure—is to use cloud compute and storage resources to create a custom file-share solution. At first glance, this setup is appealing due to the high degree of control that it offers over the type of storage used, access protocols, file system features, and more. However, these solutions very often run into complications.

## Reliable Access

The first major issue with a custom file service solution is ensuring reliable access to the files and guaranteeing cross-site durability. A file service is built on cloud compute and block storage. Most block-level cloud storage solutions provide redundancy within a single location; however, extra steps need to be taken to ensure durability beyond a single site. Asynchronous replication or copying the data on a schedule always means that the files being stored are susceptible to data loss.

## Manual Administration

The second issue is the level of manual administration that these solutions require. For example, allocating a new file system requires allocating new storage, mounting it to the compute nodes that will serve out the data, and potentially initializing the new share with existing data. If the file system needs to grow, this growth must be handled manually. If the performance of the underlying disks needs to be upgraded, the allocation of the new storage and migration of existing files need to be taken care of while still trying to minimize downtime.

## Complexity Over Time

The third issue is the complexity of managing the storage over time as the deployment grows. Cloud administrators working with Linux file shares need to maintain uninterrupted access to the files, provide backup or snapshot facilities, allow test copies of the data to be created, and much more. Growth of data is essentially never-ending, so scale becomes a considerable challenge with these manually built environments. Providing robust support for this kind of functionality requires a high level of technical expertise.

## Performance Constraints

Finally, achieving high performance out of a manually built file service is extremely challenging, if not impossible. The file server is at the mercy of the compute node and the block-level storage attached to it. Achieving this performance requires high-performance compute nodes with high-speed volumes attached. This is costly, and if you have any requirements on protecting against data loss, it is impossible to get the on-premises-like performance that your critical applications often require.

## A Use Case - GlusterFS

Let's look at one such solution. The GlusterFS file service can distribute files across both on-premises and virtual devices, offering scalability and reliability. GlusterFS can be used for Azure, AWS, and GCP. It is capable of making connections via NFS, SMB, and iSCSI, with the correct drivers and add-ons in place.

Gluster storage configurations offer a range of options for file sharing, including striped, replicated, dispersed, distributed, and a

number of combinations of those. It also has the ability to take snapshots, as well as to clone snapshots.

As an open-source technology, the setup and configuration for GlusterFS are completely up to the user. This can be a considerable task that may require a large investment in time and money. Problems that arise through use of the platform require maintenance and active resolution. Although GlusterFS is an open-source technology, support for it is commercially available through Red Hat.

## “ROLL-YOUR-OWN” AT A GLANCE

### Benefits:

- Full control over the deployment, including storage disks and access protocols
- Suitable for small deployments

### Considerations:

- Manual administration required
- Complex to support over time
- No built-in cross-site data redundancy
- Limited performance based on compute and storage resources



# Azure NetApp Files

Azure NetApp Files is Microsoft's fully managed, high-performance, cloud-native file service. Available directly in the Azure portal, it is delivered and supported by Microsoft but built on NetApp technology. The service delivers both Linux NFSv3 (with NFSv4 coming imminently, as of this writing) and Windows (SMB) protocol support, advanced data services, and is integrated directly into the Azure service infrastructure (including AzureCLI).


Users can select one of three service levels: Standard, Premium, or Ultra, with the following "on-premises-like" performance levels:

- Standard service level:  
16MB of throughput per TB
- Premium service level:  
64MB of throughput per TB
- Extreme service level:  
128MB of throughput per TB

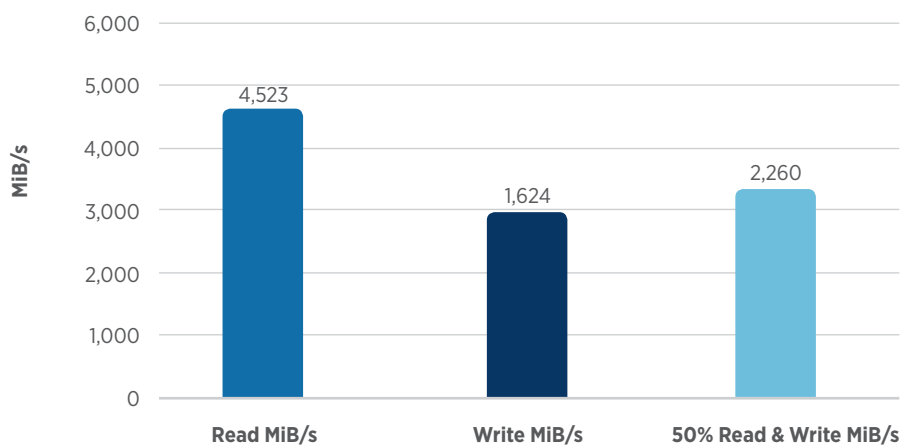
## Unprecedented HA Performance – Benchmark-Tested in Labs

High-performance service levels come with the ability to spin up 100TB of storage in under 10 seconds, each with built-in data protection. The Azure NetApp Files performance levels operate at sub-millisecond latencies – unprecedented across any cloud file service – enabling companies to move workloads to the cloud that would never have been possible before.

**8KiB I/O - SINGLE VOLUME LIMITS  
TWELVE D32s V3 VIRTUAL MACHINES**



**64KiB THROUGHPUT TEST - SINGLE VOLUME LIMITS  
TWELVE D32s V3 VIRTUAL MACHINES**



On top of the performance, with Azure NetApp Files, users no longer have to worry about storage management. The service takes care of all the setup, configuration, updates, performance, and service levels. The service is built on the NetApp infrastructure that has served enterprise customers for more than 25 years – faster and better than ever.


### NetApp Snapshot™ Technology

For data protection, no solution supports a file service better than NetApp Snapshot technology. Snapshot copies give Azure NetApp Files the ability to provide companies with point-in-time backups that can be created instantly and do not add to your data footprint in Azure, which can save both time and money while ensuring that important data protection goals are maintained. You also get powerful, high-speed data copies that can be used for a number of purposes, such as building test environments and restoring systems in the event of accidental or malicious data loss events.

### Encryption of Data

Azure NetApp Files comes with built-in, always-on protection via encryption of data at rest, encryption in-transit for SMB connections and over VPNs, as well as high availability, and NetApp reliability.

- **Fully managed service.** Data management is handled completely by Microsoft, not the customer.
- **Scalability and performance.** Spin up to 100TB at a time of extremely high-performance storage in just seconds.
- **Multiple protocols.** Support for SMB, NFSv3, and (coming soon) NFSv4 file shares.
- **Compatibility.** Supports shared file access across Linux, UNIX, and Windows operating systems for greater host client data operability.



- **Integratable.** Complete integration with file directory metadata, keeping domain credentials, access and authentication, and group memberships, including full compatibility with Microsoft Active Directory.
- **Performance levels.** Provides high performance at the service level of your choice.
- **Data protection.** Data corruption or loss can be prevented with efficient, automatic data Snapshot copies.
- **Migratable.** Import data from on-premises and other storage repositories via Cloud Sync.
- **Business continuity.** High availability ensures business continuity with no data loss (RPO = 0) and short recovery times (RTO < 60 secs).
- **Automation.** Schedule tasks directly via the Azure CLI to meet file-share demands with automation and orchestration capabilities.

## AZURE NetApp FILES AT A GLANCE

### Benefits:

- Managed by Microsoft as an Azure service
- No special skills – spin up HA Linux environment in minutes
- Supports multiple protocols with no code changes
- Maximums
  - Volume size 100TB
  - File Size 16TB
- Suitable for all deployments – multiple service levels
- Secure, encrypted

### Considerations:

- Extreme volume sizes over 100TB not supported
- Small volume sizes under 100GB not supported
- Focused global region support



# Conclusion

Linux file services in the cloud come with a range of options, but users should be smart about how they move ahead with a file service. It's important to make sure that the file service has the performance and features to meet their organization's needs. Most important in this case are scalability, data protection, and compatibility, all of which Azure NetApp Files can provide better than any other option.

To get the most out of the cloud, [sign up today for Azure NetApp Files.](#)





# References

<sup>1</sup> <https://www.zdnet.com/article/microsoft-developer-reveals-linux-is-now-more-used-on-azure-than-windows-server/>

<sup>2</sup> <https://www.idc.com/getdoc.jsp?containerId=US43753318>

<sup>3</sup> <https://www.computerweekly.com/news/2240242478/Satya-Nadella-Every-business-will-be-a-software-business>

<sup>4</sup> <https://www.statista.com/statistics/475768/cloud-applications-market-cagr-by-segment/>

<sup>5</sup> <https://www.statista.com/study/15293/cloud-computing-statista-dossier/>