

Network Attached Storage Solutions for AWS

Selecting storage solutions for NAS
on Amazon Web Services

December 2018



© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Abstract	4
Introduction	5
Traditional Enterprise NAS Solutions	5
NetApp Cloud Volumes ONTAP	5
SoftNAS	10
Scale-Out NAS	11
Weka IO Matrix	12
Qumulo File Fabric	14
Gateway NAS Solutions	15
Panzura	16
CTERA	17
Nasuni	19
Conclusion	21
Contributors	21
Document Revisions	21

Abstract

There are many reasons a customer may choose to use a Technology Partner Network Attached Storage (NAS) solution for storage workloads on Amazon Web Services. Reasons include familiarity with existing storage management solutions, compatibility of replication and recovery workflows, application feature dependencies, and established procedures around monitoring and protection. This paper discusses Technology Partner solutions for NAS and highlights each solution's strengths around performance, use cases, availability, protocol, and application support.

Introduction

File shares, often called network attached storage (NAS), support use cases for simultaneous access by hosts to a shared data set. This shared data access method can be used for a variety of use cases, from database repositories to home directories. Because of the ubiquity of use cases, it's common to share file data in organizations across different types of computers. Therefore, multiple protocols - for both Windows and Unix or Linux - are usually supported for the same data set.

The most common protocols are Network File System (NFS) and Server Message Block (SMB) (synonymous with Common Internet File System, or CIFS). NFS has been the standard for Unix platforms since 1984 and is prevalent in Linux environments. CIFS is based on Microsoft's SMB, both of which are the standards for Windows hosts.

Aside from these core connectivity and compatibility standards, enterprise NAS solutions today have built robust performance, replication, tiering, backup, and data protection features that users have come to rely on for their shared data pool.

In this paper, we will cover partner NAS solutions that fall into three categories: traditional enterprise, scale-out, and gateway solutions.

Traditional Enterprise NAS Solutions

Partner solutions in this category provide enterprise functionality such as read-writable snapshots, replication, and tiering. Customers can extend their current processes and management used today on premises to now run on AWS. Traditional Enterprise solutions also support access from both NFS and CIFS/SMB to the same file system.

NetApp Cloud Volumes ONTAP

The NetApp solution for AWS is a fit for customers looking to migrate and protect their on-premises NetApp data to AWS. Additionally, NetApp Cloud Volumes ONTAP can be used when customers need support for NetApp-based applications and workloads on AWS without modifying the protocol interface or known processes and administrative practices.

ARCHITECTURAL OVERVIEW

Figure 1 below shows a NetApp Cloud Volumes ONTAP High Availability (HA) cluster pair, representing the maximum cluster size for Cloud Volumes ONTAP. Along with the two nodes in separate Availability Zones (AZs), there is a mediator instance in a third AZ. The mediator ensures that the cluster nodes are alive and communicate a failure to a surviving node. Lastly, the Cloud Manager is the user interface for configuration and management of the data cluster.

Cloud Volumes ONTAP utilizes Amazon Elastic Block Store (EBS) for the primary disk tier. In an HA configuration, writes to Cluster Node A are synchronously copied to Cluster Node B. This ensures data consistency across both nodes, necessary for failover in the event of Amazon Elastic Compute Cloud (EC2) instance failure. Data protection is increased if nodes are deployed in distinct availability zones. Overall resiliency is improved by in an HA configuration by protecting against failures of an instance, Amazon EBS volumes, and even an availability zone.

Pro Tip: Even though a synchronous write configuration exposes a small amount of additional latency for write operations, production deployments requiring HA configuration make it a necessity. Deploying CIFS, NFS, or iSCSI volumes in a single node configuration presents a risk of reduced availability in the event of an Amazon EC2 instance failure and data loss in the event of underlying Amazon EBS volume failure.

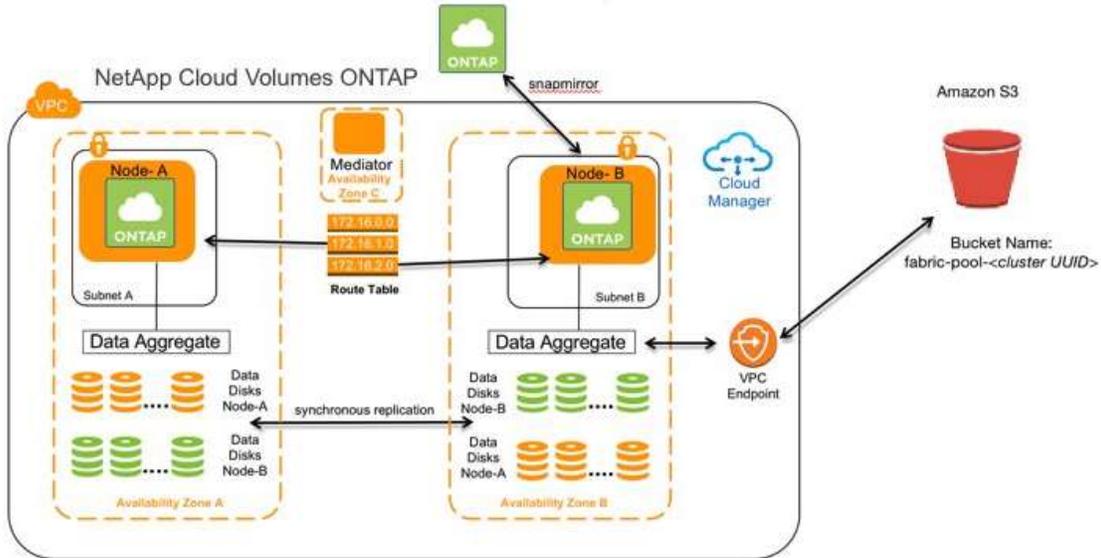


Figure 1: NetApp Cloud Volumes ONTAP HA cluster pair

Along with Amazon EBS as the primary underlying storage tier, volumes and aggregates also can tier data to Amazon Simple Storage Service (S3) according to a user defined policy in NetApp OnCommand Cloud Manager. By setting a tiering policy, a new Amazon S3 bucket will be created (or optionally use an existing bucket) and cold data will be abstracted from the Amazon EBS block device and moved to Amazon S3. In most cases, to deploy an operationally efficient architecture with Cloud Volumes ONTAP, tiering to Amazon S3 is a crucial configuration option. This is especially true in the case of a NetApp SnapMirror target as shown in the overview diagram. The reason being, in an HA configuration Amazon EBS storage is provisioned twice, but tiered data to Amazon S3 is a single copy known by both cluster nodes. Amazon S3 already is set up to protect data across an entire region, so having a single copy is both cost-effective and durable. NetApp SnapMirror on both nodes is aware of the tiering to Amazon S3.

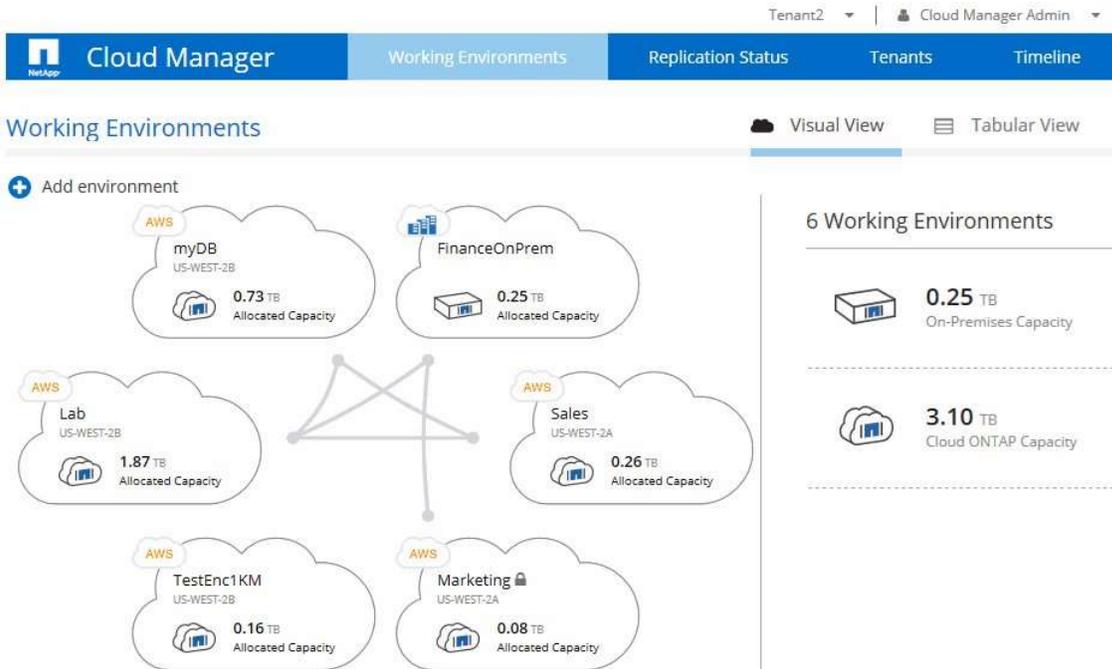


Figure 2: NetApp OnCommand Cloud Manager GUI

Pro Tip: When extending your configuration to utilize Amazon S3, you must allocate enough Amazon EBS storage to cover new writes to the file share. Once written, Cloud Volumes ONTAP will shuffle cold data off to Amazon S3 to maintain the policy threshold for tiering cold data. In order to make sure tiering continues without incident, new writes need enough capacity for Cloud Volume ONTAP to reconcile it with old data in the background, often during periods of reduced activity. In the current version of Cloud Volumes ONTAP 9.4, there is another restriction: you cannot tier writes to Amazon S3 until they age for a minimum of two days.

Use Case: Moving On-Premises NetApp Data to AWS with SnapMirror and SnapVault

SnapMirror and SnapVault are supported for NetApp Cloud Volumes ONTAP. The most common use case is replicating the data from an on-premises NetApp system to a cluster on AWS utilizing NetApp SnapMirror.

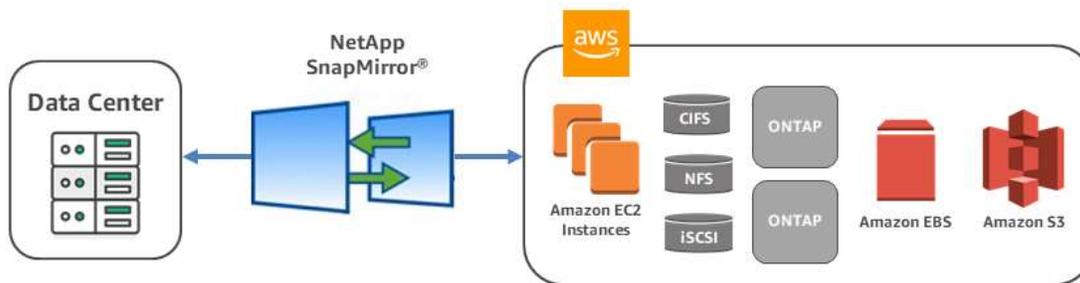


Figure 3: NetApp SnapMirror Architecture

Note: Double check for ONTAP versions for two-way replication. Some older versions of ONTAP on-premises may not be suitable targets.

To save costs on the cloud, it is highly recommend you employ cloud tiering to Amazon S3 for the SnapMirror replicated data. The tiering function in Cloud Volumes ONTAP recognizes incoming data from SnapMirror as backup and classifies it automatically to “cold” and eligible to move directly to Amazon S3. This results in significant savings in AWS storage costs.

Use Case: Archiving On-Premises NetApp Files to AWS with FabricPool

Similar to cloud tiering in Cloud Volumes ONTAP, NetApp FabricPool allows you to set a policy on your on-premises NetApp filer to tier data to Amazon S3. In ONTAP 9.2 and 9.3, FabricPool supports tiering NetApp snapshot copies and data in data protection volumes to Amazon S3 by using the Snapshot-Only and Backup volume tiering policies.

In ONTAP 9.4, an auto tiering policy was added. When a NetApp volume tiering policy is set to Auto, all cold blocks in the volume are moved to Amazon S3. By default, it takes 31 days for inactive blocks to become cold, but auto tiering is user configurable per volume.

Pro Tip: You can enable inactive data reporting (IDR) to determine the amount of inactive (cold) data that can be tiered from an aggregate. This will provide a clear estimate of the data sent to Amazon S3 with the Auto policy. To enable IDR:

```
storage aggregate modify -aggregate <name>
-is-inactive-data-reporting-enabled true
```

Use Case: Development and Test with NetApp Snapshots and FlexClones on AWS

NetApp read/write snapshot capabilities are available via Cloud Volumes ONTAP. This use case is ideal for situations in which you need multiple copies of identical datasets or temporary copies of a dataset (e.g. testing an application against a production dataset). NetApp FlexClones are space-efficient copies that share data blocks with parent volumes so they consume no storage except what is required for metadata until changes are written to the copy. This methodology also provides copies of even large volumes to be instantly available to provision to an alternate host.

SoftNAS

SoftNAS is a software-based NAS solution for storage administrators looking to replace existing enterprise NAS solutions with a lower-cost alternative. The SoftNAS solution provides enterprise features such as snapshots, replication, migration, and comprehensive protocol support and can be provisioned on an on-premises physical or virtual machine as well as on an AWS instance. To allow for cost optimization, SoftNAS does not require any specialized hardware and can create SoftNAS S3 Cloud Disk utilizing Amazon S3.

ARCHITECTURAL OVERVIEW

As with any storage architecture that utilizes Amazon EBS, it is important for data durability to deploy the solution in a Multi-AZ HA configuration. SoftNAS is designed on the Zettabyte File System (ZFS).

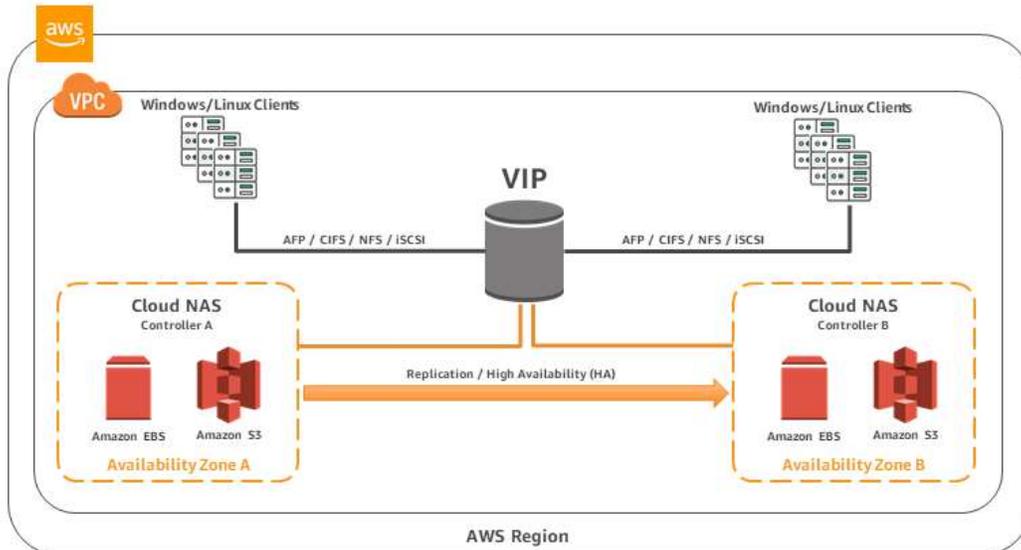


Figure 4: SoftNAS Architecture

The HA cluster configuration spans AWS AZs and write operations are replicated asynchronously from Cloud NAS Controller A to Controller B. Because it is asynchronous, you will see better latency performance on write operations at the cost of potentially losing transactions if a failover is required.

Windows and Linux instances can access data through Apple Filing Protocol (AFP), CIFS, NFS, and iSCSI. When setting up a cluster, consider the instance type that best suits your performance goals. SoftNAS provides guidance here with standard, medium, and workload-specific examples.

Use Case: Scale to Multi-Petabytes with Cloud Disks

SoftNAS Cloud provides support for a feature known as SoftNAS S3 Cloud Disks. Here, Amazon S3 storage is presented as block devices. By using Amazon S3 storage in this way, the SoftNAS Cloud can scale to much larger capacity than traditional block devices. Each cloud disk holds up to four petabytes (PB) of data and you can stripe across cloud disks for maximum scale.

Each SoftNAS S3 Cloud Disk occupies a single Amazon S3 bucket in AWS. For best performance, choose the same AWS Region for both the SoftNAS Cloud Amazon EC2 instance and its Amazon S3 buckets. SoftNAS Cloud storage pools and volumes using cloud disks have deduplication, compression, caching, and storage snapshots available, as well as support shared access through NFS, CIFS, AFP, and iSCSI.

Pro Tip: When you use a Cloud Disk, use a block device local to the SoftNAS Cloud virtual appliance as a read cache to improve IOPS and performance for reads.

Scale-Out NAS

Scale-Out NAS partner solutions focus on high performance workloads and scale-out capacity. These solutions often fit well with high performance computing (HPC), media, life sciences, and financial services. Scale-Out NAS solutions possess capabilities, such as advanced analytics and automated tiering, to control costs and right-size for AWS.

Weka IO Matrix

If performance is the main criteria for your AWS-based NAS solution, Weka IO should be considered. Weka IO has a unique capability to shut down the cluster and persist the data to Amazon S3. The cluster can then be re-instantiated with the Amazon S3 data with a parallel copy process when new jobs need to be processed. For HPC workloads in particular, Weka IO's use of Amazon S3 lowers the cost of performance compute and storage resources on AWS.

Weka IO has a product called Matrix which is a scale-out NAS platform. Matrix was designed on the AWS Cloud with optimization in mind and can run on AWS as well as on-premises. MatrixFS, which is the underlying file system of the Matrix NAS, is a distributed parallel file system. Matrix includes many features you would expect from a Scale-Out NAS system, including a Global Name Space, support for NAS protocols like NFS and SMB, and the ability to scale both performance and capacity. Some key differentiators for Matrix include linear performance scalability, seamless tiering to Amazon S3, snapshots to Amazon S3 (including the entire cluster), and use of a native client to provide true parallel processing across multiple nodes (currently only available for Linux).

ARCHITECTURAL OVERVIEW

On AWS, Matrix runs on either “R” or “I” instance families, both which include instance storage that is used as the primary performance storage tier for the cluster. Matrix requires a minimum of six nodes to have a fully formed cluster. Matrix software runs inside a container, so it can either be installed on instances by itself or along with other applications, which generally would be the application utilizing the file system. Matrix has configurable redundancy from N+2 to N+4 and can scale from the base six nodes up to 1,000 of nodes. MatrixFS distributes all data and metadata across the nodes in the cluster. Matrix is designed to have an active tier of data on instance storage and an inactive tier on Amazon S3. While tiering is optional, it is recommend that most implementations use the Amazon S3 tier, which can be enabled any time after the cluster is running.



Figure 5: Weka IO Node Cluster

Use Case: Persist HPC Data to Amazon S3 to Lower Costs

As mentioned previously, Matrix can run both on AWS and on premises. In either configuration, Matrix has policy-driven tiering to Amazon S3. Because Matrix utilizes instance storage for the primary performance tier, Amazon S3 tiering for lower use data provides significant cost savings for operational efficiency. Matrix also has a snapshot functionality, which copies data to Amazon S3 that doesn't already reside there as a result of tiering. The snapshot functionality enables Matrix to run in a hybrid mode with sharing of data between on premises and a cluster run on Amazon EC2 or between clusters running in different regions. This can be used for both disaster recovery (DR) and bursting use cases. A file system can utilize snapshots from a cluster running on premises and a cluster can be spun up on AWS on demand utilizing that data. Multiple nodes retrieve data from Amazon S3 in a parallel fashion, creating a cost-effective solution and that is still highly performant. Therefore, data only has to reside in higher cost instance storage during actual compute job generation. When a large-scale compute job completes, data can be efficiently moved to Amazon S3 and compute nodes with instance storage can be shut down to save costs.

Use Case: HPC

Due to Matrix's parallel distributed nature, it offers performance for the most demanding workloads at very low latency, on both NFS and SMB. The solution has shown to provide linear scaling performance when adding additional Amazon EC2 instances to the cluster on AWS. To get the best single-client performance, Matrix offers a client that can run on a Linux Amazon EC2 instances. The client provides a POSIX compatible mount on the instance that offers the highest level of single-client performance, resulting in multiple GB/s.

Qumulo File Fabric

Qumulo File Fabric (QF2) is a scale-out file system and is a great fit for storage administrators looking for high performance coupled with real-time statistics on how storage is being used at very large scale. QF2 also is a fit for customers with an existing on-premises footprint for render processing that requires elasticity to extend the capacity of the render farm to AWS.

ARCHITECTURAL OVERVIEW

QF2 can run on-premises or on AWS and has a distributed architecture where many individual computing nodes work together to form a cluster with scalable performance and capacity under a single, unified file system. Clusters start at four instances and can scale to 1,000 instances. Clients can access the filesystem through NFS or SMB protocols so that both Windows and Linux hosts can access the same file share.

This scale-out architecture not only aggregates performance, but also provides the ability to give real-time views of data usage and up-to-the-minute file system analytics on billions of files and directories. The use of aggregated metadata enables this view into the storage even as it scales out. Trends for capacity and performance for the cluster are visible down to file-level granularity. The analytics also give administrators immediate access to information needed to diagnose problems as well as enable proactive planning for capacity and performance tiers regardless of the scale.

This instant view into cluster use can work in conjunction with the real-time quota feature. If an administrator sees concerning usage patterns, he or she can apply a quota for that user which will take effect immediately. Each QF2 cluster also can work together to form a globally distributed storage fabric tied together with continuous replication relationships both on AWS and on the customer premises.

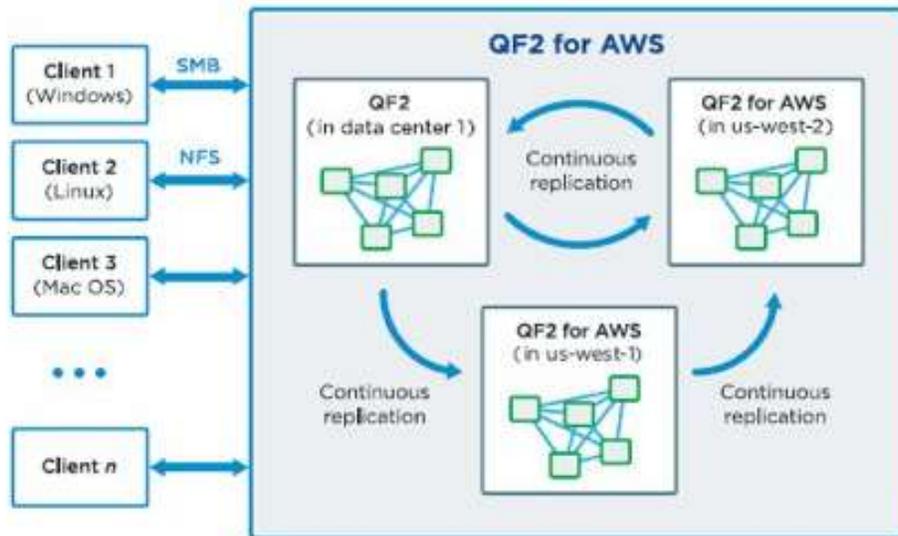


Figure 6: Qumulo QF2 for AWS Architecture

QF2 snapshots let system administrators capture the state of a file system or directory at a given point in time. QF2 does not currently limit the number of concurrent snapshots and there is no data movement overhead creating point in time copies.

QF2 clusters on AWS have utility pricing that is based on hours of use, capacity, and performance. QF2 can be used for free in non-clustered, standalone mode on AWS.

Use Case: Rendering Jobs to AWS Cloud

By utilizing the continuous replication feature in QF2, you can stitch together your namespace from both on-premises and Amazon EC2 instances. With QF2 for AWS, rendering can be done with the on-demand scale and elasticity of Amazon EC2 instances running with Qumulo. This avoids over-provisioning for individual projects, rental fees, and time spent deploying permanent or temporary infrastructure on site so studios can make tight deadlines more efficiently.

Gateway NAS Solutions

Partner Gateway NAS Solutions enable you to store data on Amazon S3, while maintaining the same front-end protocols used with traditional applications and workloads, including SMB and NFS. Gateway Solutions also possess file

sync and share capabilities that enable collaboration among geographically dispersed users and data in Amazon S3.

Panzura

Panzura Freedom is a good fit for administrators that have an existing, heterogeneous footprint of NAS storage solutions and would like to unify dispersed data solutions in multiple locations to AWS for collaboration and archive. Data centralized in Amazon S3 will appear at multiple sites as a single global namespace, reducing the storage requirements of on-premises and edge locations.

ARCHITECTURAL OVERVIEW

Panzura Freedom is based on the Panzura Cloud File System (PCFS). PCFS enables transparent file movement and management across multiple geographic locations and Amazon S3. You can install Freedom on a virtual machine, physical device, or Amazon EC2 instance. Once installed, both Linux and Windows clients will have access to the data through NFS and SMB protocols in a single global namespace. This includes global file locking capabilities to enable collaboration from geographically dispersed users to the same file system. Deduplication and compression also can be applied as data is distributed throughout the Panzura Cloud File System.

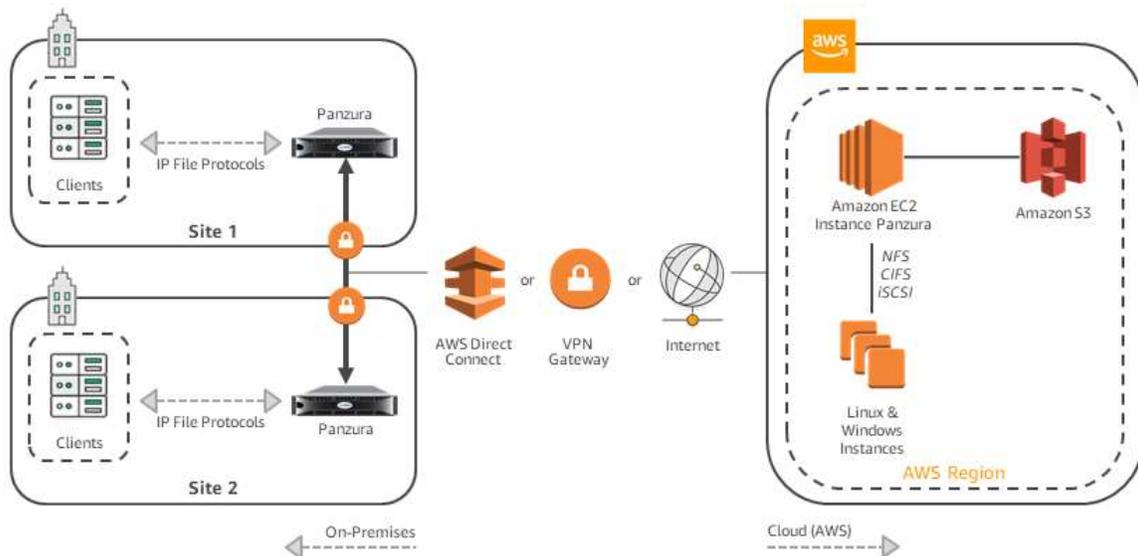


Figure 7: Panzura Architecture

Freedom Use Cases - Archive On-Premises Data to AWS

Freedom archive is installed at the customer data center or remote branch office and will support local clients through SMB and NFS. Data that is actively accessed is cached at the local site while 100% of the data resides in Amazon S3. This provides both an archive for less frequently accessed data, as well as a disaster recover copy on Amazon S3. Included in recovery are integrated snapshots, which provide a specific point-in-time of recovery in the Panzura Cloud File System.

Freedom Use Cases - Centralize Remote Data to Amazon S3

The capabilities in the Freedom archive use case extend to support multi-site NAS. With multi-site NAS, file data can be centralized in Amazon S3 and accessed at the remote or branch offices without any need for modifying existing applications or workflows using NAS. Data files are synchronized to Amazon S3 and active data is left in cache at the remote and on-premises sites. Files are synchronized immediately to Amazon S3 when updated, and Panzura uses compression and deduplication to increase transfer efficiency. Because all data files are synchronized to Amazon S3, data can be used for purposes of disaster recovery for the multiple distributed sites.

Freedom Use Cases - Collaboration from Multiple Sites on Amazon S3

For multi-site scenarios, Panzura can provide read and write access to files. Using global file locking, users in dispersed sites can collaborate and make updates to shared files stored in Amazon S3. This collaboration use case leverages transfer efficiency with deduplication and compression while utilizing automated caching so only active files are stored at the edge site for faster access.

CTERA

CTERA is a comprehensive solution for consolidating data from multiple locations to AWS. This includes local on-premises NAS storage delivered by the CTERA gateway, as well as mobile devices and endpoints using secure data transfer to AWS from the CTERA host agent without the need for a gateway.

ARCHITECTURAL OVERVIEW

The CTERA architecture provides secure data transfer of your endpoints whether they are located in a datacenter, remote location, or completely mobile. CTERA Portal connects to the following devices: CTERA Cloud Storage

Gateway, CTERA Agents, and CTERA Mobile, supporting desktops, servers, and mobile endpoints. The CTERA Cloud Storage Gateway is a physical or virtual appliance that provide local storage, cloud storage, data protection, and collaboration capabilities.

CTERA Cloud Storage Gateways feature a full set of NAS capabilities and comprehensive backup and file sync and share functionalities, utilizing on-premises storage capabilities for speed and local sharing, while taking advantage of cloud storage for off-site backup, universal access, file sharing, and folder synchronization. By synchronizing to Amazon S3 storage, silos of off premises and branch office data can be centralized and managed for better data visibility, analytics, and collaboration.

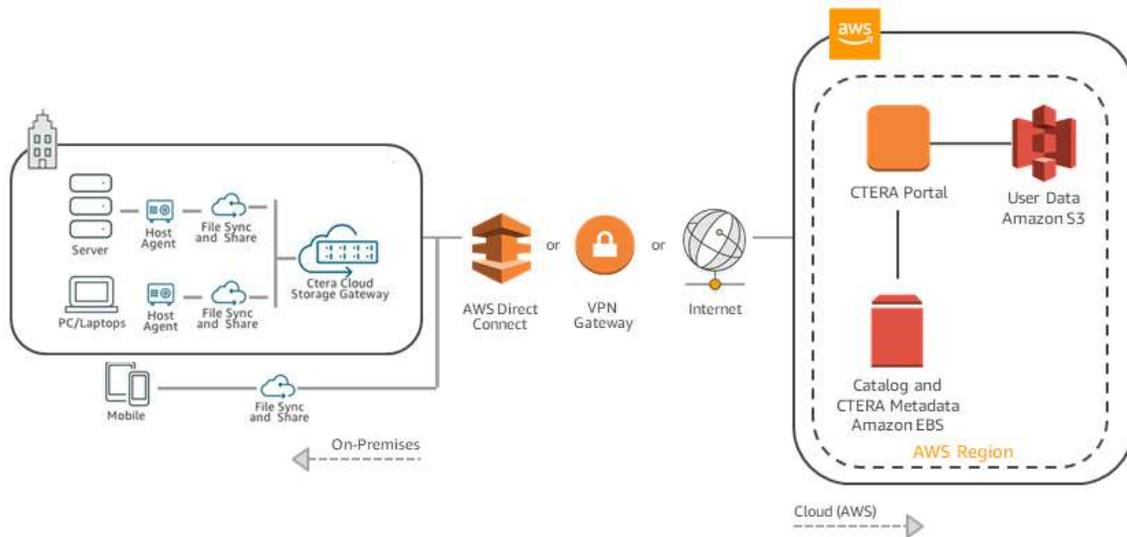


Figure 8: CTERA Architecture

Use Case: Global File Collaboration for NAS and Mobile Devices to AWS

Utilizing the file sync agent with CTERA, you can select directories and files to automatically synchronize to the CTERA portal running on AWS. The agent will also provide a virtual cloud drive on the desktop, creating a simple target for securely transferring data to Amazon S3 through the CTERA Portal. Data in Amazon S3 then can be accessed by multiple users but is controlled and secured by the CTERA Portal running on Amazon EC2. Amazon S3 becomes a centralized data repository for collaboration on files and data across the enterprise.

Use Case: Primary Storage at the Remote or Branch Office with Sync to Amazon S3

CTERA Cloud Storage Gateway can provide primary storage for both CIFS/SMB and NFS at on premises or remote sites. Data updates are automatically synchronized according to policy to Amazon S3 for backup, recovery, or in-cloud workload use. The Cloud Storage Gateway sends data to Amazon S3 encrypted, deduplicated, and compressed to minimize the time and amount of data required to complete the synchronization.

Nasuni

Nasuni provides a Gateway Solution, distinguished by its global locking feature used during global file sharing. Nasuni global file locking ensures that data can be written from many global sources by one client at a time and should be utilized when single-writer consistency is required for file collaboration.

ARCHITECTURAL OVERVIEW

At the Nasuni core is a global file system, UniFS. File system data and metadata are stored locally on a Nasuni Edge Appliance, also referred to as “Filer,” with all data persisted in Amazon S3. When the file system changes at the local site, the changes are sent to Amazon S3 by the Nasuni Filer. The updates can be controlled by user settings for Quality of Service (QoS) and scheduled snapshots. Nasuni also will track file versions and the location of those versions for creating, updating, deleting, and moving files. In addition to files, volume operations also can have version history for operations such as creating, deleting, and moving.

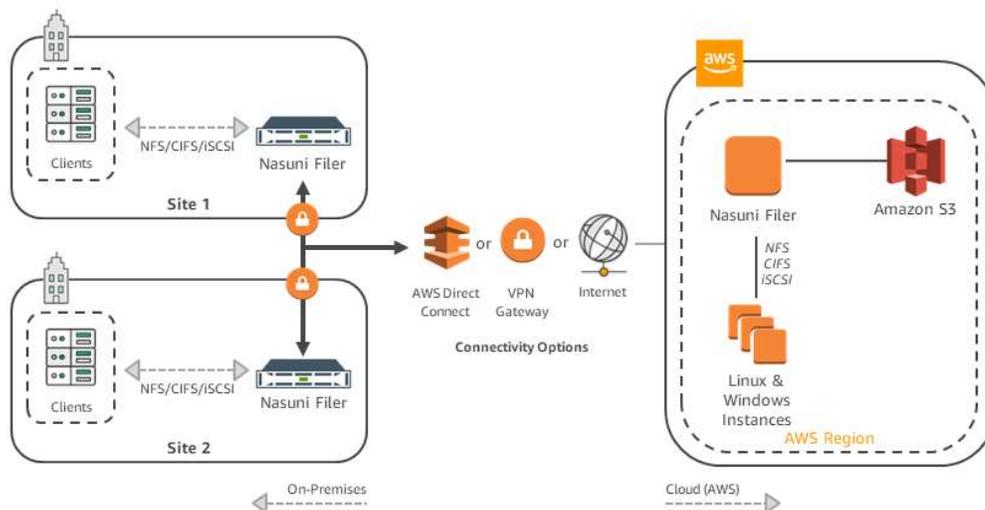


Figure 9: Nasuni Architecture

Synchronization also can occur independently of snapshots if the global file locking is enabled for a volume. The Global Locking feature prevents conflicts when two or more users attempt to change the same file on different Nasuni Filers. If you enable the Global Locking feature for a directory, any files in that directory can only be changed by one user at a time. In this case, the synchronization will occur when the file releases the file lock. Global file lock then presents a consistent version of user data across multiple sites and geographies.

Use Case: Global File Collaboration with Global Lock on AWS

Global File Locking ensures that no two users in any site can change data simultaneously. Nasuni Global Lock is recommended when the consistency of the data is paramount and users want to avoid any update race conditions. Nasuni Filers can be installed on a VM or physical device on premises or on an Amazon EC2 instance. Frequently accessed data will be cached on Nasuni Filers and persisted to Amazon S3. Both Windows and Linux hosts can access the data in multiple locations from the Nasuni Filer with support for NFS, SMB, HTTP/REST, and SFTP.

Use Case: Remote Branch Data Protection to Amazon S3

Protect remote office or branch office data by persisting snapshots to Amazon S3 and add versioning for updates. Using snapshots, you can restore a file or folder (for a CIFS or NFS volume, or FTP/SFTP directory) or an entire volume (for an iSCSI volume) from any location. Snapshots can be scheduled to account for bandwidth or recovery point objective (RPO) requirements and snapshots are stored in Amazon S3 with eleven 9s of durability.

Pro Tip: Enabling Global Locking can have an impact on performance, so do not enable Global Locking if users do not need to collaborate on the same files in different locations. Also, for large data loads, only enable Global Lock after the initial data seeding.

Pro Tip: If two Nasuni Filers both have Global Locking enabled for the same directory and a file is deleted or removed in the directory on just one of the Nasuni Filers, the data may still persist to the other Nasuni Filer.

Conclusion

There is a wide variety of partner NAS solutions available on AWS. Each solution has unique characteristics to leverage the elasticity, scale, and economics of persisting data to AWS Cloud storage. These solutions also offer customers features such as collaboration, multi-protocol access, instant point-in-time copies, and replication. For each case, partner solutions enable the transition and movement of data and application workloads to AWS.

Contributors

The following individuals and organizations contributed to this document:

- Peter Kisich, Storage Partner Solutions Architect, AWS
- Henry Axelrod, Storage Partner Solutions Architect, AWS
- Isaiah Weiner, Senior Manager, Solutions Architecture, AWS

Document Revisions

Date	Description
December 2018	First publication
