



Data Migration and Data Protection

NetApp Solutions

NetApp
September 19, 2024

Table of Contents

- Data Migration and Data Protection 1
 - ONTAP cyber vault..... 1
 - Data Migration 7
 - Data Protection 91
 - Security 92

Data Migration and Data Protection

ONTAP cyber vault

ONTAP cyber vault overview

The primary driving threat that necessitates the implementation of cyber vaulting is the growing prevalence and evolving sophistication of cyber-attacks, particularly ransomware and data breaches. [With a rise in phishing](#) and ever more sophisticated methods of credential stealing, credentials used to begin a ransomware attack could then be used to access infrastructure systems. In these cases, even hardened infrastructure systems are at risk of attack. The only defense to a compromised system is to have your data protected and isolated in a cyber vault.



Beginning in July 2024, content from technical reports previously published as PDFs has been integrated with ONTAP product documentation. In addition, new technical reports (TRs) such as this document will no longer be getting TR numbers.

What is cyber vaulting?

Cyber vaulting is a specific data protection technique that involves storing critical data in an isolated environment, separate from the primary IT infrastructure.

"Air gapped", **immutable** and **indelible** data repository that is immune to threats affecting the main network, such as malware, ransomware, or even insider threats. Cyber vaulting can be achieved with **immutable** and **indelible** snapshots.

Air-gapping backups that use traditional methods involve creating space and physically separating the primary and secondary media. By moving the media offsite and/or severing connectivity, bad actors have no access to the data. This protects the data but can lead to slower recovery times.

NetApp's approach to cyber vaulting

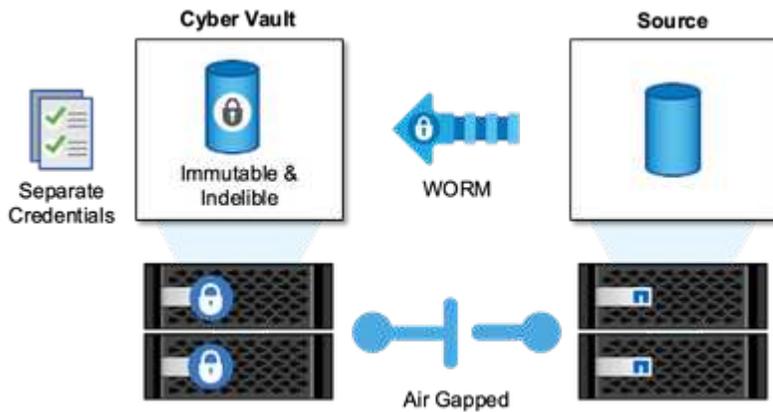
Key features of NetApp reference architecture for cyber vaulting include:

- Secure, isolated storage infrastructure (e.g., air gapped storage systems)
- Copies of the data must be both **immutable** and **indelible** without exception
- Strict access controls and multi-factor authentication
- Rapid data restoration capabilities

You can use NetApp storage with ONTAP as an air-gapped cyber vault by leveraging [SnapLock Compliance to WORM-protect Snapshot copies](#). You can perform all the basic SnapLock Compliance tasks on the Cyber vault. Once configured, Cyber vault volumes are automatically protected, eliminating the need to manually commit the Snapshot copies to WORM. More information on logical air-gapping can be found in this [blog](#)

SnapLock Compliance is used to comply with the Banking and Financial regulations SEC 70-a-4(f), FINRA 4511(c), and CFTC 1.31(c)-(d). It has been certified by Cohasset Associates to adhere to these regulations (audit report available upon request). By using SnapLock Compliance with this certification you get a hardened mechanism for air gapping of your data that is relied upon by the largest financial institutions in the world to

ensure both retention and retrieval of banking records.



Cyber vault ONTAP terminology

These are the terms commonly used in cyber vault architectures.

Autonomous Ransomware Protection (ARP) - Autonomous Ransomware Protection (ARP) feature uses workload analysis in NAS (NFS and SMB) environments to proactively, and in real time, detect and warn about abnormal activity that might indicate a ransomware attack. When an attack is suspected, ARP also creates new Snapshot copies, in addition to existing protection from scheduled Snapshot copies. For more information, see the [ONTAP documentation on Autonomous Ransomware Protection](#)

Airgap (Logical) - You can configure NetApp storage with ONTAP as a logical air-gapped cyber vault by leveraging [SnapLock Compliance to WORM-protect Snapshot copies](#)

Airgap (Physical) - A physical airgapped system has no network connectivity to it. Using tape backups, you can move the images to another location. The SnapLock Compliance logical air gap is just as robust as a physical airgapped system.

Bastion host - A dedicated computer on an isolated network, configured to withstand attacks.

Immutable Snapshot copies - Snapshot copies that are not able to be modified, without exception (including a support organization or the ability to low level format the storage system).

Indelible Snapshot copies - Snapshot copies that are not able to be deleted, without exception (including a support organization or the ability to low level format the storage system).

Tamperproof Snapshot copies - Tamperproof Snapshot copies use the SnapLock Compliance clock feature to lock Snapshot copies for a specified period. These locked snapshots can not be deleted by any user or NetApp support. You can use locked Snapshot copies to recover data if a volume is compromised by a ransomware attack, malware, hacker, rogue administrator or accidental deletion. For more information, see the [ONTAP documentation on Tamperproof Snapshot copies](#)

SnapLock - SnapLock is a high-performance compliance solution for organizations that use WORM storage to retain files in unmodified form for regulatory and governance purposes. For more information, see the [ONTAP documentation on SnapLock](#).

SnapMirror - SnapMirror is disaster recovery replication technology, designed to efficiently replicate data. SnapMirror can create a mirror (or exact copy of the data), vault (a copy of the data with longer Snapshot copy retention), or both to a secondary system, on premises or in the cloud. These copies can be used for many different purposes such as a disaster, bursting to the cloud, or a cyber vault (when using the vault policy and

locking the vault). For more information, see the [ONTAP documentation on SnapMirror](#)

SnapVault - In ONTAP 9.3 SnapVault was deprecated in favor of configuring SnapMirror using the vault or mirror-vault policy. This is term, while still used, has been deprecated as well. For more information, see the [ONTAP documentation on SnapVault](#).

Cyber vault sizing with ONTAP

Sizing a cyber vault requires understanding how much data that will need to be restored in a given Recovery Time Objective (RTO). Many factors play into properly designing a right sized cyber vault solution.

Sizing considerations

1. What are the source platform models (FAS v AFF A-Series v AFF C-Series)?
2. What is the bandwidth and latency between the source and cyber vault?
3. How large are the file sizes and how many files?
4. What is your recovery time objective?
5. How much data do you need to be recovered within the RTO?
6. How many SnapMirror fan-in relationships will the cyber vault be ingesting?
7. Will there be single or multiple recoveries happening at the same time?
8. Will those multiple recoveries be happening to the same primary?
9. Will SnapMirror be replicating to the vault during a recovery from a vault?

Creating a cyber vault with ONTAP

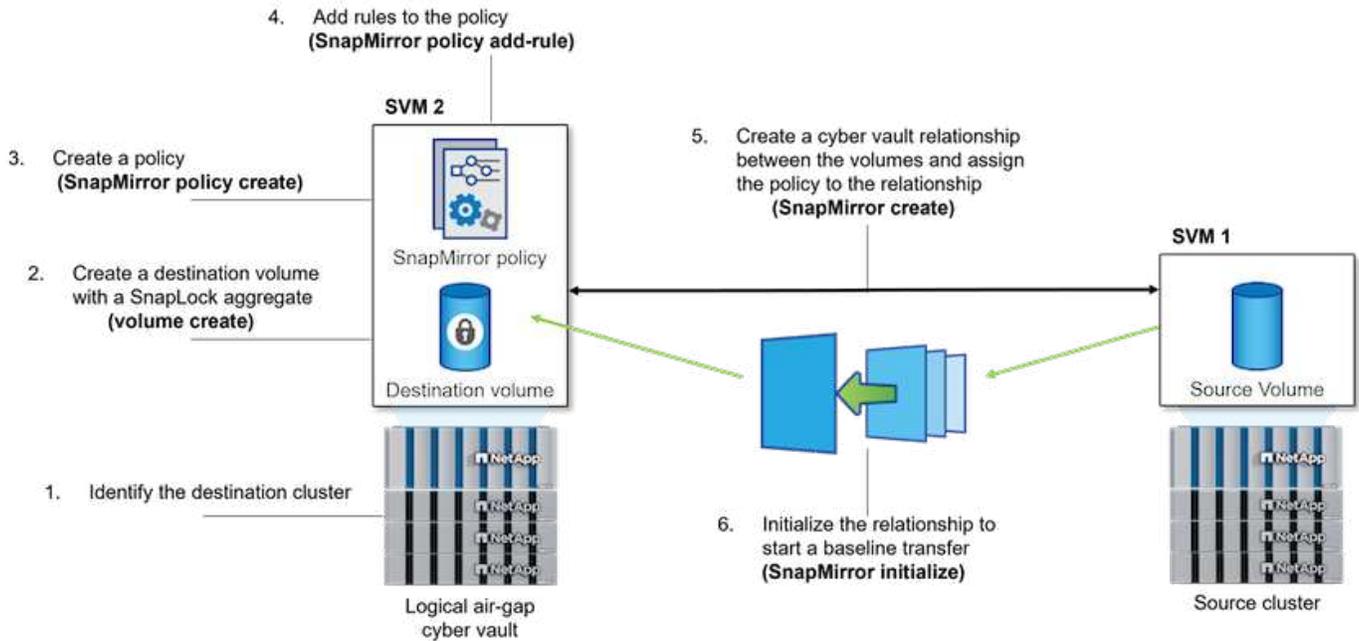
The steps below will assist with the creation of a cyber vault with ONTAP.

Before you begin

- The source cluster must be running ONTAP 9 or later.
- The source and destination aggregates must be 64-bit.
- The source and destination volumes must be created in peered clusters with peered SVMs. For more information, see [Cluster Peering](#).
- If volume autogrow is disabled, the free space on the destination volume must be at least five percent more than the used space on the source volume.

About this task

The following illustration shows the procedure for initializing a SnapLock Compliance vault relationship:



Steps

1. Identify the destination array to become the cyber vault to receive the air gapped data.
2. On the destination array, to prepare the cyber vault, [install the ONTAP One license](#), [initialize the Compliance Clock](#), and, if you are using an ONTAP release earlier than 9.10.1, [create a SnapLock Compliance aggregate](#).
3. On the destination array, create a SnapLock Compliance destination volume of type DP:

```
volume create -vserver SVM_name -volume volume_name -aggregate aggregate_name
-snaplock-type compliance|enterprise -type DP -size size
```

4. Beginning with ONTAP 9.10.1, SnapLock and non-SnapLock volumes can exist on the same aggregate; therefore, you are no longer required to create a separate SnapLock aggregate if you are using ONTAP 9.10.1. You use the volume `-snaplock-type` option to specify a Compliance type. In ONTAP releases earlier than ONTAP 9.10.1, the SnapLock mode, Compliance is inherited from the aggregate. Version-flexible destination volumes are not supported. The language setting of the destination volume must match the language setting of the source volume.

The following command creates a 2 GB SnapLock Compliance volume named `dstvolB` in SVM2 on the aggregate `node01_aggr`:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate node01_aggr
-snaplock-type compliance -type DP -size 2GG
```

5. On the destination cluster, to create the airgap, set the default retention period, as described in [Set the default retention period](#).
A SnapLock volume that is a vault destination has a default retention period assigned to it. The value for this period is initially set to a minimum of 0 years and maximum of 30 years for SnapLock Compliance volumes. Each NetApp Snapshot copy is committed with this default retention period at first. The default-retention-period must be changed. The retention period can be extended later, if needed, but never shortened. For more information, see [Set retention time overview](#).
6. [Create a new replication relationship](#) between the non-SnapLock source and the new SnapLock destination you created in Step 3.

This example creates a new SnapMirror relationship with destination SnapLock volume dstvolB using a policy of XDPDefault to vault Snapshot copies labeled daily and weekly on an hourly schedule:

```
cluster2::> snapmirror create -source-path SVM1:srcvolA -destination-path SVM2:dstvolB -vserver SVM2 -policy XDPDefault -schedule hourly
```

[Create a custom replication policy](#) or a [custom schedule](#) if the available defaults are not suitable.

7. On the destination SVM, initialize the SnapVault relationship created in Step 5:

```
snapmirror initialize -destination-path destination_path
```

8. The following command initializes the relationship between the source volume srcvolA on SVM1 and the destination volume dstvolB on SVM2:

```
cluster2::> snapmirror initialize -destination-path SVM2:dstvolB
```

9. After the relationship is initialized and idle, use the snapshot show command on the destination to verify the SnapLock expiry time applied to the replicated Snapshot copies.

This example lists the Snapshot copies on volume dstvolB that have the SnapMirror label and the SnapLock expiration date:

```
cluster2::> snapshot show -vserver SVM2 -volume dstvolB -fields snapmirror-label, snaplock-expiry-time
```

Cyber vault hardening

These are the additional recommendations to harden an ONTAP cyber vault. Please consult the ONTAP hardening guide below for more recommendations and procedures.

Cyber vault hardening recommendations

- Isolate the cyber vault's management planes
- Do not enable data LIFs on the destination cluster as they are an additional attack vector
- On the destination cluster, limit intercluster LIF access to the source cluster with a service policy
- Segment the management LIF on the destination cluster for limited access with a service policy and a bastion host
- Restrict all data traffic from the source cluster to the cyber vault to allow only the ports required for SnapMirror traffic
- Where possible, disable any unneeded management access methods within ONTAP to decrease the attack surface
- Enable audit logging and remote log storage
- Enable multi-admin verification and require verification from an admin outside your regular storage administrators (e.g. CISO staff)
- Implement role-based access controls
- Require administrative multifactor authentication for System Manager and ssh
- Use token based authentication for scripts and REST API calls

Please refer to the [ONTAP hardening guide](#), [Multi-admin verification overview](#) and [ONTAP multifactor authentication guide](#) for how to accomplish these hardening steps.

Cyber vault interoperability

ONTAP hardware and software can be used to create a cyber vault configuration.

ONTAP hardware recommendations

All ONTAP unified physical arrays can be used for a cyber vault implementation.

- FAS hybrid storage offers the most cost-efficient solution.
- AFF C-Series offers the most efficient power consumption and density.
- AFF A-Series is the highest performing platform offering the best RTO. With the recent announcement of our latest AFF A-Series, this platform will offer the best storage efficiency without performance compromise.

ONTAP software recommendations

Beginning with ONTAP 9.14.1, you can specify retention periods for specific SnapMirror labels in the SnapMirror policy of the SnapMirror relationship so that the replicated Snapshot copies from the source to the destination volume are retained for the retention-period specified in the rule. If no retention period is specified, the default-retention-period of the destination volume is used.

Beginning with ONTAP 9.13.1, you can instantaneously restore a locked Snapshot copy on the destination SnapLock volume of a SnapLock vault relationship by creating a FlexClone with the snaplock-type option set to "non-snaplock" and specifying the Snapshot copy as the "parent-snapshot" when executing the volume clone creation operation. Learn more about [creating a FlexClone volume with a SnapLock type](#).

MetroCluster configuration

For MetroCluster configurations, you should be aware of the following:

- You can create a SnapVault relationship only between sync-source SVMs, not between a sync-source SVM and a sync-destination SVM.
- You can create a SnapVault relationship from a volume on a sync-source SVM to a data-serving SVM.
- You can create a SnapVault relationship from a volume on a data-serving SVM to a DP volume on a sync-source SVM.

Cyber vault resources

To learn more about the information described in this cyber vault information, refer to the following additional information and security concepts.

- [NetApp Cyber vaulting: Multilayered Data Protection Solutions Brief](#)
- [NetApp Earns AAA Rating for Industry-First AI-Driven On-Box Ransomware Detection Solution](#)
- [Elevate cyber resilience with the most secure storage on the planet](#)
- [ONTAP security hardening guide](#)
- [NetApp Zero Trust](#)

- [NetApp Cyber Resilience](#)
- [NetApp Data Protection](#)
- [Cluster and SVM peering overview with the CLI](#)
- [SnapVault archiving](#)

Data Migration

Best-Practice Guidelines for NetApp XCP

TR-4863: Best-Practice Guidelines for NetApp XCP - Data Mover, File Migration, and Analytics

Karthikeyan Nagalingam, NetApp

This document provides NetApp XCP best-practice guidelines and a test scenario-based solution. These best practices cover the migration workflow for on-premises as well as cloud, file-system analytics, troubleshooting, and performance tuning of XCP. The test-scenario section covers customer use cases and their requirements, the NetApp solution using XCP, and benefits to the customer.

NetApp XCP

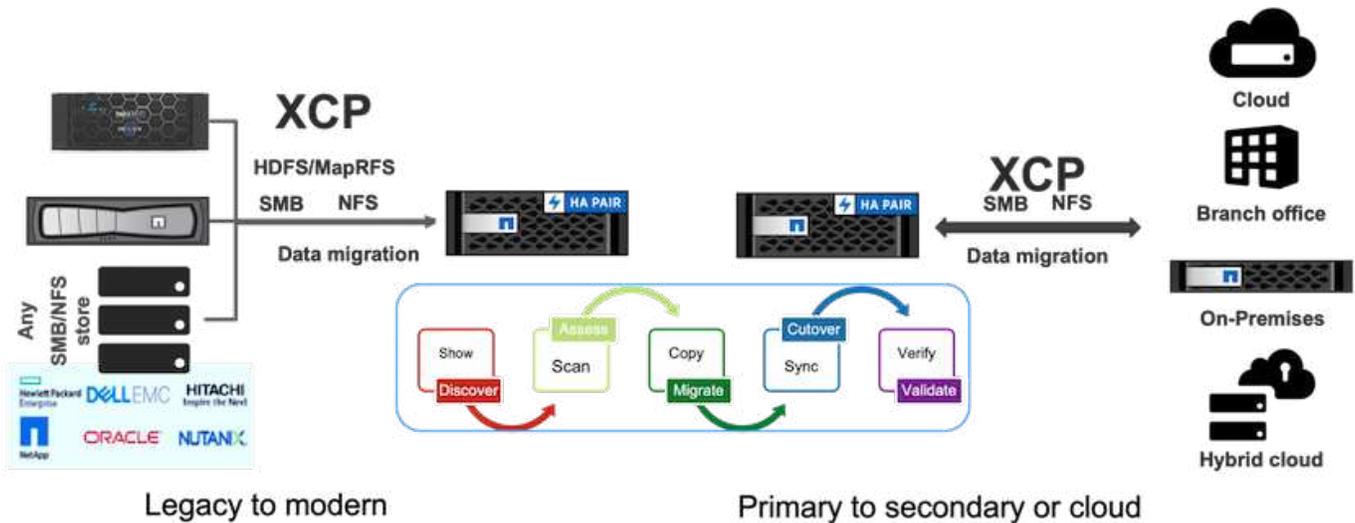
NetApp XCP transfers data by using multithreads and customizable features. It is designed for three major use cases: data move or migration, file-system analytics, and fast directory tree deletion.

Data move or migration

NetApp XCP transfers data from any NAS to NetApp NAS. This process consists of four major operations: scan, copy, sync, and verify. There are some additional features that help the data monitoring and transfer:

- **Scan.** Provides a high-level layout of NAS and MapR/HDFS data.
- **Copy.** Performs a baseline data transfer.
- **Sync.** Performs the incremental data transfer.
- **Verify.** Performs a thorough verification of the target.
- **Show (optional).** Discovers NAS shares.

The following figure illustrates XCP data migration and replication operations.



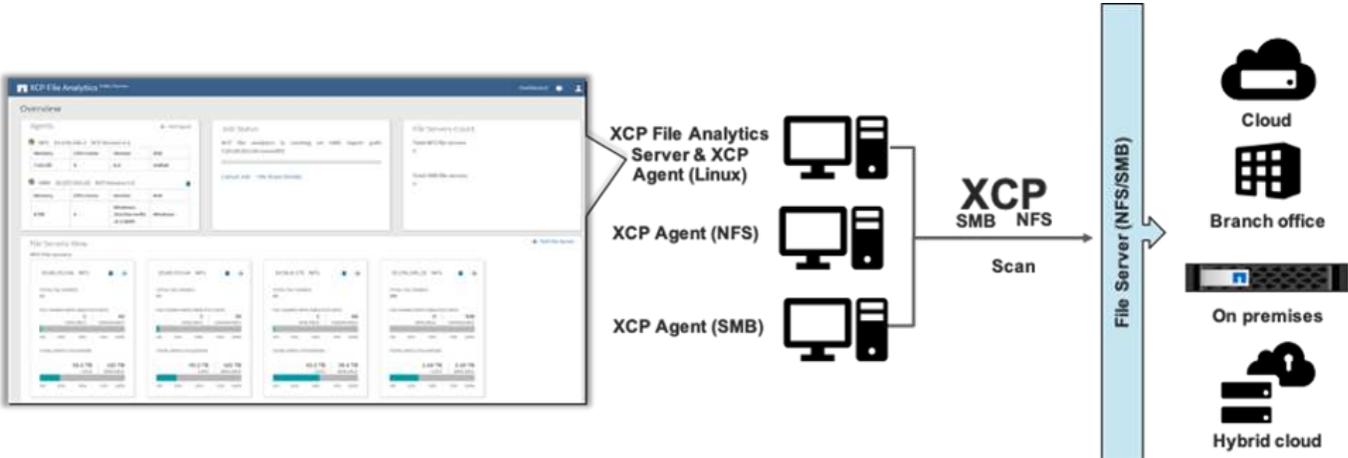
File-system analytics

NetApp XCP natively enables you to identify, scrutinize, and analyze unstructured data to improve insights—a key requirement for enterprise customers who want to use those insights for better planning, to operationalize high-value digital assets, and for data governance through reporting and assessment.

Customers that deal with sensitive data can use NetApp XCP to answer typical operational questions, such as the following:

- Where is my data?
- How much data and what types of files do we have?
- What data is actively used and how much is dormant?

The following figure illustrates NetApp XCP file analytics communication from the GUI.



Delete

It can be very challenging for storage teams and Electronic Design Automation (EDA) workloads to clean up large directories, whether it's stale data or test data that needs to be cleaned to recover storage space. XCP provides a fast delete functionality that can delete a complete directory tree. The NetApp XCP Delete function removes files and folders from a given NAS path. You can leverage the match filters to delete a specific set of files and folders. For a large number of files and folders, you can use the Force option, which does not require

a confirmation to delete.

Live Source Migration support

Live Source Migration support included in XCP 1.7 allows migration from a data source that is in active use (read and write activity). XCP leaves out files that are being used during the migration job, such as copy and sync running, and skipped files information is captured in the XCP log.

This feature supports changes on the source but does not support changes on the destination. During migration, the destination should not be active. Live Source Migration support is only available for NFS migrations.



No special settings are required for Live Source Migrations.

Prerequisites for XCP

Before you deploy NetApp XCP, the following prerequisites must be met:

1. Verify the NFS ports used by the NFS server by running the following command:

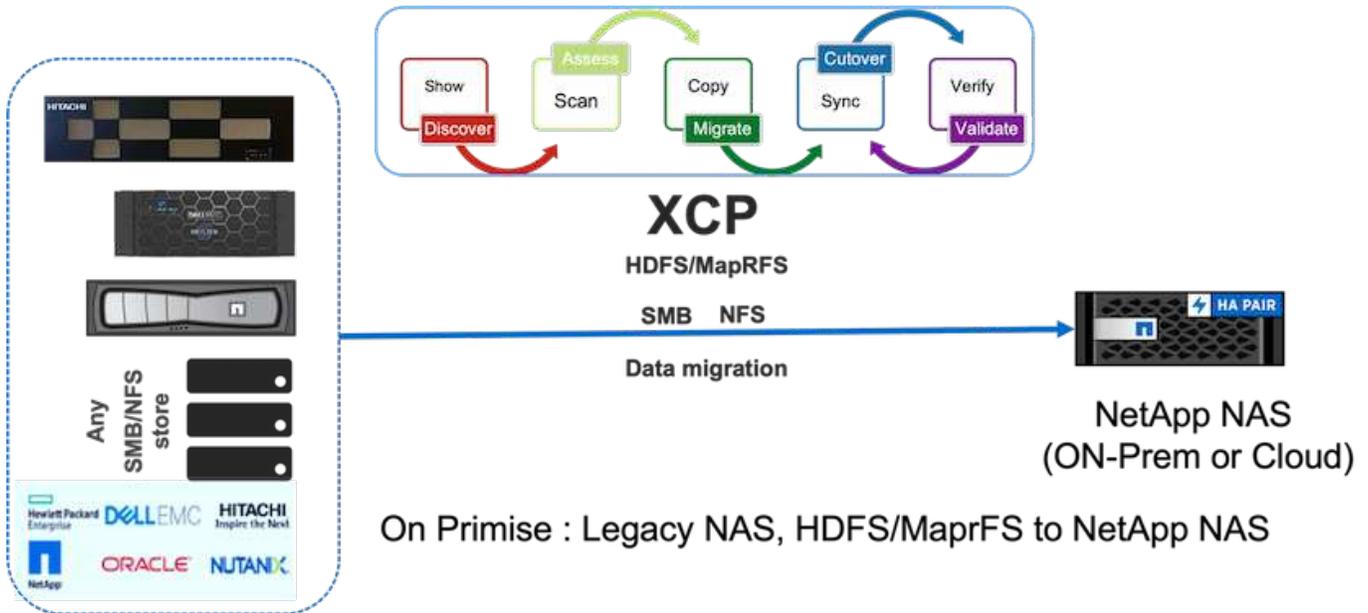
```
rpcinfo -p < NAS IP or on-prem nfs data LIF ip >
```

2. To access the location where you execute the XCP operations, such as on-premises or cloud instances (for example, Azure, AWS, or Google virtual machine [VM] instances), open the firewall ports for the NFS ports.
3. Verify that the NFS port is accessible from the XCP server by using the telnet command `<on-prem nfs data LIF ip or NAS IP > 2049`. The default port is 2049. If your environment has a different port, use that IP.
4. For NFS, verify that the shares are accessible from the XCP server by using the `showmount -e < NAS IP >` command.
5. Increase the number of inodes on the destination volume to more than the file count (number of files) on the source files.
6. Download the XCP license from the [NetApp XCP License Portal](#).
 - a. You must have a NetApp account in mysupport.netapp.com or you can register for free.
 - b. Download the license and have it ready.
7. Create one NFS share on-premises for each Azure NetApp volume or for the Cloud Volume Service (premium service level) in cloud for the XCP catalog.
8. Create an NAS volume and configure the share for the data destination.
9. For multiple XCP instances, you must have one or more servers or cloud instances to transfer the data from multiple source folders or files to the destination.
10. The `maxdir` size (default is 308MB) defines the maximum file count (approximately one million) in a single folder. Increase the `maxdir` size value to increase the file count. Increasing the value has an effect on additional CPU cycles.
11. In the cloud, NetApp recommends that you have ExpressRoute (Azure), Direct Connect (AWS), or Cloud Interconnect (GCP) between on-premises and cloud.

Migration workflow

Migration has different phases to follow for better planning and completion of the migration. To migrate data from third-party NAS storage or directly attached NAS exported storage using NetApp XCP, follow the migration guidelines provided in this section.

The following figure illustrates the migration workflow from any NAS to NetApp NAS.



On-premises

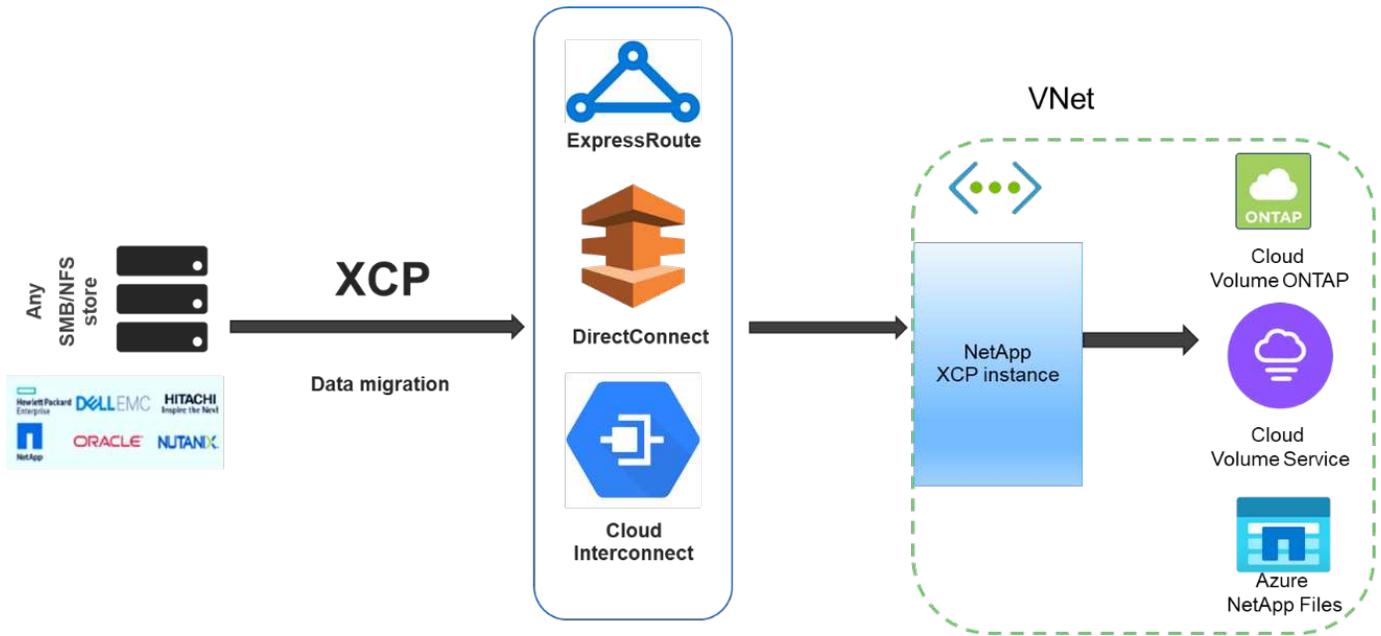
The migration workflow from any NAS to NetApp NAS includes the following steps:

1. Discover the NAS shares and data.
2. Scan the data and produce a report to find the layout of the data.
3. Create a baseline by running the XCP Copy command. For faster migrations, select more XCP instances and split the workload at the subfolder level to initiate parallel migration jobs.
4. For incremental updates, use XCP sync until the change rate is low for the cutover window.
5. Mark the source as read-only to perform a final sync by running the XCP sync command to complete the migration.
6. To verify that the data transferred correctly, compare the source and destination by running the `xcp verify` command.

Cloud

For the cloud, you can follow a similar on-premises migration workflow if the connectivity between on-premises and the cloud is direct connect (AWS), ExpressRoute (Azure), or cloud interconnect (GCP).

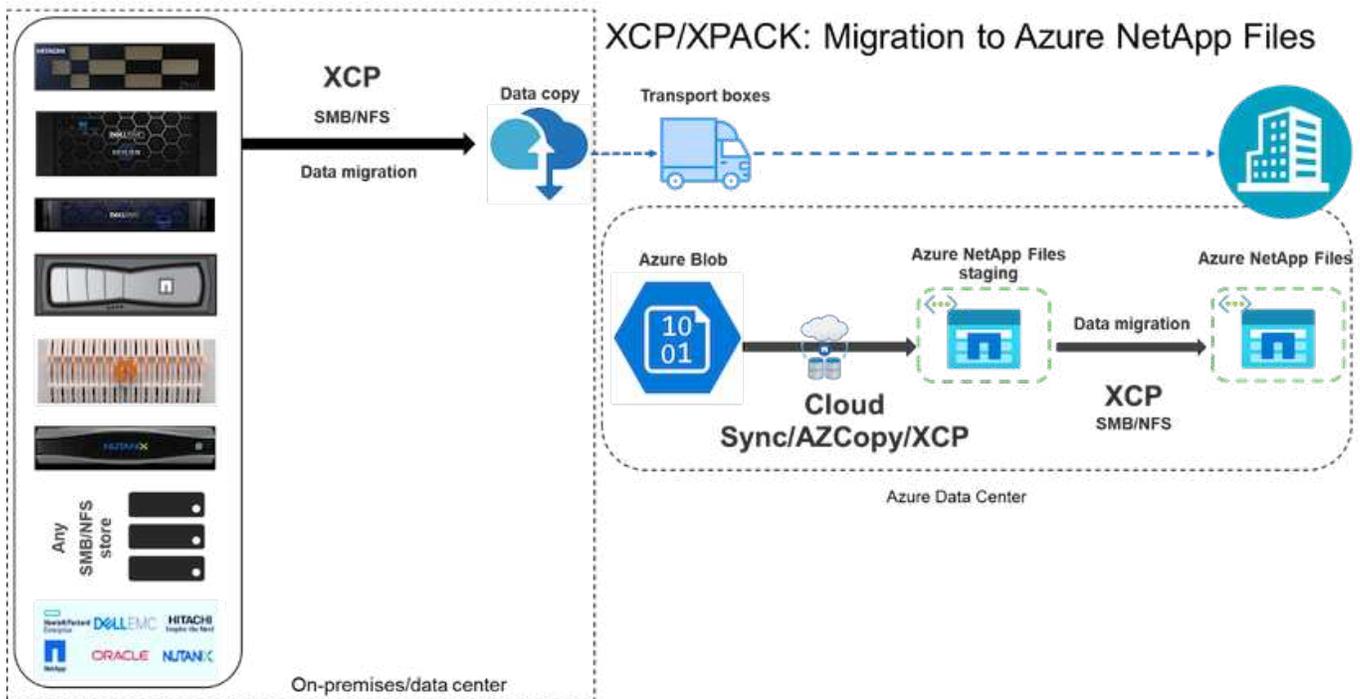
The following figure illustrates the migration workflow from on-premises to the cloud.



Data migration from any storage to cloud

If there is no direct internet connection between on-premises and the cloud, you must transfer the data from on-premises to the cloud through an offline data transport method such as truck. Each cloud service provider has a different method with different terminology to move data to their data center.

The following figure depicts the data mover solution for on-premises to Azure without ExpressRoute.

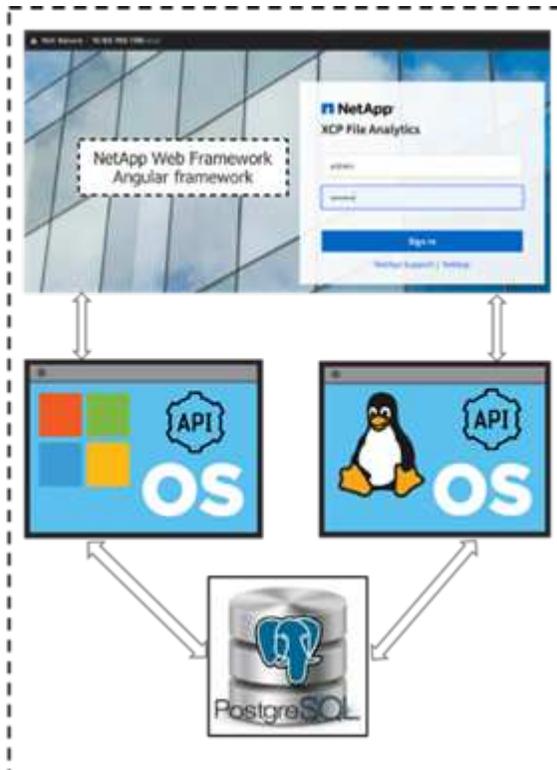


You can use a similar architecture with the respective components from the various cloud service providers.

File analytics

The NetApp XCP file analytics GUI helps to run file system scans by using XCP at the back end and visualizing statistics such as graphs and views for any NAS (NFS, SMB) file system. Starting in 1.6, XCP can be run as a service with the help of simple deployment steps by using the Configure and systemctl options. The XCP Configure option guides you to install and configure Postgres and a web server as well as collect credentials. The systemctl option runs XCP as a service for REST API communications from the GUI.

The following figure illustrates the XCP file analytics flow.



For more information about the high-level architecture of XCP file analytics, GUI-based dashboard views such as stats view, and file distribution view details, see the blog post [NetApp XCP 1.6 Delivers Open File Analytics and Infrastructure Improvements](#).

There is a limited GUI in XCP 1.6 for customized graphs. To create the required graphs, you can use the CLI to run the `xcp scan` command with matching filters. See the following examples.

1. Generate a list of files modified beyond a year by using `xcp scan` and the `-match` filter with the space consumed.

```

[root@ch-vm-cent7-2 linux]# ./xcp scan -match "modified > 1*year" -l -q
192.168.89.110:/ifs/data_for_analysis > modified_morethan_year
XCP 1.6P1; (c) 2020 NetApp, Inc.; Licensed to Karthikeyan Nagalingam
[NetApp Inc] until Wed Sep  9 13:19:35 2020

xcp: WARNING: CPU count is only 1!

Filtered: 1 did not match

Xcp command : xcp scan -match modified > 1*year -l -q
192.168.89.110:/ifs/data_for_analysis
5,055 scanned, 5,054 matched, 0 error
Speed          : 1.10 MiB in (510 KiB/s), 110 KiB out (49.5 KiB/s)
Total Time    : 2s.
STATUS        : PASSED
[root@ch-vm-cent7-2 linux]#
[root@ch-vm-cent7-2 linux]# cat modified_morethan_year
rwxr-xr-x --- 7056 503          0      512  7y99d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/6/_SUCCESS
rwxr-xr-x --- 7056 503        270 8.50KiB  7y99d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/6/part-r-00000
rw-r--r-- --- 7056 503          0      512  7y58d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/6/SUCCESS.crc
rw-r--r-- --- 7056 503        270 8.50KiB  7y99d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/6/out_original
rw-r--r-- --- 7056 503        270 8.50KiB  7y99d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/6/out_sorted
rwxr-xr-x --- 7056 503          0      512  7y99d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/2/_SUCCESS
rwxr-xr-x --- 7056 503         90 8.50KiB  7y99d
data_for_analysis/benchmarks/benchmarks/udf_TOBAGandTOTUPLE_7_benchmark.
out/2/part-r-00000
...
< console output removed due o page space size >
...

```

2. Find the space used by files that are more than one year old.

```

[root@ch-vm-cent7-2 linux]# ./xcp -du -match "modified > 1*year"

```

```

192.168.89.110:/ifs/data_for_analysis/
XCP 1.6.1; (c) 2020 NetApp, Inc.; Licensed to Karthikeyan Nagalingam
[NetApp Inc] until Wed Sep  9 13:19:35 2020
xcp: WARNING: CPU count is only 1!
52.5KiB
data_for_analysis/benchmarks/benchmarks/Macro_Scope_1_benchmark.out
28.5KiB
data_for_analysis/benchmarks/benchmarks/CollectedGroup_6_benchmark.out
28.5KiB data_for_analysis/benchmarks/benchmarks/Foreach_11_benchmark.out
153KiB
data_for_analysis/benchmarks/benchmarks/SecondarySort_9_benchmark.out
412KiB
data_for_analysis/benchmarks/benchmarks/CoGroupFlatten_6_benchmark.out
652KiB data_for_analysis/benchmarks/benchmarks/Iterator_1_benchmark.out
652KiB
data_for_analysis/benchmarks/benchmarks/LoaderDefaultDir_1_benchmark.out
652KiB data_for_analysis/benchmarks/benchmarks/Order_4_benchmark.out
28.5KiB
data_for_analysis/benchmarks/benchmarks/MapPartialAgg_4_benchmark.out/2
28.5KiB
data_for_analysis/benchmarks/benchmarks/CastScalar_11_benchmark.out/2
1.29MiB data_for_analysis/benchmarks/benchmarks/Order_18_benchmark.out
652KiB
data_for_analysis/benchmarks/benchmarks/FilterBoolean_5_benchmark.out
20.5KiB
data_for_analysis/benchmarks/benchmarks/Macro_DefinitionAndInline_5_benc
hmark.out/2
628KiB data_for_analysis/benchmarks/benchmarks/Types_29_benchmark.out
...
< console output removed due o page space size >
...
3.18MiB data_for_analysis/benchmarks/benchmarks/hadoop10
340KiB data_for_analysis/benchmarks/benchmarks/Split_5_benchmark.out
5.90GiB data_for_analysis/benchmarks/benchmarks
6.56GiB data_for_analysis/benchmarks
6.56GiB data_for_analysis

Filtered: 488 did not match

Xcp command : xcp -du -match modified > 1*year
192.168.89.110:/ifs/data_for_analysis/
Stats          : 5,055 scanned, 4,567 matched
Speed          : 1.10 MiB in (1.36 MiB/s), 110 KiB out (135 KiB/s)
Total Time    : 0s.
STATUS        : PASSED
[root@ch-vm-cent7-2 linux]#

```

3. Find the total size and graphical view of data that was modified more than one year ago.

```
[root@ch-vm-cent7-2 linux]# ./xcp -stats -match "modified > 1*year"
-html 192.168.89.110:/ifs/data_for_analysis/ >
modified_morethan_year_stats.html
XCP 1.6.1; (c) 2020 NetApp, Inc.; Licensed to Karthikeyan Nagalingam
[NetApp Inc] until Wed Sep 9 13:19:35 2020

xcp: WARNING: CPU count is only 1!

Xcp command : xcp -stats -match modified > 1*year -html
192.168.89.110:/ifs/data_for_analysis/
Stats       : 5,055 scanned, 4,567 matched
Speed       : 1.10 MiB in (919 KiB/s), 110 KiB out (89.1 KiB/s)
Total Time  : 1s.
STATUS      : PASSED
[root@ch-vm-cent7-2 linux]#
```

The following report is a custom example scan of files that were modified more than one year ago.

Command scan 192.168.89.110:/ifs/data_for_analysis

Options '-stats': True, '-match': 'modified > 1*year'

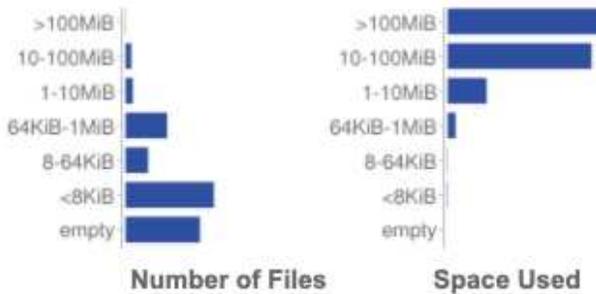
Unreadable directories None Unreadable files None

Filters: Unmatched None

Summary 5,055 scanned, 4,567 matched, 1.10 MiB in (924 KiB/s), 110 KiB out (89.7 KiB/s), 1s.

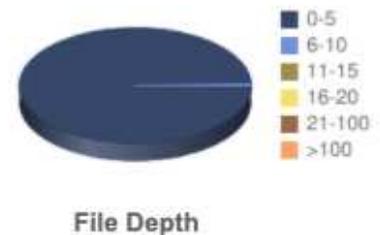
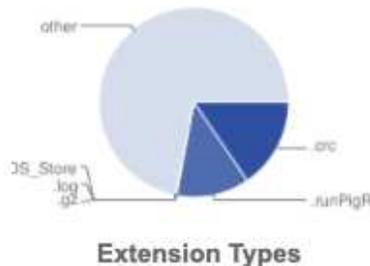
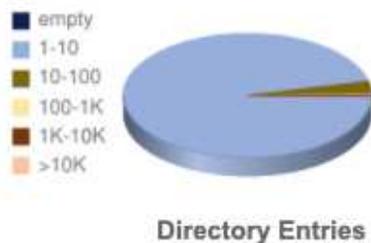
	Count	Used	Avg	Max
All File Types	4,567	6.56 GiB	Name Length 14	52
Regular Files	3,894	6.56 GiB	File Size 1.72 MiB	678 MiB
Directories	673	2.75 MiB	Directory Entries 7	1,463
Symlinks	None	0	File Depth 3	6
Specials	None	0		

7056 4,567
Top 5 File Owners



7056 6.56 GiB
Top 5 Space Users

Dedupe Estimate: N/A Zero Blocks: N/A Hard Links: 0 Extensions: 6 Groups: 1 Users: 1



Deployment steps

This section covers deployment steps for NetApp XCP for data transfer.

Test bed details

The following table provides the details of the test bed that was used for this deployment and performance validation.

Solution components	Details
XCP version 1.7	<ul style="list-style-type: none">• One Linux server - Linux (RHEL 7.9 or RHEL 8)• One Windows server – Windows Server 2019 standard
NetApp AFF storage array HA pair for the source volume	<ul style="list-style-type: none">• AFF8080• NetApp ONTAP 9• NFS protocol
NetApp AFF storage array HA pair for destination volume	<ul style="list-style-type: none">• AFF A800• ONTAP 9• NFS protocol
Fujitsu PRIMERGY RX2540 server	Each equipped with: <ul style="list-style-type: none">* 48 CPUs* Intel Xeon* 256GB physical memory* 10GbE dual port
Networking	10GbE

Deployment steps - NAS

To deploy NetApp XCP for data transfer, first install and activate the XCP software on the destination location. You can review the details in the [NetApp XCP User Guide](#). To do so, complete the following steps:

1. Meet the prerequisites as detailed in the section “[Prerequisites for XCP.](#)”
2. Download the XCP software from the [NetApp XCP \(Downloads\) page](#).
3. Copy the downloaded XCP tar files to the XCP server.

```
# scp Documents/OneDrive\ -\ NetApp\  
Inc/XCP/software/1.6.1/NETAPP_XCP_1.6.1.tgz  
mailto:root@10.63.150.53:/usr/src
```

4. Untar the tarfile.

```
[root@mastr-53 src]# tar -zxvf NETAPP_XCP_1.6.1.tgz
```

5. Download the license from <https://xcp.netapp.com/license/xcp.xwic> and copy to the XCP server.
6. Activate the license.

```
[root@mastr-53 linux]# ./xcp activate
[root@mastr-53 src]# cp license /opt/NetApp/xFiles/xcp/license
[root@mastr-53 src]# cd /usr/src/xcp/linux/
[root@mastr-53 linux]# ./xcp activate
```

7. Find the source NFS port and destination NFS server. The default port is 2049.

```
[root@mastr-53 ~]# rpcinfo -p 10.63.150.213
[root@mastr-53 ~]# rpcinfo -p 10.63.150.63
```

8. Check the NFS connection. Check the NFS server (for both source and destination) by using telnet to the NFS server port.

```
[root@mastr-53 ~]# telnet 10.63.150.127 2049
[root@mastr-53 ~]# telnet 10.63.150.63 2049
```

9. Configure the catalog.

a. Create an NFS volume and export NFS for the XCP catalog. You can also leverage the operating system NFS export for XCP catalog.

```
A800-Node1-2::> volume create -vserver Hadoop_SVM -volume xcpcatalog
-aggregate aggr_Hadoop_1 -size 50GB -state online -junction-path
/xcpcatalog -policy default -unix-permissions ---rwxr-xr-x -type RW
-snapshot-policy default -foreground true
A800-Node1-2::> volume mount -vserver Hadoop_SVM -volume
xcpcatalog_vol -junction-path /xcpcatalog
```

b. Check the NFS export.

```
[root@mastr-53 ~]# showmount -e 10.63.150.63 | grep xcpca
/xcpcatalog (everyone)
```

c. Update xcp.ini.

```
[root@mastr-53 ~]# cat /opt/NetApp/xFiles/xcp/xcp.ini
# Sample xcp config
[xcp]
catalog = 10.63.150.64:/xcpcatalog

[root@mastr-53 ~]#
```

10. Find the source NAS exports by using `xcp show`. Look for:

```
== NFS Exports ==  
== Attributes of NFS Exports ==
```

```
[root@mastr-53 linux]# ./xcp show 10.63.150.127  
== NFS Exports ==  
<check here>  
== Attributes of NFS Exports ==  
<check here>
```

11. (Optional) Scan the source NAS data.

```
[root@mastr-53 linux]# ./xcp scan -newid xcpscantest4 -stats  
10.63.150.127:/xcpsrc_vol
```

Scanning the source NAS data helps you understand the data layout and find any potential issues for migration. The XCP scanning operation time is proportional to the number of files and the directory depth. You can skip this step if you are familiar with your NAS data.

12. Check the report created by `xcp scan`. Search mainly for unreadable folders and unreadable files.

```
[root@mastr-53 linux]# mount 10.63.150.64:/xpcatalog /xpcatalog  
base) nkarthik-mac-0:~ karthikeyannagalingam$ scp -r  
root@10.63.150.53:/xpcatalog/catalog/indexes/xcpscantest4  
Documents/OneDrive\ -\ NetApp\ Inc\XCP\customers\reports/
```

13. (Optional) Change the inode. View the number of inodes and modify the number based on the number of files to migrate or copy for both catalog and destination volumes (if required).

```
A800-Node1-2::> volume show -volume xpcatalog -fields files,files-used  
A800-Node1-2::> volume show -volume xcpdest -fields files,files-used  
A800-Node1-2::> volume modify -volume xpcatalog -vserver A800-Node1_vs1  
-files 2000000  
Volume modify successful on volume xpcatalog of Vserver A800-Node1_vs1.  
  
A800-Node1-2::> volume show -volume xpcatalog -fields files,files-used
```

14. Scan the destination volume.

```
[root@mastr-53 linux]# ./xcp scan -stats 10.63.150.63:/xcpdest
```

15. Check the source and destination volume space.

```
[root@mastr-53 ~]# df -h /xcpsrc_vol
[root@mastr-53 ~]# df -h /xcpdest/
```

16. Copy the data from source to destination by using `xcp copy` and check the summary.

```
[root@mastr-53 linux]# ./xcp copy -newid create_Sep091599198212
10.63.150.127:/xcpsrc_vol 10.63.150.63:/xcpdest
<command inprogress results removed>
Xcp command : xcp copy -newid create_Sep091599198212 -parallel 23
10.63.150.127:/xcpsrc_vol 10.63.150.63:/xcpdest
Stats          : 9.07M scanned, 9.07M copied, 118 linked, 9.07M indexed,
173 giants
Speed          : 1.57 TiB in (412 MiB/s), 1.50 TiB out (392 MiB/s)
Total Time    : 1h6m.
STATUS        : PASSED
[root@mastr-53 linux]#
```



By default, XCP creates seven parallel processes to copy the data. This can be tuned.



NetApp recommends that the source volume be read only. In real time, the source volume is a live, active file system. The `xcp copy` operation might fail because NetApp XCP does not support a live source that is continuously changed by an application.

For Linux, XCP requires an Index ID because XCP Linux performs cataloging.

17. (Optional) Check the inodes on the destination NetApp volume.

```
A800-Node1-2::> volume show -volume xcpdest -fields files,files-used
vserver          volume  files    files-used
-----
A800-Node1_vs1  xcpdest 21251126 15039685

A800-Node1-2::>
```

18. Perform the incremental update by using `xcp sync`.

```
[root@mastr-53 linux]# ./xcp sync -id create_Sep091599198212
Xcp command : xcp sync -id create_Sep091599198212
Stats       : 9.07M reviewed, 9.07M checked at source, no changes, 9.07M
reindexed
Speed       : 1.73 GiB in (8.40 MiB/s), 1.98 GiB out (9.59 MiB/s)
Total Time  : 3m31s.
STATUS     : PASSED
```

For this document, to simulate real-time, the one million files in the source data were renamed, and then the updated files were copied to the destination by using `xcp sync`. For Windows, XCP needs both source and destination paths.

19. Validate data transfer. You can validate that the source and destination have the same data by using `xcp verify`.

```
Xcp command : xcp verify 10.63.150.127:/xcpsrc_vol 10.63.150.63:/xcpdest
Stats       : 9.07M scanned, 9.07M indexed, 173 giants, 100% found
(6.01M have data), 6.01M compared, 100% verified (data, attrs, mods)
Speed       : 3.13 TiB in (509 MiB/s), 11.1 GiB out (1.76 MiB/s)
Total Time  : 1h47m.
STATUS     : PASSED
```

XCP documentation provides multiple options (with examples) for the `scan`, `copy`, `sync`, and `verify` operations. For more information, see the [NetApp XCP User Guide](#).



Windows customers should copy the data by using access control lists (ACLs). NetApp recommends using the command `xcp copy -acl -fallbackuser\. To maximum performance, considering the source volume that has SMB data with ACL and the data accessible by both NFS and SMB, the target must be an NTFS volume. Using XCP (NFS version), copy the data from the Linux server and execute the XCP (SMB version) sync with the -acl and -nodata options from the Windows server to copy the ACLs from source data to the target SMB data.`

For detailed steps, see [Configuring 'Manage Auditing and Security Log' Policy](#).

Deployment steps - HDFS/MapRFS data migration

In this section, we discuss the new XCP feature called Hadoop Filesystem Data Transfer to NAS, which migrates data from HDFS/MapRFS to NFS and vice versa.

Prerequisites

For the MapRFS/HDFS feature, you must perform the following procedure in a non-root user environment. Normally the non-root user is `hdfs`, `mapr`, or a user who has permission to make changes in the HDFS and MapRFS filesystem.

1. Set the CLASSPATH, HADOOP_HOME, NHDFS_LIBJVM_PATH, LD_LIBRARY_PATH, and NHDFS_LIBHDFS_PATH variables in the CLI or the .bashrc file of the user along with the xcp command.
 - NHDFS_LIBHDFS_PATH points to the libhdfs.so file. This file provides HDFS APIs to interact and manipulate the HDFS/MapRFS files and filesystem as a part of the Hadoop distribution.
 - NHDFS_LIBJVM_PATH points to the libjvm.so file. This is a shared JAVA virtual machine library in the jre location.
 - CLASSPATH points to all jars files using (Hadoop classpath -glob) values.
 - LD_LIBRARY_PATH points to the Hadoop native library folder location.

See the following sample based on a Cloudera cluster.

```
export CLASSPATH=$(hadoop classpath --glob)
export LD_LIBRARY_PATH=/usr/java/jdk1.8.0_181-
cloudera/jre/lib/amd64/server/
export HADOOP_HOME=/opt/cloudera/parcels/CDH-6.3.4-
1.cdh6.3.4.p0.6751098/
#export HADOOP_HOME=/opt/cloudera/parcels/CDH/
export NHDFS_LIBJVM_PATH=/usr/java/jdk1.8.0_181-
cloudera/jre/lib/amd64/server/libjvm.so
export NHDFS_LIBHDFS_PATH=$HADOOP_HOME/lib64/libhdfs.so
```

In this release, we support XCP scan, copy, and verify operations and data migration from HDFS to NFS. You can transfer data from a data lake cluster single worker node and multiple worker nodes. In the 1.8 release, root and non-root users can perform data migration.

Deployment steps - Non-root user migrates HDFS/MaprFS data to NetApp NFS

1. Follow the same steps mentioned from 1-9 steps from steps for deployment section.
2. In the following example, the user migrates data from HDFS to NFS.
 - a. Create a folder and files (using `hadoop fs -copyFromLocal`) in HDFS.

```

[root@n138 ~]# su - tester -c 'hadoop fs -mkdir
/tmp/testerfolder_src/util-linux-2.23.2/mohankarthikhdfs_src'
[root@n138 ~]# su - tester -c 'hadoop fs -ls -d
/tmp/testerfolder_src/util-linux-2.23.2/mohankarthikhdfs_src'
drwxr-xr-x   - tester supergroup           0 2021-11-16 16:52
/tmp/testerfolder_src/util-linux-2.23.2/mohankarthikhdfs_src
[root@n138 ~]# su - tester -c "echo 'testfile hdfs' >
/tmp/a_hdfs.txt"
[root@n138 ~]# su - tester -c "echo 'testfile hdfs 2' >
/tmp/b_hdfs.txt"
[root@n138 ~]# ls -ltrah /tmp/*_hdfs.txt
-rw-rw-r-- 1 tester tester 14 Nov 16 17:00 /tmp/a_hdfs.txt
-rw-rw-r-- 1 tester tester 16 Nov 16 17:00 /tmp/b_hdfs.txt
[root@n138 ~]# su - tester -c 'hadoop fs -copyFromLocal
/tmp/*_hdfs.txt hdfs:///tmp/testerfolder_src/util-linux-
2.23.2/mohankarthikhdfs_src'
[root@n138 ~]#

```

b. Check permissions in the HDFS folder.

```

[root@n138 ~]# su - tester -c 'hadoop fs -ls
hdfs:///tmp/testerfolder_src/util-linux-2.23.2/mohankarthikhdfs_src'
Found 2 items
-rw-r--r--   3 tester supergroup           14 2021-11-16 17:01
hdfs:///tmp/testerfolder_src/util-linux-
2.23.2/mohankarthikhdfs_src/a_hdfs.txt
-rw-r--r--   3 tester supergroup           16 2021-11-16 17:01
hdfs:///tmp/testerfolder_src/util-linux-
2.23.2/mohankarthikhdfs_src/b_hdfs.txt

```

c. Create a folder in NFS and check permissions.

```

[root@n138 ~]# su - tester -c 'mkdir
/xcpsrc_vol/mohankarthiknfs_dest'
[root@n138 ~]# su - tester -c 'ls -l
/xcpsrc_vol/mohankarthiknfs_dest'
total 0
[root@n138 ~]# su - tester -c 'ls -d
/xcpsrc_vol/mohankarthiknfs_dest'
/xcpsrc_vol/mohankarthiknfs_dest
[root@n138 ~]# su - tester -c 'ls -ld
/xcpsrc_vol/mohankarthiknfs_dest'
drwxrwxr-x 2 tester tester 4096 Nov 16 14:32
/xcpsrc_vol/mohankarthiknfs_dest
[root@n138 ~]#

```

d. Copy the files from HDFS to NFS using XCP, and check permissions.

```

[root@n138 ~]# su - tester -c '/usr/src/hdfs_nightly/xcp/linux/xcp
copy -chown hdfs:///tmp/testerfolder_src/util-linux-
2.23.2/mohankarthikhdfs_src/
10.63.150.126:/xcpsrc_vol/mohankarthiknfs_dest'
XCP Nightly_dev; (c) 2021 NetApp, Inc.; Licensed to Karthikeyan
Nagalingam [NetApp Inc] until Wed Feb 9 13:38:12 2022

xcp: WARNING: No index name has been specified, creating one with
name: autoname_copy_2021-11-16_17.04.03.652673

Xcp command : xcp copy -chown hdfs:///tmp/testerfolder_src/util-
linux-2.23.2/mohankarthikhdfs_src/
10.63.150.126:/xcpsrc_vol/mohankarthiknfs_dest
Stats          : 3 scanned, 2 copied, 3 indexed
Speed          : 3.44 KiB in (650/s), 80.2 KiB out (14.8 KiB/s)
Total Time    : 5s.
STATUS        : PASSED
[root@n138 ~]# su - tester -c 'ls -l
/xcpsrc_vol/mohankarthiknfs_dest'
total 0
-rw-r--r-- 1 tester supergroup 14 Nov 16 17:01 a_hdfs.txt
-rw-r--r-- 1 tester supergroup 16 Nov 16 17:01 b_hdfs.txt
[root@n138 ~]# su - tester -c 'ls -ld
/xcpsrc_vol/mohankarthiknfs_dest'
drwxr-xr-x 2 tester supergroup 4096 Nov 16 17:01
/xcpsrc_vol/mohankarthiknfs_dest
[root@n138 ~]#

```

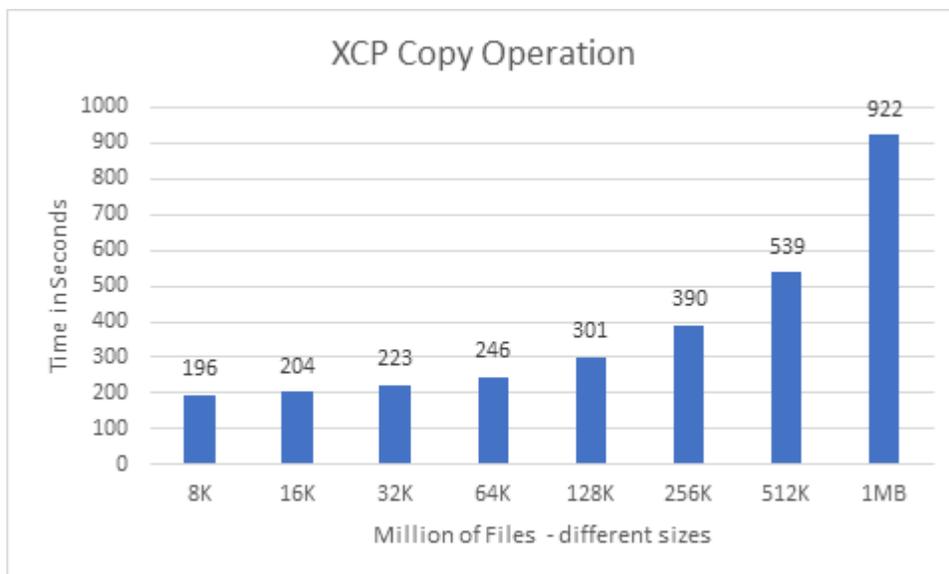
Sizing guidelines

This section provides the approximate time to perform the XCP copy and XCP sync operations with a different file size of one million files for NFS.

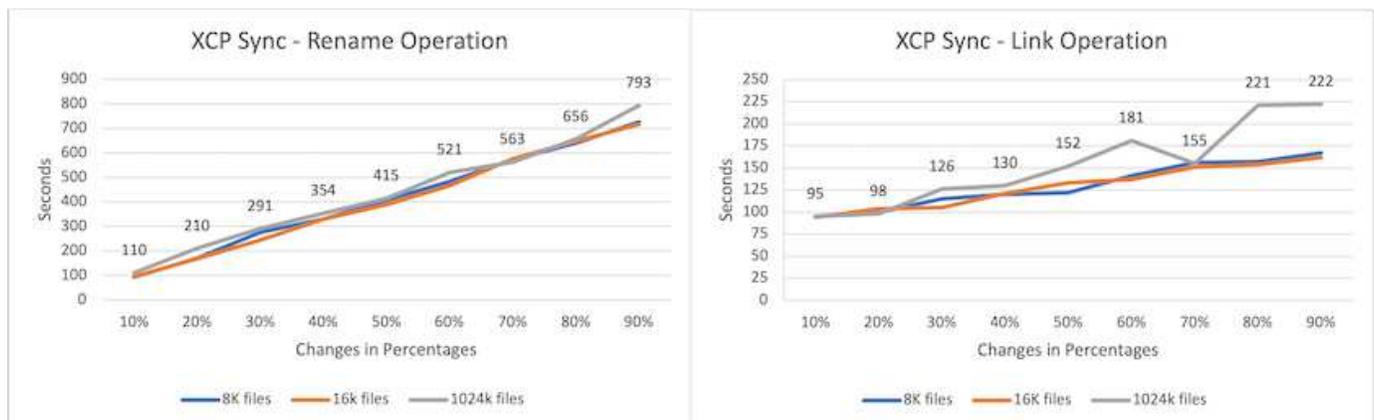
Time estimate based on testing

The tests for the XCP copy and sync operations used the same test bed that was used for deployment. One million files of three sets of 8K, 16K, and 1MB files were created and the changes were performed in real time. The XCP sync function performed the differential incremental updates from the source to the target at the file level. The incremental update operation is one or more of these four operations: rename existing files and folders, append data to existing files, delete files and folders, and include additional hard, soft, and multilinks. For test purposes, we focused on the rename, append, delete, and links operations. In other words, the modification operations such as rename, append, and delete were performed at a change rate of 10% to 90% on one million files.

The following figure shows the results of the XCP copy operation.



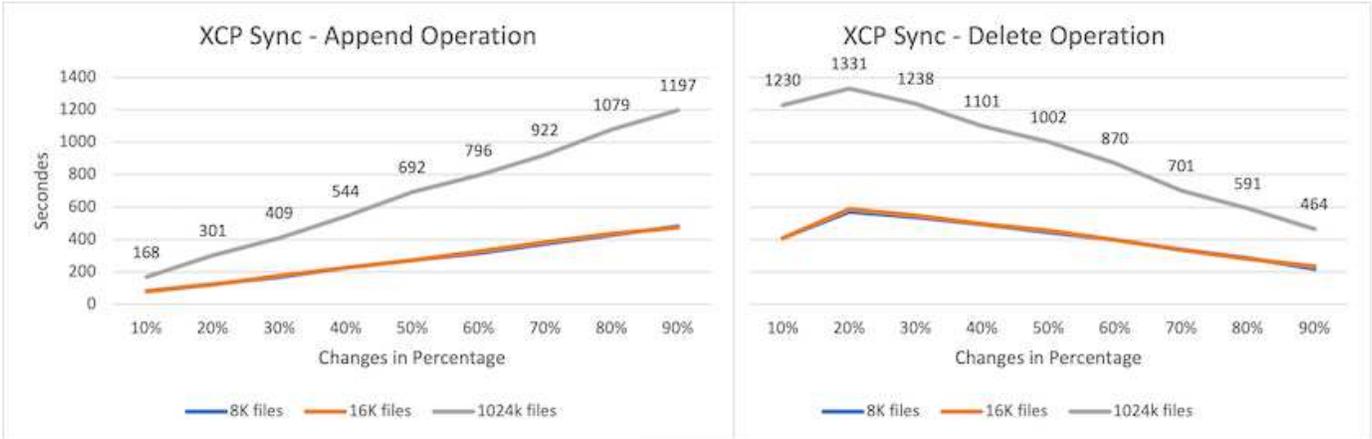
The following figure shows the results of the XCP Sync rename and link operations.



The file size is not proportional to the xcp sync completion time for transferring the renamed source files; the graphs are linear.

The link types are soft links, hard links, and multi-links. Soft links are considered normal files. The size of the files is not relevant for the time to complete the XCP sync operation.

The following figures show the results of the XCP sync append and delete operations.

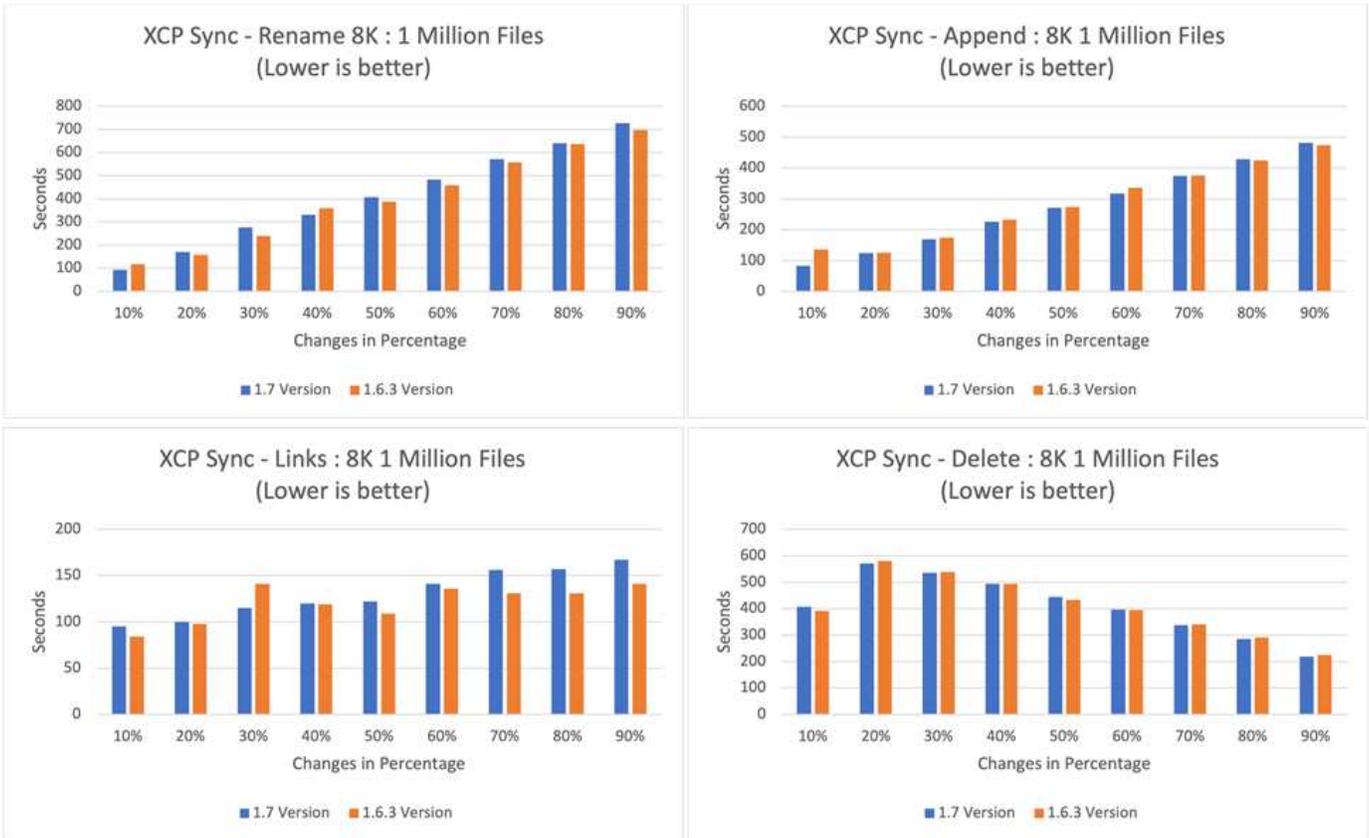


For the append and delete operations, large file sizes take more time compared to small file sizes. The time to complete the operation is linear to the percentage of append and delete changes.

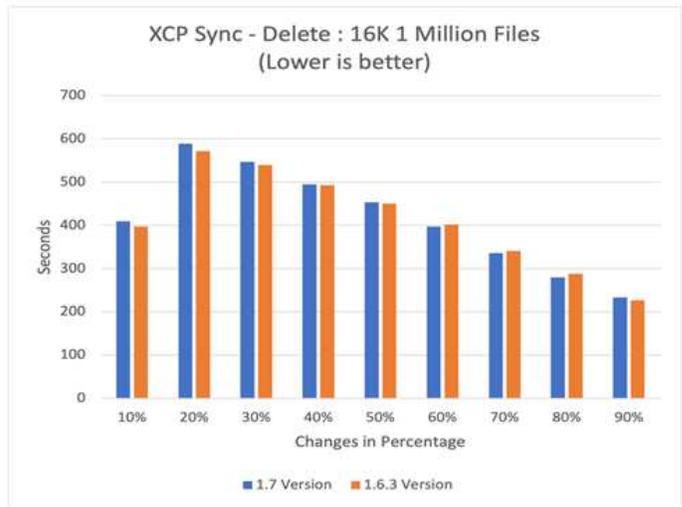
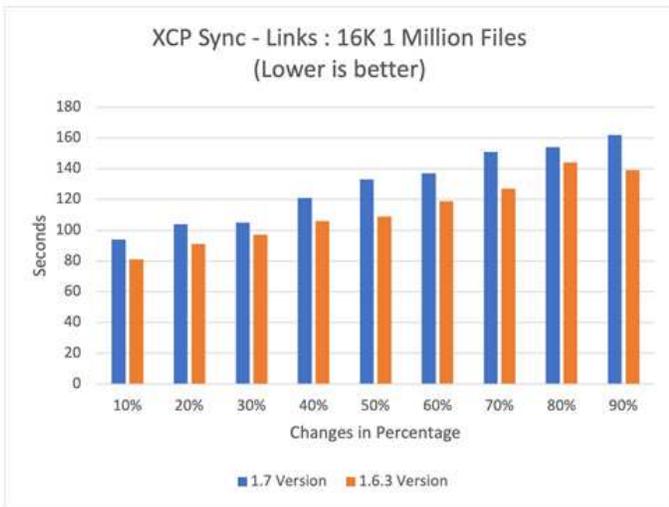
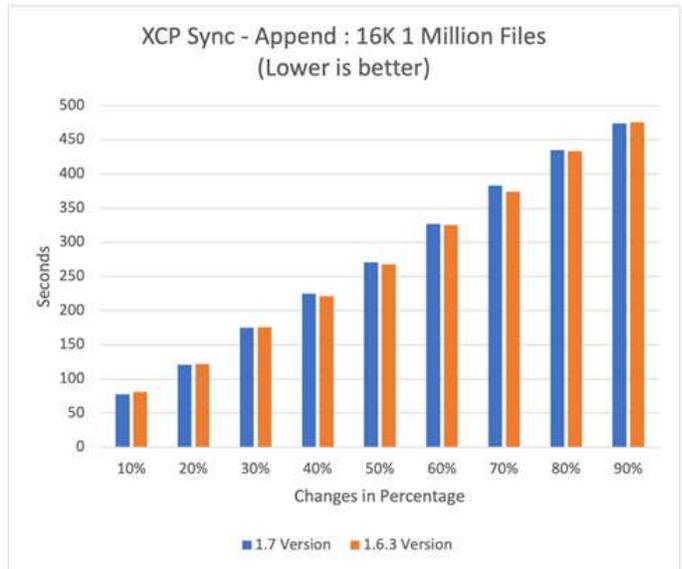
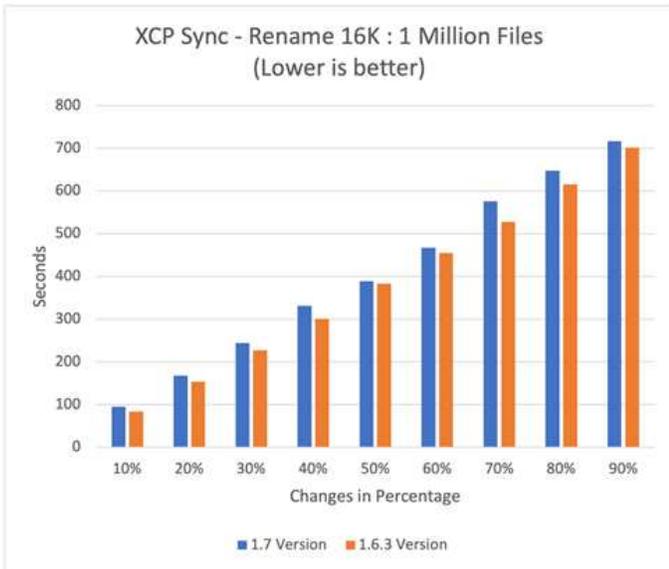
Comparing XCP 1.6.1 to XCP 1.5

Compared to previous versions, XCP 1.6.3 and 1.7 provides improved performance. The following section shows a sync performance comparison between XCP 1.6.3 and 1.7 for 8K, 16K, and 1MB sizes of one million files.

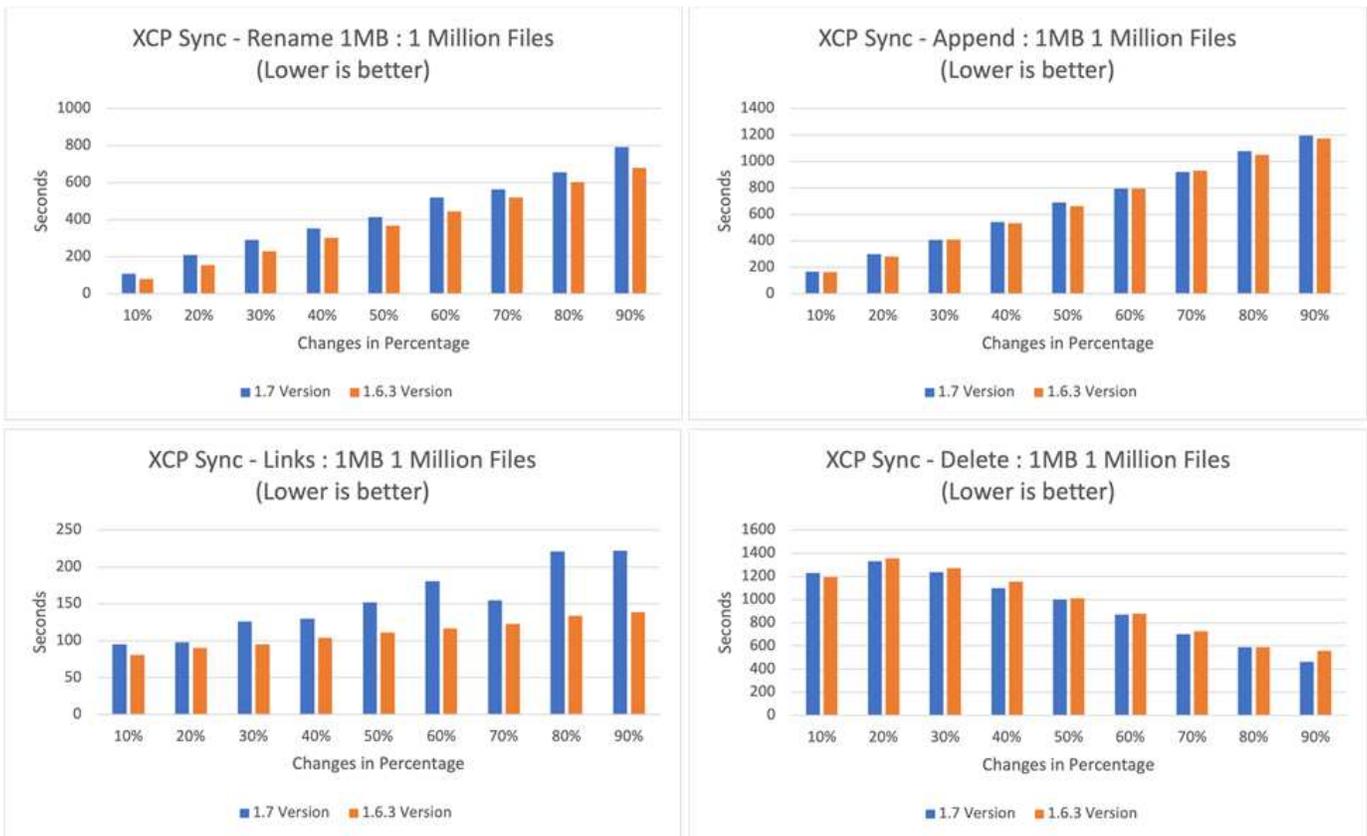
The following figures shows the results of the XCP sync performance for XCP 1.6.3 versus 1.7 (with an 8K size of one million files).



The following figure shows the results of the XCP sync performance for XCP 1.6.1 versus 1.5 (with a 16K size of one million files).



The following figure shows the results of the XCP sync performance for XCP 1.6.1 versus 1.5 with a 1MB size of one million files.



On average, the XCP 1.7 performance improved on or was similar to XCP 1.6.3 for the `xcp_sync` differential incremental update—rename, append, link, and delete operations with a 1MB size of one million files.

Based on this performance validation, NetApp recommends using XCP 1.7 for your data migration on-premises and in the cloud.

Performance tuning

This section provides some of the tuning parameters that help to improve the performance of XCP operations:

- For better scaling and to distribute the workload across multiple XCP instances, split the subfolders for each XCP instance for the migration and data transfer.
- XCP can use maximum CPU resources—the more the CPU cores, the better the performance. Therefore, you should have more CPUs in the XCP server. We lab tested 128GB RAM and 48x core CPUs, which provided better performance than 8x CPUs and 8GB RAM.
- XCP copy with the `-parallel` option is based on the number of CPUs. The default number of parallel threads (seven) is sometimes sufficient for most XCP data transfer and migration operations. For XCP Windows by default, the number of parallel processes is equal to the number of CPUs. The maximum number for the `-parallel` option should be less than or equal to the number of cores.
- 10GbE is a good start for data transfer. However, we tested with 25GbE and 100GbE, which provided better data transfer and are recommended for large file-size data transfer.
- For Azure NetApp Files, the performance varies based on the service level. For more information, see the following table, which shows Azure NetApp Files service levels and performance details.

Service level	Standard	Premium	Ultra
Throughput	16MBps/terabyte (TB)	64MBps/TB	128MBps/TB
Workload types	General purpose file shares, email, and web	BMs, databases, and applications	Latency-sensitive applications
Performance explained	Standard performance: 1,000 IOPS per TB (16K I/O) and 16MBps/TB	Premium performance – 4,000 IOPS per TB (16k I/O) and 64MBps/TB	Extreme performance: 8,000 IOPS per TB (16k I/O) and 128MBps/TB

You must choose the right service level based on the throughput and workload types. Most customers start with the Premium level and change the service level based on the workload.

Customer scenarios

Overview

This section describes customer scenarios and their architectures.

Data lake to ONTAP NFS

This use case is based on the largest financial customer proof of concept (CPOC) that we have done. Historically, we used the NetApp In-Place Analytics Module (NIPAM) to move analytics data to NetApp ONTAP AI. However, because of recent enhancements and the improved performance of NetApp XCP as well as the unique NetApp data mover solution approach, we reran the data migration using NetApp XCP.

Customer challenges and requirements

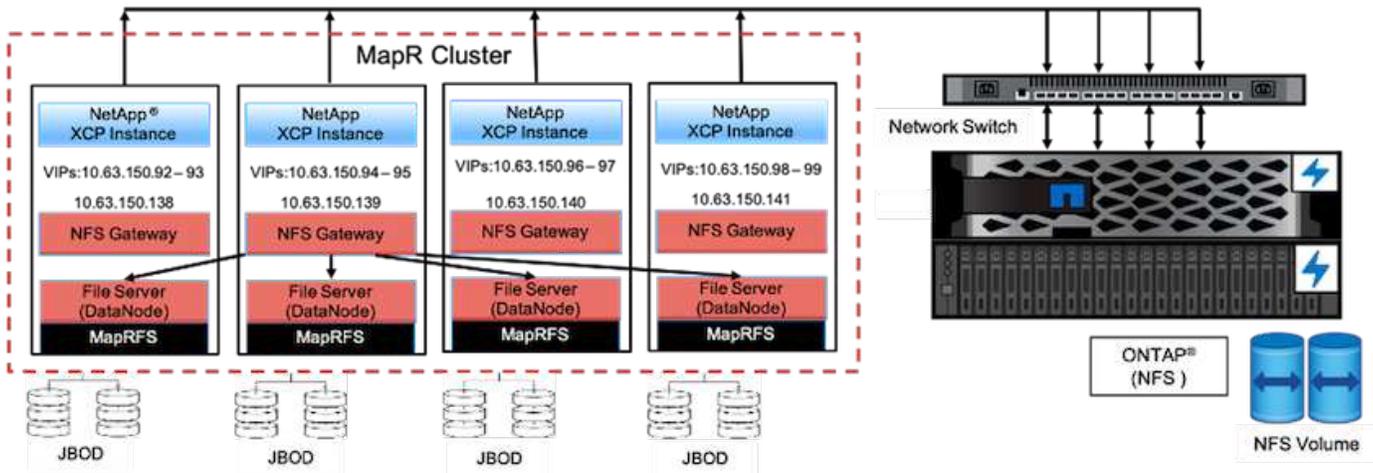
Customer challenges and requirements that are worth noting include the following:

- Customers have different types of data, including structured, unstructured, and semistructured data, logs, and machine-to-machine data in data lakes. AI systems require all these types of data to process for prediction operations. When data is in a data lake-native file system, it is difficult to process.
- The customer's AI architecture is not able to access data from Hadoop Distributed File System (HDFS) and Hadoop Compatible File System (HCFS), so the data is not available to AI operations. AI requires data in an understandable file system format such as NFS.
- Some special processes are required to move data from the data lake because of the large amount of data and high-throughput, and a cost-effective method is required to move the data to the AI system.

Data mover solution

In this solution, the MapR File System (MapR-FS) is created from local disks in the MapR cluster. The MapR NFS Gateway is configured on each data node with virtual IPs. The file server service stores and manages the MapR-FS data. NFS Gateway makes Map-FS data accessible from the NFS client through the virtual IP. An XCP instance is running on each MapR data node to transfer the data from the Map NFS Gateway to NetApp ONTAP NFS. Each XCP instance transfers a specific set of source folders to the destination location.

The following figure illustrates the NetApp data mover solution for MapR cluster using XCP.



For detailed customer use cases, recorded demos, and test results, see the [Using XCP to Move Data from a Data Lake and High-Performance Computing to ONTAP NFS](#) blog.

For detailed steps on moving MapR-FS data into ONTAP NFS by using NetApp XCP, see Appendix B in [TR-4732: Big Data Analytics Data to Artificial Intelligence](#).

High-performance computing to ONTAP NFS

This use case is based on requests from field organizations. Some NetApp customers have their data in a high-performance computing environment, which provides data analytics for training models and enables research organizations to gain insight and understanding of large amount of digital data. NetApp field engineers need a detailed procedure to extract the data from IBM’s GPFS to NFS. We used NetApp XCP to migrate the data from GPFS to NFS so that GPUs can process the data. AI typically processes data from a network file system.

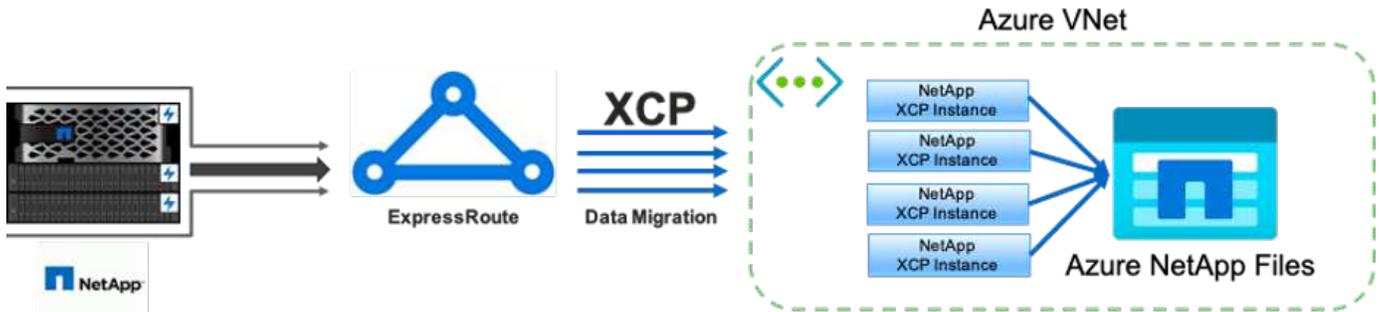
For more information about the high-performance computing to ONTAP NFS use case, a recorded demo, and test results, see the [Using XCP to Move Data from a Data Lake and High-Performance Computing to ONTAP NFS](#) blog.

For detailed steps on moving MapR-FS data into ONTAP NFS by using NetApp XCP, see Appendix A: GPFS to NFS—Detailed Steps in <https://docs.netapp.com/us-en/netapp-solutions/data-analytics/bda-ai-introduction.html>.

Using the XCP Data Mover to migrate millions of small files to flexible storage

This use case is based on the largest NetApp tourism industry customer for on-premises-to-cloud data migration. Because COVID-19 has reduced demand in the travel industry, customers want to save capital expenses on high-end storage in their on-premises environment for the demand pricing application. This customer has a tight SLA to migrate millions of small files to the cloud.

The following figure depicts data migration from on-premises to Azure NetApp Files for small files.



For more information, see the [NetApp XCP Data Mover Solution: On Premises to Cloud](#) blog.

Using the XCP Data Mover to migrate large files

This use case is based on a television network customer. The customer wanted to migrate Oracle Recovery Manager (RMAN) backup files to the cloud and run the Oracle E-Business Suite (EBS) application by using Azure NetApp Files with Pacemaker software. The customer also wanted to migrate their database backup files to on-demand cloud storage and transfer large files (in the range of 25GB to 50GB each) to Azure.

The following figure illustrates the data migration from on-premises to Azure NetApp Files for large files.

For more information, see the [NetApp XCP Data Mover Solution: On Premises to Cloud](#) blog.

Duplicate files

NetApp received a request to find duplicate files from a single volume or multiple volumes. NetApp provided the following solution.

For single volume, run the following commands:

```

[root@mastr-51 linux]# ./xcp -md5 -match 'type==f and nlinks==1 and size
!= 0' 10.63.150.213:/common_volume/nfsconnector_hw_cert/ | sort | uniq -cd
--check-chars=32
XCP 1.5; (c) 2020 NetApp, Inc.; Licensed to Calin Salagean [NetApp Inc]
until Mon Dec 31 00:00:00 2029

176,380 scanned, 138,116 matched, 138,115 summed, 10 giants, 61.1 GiB in
(763 MiB/s), 172 MiB out (2.57 MiB/s), 1m5s

Filtered: 38264 did not match
176,380 scanned, 138,116 matched, 138,116 summed, 10 giants, 62.1 GiB in
(918 MiB/s), 174 MiB out (2.51 MiB/s), 1m9s.
    3 00004964ca155eca1a71d0949c82e37e
nfsconnector_hw_cert/grid_01082017_174316/0/hadoopqe/accumulo/shell/pom.xml
1
    2 000103fbed06d8071410c59047738389
nfsconnector_hw_cert/usr_hdp/2.5.3.0-37/hive2/doc/examples/files/dim-
data.txt
    2 000131053a46d67557d27bb678d5d4a1
nfsconnector_hw_cert/grid_01082017_174316/0/log/cluster/mahout_1/artifacts
/classifier/20news_reduceddata/20news-bydate-test/alt.atheism/53265

```

For multiple volumes, run the following commands:

```

[root@mastr-51 linux]# cat multiplevolume_duplicate.sh
#!/usr/bin/bash

#user input
JUNCTION_PATHS='/nc_volume1 /nc_volume2 /nc_volume3 /oplogarchivevolume'
NFS_DATA_LIF='10.63.150.213'

#xcp operation
for i in $JUNCTION_PATHS
do
echo "start - $i" >> /tmp/duplicate_results
/usr/src/xcp/linux/xcp -md5 -match 'type==f and nlinks==1 and size != 0'
${NFS_DATA_LIF}:$i | sort | uniq -cd --check-chars=32 | tee -a
/tmp/duplicate_results
echo "end - $i" >> /tmp/duplicate_results
done

[root@mastr-51 linux]# nohup bash +x multiplevolume_duplicate.sh &
[root@mastr-51 linux]# cat /tmp/duplicate_results

```

Specific date-based scan and copy of data

This solution is based on a customer who needs to copy data based on a specific date. Verify the following details:

Created a file in Y: and checked the scan command to list them.

```
c:\XCP>dir Y:\karthik_test
Volume in drive Y is from
Volume Serial Number is 80F1-E201

Directory of Y:\karthik_test

05/26/2020  02:51 PM    <DIR>          .
05/26/2020  02:50 PM    <DIR>          ..
05/26/2020  02:51 PM                2,295 testfile.txt
                1 File(s)                2,295 bytes
                2 Dir(s)          658,747,392 bytes free
```

```
c:\XCP>
```

```
c:\XCP>xcp scan -match "strfttime(ctime,'%Y-%m-%d')>'2020-05-01'" -fmt
"'{},{}'format(iso(mtime),name)" Y:\
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to Calin Salagean [NetApp
Inc] until Mon Dec 31 00:00:00 2029
```

It appears that you are not running XCP as Administrator. To avoid access issues please run XCP as Administrator.

```
2020-05-26_14:51:13.132465,testfile.txt
2020-05-26_14:51:00.074216,karthik_test
```

```
xcp scan -match strfttime(ctime,'%Y-%m-%d')>'2020-05-01' -fmt
'"{},{}'format(iso(mtime),name) Y:\ : PASSED
30,205 scanned, 2 matched, 0 errors
Total Time : 4s
STATUS : PASSED
```

Copy the files based on date (2020 YearMay month first date) from Y: to Z:

```
c:\XCP>xcp copy -match "strfttime(ctime,'%Y-%m-%d')>'2020-05-01'" Y:
Z:\dest_karthik
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to Calin Salagean [NetApp
Inc] until Mon Dec 31 00:00:00 2029
```

It appears that you are not running XCP as Administrator. To avoid access

```
issues please run XCP as Administrator.
```

```
30,205 scanned, 3 matched, 0 copied, 0 errors, 5s
xcp copy -match strftime(ctime,'%Y-%m-%d')>'2020-05-01' Y: Z:\dest_karthik
: PASSED
30,205 scanned, 3 matched, 2 copied, 0 errors
Total Time : 6s
STATUS : PASSED
```

```
c:\XCP>
```

Check the destination Z:

```
c:\XCP>dir Z:\dest_karthik\karthik_test
Volume in drive Z is to
Volume Serial Number is 80F1-E202
```

Directory of Z:\dest_karthik\karthik_test

```
05/26/2020  02:51 PM    <DIR>          .
05/26/2020  02:50 PM    <DIR>          ..
05/26/2020  02:51 PM                2,295 testfile.txt
                1 File(s)          2,295 bytes
                2 Dir(s)      659,316,736 bytes free
```

```
c:\XCP>
```

Creating a CSV file from SMB/CIFS share

The following command dumps data in the CSV format. You can sum up the size column to get the total size of the data.

```
xcp scan -match "((now-x.atime) / 3600) > 31*day" -fmt "'{ }, { }, { },
{ }'.format(relpath, name, strftime(x.atime, '%y-%m-%d-%H:%M:%S'),
humanize_size(size))" -preserve-atime >file.csv
```

The output should look similar to this example:

```
erase\report_av_fp_cdot_crosstab.csvreport_av_fp_cdot_crosstab.csv20-01-
29-10:26:2449.6MiB
```

To scan up to the depth of three subdirectories and provide the result in sorting order, run the `xcp -du` command and dump the size at each directory level up to the depth of three subdirectories.

```
./xcp scan -du -depth 3 NFS_Server_IP:/source_vol
```

To sort, dump the information to a CSV file and sort the information.

```
xcp scan -match "type == d" -depth 3 -fmt "'{}, {}, {}, {}'.format(name, relpath, size)" NFS_Server_IP:/share > directory_report.csv
```

This is a custom report that uses the `-fmt` command. It scans all the directories and dumps the name of the directory, path, and size of directory into a CSV file. You can sort the size column from the spreadsheet application.

Data migration from 7-Mode to ONTAP

This section provides detailed steps for migrating data from NetApp Data ONTAP operating in 7-Mode to ONTAP.

Transitioning 7-Mode NFSv3 storage to ONTAP for NFS data

This section provides the step-by-step procedure in the following table for transitioning a source 7-Mode NFSv3 export to an ONTAP system.

NetApp assumes that the source 7-Mode NFSv3 volume is exported and mounted on the client system and that XCP is already installed on a Linux system.

1. Verify that the target ONTAP system is healthy.

```

CLUSTER::> cluster show
Node                Health  Eligibility
-----
CLUSTER-01          true   true
CLUSTER-02          true   true
2 entries were displayed.
CLUSTER::> node show
Node      Health Eligibility Uptime           Model      Owner      Location
-----
CLUSTER-01
           true  true           78 days 21:01 FAS8060
           RTP
CLUSTER-02
           true  true           78 days 20:50 FAS8060
           RTP
2 entries were displayed.
CLUSTER::> storage failover show
Node      Partner      Takeover
Possible State Description
-----
CLUSTER-01  CLUSTER-02  true      Connected to CLUSTER-02
CLUSTER-02  CLUSTER-01  true      Connected to CLUSTER-01
2 entries were displayed.

```

2. Verify that at least one nonroot aggregate exists on the target system. The aggregate is normal.

```

CLUSTER::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
aggr0          368.4GB   17.85GB   95% online    1 CLUSTER-01
raid_dp,

normal
aggr0_CLUSTER_02_0
                368.4GB   17.85GB   95% online    1 CLUSTER-02
raid_dp,

normal
source         1.23TB    1.10TB   11% online    6 CLUSTER-01
raid_dp,

normal
3 entries were displayed.

```

If there is no data aggregate, create a new one using the `storage aggr create` command.

3. Create a storage virtual machine (SVM) on the target cluster system.

```

CLUSTER::> vserver create -vserver dest -rootvolume dest_root -aggregate
poc -rootvolume-security-style mixed
[Job 647] Job succeeded:
Vserver creation completed
Verify the security style and language settings of the source

Verify that the SVM was successfully created.
CLUSTER::> vserver show -vserver dest
                                Vserver: dest
                                Vserver Type: data
                                Vserver Subtype: default
                                Vserver UUID: 91f6d786-0063-11e5-b114-
00a09853a969
                                Root Volume: dest_root
                                Aggregate: poc
                                NIS Domain: -
                                Root Volume Security Style: mixed
                                LDAP Client: -
                                Default Volume Language Code: C.UTF-8
                                Snapshot Policy: default
                                Comment:
                                Quota Policy: default
                                List of Aggregates Assigned: -
                                Limit on Maximum Number of Volumes allowed: unlimited
                                Vserver Admin State: running
                                Vserver Operational State: running
                                Vserver Operational State Stopped Reason: -
                                Allowed Protocols: nfs, cifs, fcp, iscsi, ndmp
                                Disallowed Protocols: -
                                Is Vserver with Infinite Volume: false
                                QoS Policy Group: -
                                Config Lock: false
                                IPspace Name: Default

```

4. Remove the FCP, iSCSI, NDMP, and CIDS protocols from the target SVM.

```

CLUSTER::> vserver remove-protocols -vserver dest -protocols
fcp,iscsi,ndmp,cifs

```

Verify that NFS is the allowed protocol for this SVM.

```

CLUSTER::> vserver show -vserver dest -fields allowed-protocols
vserver allowed-protocols
-----
dest      nfs

```

5. Create a new read-write data volume on the destination SVM. Verify that the security style, language settings, and capacity requirements match the source volume.

```

CLUSTER::> vol create -vserver dest -volume dest_nfs -aggregate poc
-size 150g -type RW -state online -security-style mixed
[Job 648] Job succeeded: Successful

```

6. Create a data LIF to serve NFS client requests.

```

CLUSTER::> network interface create -vserver dest -lif dest_lif -address
10.61.73.115 -netmask 255.255.255.0 -role data -data-protocol nfs -home
-node CLUSTER-01 -home-port e0i

```

Verify that the LIF was successfully created.

```

CLUSTER::> network interface show -vserver dest

```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
dest	dest_lif	up/up	10.61.73.113/24	CLUSTER-01	e0i
true					

7. Create a static route with the SVM, if required.

```

CLUSTER::> network route create -vserver dest -destination 0.0.0.0/0
-gateway 192.168.100.111

```

Verify that the route was successfully created.

```

CLUSTER::> network route show -vserver source
Vserver          Destination      Gateway          Metric
-----
dest
                0.0.0.0/0       10.61.73.1      20

```

8. Mount the target NFS data volume in the SVM namespace.

```

CLUSTER::> volume mount -vserver dest -volume dest_nfs -junction-path
/dest_nfs -active true

```

Verify that the volume is successfully mounted.

```

CLUSTER::> volume show -vserver dest -fields junction-path
vserver volume  junction-path
-----
dest    dest_nfs /dest_nfs
dest    dest_root
        /
2 entries were displayed.

```

You can also specify volume mount options (junction path) with the `volume create` command.

9. Start the NFS service on the target SVM.

```

CLUSTER::> vserver nfs start -vserver dest

```

Verify that the service is started and running.

```

CLUSTER::> vserver nfs status
The NFS server is running on Vserver "dest".
CLUSTER::> nfs show
Vserver: dest
    General Access: true
                v3: enabled
                v4.0: disabled
                4.1: disabled
                UDP: enabled
                TCP: enabled
    Default Windows User: -
    Default Windows Group: -

```

10. Verify that the default NFS export policy was applied to the target SVM.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver          Policy Name
-----
dest             default
```

11. If required, create a new custom export policy for the target SVM.

```
CLUSTER::> vserver export-policy create -vserver dest -policyname
xcpexportpolicy
```

Verify that the new custom export policy was successfully created.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver          Policy Name
-----
dest             default
dest             xcpexportpolicy
2 entries were displayed.
```

12. Modify the export policy rules to allow access to NFS clients.

```
CLUSTER::> export-policy rule modify -vserver dest -ruleindex 1
-policyname xcpexportpolicy -clientmatch 0.0.0.0/0 -rorule any -rwrule
any -anon 0
Verify the policy rules have modified
CLUSTER::> export-policy rule show -instance
                                Vserver: dest
                                Policy Name: xcpexportpolicy
                                Rule Index: 1
                                Access Protocol: nfs3
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
                                RO Access Rule: none
                                RW Access Rule: none
User ID To Which Anonymous Users Are Mapped: 65534
                                Superuser Security Types: none
                                Honor SetUID Bits in SETATTR: true
                                Allow Creation of Devices: true
```

13. Verify that the client is allowed access to the volume.

```
CLUSTER::> export-policy check-access -vserver dest -volume dest_nfs
-client-ip 10.61.82.215 -authentication-method none -protocol nfs3
-access-type read-write
```

Path	Policy	Policy Owner	Policy Owner Type	Rule Index
Access				
-----	-----	-----	-----	-----
/	xcpexportpolicy	dest_root	volume	1
read				
/dest_nfs	xcpexportpolicy	dest_nfs	volume	1
read-write				

2 entries were displayed.

14. Connect to the Linux NFS server. Create a mount point for the NFS exported volume.

```
[root@localhost /]# cd /mnt
[root@localhost mnt]# mkdir dest
```

15. Mount the target NFSv3 exported volume at this mount point.



The NFSv3 volumes should be exported but not necessarily mounted by the NFS server. If they can be mounted, the XCP Linux host client mounts these volumes.

```
[root@localhost mnt]# mount -t nfs 10.61.73.115:/dest_nfs /mnt/dest
```

Verify that the mount point was successfully created.

```
[root@localhost /]# mount | grep nfs
10.61.73.115:/dest_nfs on /mnt/dest type nfs
(rw,relatime,vers=3,rsize=65536,wsiz=65536,namlen=255,hard,proto=tcp,tim=
meo=600,retrans=2,sec=sys,mountaddr=10.61.82.215,mountvers=3,mountport=4
046,mountproto=udp,local_lock=none,addr=10.61.73.115)
```

16. Create a test file on the NFS exported mount point to enable read-write access.

```
[root@localhost dest]# touch test.txt
Verify the file is created
[root@localhost dest]# ls -l
total 0
-rw-r--r-- 1 root bin 0 Jun  2 03:16 test.txt
```



After the read-write test is complete, delete the file from the target NFS mount point.

17. Connect to the Linux client system in which XCP is installed. Browse to the XCP install path.

```
[root@localhost ~]# cd /linux/
[root@localhost linux]#
```

18. Query the source 7-Mode NFSv3 exports by running the `xcp show` command on the XCP Linux client host system.

```
[root@localhost]# ./xcp show 10.61.82.215
== NFS Exports ==
Mounts  Errors  Server
      4      0 10.61.82.215
      Space   Files   Space   Files
      Free    Free    Used    Used Export
23.7 GiB 778,134 356 KiB   96 10.61.82.215:/vol/nfsvol1
17.5 GiB 622,463 1.46 GiB 117 10.61.82.215:/vol/nfsvol
328 GiB 10.8M 2.86 GiB 7,904 10.61.82.215:/vol/vol0/home
328 GiB 10.8M 2.86 GiB 7,904 10.61.82.215:/vol/vol0
== Attributes of NFS Exports ==
drwxr-xr-x --- root wheel 4KiB 4KiB 2d21h 10.61.82.215:/vol/nfsvol1
drwxr-xr-x --- root wheel 4KiB 4KiB 2d21h 10.61.82.215:/vol/nfsvol
drwxrwxrwx --t root wheel 4KiB 4KiB 9d22h 10.61.82.215:/vol/vol0/home
drwxr-xr-x --- root wheel 4KiB 4KiB 4d0h 10.61.82.215:/vol/vol0
3.89 KiB in (5.70 KiB/s), 7.96 KiB out (11.7 KiB/s), 0s.
```

19. Scan the source NFSv3 exported paths and print the statistics of their file structure.

NetApp recommends putting the source NFSv3 exports in read-only mode during `xcp scan`, `copy`, and `sync` operations.

```

[root@localhost /]# ./xcp scan 10.61.82.215:/vol/nfsvol
nfsvol
nfsvol/n5000-uk9.5.2.1.N1.1.bin
nfsvol/821_q_image.tgz
nfsvol/822RC2_q_image.tgz
nfsvol/NX5010_12_node_RCF_v1.3.txt
nfsvol/n5000-uk9-kickstart.5.2.1.N1.1.bin
nfsvol/NetApp_CN1610_1.1.0.5.stk
nfsvol/glibc-common-2.7-2.x86_64.rpm
nfsvol/glibc-2.7-2.x86_64.rpm
nfsvol/rhel-server-5.6-x86_64-dvd.iso.filepart
nfsvol/xcp
nfsvol/xcp_source
nfsvol/catalog
23 scanned, 7.79 KiB in (5.52 KiB/s), 1.51 KiB out (1.07 KiB/s), 1s.

```

20. Copy the source 7-Mode NFSv3 exports to NFSv3 exports on the target ONTAP system.

```

[root@localhost /]# ./xcp copy 10.61.82.215:/vol/nfsvol
10.61.73.115:/dest_nfs
 44 scanned, 39 copied, 264 MiB in (51.9 MiB/s), 262 MiB out (51.5
MiB/s), 5s
 44 scanned, 39 copied, 481 MiB in (43.3 MiB/s), 479 MiB out (43.4
MiB/s), 10s
 44 scanned, 40 copied, 748 MiB in (51.2 MiB/s), 747 MiB out (51.3
MiB/s), 16s
 44 scanned, 40 copied, 1.00 GiB in (55.9 MiB/s), 1.00 GiB out (55.9
MiB/s), 21s
 44 scanned, 40 copied, 1.21 GiB in (42.8 MiB/s), 1.21 GiB out (42.8
MiB/s), 26s
Sending statistics...
44 scanned, 43 copied, 1.46 GiB in (47.6 MiB/s), 1.45 GiB out (47.6
MiB/s), 31s.

```

21. After the copy is finished, verify that the source and destination NFSv3 exports have identical data. Run the `xcp verify` command.

```
[root@localhost /]# ./xcp verify 10.61.82.215:/vol/nfsvol
10.61.73.115:/dest_nfs
44 scanned, 44 found, 28 compared, 27 same data, 2.41 GiB in (98.4
MiB/s), 6.25 MiB out (255 KiB/s), 26s
44 scanned, 44 found, 30 compared, 29 same data, 2.88 GiB in (96.4
MiB/s), 7.46 MiB out (249 KiB/s), 31s
44 scanned, 100% found (43 have data), 43 compared, 100% verified (data,
attrs, mods), 2.90 GiB in (92.6 MiB/s), 7.53 MiB out (240 KiB/s), 32s.
```

If `xcp verify` finds differences between the source and destination data, then the error `no such file or directory` is reported in the summary. To fix that issue, run the `xcp sync` command to copy the source changes to the destination.

22. Before and during the cutover, run `verify` again. If the source has new or updated data, then perform incremental updates. Run the `xcp sync` command.

```
For this operation, the previous copy index name or number is required.
[root@localhost /]# ./xcp sync -id 3
Index: {source: '10.61.82.215:/vol/nfsvol', target:
'10.61.73.115:/dest_nfs1'}
64 reviewed, 64 checked at source, 6 changes, 6 modifications, 51.7 KiB
in (62.5 KiB/s), 22.7 KiB out (27.5 KiB/s), 0s.
xcp: sync '3': Starting search pass for 1 modified directory...
xcp: sync '3': Found 6 indexed files in the 1 changed directory
xcp: sync '3': Rereading the 1 modified directory to find what's new...
xcp: sync '3': Deep scanning the 1 directory that changed...
11 scanned, 11 copied, 12.6KiB in (6.19KiBps), 9.50 KiB out (4.66KiBps),
2s.
```

23. To resume a previously interrupted copy operation, run the `xcp resume` command.

```

[root@localhost /]# ./xcp resume -id 4
Index: {source: '10.61.82.215:/vol/nfsvol', target:
'10.61.73.115:/dest_nfs7'}
xcp: resume '4': WARNING: Incomplete index.
xcp: resume '4': Found 18 completed directories and 1 in progress
106 reviewed, 24.2 KiB in (30.3 KiB/s), 7.23 KiB out (9.06 KiB/s), 0s.
xcp: resume '4': Starting second pass for the in-progress directory...
xcp: resume '4': Found 3 indexed directories and 0 indexed files in the
1 in-progress directory
xcp: resume '4': In progress dirs: unindexed 1, indexed 0
xcp: resume '4': Resuming the 1 in-progress directory...
 20 scanned, 7 copied, 205 MiB in (39.6 MiB/s), 205 MiB out (39.6
MiB/s), 5s
 20 scanned, 14 copied, 425 MiB in (42.1 MiB/s), 423 MiB out (41.8
MiB/s), 11s
 20 scanned, 14 copied, 540 MiB in (23.0 MiB/s), 538 MiB out (23.0
MiB/s), 16s
 20 scanned, 14 copied, 721 MiB in (35.6 MiB/s), 720 MiB out (35.6
MiB/s), 21s
 20 scanned, 15 copied, 835 MiB in (22.7 MiB/s), 833 MiB out (22.7
MiB/s), 26s
 20 scanned, 16 copied, 1007 MiB in (34.3 MiB/s), 1005 MiB out (34.3
MiB/s), 31s
 20 scanned, 17 copied, 1.15 GiB in (33.9 MiB/s), 1.15 GiB out (33.9
MiB/s), 36s
 20 scanned, 17 copied, 1.27 GiB in (25.5 MiB/s), 1.27 GiB out (25.5
MiB/s), 41s
 20 scanned, 17 copied, 1.45 GiB in (36.1 MiB/s), 1.45 GiB out (36.1
MiB/s), 46s
 20 scanned, 17 copied, 1.69 GiB in (48.7 MiB/s), 1.69 GiB out (48.7
MiB/s), 51s
Sending statistics...
20 scanned, 20 copied, 21 indexed, 1.77 GiB in (33.5 MiB/s), 1.77 GiB
out (33.4 MiB/s), 54s.

```

After `resume` finishes copying files, run `verify` again so that the source and destination storage have identical data.

24. The NFSv3 client host needs to unmount the source NFSv3 exports provisioned from the 7-Mode storage and mounts the target NFSv3 exports from ONTAP. Cutover requires an outage.

Transitioning 7-Mode volume Snapshot copies to ONTAP

This section covers the procedure for transitioning a source 7-Mode volume NetApp Snapshot copy to ONTAP.



NetApp assumes that the source 7-Mode volume is exported and mounted on the client system and that XCP is already installed on a Linux system. A Snapshot copy is a point-in-time image of a volume that records incremental changes since the last Snapshot copy. Use the `-snap` option with a 7-Mode system as the source.

Warning: Keep the base Snapshot copy. Do not delete the base Snapshot copy after the baseline copy is complete. The base Snapshot copy is required for further sync operations.

1. Verify that the target ONTAP system is healthy.

```
CLUSTER::> cluster show
Node                Health  Eligibility
-----
CLUSTER-01          true   true
CLUSTER-02          true   true
2 entries were displayed.
CLUSTER::> node show
Node      Health Eligibility Uptime      Model      Owner      Location
-----
CLUSTER-01
           true  true      78 days 21:01 FAS8060
           CLUSTER-02
           true  true      78 days 20:50 FAS8060
           CLUSTER-01
           true  true      78 days 20:50 FAS8060
2 entries were displayed.
CLUSTER::> storage failover show
Node      Partner      Takeover
Possible State Description
-----
CLUSTER-01  CLUSTER-02  true      Connected to CLUSTER-02
CLUSTER-02  CLUSTER-01  true      Connected to CLUSTER-01
2 entries were displayed.
```

2. Verify that at least one nonroot aggregate exists on the target system. The aggregate is normal.

```

CLUSTER::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
aggr0          368.4GB   17.85GB   95% online    1 CLUSTER-01
raid_dp,

normal
aggr0_CLUSTER_02_0
                368.4GB   17.85GB   95% online    1 CLUSTER-02
raid_dp,

normal
source         1.23TB    1.10TB   11% online    6 CLUSTER-01
raid_dp,

normal
3 entries were displayed.

```

If there is no data aggregate, create a new one using the `storage aggr create` command.

3. Create an SVM on the target cluster system.

```

CLUSTER::> vserver create -vserver dest -rootvolume dest_root -aggregate
poc -rootvolume-security-style mixed
[Job 647] Job succeeded:
Vserver creation completed
Verify the security style and language settings of the source

Verify that the SVM was successfully created.
CLUSTER::> vserver show -vserver dest
                                Vserver: dest
                                Vserver Type: data
                                Vserver Subtype: default
                                Vserver UUID: 91f6d786-0063-11e5-b114-
00a09853a969
                                Root Volume: dest_root
                                Aggregate: poc
                                NIS Domain: -
                                Root Volume Security Style: mixed
                                LDAP Client: -
                                Default Volume Language Code: C.UTF-8
                                Snapshot Policy: default
                                Comment:
                                Quota Policy: default
                                List of Aggregates Assigned: -
                                Limit on Maximum Number of Volumes allowed: unlimited
                                Vserver Admin State: running
                                Vserver Operational State: running
                                Vserver Operational State Stopped Reason: -
                                Allowed Protocols: nfs, cifs, fcp, iscsi, ndmp
                                Disallowed Protocols: -
                                Is Vserver with Infinite Volume: false
                                QoS Policy Group: -
                                Config Lock: false
                                IPspace Name: Default

```

4. Remove the FCP, iSCSI, NDMP, and CIFS protocols from the target SVM.

```

CLUSTER::> vserver remove-protocols -vserver dest -protocols
fcp,iscsi,ndmp,cifs
Verify that NFS is the allowed protocol for this SVM.
CLUSTER::> vserver show -vserver dest -fields allowed-protocols
vserver allowed-protocols
-----
dest      nfs

```

5. Create a new read-write data volume on the destination SVM. Verify that the security style, language settings, and capacity requirements match the source volume.

```
CLUSTER::> vol create -vserver dest -volume dest_nfs -aggregate poc
-size 150g -type RW -state online -security-style mixed
[Job 648] Job succeeded: Successful
```

6. Create a data LIF to serve NFS client requests.

```
CLUSTER::> network interface create -vserver dest -lif dest_lif -address
10.61.73.115 -netmask 255.255.255.0 -role data -data-protocol nfs -home
-node CLUSTER-01 -home-port e01
```

Verify that the LIF was successfully created.

```
CLUSTER::> network interface show -vserver dest
```

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
dest	dest_lif	up/up	10.61.73.113/24	CLUSTER-01	e0i
true					

7. If required, create a static route with the SVM.

```
CLUSTER::> network route create -vserver dest -destination 0.0.0.0/0
-gateway 192.168.100.111
```

Verify that the route was successfully created.

```
CLUSTER::> network route show -vserver source
```

Vserver	Destination	Gateway	Metric
dest	0.0.0.0/0	10.61.73.1	20

8. Mount the target NFS data volume in the SVM namespace.

```
CLUSTER::> volume mount -vserver dest -volume dest_nfs -junction-path
/dest_nfs -active true
```

Verify that the volume was successfully mounted.

```
CLUSTER::> volume show -vserver dest -fields junction-path
vserver volume    junction-path
-----
dest    dest_nfs /dest_nfs
dest    dest_root
          /
2 entries were displayed.
```

You can also specify the volume mount options (junction path) with the `volume create` command.

9. Start the NFS service on the target SVM.

```
CLUSTER::> vserver nfs start -vserver dest
```

Verify that the service is started and running.

```
CLUSTER::> vserver nfs status
The NFS server is running on Vserver "dest".
CLUSTER::> nfs show
Vserver: dest
      General Access: true
                v3: enabled
                v4.0: disabled
                4.1: disabled
                UDP: enabled
                TCP: enabled
      Default Windows User: -
      Default Windows Group: -
```

10. Verify that the default NFS export policy is applied to the target SVM.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver          Policy Name
-----
dest             default
```

11. If required, create a new custom export policy for the target SVM.

```
CLUSTER::> vserver export-policy create -vserver dest -policyname
xcpexportpolicy
```

Verify that the new custom export policy was successfully created.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver          Policy Name
-----
dest             default
dest             xcpexportpolicy
2 entries were displayed.
```

12. Modify the export policy rules to allow access to NFS clients on the target system.

```
CLUSTER::> export-policy rule modify -vserver dest -ruleindex 1
-policyname xcpexportpolicy -clientmatch 0.0.0.0/0 -rorule any -rwrule
any -anon 0
Verify the policy rules have modified
CLUSTER::> export-policy rule show -instance
                Vserver: dest
                Policy Name: xcpexportpolicy
                Rule Index: 1
                Access Protocol: nfs3
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
                RO Access Rule: none
                RW Access Rule: none
User ID To Which Anonymous Users Are Mapped: 65534
                Superuser Security Types: none
                Honor SetUID Bits in SETATTR: true
                Allow Creation of Devices: true
```

13. Verify that the client has access to the target volume.

```
CLUSTER::> export-policy check-access -vserver dest -volume dest_nfs
-client-ip 10.61.82.215 -authentication-method none -protocol nfs3
-access-type read-write
```

Path	Policy	Policy Owner	Policy Owner Type	Rule Index
Access				
-----	-----	-----	-----	-----
/	xcpexportpolicy	dest_root	volume	1
read				
/dest_nfs	xcpexportpolicy	dest_nfs	volume	1
read-write				

2 entries were displayed.

14. Connect to the Linux NFS server. Create a mount point for the NFS exported volume.

```
[root@localhost /]# cd /mnt
[root@localhost mnt]# mkdir dest
```

15. Mount the target NFSv3 exported volume at this mount point.



The NFSv3 volumes should be exported but not necessarily mounted by the NFS server. If they can be mounted, the XCP Linux host client mounts these volumes.

```
[root@localhost mnt]# mount -t nfs 10.61.73.115:/dest_nfs /mnt/dest
```

Verify that the mount point was successfully created.

```
[root@localhost /]# mount | grep nfs
10.61.73.115:/dest_nfs on /mnt/dest type nfs
```

16. Create a test file on the NFS exported mount point to enable read-write access.

```
[root@localhost dest]# touch test.txt
Verify the file is created
[root@localhost dest]# ls -l
total 0
-rw-r--r-- 1 root bin 0 Jun  2 03:16 test.txt
```



After the read-write test is complete, delete the file from the target NFS mount point.

17. Connect to the Linux client system in which XCP is installed. Browse to the XCP install path.

```
[root@localhost ~]# cd /linux/  
[root@localhost linux]#
```

18. Query the source 7-Mode NFSv3 exports by running the `xcp show` command on the XCP Linux client host system.

```
[root@localhost]# ./xcp show 10.61.82.215  
== NFS Exports ==  
Mounts  Errors  Server  
      4      0 10.61.82.215  
Space   Files   Space   Files  
Free    Free    Used    Used Export  
23.7 GiB 778,134 356 KiB   96 10.61.82.215:/vol/nfsvol1  
17.5 GiB 622,463 1.46 GiB 117 10.61.82.215:/vol/nfsvol  
328 GiB 10.8M 2.86 GiB 7,904 10.61.82.215:/vol/vol0/home  
328 GiB 10.8M 2.86 GiB 7,904 10.61.82.215:/vol/vol0  
== Attributes of NFS Exports ==  
drwxr-xr-x --- root wheel 4KiB 4KiB 2d21h 10.61.82.215:/vol/nfsvol1  
drwxr-xr-x --- root wheel 4KiB 4KiB 2d21h 10.61.82.215:/vol/nfsvol  
drwxrwxrwx --t root wheel 4KiB 4KiB 9d22h 10.61.82.215:/vol/vol0/home  
drwxr-xr-x --- root wheel 4KiB 4KiB 4d0h 10.61.82.215:/vol/vol0  
3.89 KiB in (5.70 KiB/s), 7.96 KiB out (11.7 KiB/s), 0s.
```

19. Scan the source NFSv3 exported paths and print the statistics of their file structure.

NetApp recommends putting the source NFSv3 exports in read-only mode during `xcp scan`, `copy`, and `sync` operations. In `sync` operation, you must pass the `-snap` option with a corresponding value.

```

[root@localhost /]# ./xcp scan 10.61.82.215:/vol/nfsvol/.snapshot/snap1
nfsvol
nfsvol/n5000-uk9.5.2.1.N1.1.bin
nfsvol/821_q_image.tgz
nfsvol/822RC2_q_image.tgz
nfsvol/NX5010_12_node_RCF_v1.3.txt
nfsvol/n5000-uk9-kickstart.5.2.1.N1.1.bin
nfsvol/catalog
23 scanned, 7.79 KiB in (5.52 KiB/s), 1.51 KiB out (1.07 KiB/s), 1s.
[root@scspr1202780001 vol_acl4]# ./xcp sync -id 7msnap1 -snap
10.236.66.199:/vol/nfsvol/.snapshot/snap10
(show scan and sync)

```

20. Copy the source 7-Mode NFSv3 snapshot (base) to NFSv3 exports on the target ONTAP system.

```

[root@localhost /]# /xcp copy 10.61.82.215:/vol/nfsvol/.snapshot/snap1
10.61.73.115:/dest_nfs
44 scanned, 39 copied, 264 MiB in (51.9 MiB/s), 262 MiB out (51.5
MiB/s), 5s
44 scanned, 39 copied, 481 MiB in (43.3 MiB/s), 479 MiB out (43.4
MiB/s), 10s
44 scanned, 40 copied, 748 MiB in (51.2 MiB/s), 747 MiB out (51.3
MiB/s), 16s
44 scanned, 40 copied, 1.00 GiB in (55.9 MiB/s), 1.00 GiB out (55.9
MiB/s), 21s
44 scanned, 40 copied, 1.21 GiB in (42.8 MiB/s), 1.21 GiB out (42.8
MiB/s), 26s
Sending statistics...
44 scanned, 43 copied, 1.46 GiB in (47.6 MiB/s), 1.45 GiB out (47.6
MiB/s), 31s.

```



Keep this base snapshot for further sync operations.

21. After copy is complete, verify that the source and destination NFSv3 exports have identical data. Run the `xcp verify` command.

```
[root@localhost /]# ./xcp verify 10.61.82.215:/vol/nfsvol
10.61.73.115:/dest_nfs
44 scanned, 44 found, 28 compared, 27 same data, 2.41 GiB in (98.4
MiB/s), 6.25 MiB out (255 KiB/s), 26s
44 scanned, 44 found, 30 compared, 29 same data, 2.88 GiB in (96.4
MiB/s), 7.46 MiB out (249 KiB/s), 31s
44 scanned, 100% found (43 have data), 43 compared, 100% verified (data,
attrs, mods), 2.90 GiB in (92.6 MiB/s), 7.53 MiB out (240 KiB/s), 32s.
```

If `verify` finds differences between the source and destination data, then the error `no such file or directory` is reported in the summary. To fix that issue, run the `xcp sync` command to copy the source changes to the destination.

22. Before and during the cutover, run `verify` again. If the source has new or updated data, then perform incremental updates. If there are incremental changes, create a new Snapshot copy for these changes and pass that snapshot path with the `-snap` option for sync operations.

Run the `xcp sync` command with the `-snap` option and snapshot path.

```
[root@localhost /]# ./xcp sync -id 3
Index: {source: '10.61.82.215:/vol/nfsvol/.snapshot/snap1', target:
'10.61.73.115:/dest_nfs1'}
64 reviewed, 64 checked at source, 6 changes, 6 modifications, 51.7 KiB
in (62.5
KiB/s), 22.7 KiB out (27.5 KiB/s), 0s.
xcp: sync '3': Starting search pass for 1 modified directory...
xcp: sync '3': Found 6 indexed files in the 1 changed directory
xcp: sync '3': Rereading the 1 modified directory to find what's new...
xcp: sync '3': Deep scanning the 1 directory that changed...
11 scanned, 11 copied, 12.6 KiB in (6.19 KiB/s), 9.50 KiB out (4.66
KiB/s), 2s..
```



For this operation, the base snapshot is required.

23. To resume a previously interrupted copy operation, run the `xcp resume` command.

```

[root@scspr1202780001 534h_dest_vol]# ./xcp resume -id 3
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxxxx [NetApp Inc]
until Mon Dec 31 00:00:00 2029
xcp: Index: {source: '10.61.82.215:/vol/nfsvol',/.snapshot/snap1,
target: 10.237.160.55:/dest_vol}
xcp: resume '7msnap_res1': Reviewing the incomplete index...
xcp: diff '7msnap_res1': Found 143 completed directories and 230 in
progress
39,688 reviewed, 1.28 MiB in (1.84 MiB/s), 13.3 KiB out (19.1 KiB/s),
0s.
xcp: resume '7msnap_res1': Starting second pass for the in-progress
directories...
xcp: resume '7msnap_res1': Resuming the in-progress directories...
xcp: resume '7msnap_res1': Resumed command: copy {-newid:
u'7msnap_res1'}
xcp: resume '7msnap_res1': Current options: {-id: '7msnap_res1'}
xcp: resume '7msnap_res1': Merged options: {-id: '7msnap_res1', -newid:
u'7msnap_res1'}
xcp: resume '7msnap_res1': Values marked with a * include operations
before resume
 68,848 scanned*, 54,651 copied*, 39,688 indexed*, 35.6 MiB in (7.04
MiB/s), 28.1 MiB out (5.57 MiB/s), 5s

```

24. The NFSv3 client host must unmount the source NFSv3 exports provisioned from the 7-Mode storage and mount the target NFSv3 exports from ONTAP. This cutover requires an outage.

Migrating ACLv4 from NetApp 7-Mode to a NetApp storage system

This section covers the step-by-step procedure for transitioning a source NFSv4 export to an ONTAP system.



NetApp assumes that the source NFSv4 volume is exported and mounted on the client system and that XCP is already installed on a Linux system. The source should be a NetApp 7-Mode system that support ACLs. ACL migration is supported from NetApp to NetApp only. To copy files with a special character in the name, make sure the source and destination support UTF-8 encoded language.

Prerequisites for migrating a source NFSv4 export to ONTAP

Before you migrate a source NFSv4 export to ONTAP, the following prerequisites must be met:

- The destination system must have NFSv4 configured.
- The NFSv4 source and target must be mounted on the XCP host. Select NFS v4.0 to match the source and target storage and verify that the ACLs are enabled on the source and target system.
- XCP requires the source/target path to be mounted on the XCP host for ACL processing. In the following example, `vol1 (10.63.5.56:/vol1)` is mounted on the `/mnt/vol1` path:

```

[root@localhost ~]# df -h
Filesystem                                Size  Used
Avail Use% Mounted on
10.63.5.56:/vol1                          973M  4.2M
969M   1% /mnt/vol1
[root@localhost ~]# ./xcp scan -l -acl4 10.63.5.56:/vol1/
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Sun Mar 31 00:00:00 2029
drwxr-xr-x --- root root 4KiB 4KiB 23h42m vol1
rw-r--r-- --- root root   4    0 23h42m vol1/DIR1/FILE
drwxr-xr-x --- root root 4KiB 4KiB 23h42m vol1/DIR1/DIR11
drwxr-xr-x --- root root 4KiB 4KiB 23h42m vol1/DIR1
rw-r--r-- --- root root   4    0 23h42m vol1/DIR1/DIR11/FILE
drwxr-xr-x --- root root 4KiB 4KiB 23h42m vol1/DIR1/DIR11/DIR2
rw-r--r-- --- root root   4    0 23h42m vol1/DIR1/DIR11/DIR2/FILE
drwxr-xr-x --- root root 4KiB 4KiB 17m43s vol1/DIR1/DIR11/DIR2/DIR22
8 scanned, 8 getacls, 1 v3perm, 7 acls, 3.80 KiB in (3.86 KiB/s), 1.21 KiB
out (1.23 KiB/s), 0s.

```

Subdirectories options

The two options to work with subdirectories are as follows:

- For XCP to work on a subdirectory (`/vol1/DIR1/DIR11`), mount the complete path (`10.63.5.56:/vol1/DIR1/DIR11`) on the XCP host.

If the complete path is not mounted, XCP reports the following error:

```

[root@localhost ~]# ./xcp scan -l -acl4 10.63.5.56:/vol1/DIR1/DIR11
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Sun Mar 31 00:00:00 2029
xcp: ERROR: For xcp to process ACLs, please mount
10.63.5.56:/vol1/DIR1/DIR11 using the OS nfs4 client.

```

- Use the subdirectory syntax (`mount: subdirectory/mtree/.snapshot`), as shown in the example below:

```
[root@localhost ~]# ./xcp scan -l -acl4 10.63.5.56:/vol1:/DIR1/DIR11
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Sun Mar 31 00:00:00 2029
drwxr-xr-x --- root root 4KiB 4KiB 23h51m DIR11
rw-r--r-- --- root root 4 0 23h51m DIR11/DIR2/FILE
drwxr-xr-x --- root root 4KiB 4KiB 26m9s DIR11/DIR2/DIR22
rw-r--r-- --- root root 4 0 23h51m DIR11/FILE
drwxr-xr-x --- root root 4KiB 4KiB 23h51m DIR11/DIR2
5 scanned, 5 getacls, 5 acls, 2.04 KiB in (3.22 KiB/s), 540 out (850/s),
0s.
```

Complete the following steps to migrate ACLv4 from NetApp 7-Mode to a NetApp storage system.

1. Verify that the target ONTAP system is healthy.

```
CLUSTER::> cluster show
Node           Health Eligibility
-----
CLUSTER-01    true  true
CLUSTER-02    true  true
2 entries were displayed.
CLUSTER::> node show
Node           Health Eligibility Uptime           Model           Owner           Location
-----
CLUSTER-01
           true  true           78 days 21:01 FAS8060           RTP
CLUSTER-02
           true  true           78 days 20:50 FAS8060           RTP
2 entries were displayed.
CLUSTER::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
CLUSTER-01    CLUSTER-02    true    Connected to CLUSTER-02
CLUSTER-02    CLUSTER-01    true    Connected to CLUSTER-01
2 entries were displayed.
```

2. Verify that at least one nonroot aggregate exists on the target system. The aggregate is normal.

```

CLUSTER::> storage aggregate show
Aggregate      Size Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
aggr0          368.4GB   17.85GB   95% online    1 CLUSTER-01
raid_dp,

normal
aggr0_CLUSTER_02_0
              368.4GB   17.85GB   95% online    1 CLUSTER-02
raid_dp,

normal
source         1.23TB    1.10TB   11% online    6 CLUSTER-01
raid_dp,

normal
3 entries were displayed.

```

If there is no data aggregate, create a new one using the `storage aggr create` command.

3. Create an SVM on the target cluster system.

```

CLUSTER::> vservers create -vservers dest -rootvolume dest_root -aggregate
poc -rootvolume-security-style mixed
[Job 647] Job succeeded:
Vserver creation completed
Verify the security style and language settings of the source

```

Verify that the SVM was successfully created.

```

CLUSTER::> vserver show -vserver dest
                Vserver: dest
                Vserver Type: data
                Vserver Subtype: default
                Vserver UUID: 91f6d786-0063-11e5-b114-
00a09853a969
                Root Volume: dest_root
                Aggregate: poc
                NIS Domain: -
                Root Volume Security Style: mixed
                LDAP Client: -
                Default Volume Language Code: C.UTF-8
                Snapshot Policy: default
                Comment:
                Quota Policy: default
                List of Aggregates Assigned: -
                Limit on Maximum Number of Volumes allowed: unlimited
                Vserver Admin State: running
                Vserver Operational State: running
                Vserver Operational State Stopped Reason: -
                Allowed Protocols: nfs, cifs, fcp, iscsi, ndmp
                Disallowed Protocols: -
                Is Vserver with Infinite Volume: false
                QoS Policy Group: -
                Config Lock: false
                IPspace Name: Default

```

4. Remove the FCP, iSCSI, NDMP, and CIFS protocols from the target SVM.

```

CLUSTER::> vserver remove-protocols -vserver dest -protocols
fcp,iscsi,ndmp,cifs

```

Verify that NFS is the allowed protocol for this SVM.

```

CLUSTER::> vserver show -vserver dest -fields allowed-protocols
vserver allowed-protocols
-----
dest      nfs

```

5. Create a new read-write data volume on the destination SVM. Verify that the security style, language settings, and capacity requirements match the source volume.

```
CLUSTER::> vol create -vserver dest -volume dest_nfs -aggregate poc
-size 150g -type RW -state online -security-style mixed
[Job 648] Job succeeded: Successful
```

6. Create a data LIF to serve NFS client requests.

```
CLUSTER::> network interface create -vserver dest -lif dest_lif -address
10.61.73.115 -netmask 255.255.255.0 -role data -data-protocol nfs -home
-node CLUSTER-01 -home-port e0i
```

Verify that the LIF was successfully created.

```
CLUSTER::> network interface show -vserver dest
```

Current Is	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Port
dest	dest_lif	up/up	10.61.73.113/24	CLUSTER-01	e0i

7. If required, create a static route with the SVM.

```
CLUSTER::> network route create -vserver dest -destination 0.0.0.0/0
-gateway 192.168.100.111
```

Verify that the route was successfully created.

```
CLUSTER::> network route show -vserver source
```

Vserver	Destination	Gateway	Metric
dest	0.0.0.0/0	10.61.73.1	20

8. Mount the target NFS data volume in the SVM namespace.

```
CLUSTER::> volume mount -vserver dest -volume dest_nfs -junction-path
/dest_nfs -active true
```

Verify that the volume was successfully mounted.

```
CLUSTER::> volume show -vserver dest -fields junction-path
vserver volume    junction-path
-----
dest    dest_nfs /dest_nfs
dest    dest_root
        /
2 entries were displayed.
```

You can also specify the volume mount options (junction path) with the `volume create` command.

9. Start the NFS service on the target SVM.

```
CLUSTER::> vserver nfs start -vserver dest
```

Verify that the service is started and running.

```
CLUSTER::> vserver nfs status
The NFS server is running on Vserver "dest".
CLUSTER::> nfs show
Vserver: dest
      General Access: true
                v3: enabled
                v4.0: enabled
                4.1: disabled
                UDP: enabled
                TCP: enabled
      Default Windows User: -
      Default Windows Group: -
```

10. Check that the default NFS export policy is applied to the target SVM.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver          Policy Name
-----
dest             default
```

11. If required, create a new custom export policy for the target SVM.

```
CLUSTER::> vserver export-policy create -vserver dest -policyname
xcpexportpolicy
```

Verify that the new custom export policy was successfully created.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver          Policy Name
-----
dest             default
dest             xcpexportpolicy
2 entries were displayed.
```

12. Modify the export policy rules to allow access to NFS clients.

```
CLUSTER::> export-policy rule modify -vserver dest -ruleindex 1
-policyname xcpexportpolicy -clientmatch 0.0.0.0/0 -rorule any -rwrule
any -anon 0
```

Verify that the policy rules have been modified.

```
CLUSTER::> export-policy rule show -instance
Vserver: dest
Policy Name: xcpexportpolicy
Rule Index: 1
Access Protocol: nfs3
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
RO Access Rule: none
RW Access Rule: none
User ID To Which Anonymous Users Are Mapped: 65534
Superuser Security Types: none
Honor SetUID Bits in SETATTR: true
Allow Creation of Devices: true
```

13. Verify that the client is allowed access to the volume.

```
CLUSTER::> export-policy check-access -vserver dest -volume dest_nfs
-client-ip 10.61.82.215 -authentication-method none -protocol nfs3
-access-type read-write
```

Path	Policy	Policy Owner	Policy Owner Type	Rule Index
Access				
-----	-----	-----	-----	-----
/	xcpexportpolicy	dest_root	volume	1
read				
/dest_nfs	xcpexportpolicy	dest_nfs	volume	1
read-write				

2 entries were displayed.

14. Connect to the Linux NFS server. Create a mount point for the NFS exported volume.

```
[root@localhost /]# cd /mnt
[root@localhost mnt]# mkdir dest
```

15. Mount the target NFSv4 exported volume at this mount point.



The NFSv4 volumes should be exported but not necessarily mounted by the NFS server. If they can be mounted, the XCP Linux host client mounts these volumes.

```
[root@localhost mnt]# mount -t nfs4 10.63.5.56:/vol1 /mnt/vol1
```

Verify that the mount point was successfully created.

```
[root@localhost mnt]# mount | grep nfs
10.63.5.56:/vol1 on /mnt/vol1 type nfs4
(rw,relatime,vers=4.0,rsize=65536,wsiz=65536,namlen=255,hard,proto=tcp,
timeo=600,
retrans=2,sec=sys,clientaddr=10.234.152.84,local_lock=none,addr=10.63.5.
56)
```

16. Create a test file on the NFS exported mount point to enable read-write access.

```
[root@localhost dest]# touch test.txt
```

Verify the file is created.

```
[root@localhost dest]# ls -l
total 0
-rw-r--r-- 1 root bin 0 Jun  2 03:16 test.txt
```



After the read-write test is complete, delete the file from the target NFS mount point.

17. Connect to the Linux client system in which XCP is installed. Browse to the XCP install path.

```
[root@localhost ~]# cd /linux/
[root@localhost linux]#
```

18. Query the source NFSv4 exports by running the `xcp show` command on the XCP Linux client host system.

```

root@localhost]# ./xcp show 10.63.5.56
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxx [NetApp Inc] until
Mon Dec 31 00:00:00 2029
getting pmap dump from 10.63.5.56 port 111...
getting export list from 10.63.5.56...
sending 6 mounts and 24 nfs requests to 10.63.5.56...
== RPC Services ==
'10.63.5.56': UDP rpc services: MNT v1/2/3, NFS v3, NLM v4, PMAP v2/3/4,
STATUS v1
'10.63.5.56': TCP rpc services: MNT v1/2/3, NFS v3/4, NLM v4, PMAP
v2/3/4, STATUS v1
== NFS Exports ==
Mounts  Errors  Server
      6      0 10.63.5.56
      Space   Files      Space   Files
      Free    Free      Used    Used Export
94.7 MiB  19,883   324 KiB   107 10.63.5.56:/
971 MiB   31,023   2.19 MiB    99 10.63.5.56:/vol2
970 MiB   31,024   2.83 MiB    98 10.63.5.56:/vol1
9.33 GiB  310,697   172 MiB    590 10.63.5.56:/vol_005
43.3 GiB   1.10M   4.17 GiB   1.00M 10.63.5.56:/vol3
36.4 GiB   1.10M  11.1 GiB   1.00M 10.63.5.56:/vol4
== Attributes of NFS Exports ==
drwxr-xr-x --- root root 4KiB 4KiB 6d2h 10.63.5.56:/
drwxr-xr-x --- root root 4KiB 4KiB 3d2h 10.63.5.56:/vol2
drwxr-xr-x --- root root 4KiB 4KiB 3d2h 10.63.5.56:/vol1
drwxr-xr-x --- root root 4KiB 4KiB 9d2h 10.63.5.56:/vol_005
drwxr-xr-x --- root root 4KiB 4KiB 9d4h 10.63.5.56:/vol3
drwxr-xr-x --- root root 4KiB 4KiB 9d4h 10.63.5.56:/vol4
6.09 KiB in (9.19 KiB/s), 12.2 KiB out (18.3 KiB/s), 0s.

```

19. Scan the source NFSv4 exported paths and print the statistics of their file structure.

NetApp recommends putting the source NFSv4 exports in read-only mode during `xcp scan`, `copy`, and `sync` operations.

```

[root@localhost]# ./xcp scan -acl4 10.63.5.56:/vol1
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxx [NetApp Inc] until
Mon Dec 31 00:00:00 2029
vol1
vol1/test/f1
vol1/test
3 scanned, 3 getacls, 3 v3perms, 1.59 KiB in (1.72 KiB/s), 696 out
(753/s), 0s.

```

20. Copy source NFSv4 exports to NFSv4 exports on the target ONTAP system.

```
[root@localhost]# ./xcp copy -acl4 -newid id1 10.63.5.56:/vol1
10.63.5.56:/vol2
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxx [NetApp Inc] until
Mon Dec 31 00:00:00 2029
3 scanned, 2 copied, 3 indexed, 3 getacls, 3 v3perms, 1 setacl, 14.7 KiB
in (11.7 KiB/s), 61 KiB out (48.4 KiB/s), 1s..
```

21. After copy is complete, verify that the source and destination NFSv4 exports have identical data. Run the `xcp verify` command.

```
[root@localhost]# ./xcp verify -acl4 -noid 10.63.5.56:/vol1
10.63.5.56:/vol2
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxx [NetApp Inc] until
Mon Dec 31 00:00:00 2029
3 scanned, 100% found (0 have data), 100% verified (data, attrs, mods,
acls), 6 getacls, 6 v3perms, 2.90 KiB in (4.16 KiB/s), 2.94 KiB out
(4.22 KiB/s), 0s.
```

If `verify` finds differences between the source and destination data, then the error `no such file or directory` is reported in the summary. To fix that issue, run the `xcp sync` command to copy the source changes to the destination.

22. Before and during the cutover, run `verify` again. If the source has new or updated data, then perform incremental updates. Run the `xcp sync` command.

```
[root@ root@localhost]# ./xcp sync -id id1
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxx [NetApp Inc] until
Mon Dec 31 00:00:00 2029
xcp: Index: {source: 10.63.5.56:/vol1, target: 10.63.5.56:/vol2}
3 reviewed, 3 checked at source, no changes, 3 reindexed, 25.6 KiB in
(32.3 KiB/s), 23.3 KiB out (29.5 KiB/s), 0s.
```



For this operation, the previous copy index name or number is required.

23. To resume a previously interrupted `copy` operation, run the `xcp resume` command.

```

[root@localhost]# ./xcp resume -id id1
XCP <version>; (c) 2020 NetApp, Inc.; Licensed to xxx [NetApp Inc] until
Mon Dec 31 00:00:00 2029
xcp: Index: {source: 10.63.5.56:/vol3, target: 10.63.5.56:/vol4}
xcp: resume 'id1': Reviewing the incomplete index...
xcp: diff 'id1': Found 0 completed directories and 8 in progress
39,899 reviewed, 1.64 MiB in (1.03 MiB/s), 14.6 KiB out (9.23 KiB/s),
ls.
xcp: resume 'id1': Starting second pass for the in-progress
directories...
xcp: resume 'id1': Resuming the in-progress directories...
xcp: resume 'id1': Resumed command: copy {-acl4: True}
xcp: resume 'id1': Current options: {-id: 'id1'}
xcp: resume 'id1': Merged options: {-acl4: True, -id: 'id1'}
xcp: resume 'id1': Values marked with a * include operations before
resume
 86,404 scanned, 39,912 copied, 39,899 indexed, 13.0 MiB in (2.60
MiB/s), 78.4 KiB out (15.6 KiB/s), 5s 86,404 scanned, 39,912 copied,
39,899 indexed, 13.0 MiB in (0/s), 78.4 KiB out (0/s), 10s
1.00M scanned, 100% found (1M have data), 1M compared, 100% verified
(data, attrs, mods, acls), 2.00M getacls, 202 v3perms, 1.00M same acls,
2.56 GiB in (2.76 MiB/s), 485 MiB out (524 KiB/s), 15m48s.

```

After `resume` finishes copying files, run `verify` again so that the source and destination storage have identical data.

Transitioning 7-Mode SMB storage to ONTAP for CIFS data

This section covers the step-by-step procedure for transitioning a source 7-Mode SMB share to an ONTAP system.



NetApp assumes that the 7-Mode and ONTAP systems are SMB licensed. The destination SVM is created, the source and destination SMB shares are exported, and XCP is installed and licensed.

1. Scan the SMB shares for the files and directories.

```

C:\xcp>xcp scan -stats \\10.61.77.189\performance_SMB_home_dirs
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to xxxx xxxx[NetApp Inc]
until Mon Dec 31 00:00:00 2029
== Maximum Values ==
Size Depth Namelen Dirsize
15.6MiB 2 8 200
== Average Values ==
Size Depth Namelen Dirsize
540KiB 2 7 81
== Top File Extensions ==
.txt .tmp
5601 2200
== Number of files ==
empty <8KiB 8-64KiB 64KiB-1MiB 1-10MiB 10-100MiB >100MiB
46 6301 700 302 200 252
== Space used ==
empty <8KiB 8-64KiB 64KiB-1MiB 1-10MiB 10-100MiB >100MiB
0 6.80MiB 8.04MiB 120MiB 251MiB 3.64GiB 0
== Directory entries ==
empty 1-10 10-100 100-1K 1K-10K >10k
18 1 77 1
== Depth ==
0-5 6-10 11-15 16-20 21-100 >100
7898
== Modified ==
>1 year >1 month 1-31 days 1-24 hrs <1 hour <15 mins future
2167 56 322 5353
== Created ==
>1 year >1 month 1-31 days 1-24 hrs <1 hour <15 mins future
2171 54 373 5300
Total count: 7898
Directories: 97
Regular files: 7801
Symbolic links:
Junctions:
Special files:
Total space for regular files: 4.02GiB
Total space for directories: 0
Total space used: 4.02GiB
7,898 scanned, 0 errors, 0s

```

- Copy the files (with or without ACL) from the source to the destination SMB share. The following example shows a copy with ACL.

```

C:\xcp>xcp copy -acl -fallback-user "DOMAIN\gabi" -fallback-group
"DOMAIN\Group" \\10.61.77.189\performance_SMB_home_dirs
\\10.61.77.56\performance_SMB_home_dirs
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to xxxx xxxx[NetApp Inc]
until Mon Dec 31 00:00:00 2029
7,898 scanned, 0 errors, 0 skipped, 184 copied, 96.1MiB (19.2MiB/s), 5s
7,898 scanned, 0 errors, 0 skipped, 333 copied, 519MiB (84.7MiB/s), 10s
7,898 scanned, 0 errors, 0 skipped, 366 copied, 969MiB (89.9MiB/s), 15s
7,898 scanned, 0 errors, 0 skipped, 422 copied, 1.43GiB (99.8MiB/s), 20s
7,898 scanned, 0 errors, 0 skipped, 1,100 copied, 1.69GiB (52.9MiB/s),
25s
7,898 scanned, 0 errors, 0 skipped, 1,834 copied, 1.94GiB (50.4MiB/s),
30s
7,898 scanned, 0 errors, 0 skipped, 1,906 copied, 2.43GiB (100MiB/s),
35s
7,898 scanned, 0 errors, 0 skipped, 2,937 copied, 2.61GiB (36.6MiB/s),
40s
7,898 scanned, 0 errors, 0 skipped, 2,969 copied, 3.09GiB (100.0MiB/s),
45s
7,898 scanned, 0 errors, 0 skipped, 3,001 copied, 3.58GiB (100.0MiB/s),
50s
7,898 scanned, 0 errors, 0 skipped, 3,298 copied, 4.01GiB (88.0MiB/s),
55s
7,898 scanned, 0 errors, 0 skipped, 5,614 copied, 4.01GiB (679KiB/s),
1m0s
7,898 scanned, 0 errors, 0 skipped, 7,879 copied, 4.02GiB (445KiB/s),
1m5s
7,898 scanned, 0 errors, 0 skipped, 7,897 copied, 4.02GiB (63.2MiB/s),
1m5s

```



If there is no data aggregate, create a new one using the storage aggr create command.

3. Sync the files on the source and destination.

```

C:\xcp>xcp sync -acl -fallback-user "DOMAIN\gabi" -fallback-group
"DOMAIN\Group" \\10.61.77.189\performance_SMB_home_dirs
\\10.61.77.56\performance_SMB_home_dirs
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to xxxx xxxx[NetApp Inc]
until Mon Dec 31 00:00:00 2029
10,796 scanned, 4,002 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 5s
15,796 scanned, 8,038 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 10s
15,796 scanned, 8,505 compared, 0 errors, 0 skipped, 0 copied, 0

```

```
removed, 15s
15,796 scanned, 8,707 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 20s
15,796 scanned, 8,730 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 25s
15,796 scanned, 8,749 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 30s
15,796 scanned, 8,765 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 35s
15,796 scanned, 8,786 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 40s
15,796 scanned, 8,956 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 45s
8 XCP v1.6 User Guide © 2020 NetApp, Inc. All rights reserved.
Step Description
15,796 scanned, 9,320 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 50s
15,796 scanned, 9,339 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 55s
15,796 scanned, 9,363 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m0s
15,796 scanned, 10,019 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m5s
15,796 scanned, 10,042 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m10s
15,796 scanned, 10,059 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m15s
15,796 scanned, 10,075 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m20s
15,796 scanned, 10,091 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m25s
15,796 scanned, 10,108 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m30s
15,796 scanned, 10,929 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m35s
15,796 scanned, 12,443 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m40s
15,796 scanned, 13,963 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m45s
15,796 scanned, 15,488 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m50s
15,796 scanned, 15,796 compared, 0 errors, 0 skipped, 0 copied, 0
removed, 1m51s
```

4. Verify that the files were copied correctly.

```
C:\xcp> xcp verify \\10.61.77.189\performance_SMB_home_dirs
\\10.61.77.56\performance_SMB_home_dir
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to xxxx xxxx[NetApp Inc]
until Mon Dec 31 00:00:00 2029
8 compared, 8 same, 0 different, 0 missing, 5s
24 compared, 24 same, 0 different, 0 missing, 10s
41 compared, 41 same, 0 different, 0 missing, 15s
63 compared, 63 same, 0 different, 0 missing, 20s
86 compared, 86 same, 0 different, 0 missing, 25s
423 compared, 423 same, 0 different, 0 missing, 30s
691 compared, 691 same, 0 different, 0 missing, 35s
1,226 compared, 1,226 same, 0 different, 0 missing, 40s
1,524 compared, 1,524 same, 0 different, 0 missing, 45s
1,547 compared, 1,547 same, 0 different, 0 missing, 50s
1,564 compared, 1,564 same, 0 different, 0 missing, 55s
2,026 compared, 2,026 same, 0 different, 0 missing, 1m0s
2,045 compared, 2,045 same, 0 different, 0 missing, 1m5s
2,061 compared, 2,061 same, 0 different, 0 missing, 1m10s
2,081 compared, 2,081 same, 0 different, 0 missing, 1m15s
2,098 compared, 2,098 same, 0 different, 0 missing, 1m20s
2,116 compared, 2,116 same, 0 different, 0 missing, 1m25s
3,232 compared, 3,232 same, 0 different, 0 missing, 1m30s
4,817 compared, 4,817 same, 0 different, 0 missing, 1m35s
6,267 compared, 6,267 same, 0 different, 0 missing, 1m40s
7,844 compared, 7,844 same, 0 different, 0 missing, 1m45s
7,898 compared, 7,898 same, 0 different, 0 missing, 1m45s,cifs
```

CIFS data migration with ACLs from a source storage box to ONTAP

This section covers the step-by-step procedure for migrating CIFS data with security information from a source to a target ONTAP system.

1. Verify that the target ONTAP system is healthy.

```

C1_sti96-vsim-ucs540m_cluster::> cluster show
Node                Health  Eligibility
-----
sti96-vsim-ucs540m  true   true
sti96-vsim-ucs540n  true   true
2 entries were displayed.
C1_sti96-vsim-ucs540m_cluster::> node show
Node      Health Eligibility Uptime           Model      Owner      Location
-----
sti96-vsim-ucs540m
           true  true           15 days 21:17 SIMBOX      ahammed    sti
sti96-vsim-ucs540n
           true  true           15 days 21:17 SIMBOX      ahammed    sti
2 entries were displayed.
cluster::> storage failover show
Node      Partner      Takeover
-----
Possible State Description
-----
sti96-vsim-ucs540m
           sti96-vsim-  true   Connected to sti96-vsim-ucs540n
           ucs540n
sti96-vsim-ucs540n
           sti96-vsim-  true   Connected to sti96-vsim-ucs540m
           ucs540m
2 entries were displayed.
C1_sti96-vsim-ucs540m_cluster::>

```

2. Verify that at least one nonroot aggregate exists on the target system. The aggregate is normal.

```

cluster::*> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
aggr0_sti96_vsim_ucs540o
      7.58GB   373.3MB   95% online    1 sti96-vsim-
raid_dp,
                                ucs540o
normal
aggr0_sti96_vsim_ucs540p
      7.58GB   373.3MB   95% online    1 sti96-vsim-
raid_dp,
                                ucs540p
normal
aggr_001    103.7GB   93.63GB   10% online    1 sti96-vsim-
raid_dp,
                                ucs540p
normal
sti96_vsim_ucs540o_aggr1
      23.93GB   23.83GB    0% online    1 sti96-vsim-
raid_dp,
                                ucs540o
normal
sti96_vsim_ucs540p_aggr1
      23.93GB   23.93GB    0% online    0 sti96-vsim-
raid_dp,
                                ucs540p
normal
5 entries were displayed.

```



If there is no data aggregate, create a new one using the `storage aggr create` command.

3. Create an SVM on the target cluster system.

```
cluster::*> vserver create -vserver vs1 -rootvolume root_vs1 -aggregate
sti96_vsim_ucs540o_aggr1 -rootvolume-security-style mixed
```

Verify that the SVM was successfully created.

```
C2_sti96_vsim_ucs540o_cluster::*> vserver show -vserver vs1
      Vserver: vs1
      Vserver Type: data
      Vserver Subtype: default
      Vserver UUID: f8bc54be-d91b-11e9-b99c-
005056a7e57e
      Root Volume: root_vs1
      Aggregate: sti96_vsim_ucs540o_aggr1
      NIS Domain: NSQA-RTP-NIS1
      Root Volume Security Style: mixed
      LDAP Client: esisconfig
      Default Volume Language Code: C.UTF-8
      Snapshot Policy: default
      Data Services: data-nfs, data-cifs,
                    data-flexcache, data-iscsi
      Comment: vs1
      Quota Policy: default
      List of Aggregates Assigned: -
      Limit on Maximum Number of Volumes allowed: unlimited
      Vserver Admin State: running
      Vserver Operational State: running
      Vserver Operational State Stopped Reason: -
      Allowed Protocols: nfs, cifs, fcp, iscsi, ndmp
      Disallowed Protocols: -
      Is Vserver with Infinite Volume: false
      QoS Policy Group: -
      Caching Policy Name: -
      Config Lock: false
      Volume Delete Retention Period: 0
      IPspace Name: Default
      Foreground Process: -
      Is Msid Preserved for DR: false
      Force start required to start Destination in multiple IDP fan-out case:
false
      Logical Space Reporting: false
      Logical Space Enforcement: false
```

4. Create a new read-write data volume on the destination SVM. Verify that the security style, language settings, and capacity requirements match the source volume.

```
CLUSTER CLUSTER::> vol create -vserver vs1 -volume dest_vol -aggregate
aggr_001 -size 150g type RW -state online -security-style ntfs
```

5. Create a data LIF to serve SMB client requests.

```
CLUSTER::> network interface create -vserver vs1 -lif sti96-vsimsim-
ucs540o_data1 -address 10.237.165.87 -netmask 255.255.240.0 -role data
-data-protocol nfs,cifs -home-node sti96-vsimsim-ucs540o -home-port e0d
```

Verify that the LIF was successfully created.

```
cluster::*> network interface show -vserver vs1
Logical      Status      Network      Current
Current Is
Vserver      Interface   Admin/Oper   Address/Mask  Node         Port
Home
-----
vs1
           sti96-vsimsim-ucs540o_data1
                   up/up      10.237.165.87/20  sti96-vsimsim-ucs540o
                                                           e0d
true
```

6. If required, create a static route with the SVM.

```
Network route create -vserver dest -destination 0.0.0.0/0 -gateway
10.237.160.1
```

Verify that the route was successfully created.

```
cluster::*> network route show -vserver vs1
Vserver      Destination      Gateway      Metric
-----
vs1
           0.0.0.0/0      10.237.160.1  20
           ::/0          fd20:8b1e:b255:9155::1
                                       20
2 entries were displayed.
```

7. Mount the target data volume in the SVM namespace.

```
CLUSTER::> volume mount -vserver vs1 -volume dest_vol -junction-path
/dest_vol -active true
```

Verify that the volume is successfully mounted.

```
cluster::*> volume show -vserver vs1 -fields junction-path
vserver volume    junction-path
-----
vs1      dest_vol /dest_vol
vs1      root_vs1 /
2 entries were displayed.
Note: You can also specify the volume mount options (junction path) with
the volume create command.
```

8. Start the CIFS service on the target SVM.

```
cluster::*> vserver cifs start -vserver vs1
Warning: The admin status of the CIFS server for Vserver "vs1" is
already "up".
```

Verify that the service is started and running.

```
cluster::*>
Verify the service is started and running
C2_sti96-vs1m-ucs540o_cluster::*> cifs show
      Server          Status   Domain/Workgroup Authentication
Vserver  Name              Admin    Name              Style
-----
vs1      D60AB15C2AFC4D6 up       CTL               domain
```

9. Verify that the default export policy is applied to the target SVM.

```
CLUSTER::> vserver export-policy show -vserver dest
Vserver      Policy Name
-----
dest         default
```

If required, create a new custom export policy for the target SVM.

```
CLUSTER::> vserver export-policy create -vserver vs1 -policyname  
xcpexport
```

10. Modify the export policy rules to allow access to CIFS clients.

```
CLUSTER::> export-policy rule modify -vserver dest -ruleindex 1  
-policyname xcpexportpolicy -clientmatch 0.0.0.0/0 -rorule any -rwrule  
any -anon 0
```

Verify that the policy rules are modified.

```

cluster::*> export-policy rule show -instance
                Vserver: vs1
                Policy Name: default
                Rule Index: 1
                Access Protocol: any
List of Client Match Hostnames, IP Addresses, Netgroups, or Domains:
0.0.0.0/0
                RO Access Rule: any
                RW Access Rule: any
User ID To Which Anonymous Users Are Mapped: 65534
                Superuser Security Types: any
                Honor SetUID Bits in SETATTR: true
                Allow Creation of Devices: true
                NTFS Unix Security Options: fail
Vserver NTFS Unix Security Options: use_export_policy
                Change Ownership Mode: restricted
Vserver Change Ownership Mode: use_export_policy
                Policy ID: 12884901889
                Vserver: vs1
                Policy Name: default
                Rule Index: 2
                Access Protocol: any
List of Client Match Hostnames, IP Addresses, Netgroups, or Domains:
0:0:0:0:0:0:0:0/0
                RO Access Rule: any
                RW Access Rule: any
User ID To Which Anonymous Users Are Mapped: 65534
                Superuser Security Types: none
                Honor SetUID Bits in SETATTR: true
                Allow Creation of Devices: true
                NTFS Unix Security Options: fail
Vserver NTFS Unix Security Options: use_export_policy
                Change Ownership Mode: restricted
Vserver Change Ownership Mode: use_export_policy
                Policy ID: 12884901889
2 entries were displayed.

```

11. Verify that the client is allowed access to the volume.

```
cluster::*> export-policy check-access -vserver vs1 -volume dest_vol
-client-ip 10.234.17.81 -authentication-method none -protocol cifs
-access-type read-write
```

Path	Policy	Policy Owner	Policy Owner Type	Rule Index
/	default	root_vs1	volume	1
/dest_vol	default	dest_vol	volume	1

2 entries were displayed.

12. Connect to the Windows client system where XCP is installed. Browse to the XCP install path.

```
C:\WRSHDNT>dir c:\netapp\xcp
dir c:\netapp\xcp
Volume in drive C has no label.
Volume Serial Number is 5C04-C0C7
Directory of c:\netapp\xcp
09/18/2019 09:30 AM <DIR> .
09/18/2019 09:30 AM <DIR> ..
06/25/2019 06:27 AM 304 license
09/18/2019 09:30 AM <DIR> Logs
09/29/2019 08:45 PM 12,143,105 xcp.exe
2 File(s) 12,143,409 bytes
3 Dir(s) 29,219,549,184 bytes free
```

13. Query the source node SMB exports by running the `xcp show` command on the XCP Windows client host system.

```

C:\WRSHDNT>c:\netapp\xcp\xcp show \\10.237.165.71
c:\netapp\xcp\xcp show \\10.237.165.71
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Mon Dec 31 00:00:00 2029
  Shares   Errors   Server
         6         0         10.237.165.71
== SMB Shares ==
  Space   Space   Current
  Free    Used    Connections Share Path           Folder Path
 9.50GiB 4.57MiB 1          \\10.237.165.71\source_share C:\source_vol
94.3MiB 716KiB 0          \\10.237.165.71\ROOTSHARE   C:\
0        0        N/A       \\10.237.165.71\ipc$       N/A
94.3MiB 716KiB 0          \\10.237.165.71\c$         C:\
== Attributes of SMB Shares ==
  Share                               Types
Remark
  source_share                        DISKTREE
  test share                          DISKTREE
  test_sh                             DISKTREE
  ROOTSHARE                          DISKTREE           \"Share mapped
to top of Vserver global namespace, created bydeux_init \"
  ipc$                                PRINTQ,SPECIAL,IPC,DEVICE
  c$                                  SPECIAL
== Permissions of SMB Shares ==
  Share                               Entity
Type
  source_share                        Everyone
Allow/Full Control
  ROOTSHARE                          Everyone
Allow/Full Control
  ipc$                                Everyone
Allow/Full Control
  c$                                  Administrators
Allow/Full Control/

```

14. Run the help command for copy.

```

C:\WRSHDNT>c:\netapp\xcp\xcp help copy
c:\netapp\xcp\xcp help copy
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Mon Dec 31 00:00:00 2029
usage: xcp copy [-h] [-v] [-parallel <n>] [-match <filter>] [-preserve-
atime]
                [-acl] [-fallback-user FALLBACK_USER]
                [-fallback-group FALLBACK_GROUP] [-root]
                source target
positional arguments:
  source
  target
optional arguments:
  -h, --help            show this help message and exit
  -v                    increase debug verbosity
  -parallel <n>         number of concurrent processes (default: <cpu-
count>)
  -match <filter>       only process files and directories that match
the
                        filter (see `xcp help -match` for details)
  -preserve-atime       restore last accessed date on source
  -acl                  copy security information
  -fallback-user FALLBACK_USER
                        the name of the user on the target machine to
receive
                        the permissions of local (non-domain) source
machine
                        users (eg. domain\administrator)
  -fallback-group FALLBACK_GROUP
                        the name of the group on the target machine to
receive
                        the permissions of local (non-domain) source
machine
                        groups (eg. domain\administrators)
  -root                 copy acl for root directorytxt

```

15. On the target ONTAP system, get the list of local user and local group names that you need to provide as values for the `fallback-user` and `fallback-group` arguments path.

```

cluster::*> local-user show
(vserver cifs users-and-groups local-user show)
Vserver      User Name      Full Name
Description
-----
vs1          D60AB15C2AFC4D6\Administrator
                                                    Built-in
administrator account
C2_sti96-vs1m-ucs540o_cluster::*> local-group show
(vserver cifs users-and-groups local-group show)
Vserver      Group Name      Description
-----
vs1          BUILTIN\Administrators      Built-in Administrators
group
vs1          BUILTIN\Backup Operators      Backup Operators group
vs1          BUILTIN\Guests      Built-in Guests Group
vs1          BUILTIN\Power Users      Restricted
administrative privileges
vs1          BUILTIN\Users      All users
5 entries were displayed

```

16. To migrate the CIFS data with ACLs from the source to target, run the `xcp copy` command with the `-acl` and `-fallback-user/group` options.

For the `fallback-user/group` options, specify any user or group that can be found in Active Directory or local user/group to target system.

```

C:\WRSHDNT>c:\netapp\xcp\xcp copy -acl -fallback-user
D60AB15C2AFC4D6\Administrator -fallback-group BUILTIN\Users
\\10.237.165.79\source_share \\10.237.165.89\dest_share
c:\netapp\xcp\xcp copy -acl -fallback-user D60AB15C2AFC4D6\Administrator
-fallback-group BUILTIN\Users \\10.237.165.79\source_share
\\10.237.165.89\dest_share
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Mon Dec 31 00:00:00 2029
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 8s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 13s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 18s
ERROR failed to obtain fallback security principal "BUILTIN\Users".
Please check if the principal with the name "BUILTIN\Users" exists on
"D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\Administrator". Please check if the principal with the
name "D60AB15C2AFC4D6\Administrator" exists on "D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal "BUILTIN\Users".
Please check if the principal with the name "BUILTIN\Users" exists on
"D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal "BUILTIN\Users".
Please check if the principal with the name "BUILTIN\Users" exists on
"D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal "BUILTIN\Users".
Please check if the principal with the name "BUILTIN\Users" exists on
"D60AB15C2AFC4D6".
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 23s
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\Administrator". Please check if the principal with the
name "D60AB15C2AFC4D6\Administrator" exists on "D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\Administrator". Please check if the principal with the
name "D60AB15C2AFC4D6\Administrator" exists on "D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\Administrator". Please check if the principal with the
name "D60AB15C2AFC4D6\Administrator" exists on "D60AB15C2AFC4D6".
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 28s
753 scanned, 0 errors, 0 skipped, 249 copied, 24.0KiB (4.82KiB/s), 33s
753 scanned, 0 errors, 0 skipped, 744 copied, 54.4KiB (6.07KiB/s), 38s
753 scanned, 0 errors, 0 skipped, 746 copied, 54.5KiB (20/s), 43s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (1.23KiB/s), 44s
C:\WRSHDNT>

```

17. If xcp copy results in the error message ERROR failed to obtain fallback security principal, add the destination box in the hosts file (C:\Windows\System32\drivers\etc\hosts).

Use the following format for the NetApp storage destination box entry.

```
<data vserver data interface ip> 1 or more white spaces <cifs server name>
```

```
cluster::*> cifs show
      Server          Status      Domain/Workgroup Authentication
Vserver Name            Admin       Name          Style
-----
vs1      D60AB15C2AFC4D6 up          CTL           domain
C2_sti96-vsim-ucs540o_cluster::*> network interface show
      Logical      Status      Network      Current
Current Is
Cluster
      sti96-vsim-ucs540p_clus1
      up/up      192.168.148.136/24 sti96-vsim-ucs540p
      e0a
true
      sti96-vsim-ucs540p_clus2
      up/up      192.168.148.137/24 sti96-vsim-ucs540p
      e0b
true
vs1
      sti96-vsim-ucs540o_data1
      up/up      10.237.165.87/20   sti96-vsim-ucs540o
      e0d
true
      sti96-vsim-ucs540o_data1_inet6
      up/up      fd20:8b1e:b255:9155::583/64
      sti96-vsim-ucs540o
      e0d
true
      sti96-vsim-ucs540o_data2
      up/up      10.237.165.88/20   sti96-vsim-ucs540o
      e0e
true
10.237.165.87 D60AB15C2AFC4D6 -> destination box entry to be added in
hosts file.
```

18. If you still get the error message ERROR failed to obtain fallback security principal after adding the destination box entry in the hosts files, then the user/group does not exist in the target system.

```

C:\WRSHDNT>c:\netapp\xcp\xcp copy -acl -fallback-user
D60AB15C2AFC4D6\unknown_user -fallback-group BUILTIN\Users
\\10.237.165.79\source_share \\10.237.165.89\dest_share
c:\netapp\xcp\xcp copy -acl -fallback-user D60AB15C2AFC4D6\unknown_user
-fallback-group BUILTIN\Users \\10.237.165.79\source_share
\\10.237.165.89\dest_share
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Mon Dec 31 00:00:00 2029
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\unknown_user". Please check if the principal with the
name "D60AB15C2AFC4D6\unknown_user" exists on "D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\unknown_user". Please check if the principal with the
name "D60AB15C2AFC4D6\unknown_user" exists on "D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\unknown_user". Please check if the principal with the
name "D60AB15C2AFC4D6\unknown_user" exists on "D60AB15C2AFC4D6".
ERROR failed to obtain fallback security principal
"D60AB15C2AFC4D6\unknown_user". Please check if the principal with the
name "D60AB15C2AFC4D6\unknown_user" exists on "D60AB15C2AFC4D6".
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 5s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 10s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 15s
753 scanned, 0 errors, 0 skipped, 284 copied, 27.6KiB (5.54KiB/s), 20s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (2.44KiB/s), 22s
C:\WRSHDNT>

```

19. Use `xcp copy` to migrate CIFs data with ACLs (with or without the root folder).

Without the root folder, run the following commands:

```

C:\WRSHDNT>c:\netapp\xcp\xcp copy -acl -fallback-user
D60AB15C2AFC4D6\Administrator -fallback-group BUILTIN\Users
\\10.237.165.79\source_share \\10.237.165.89\dest_share
c:\netapp\xcp\xcp copy -acl -fallback-user
D60AB15C2AFC4D6\Administrator -fallback-group BUILTIN\Users
\\10.237.165.79\source_share \\10.237.165.89\dest_share
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Mon Dec 31 00:00:00 2029
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 5s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 10s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 15s
753 scanned, 0 errors, 0 skipped, 210 copied, 20.4KiB (4.08KiB/s), 20s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (2.38KiB/s), 22s
C:\WRSHDNT>

```

With the root folder, run the following commands:

```

C:\WRSHDNT>c:\netapp\xcp\xcp copy -acl -root -fallback-user
D60AB15C2AFC4D6\Administrator -fallback-group BUILTIN\Users
\\10.237.165.79\source_share \\10.237.165.89\dest_share
c:\netapp\xcp\xcp copy -acl -root -fallback-user
D60AB15C2AFC4D6\Administrator -fallback-group BUILTIN\Users
\\10.237.165.79\source_share \\10.237.165.89\dest_share
XCP SMB 1.6; (c) 2020 NetApp, Inc.; Licensed to XXX [NetApp Inc] until
Mon Dec 31 00:00:00 2029
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 5s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 10s
753 scanned, 0 errors, 0 skipped, 0 copied, 0 (0/s), 15s
753 scanned, 0 errors, 0 skipped, 243 copied, 23.6KiB (4.73KiB/s), 20s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (6.21KiB/s), 25s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 30s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 35s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 40s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 45s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 50s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 55s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 1m0s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (0/s), 1m5s
753 scanned, 0 errors, 0 skipped, 752 copied, 54.7KiB (817/s), 1m8s
C:\WRSHDNT>

```

Best practice guidelines and recommendations

- Use the XCP client operating system, which is IMT supported. The IMT supported client is qualified by

NetApp.

- Run XCP as a root user in the Linux operating system to perform migration. you can run the xcp command as the sudo user, but it is not supported by XCP.
- Run only one instance of XCP per client. Technically you can run multiple instanced of XCP on the same host from a different location, however this is not a supported practice. Indeed, running many instances might result in failure.
- In the current XCP version, Live Source is not supported. If the source NetApp volume is active and continuously changed by applications and users, you should take a snapshot of the source volume to perform a migration.
- It is a best practice to create a new snapshot with a different name for every incremental sync so that it is easy to create an incremental migration path based on the snapshot name in the event of failure.
- If you are performing a snapshot-based migration, it is a best practice to continue snapshot-based migration until cutover.
- If you have more than 10 million files and you have incremental data change of more than 50%, it is a best practice to use a higher core count and more memory than the minimum recommendation in the installation and administration guide.

Troubleshooting

This section provides troubleshooting guidance for data migration using NetApp XCP.

Error 1: XCP Failed with nfs3 error 70: stale filehandle Error in the xcp.log

Reason and guidance.

Mount the source folder and verify that the folder exists. If it does not exist or if it has been removed, you will receive a `stale filehandle` error, in which case, you can ignore the error.

Error 2: NetApp NFS Destination Volume Has Space, but XCP Failed with nfs3 error 28: no space left on device

Reason and guidance.

1. Check the space of the NFS destination volume by running the `df` command or check the storage.

```
root@workkr-140: USER3# df -h /xcpdest
Filesystem                Size  Used Avail Use% Mounted on
10.63.150.127:/xcpsrc_vol 4.3T  1.7T  2.6T  40% /xcpsrc_vol
```

2. Check the inodes in the storage controller.

```
A800-Node1-2::> volume show -volume xcpdest -fields files,files-used
vserver          volume  files    files-used
-----
A800-Node1_vs1  xcpdest 21251126 21251126
A800-Node1-2::>
```

3. If inode is used, increase the number of inodes by running the following command:

```
A800-Node1-2::> volume modify -volume xcpdest -vserver A800-Node1_vs1
-files 40000000
Volume modify successful on volume xcpdest of Vserver A800-Node1_vs1.
A800-Node1-2::> volume show -volume xcpdest -fields files,files-used
vserver          volume  files    files-used
-----
A800-Node1_vs1  xcpdest 39999990 21251126
A800-Node1-2::>
```

Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

- [NetApp XCP blogs](#)
- [NetApp XCP documentation](#)
- [Big data analytics data to artificial intelligence](#)

Data Protection

TR-4320: NetApp E-Series and Commvault Data Platform V11 - Reference architecture and storage best practices

Akash Gupta, NetApp
Girish Chanchlani, Commvault

TR-4320 outlines the reference architecture and best practices when using NetApp E-Series storage in a Commvault Data Platform V11 environment. Commvault and NetApp have jointly developed this reference architecture to provide guidance for Commvault Data Platform V11 deployments with NetApp E-Series storage that will accelerate time to application for this solution.

[TR-4320: NetApp E-Series and Commvault Data Platform V11 - reference architecture and storage best practices](#)

TR-4471: E-Series and EF-Series reference architecture and storage best practices with Veeam Backup & Replication 9.5

Akash Gupta, NetApp
Shawn Lieu (Americas), Stefan Renner (EMEA), and Michael Cade (Performance), Veeam

TR-4471 outlines the reference architecture and best practices when using NetApp E-Series storage in a Veeam Backup & Replication 9.5 environment.

[TR-4471: E-Series and EF-Series reference architecture and storage best practices with Veeam Backup & Replication 9.5](#)

TR-4704: Deploying Veritas NetBackup with NetApp E-Series storage

Akash Gupta and Principled Technologies, NetApp

TR-4704 describes deployment of Veritas NetBackup on NetApp E-Series storage.

[TR-4704: Deploying Veritas NetBackup with NetApp E-Series storage](#)

Security

NVA-1143: NetApp HCI - NIST security controls for FISMA with HyTrust for multitenant infrastructure - NVA design and deployment

Arvind Ramakrishnan, Abhinav Singh, NetApp

NVA-1143 describes how NetApp HCI can be designed and deployed to meet National Institute of Standards and Technology (NIST) SP 800-53 Revision 4 Security and Privacy controls, which are crucial for private cloud infrastructures and multitenant deployments.

[NVA-1143: NetApp HCI - NIST security controls for FISMA with HyTrust for multitenant infrastructure - NVA design and deployment](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.