



Technical Report

Successfully Transitioning to Clustered Data ONTAP

7-Mode Transition Tool (7MTT)
Data ONTAP 8.2.x and 8.3.x

Jay Bounds - NetApp
February 2016 | TR-4052

Abstract

The 7-Mode Transition Tool (7MTT) enables customers to migrate their 7-Mode data from the NetApp Data ONTAP® operating system to clustered Data ONTAP 8.3.2 and above. This technical report focuses on the functionality of the components that make-up 7MTT; Collect and Assess, Copy-based Transition (CBT) and Copy-Free Transition (CFT).

Data Classification

Public

TABLE OF CONTENTS

1	Overview	5
1.1	Before You Begin.....	5
1.2	Terminology	5
2	Transition Fundamentals	7
2.1	Identify	9
2.2	Design.....	11
2.3	Implement	14
3	Transition Tools.....	18
3.1	Discovery and Assessment Tools	18
3.2	Data Migration Tools Overview	19
3.3	7-Mode Transition Tool	23
3.4	Host-Based And Third-Party Tools	26
3.5	Foreign LUN Import (FLI).....	27
4	Cluster Design Considerations	27
4.1	Cluster Creation and Expansion	27
4.2	Namespaces and SVMs	27
4.3	Nondisruptive Movement	28
4.4	Cluster Lifecycle.....	28
5	Performance Assessment	29
5.1	Workload Characterization.....	29
5.2	Utilization Assessment	31
5.3	Sizing Considerations	32
6	Disaster Recovery and Backup Migration Scenarios	32
6.1	Scenario 1: Basic Disaster Recovery	33
6.2	Scenario 2: Bidirectional Disaster Recovery	34
6.3	Scenario 3: Basic Backup	35
6.4	Scenario 4: Disaster Recovery and Backup Cascade.....	36
6.5	Scenario 5: Disaster Recovery and Backup Fan-Out.....	37
6.6	Scenario 6: Disaster Recovery Fan-Out	39
6.7	Scenario 7: SnapVault-to-Disaster Recovery Cascade.....	40
6.8	Scenario 8: SnapVault Fan-In.....	41
6.9	Scenario 9: Disaster Recovery Cascade	42
6.10	Scenario 10: Qtree Replication	43

7	Qtree Migration	43
7.1	Should Qtrees Be Reorganized?	43
7.2	Scenario 1: Qtree-to-Qtree Migration Using CBT	44
7.3	Scenario 3: Qtree-to-Volume Migration	45
7.4	Scenario 4: Qtree Consolidation	46
7.5	Additional Qtree Considerations	47
8	Product and Feature Considerations	48
8.1	WAFL	48
8.2	Nondisruptive Operations and High Availability	49
8.3	Deduplication and Compression	50
8.4	Flash Cache and Flash Pool	51
8.5	FlexClone	51
8.6	Data ONTAP Networking	51
8.7	SnapMirror	52
8.8	SnapVault	53
8.9	MetroCluster	53
8.10	NDMP	53
8.11	SMTape	53
8.12	CIFS	54
8.13	NFS	54
8.14	SAN	55
8.15	OnCommand Management Products	57
8.16	NetApp Manage ONTAP API Suite	59
8.17	SnapDrive and SnapManager Products	59
9	NetApp Service Offerings	60
10	Additional Resources	62
	Version history	64
11	Contact Us	64
12	Disclaimer	64

LIST OF TABLES

Table 1)	Key SAN differences between 7-Mode and clustered Data ONTAP.	55
Table 2)	OnCommand management product versions.	57
Table 3)	SnapDrive and SnapManager minimum product versions for clustered Data ONTAP.	59

LIST OF FIGURES

Figure 1) An overview of the journey to clustered Data ONTAP.....	7
Figure 2) Transition fundamentals overview.....	8
Figure 3) Standard disruptive data migration steps.....	16
Figure 4) 7MTT decision tree overview.....	21
Figure 5) 7MTT CBT data migration process overview.....	24
Figure 6) 7MTT CFT data migration process overview.....	25
Figure 7) Hub-spoke topology for disaster recovery and backup relationships.....	33
Figure 8) Basic disaster recovery in 7-Mode.....	33
Figure 9) Basic disaster recovery in clustered Data ONTAP.....	34
Figure 10) Bidirectional disaster recovery in 7-Mode.....	34
Figure 11) Bidirectional disaster recovery in clustered Data ONTAP.....	35
Figure 12) Basic backup in 7-Mode.....	36
Figure 13) Basic backup in clustered Data ONTAP.....	36
Figure 14) Disaster recovery and backup cascade in 7-Mode.....	37
Figure 15) Disaster recovery and backup cascade on clustered Data ONTAP.....	37
Figure 16) Disaster recovery and backup fan-out in 7-Mode.....	38
Figure 17) Disaster recovery and backup fan-out in clustered Data ONTAP.....	38
Figure 18) Disaster recovery fan-out in 7-Mode.....	39
Figure 19) Disaster recovery fan-out in clustered Data ONTAP.....	39
Figure 20) SnapVault-to-disaster recovery cascade in 7-Mode.....	40
Figure 21) SnapVault-to-disaster recovery cascade in clustered Data ONTAP.....	40
Figure 22) SnapVault fan-in in 7-Mode.....	41
Figure 23) SnapVault fan-in in clustered Data ONTAP.....	41
Figure 24) Disaster recovery cascade in 7-Mode.....	42
Figure 25) Disaster recovery cascade in clustered Data ONTAP.....	42
Figure 26) Qtree-to-qtree migration overview—CBT.....	44
Figure 27) Qtree-to-volume migration using replication-based migration methods.....	45
Figure 28) Qtree-to-volume migration using copy-based migration methods.....	46
Figure 29) Qtree consolidation using replication based migration methods.....	46
Figure 30) Qtree consolidation using copy-based migration methods.....	47

1 Overview

Clustered Data ONTAP® provides enterprise-ready, unified scale-out storage on the premises or in the cloud. The storage was developed from a solid foundation of proven Data ONTAP technology and innovation. Clustered Data ONTAP is the basis for virtualized shared storage infrastructures that are architected for nondisruptive operations in a scalable architecture over the lifetime of the system.

Clustered Data ONTAP's nondisruptive operations enable you to service your infrastructure and redistribute load as needed without disrupting access to user data and applications. You can add capacity as you grow across both SAN and NAS environments, without needing to reconfigure running applications. As a result, you can start small and grow big without disruptive hardware upgrades that are commonly required by other storage vendors.

If you are an existing NetApp® customer with 7-Mode storage systems, you need to transition your 7-Mode environment to take advantage of clustered Data ONTAP capabilities. This document addresses the key knowledge that is necessary to move your 7-Mode environment to clustered Data ONTAP.

1.1 Before You Begin

Before starting your transition activity, you need to gain a fundamental understanding of clustered Data ONTAP. Core clustered Data ONTAP concepts such as namespace, storage virtual machine (SVM), and LIFs (logical interfaces) must be understood for you to effectively design and deploy a clustered Data ONTAP system. If the previously mentioned concepts are foreign to you, you need to read further to gain a fundamental understanding of clustered Data ONTAP. The best place to start learning about the core clustered Data ONTAP concepts is in [TR-3982: Clustered Data ONTAP: An Introduction](#). Additionally, a number of clustered Data ONTAP classes are available through NetApp University that cover the fundamentals of both clustered Data ONTAP and transition:

- [Clustered Data ONTAP Fundamentals](#) (web based)
- [Clustered Data ONTAP Administration](#) (instructor led)
- [NetApp Transition Fundamentals](#) (web based)
- [Planning and Implementing Transition Using the 7-Mode Transition Tool](#) (web-based)
- [Transitioning to Clustered Data ONTAP](#) (web based)

Once you have gained a fundamental understanding of clustered Data ONTAP, you are ready to start planning your transition.

1.2 Terminology

This document uses several key terms that are common in the storage industry. Although the terms are common, they can mean different things to different people. This section defines the key terms for the purposes of this document.

- **Transition:** This term refers to the end-to-end people, process, and technology activities involved in moving a 7-Mode environment to clustered Data ONTAP.
- **Data migration:** This term refers to the activity of copying, replicating or moving data from a source storage system to a destination storage system. Data migration is a subset of the activities involved in making the transition to clustered Data ONTAP.
- **Cutover:** This term encompasses the tasks involved in data migration from when clients and applications no longer communicate with the source 7-Mode storage system to when the clients and applications successfully communicate with the destination cluster.

A nondisruptive cutover involves fewer tasks than a disruptive cutover because there is no need to disconnect and reconnect clients or shut down and restart application services. These tasks make the clients depend on the destination cluster and no longer depend on the source 7-Mode storage system.

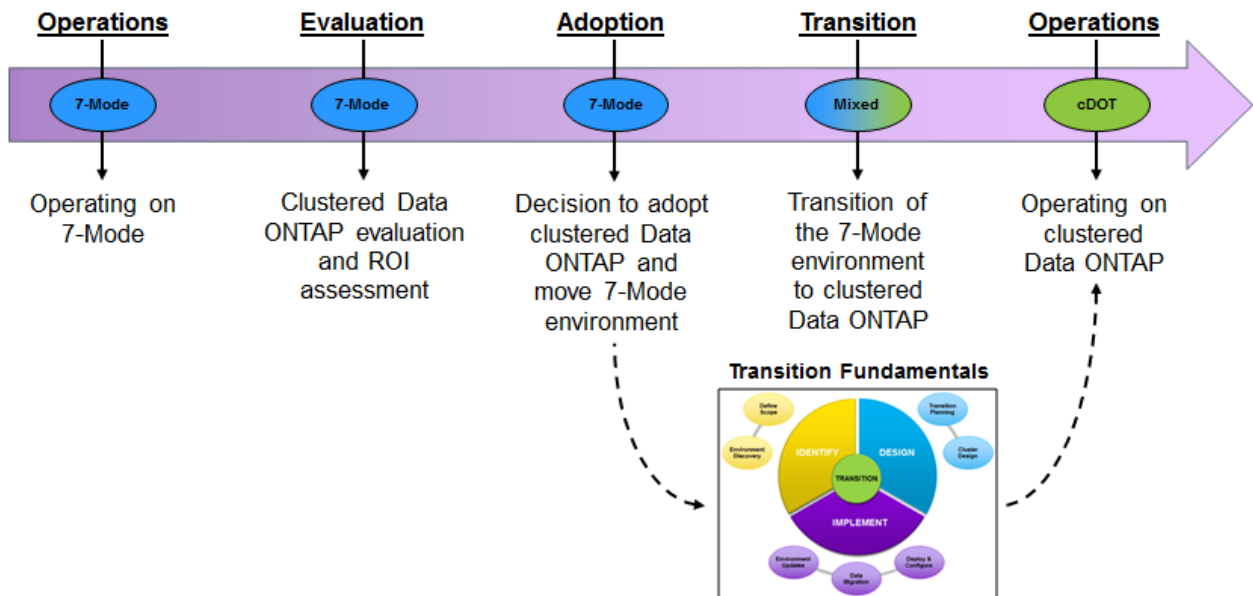
For disruptive cutovers, the tasks involved usually include:

- Disconnecting clients and shutting down application services
- An incremental data copy or replication, referred to as a final update
- Physical movement and recabling of disk shelf stacks
- Any final configuration changes to the destination cluster
- Reconnecting clients and restarting application services
- **Cutover window:** This term refers to the amount of time it takes to complete the tasks involved in cutover.
- **Copy-Based Transition:** This term indicates the method being used to move data that is based on SnapMirror technology to copy the data from the source to the destination.
- **Copy-Free Transition:** This term refers to the physical movement or reconnect of disk shelf stacks from a 7-Mode system to a clustered Data ONTAP system. No data copying or replication required.
- **Offbox:** This term refers to any technical elements that are not native features or functions of Data ONTAP. For example, a SnapMirror relationship in Data ONTAP is “onbox” a native feature of Data ONTAP. NetApp SnapDrive[®] technology for Windows[®] is installed on a Windows host and therefore is “offbox,” because it is not installed on the NetApp storage system.
- **TDP SnapMirror:** This term is used to indicate that a SnapMirror relationship’s source is 7-Mode and the destination is clustered Data ONTAP. TDP, or Transition Data Protection, refers to the type of SnapMirror relationship that is reported by Data ONTAP. A TDP SnapMirror relationship is always a volume SnapMirror relationship.
- **Operationalization:** This term is used to refer to the objective of updating people, process, and technology operating environment dependencies (on storage systems) from one operating environment to another.
- **Applications:** This term refers to the applications that are deployed in your environment; for example, Microsoft[®] Exchange or a homegrown application. Virtualized infrastructure is also grouped into applications for the purposes of transition.
- **Workloads:** This term encompasses any primary (client-generated) load that is being put on the storage system. Applications generate a workload, but workloads also include nonapplication-specific loads (for example, home directories).
- **Background workload:** This term refers to the storage system operations that are used for maintenance activities (for example, scrubs), specific product features (for example, deduplication), or backup and disaster recovery (for example, SnapMirror relationships).
- **Storage container:** This term applies to the source storage object that is being moved. This object can be an entire storage system, aggregate, volume, LUN, qtree, directory, or, in some cases, a specific file.
- **7-Mode environment:** This term indicates the grouping of all deployed 7-Mode storage systems and the technology elements that depend on those storage systems being online and serving data.
- **Operating environment:** This term refers to a specific grouping of interdependent technology. This technology can consist of both hardware and software elements. For example, my Microsoft Exchange application is deployed on a NetApp storage system, forming a dependency between those technologies (creating an operating environment).
- **Indirect data access:** This term describes how a client communicates with nodes and accesses volumes in the cluster. Indirect data access occurs when the node being contacted does not own the volume being accessed.
- **Direct data access:** This term describes how a client communicates with nodes and accesses volumes in the cluster. Direct data access occurs when the node being contacted owns the volume being accessed.

2 Transition Fundamentals

The transition fundamentals framework addresses the fundamental knowledge that is needed to holistically complete a successful transition to clustered Data ONTAP from 7-Mode. Transition begins when you decide to adopt clustered Data ONTAP and move your existing 7-Mode environment. Figure 1 shows how the transition fundamentals fit into a simplified overview of the journey to clustered Data ONTAP.

Figure 1) An overview of the journey to clustered Data ONTAP.



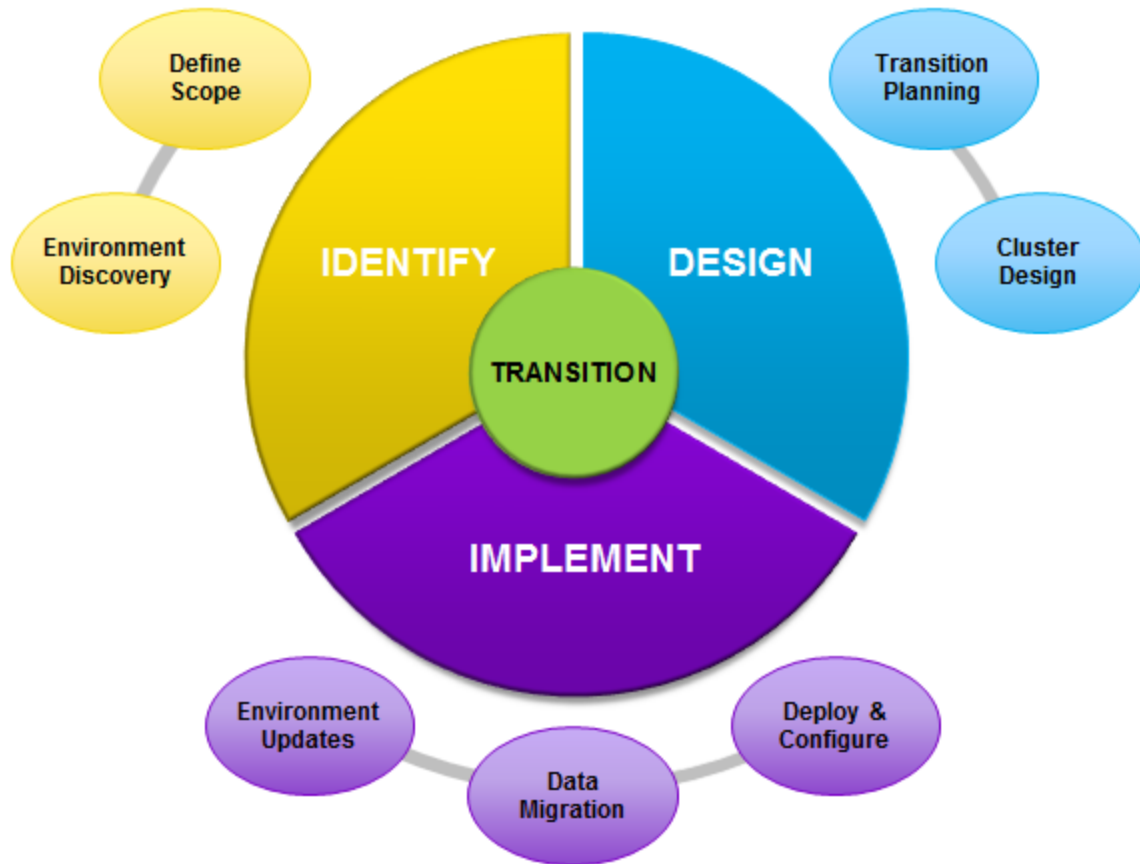
When you start your transition, your NetApp environment consists of 7-Mode storage systems (potentially with some preexisting clustered Data ONTAP systems that were deployed to address new workloads or storage needs). As the transition activity progresses, your NetApp environment is a mix of both 7-Mode and clustered Data ONTAP. By the end of the transition activity, your NetApp environment operates completely on clustered Data ONTAP.

The transition fundamentals are broken into three logical categories:

- **Identify:** This category addresses the initial business decisions that will drive the identification of which transition method(s) will be used in the NetApp environment. These decisions involve understanding which applications and workloads are deployed in your 7-Mode environment; the business requirements for the transition activity; and how to group together the systems, applications, or workloads that will move and prioritize their order.
- **Design:** This category addresses the significant activities of planning the transition activity and the cluster design and architecture. Transition planning encompasses detailed discovery of the 7-Mode environment and migration planning. Cluster design and architecture encompass the planning of which cluster capabilities are required to address the migrated workloads and future growth needs.
- **Implement:** This category addresses executing the activities involved with deployment and/or configuration of the destination cluster, migration of workloads, and data from 7-Mode to clustered Data ONTAP. The category also involves updating the operating environment so that the NetApp environment and all the technologies that depend on it work as expected.

Figure 2 displays an overview of the categories and activities that constitute the transition fundamentals.

Figure 2) Transition fundamentals overview.



Not a Process

The transition fundamentals framework is not a process. Rather, it is a logical organization of core topics that indicate which activity needs to occur and how to accomplish each task. NetApp understands that customers already have their own workflows and processes to execute projects similar to transition. The intention of the transition fundamentals is that they plug into your existing process or workflow.

Although not a process, there is a logical flow to achieving the tasks and activities involved with transition. To that end, when applying the fundamentals, start with the activities covered in Identify, move to Design, and conclude with Implement.

NetApp and its partners have a standard process they use when delivering transition-related services to customers. See the [NetApp Service Offerings](#) section of this document for more information.

Iteration

Regardless of the process or workflow you use to execute the transition activity from start to finish, it is natural that iteration of some aspects of the transition activity will occur. If you apply the transition fundamentals in the order in which they are shown in Figure 2, it is reasonable to expect that in medium to large environments, planning and executing data migrations (if not other aspects of the fundamentals as well) will be iterative.

For example, the initial transition activity and the first systems and workloads that move to clustered Data ONTAP might see all the fundamentals applied from Identify through to Implement. Data migrations that follow might simply iterate through the implement activities until all planned systems and workloads have

been transitioned to clustered Data ONTAP. NetApp expects that customers will understand how to apply the fundamentals to their existing processes and workflows to achieve their desired outcomes.

2.1 Identify

The purpose of the fundamentals contained within Identify is to determine (identify) the applications, workloads, and storage systems in your 7-Mode environment. Identify also assists in determining which process – CBT or CFT, is the method of choice to move data. Once that has been determined (or if you've already determined this), you can start planning your transition. Although there are some technical aspects to the activities in the Identify category, the key decisions to make are primarily business driven.

The fundamentals in the Identify category are:

- **Environment discovery:** The task of identifying the existing systems, workloads, applications, constituent groups, and business requirements for your 7-Mode environment.
- **Define scope:** The task of defining which components of the 7-Mode environment will move together and in what order they will move, and the task of defining whether all components can be moved in a single instance.

For a list of NetApp service offerings that can assist with the activities of the Identify fundamentals, see the [NetApp Service Offerings](#) section of this document.

Objective Summary

- A list of applications and workloads along with their dependent storage systems
- Business and cutover requirements for each application or workload
- Dividing the 7-Mode environment into one or more groupings to move to clustered Data ONTAP together
- For all defined groups, a priority order for each group, identifying which group will transition first

Environment Discovery

The task of identifying your 7-Mode environment might be simple or it might require some cross-functional communication, depending on the size of your 7-Mode environment. If you have just a few 7-Mode storage systems deployed, determining what is deployed will be straightforward. If you have hundreds of storage systems deployed, it is prudent to take the time to identify what is in the current environment.

A highly detailed technical discovery of your 7-Mode environment is not necessary at this point. More often than not, business requirements drive the decision as to what will move and by when. If simply told what is to move and to start planning, you won't need to do higher-level environment discovery and will need to look at transition planning. Part of transition planning is carrying out detailed technical discovery of your 7-Mode configuration, environment dependences, applications, and clients.

The objectives of environmental discovery are:

- A list of applications and workloads along with their dependent storage systems
- Business and cutover requirements for each application or workload

Producing a list of applications and workloads should be fairly straightforward. Establish enough information as is necessary to make an informed business decision when defining the scope of the transition. Commonly, a single 7-Mode storage system might have multiple application or workload dependencies that will be critical to capture during this activity. If your systems currently report NetApp AutoSupport™ information to NetApp, you can quickly identify which 7-Mode storage systems you have deployed by viewing the information in My AutoSupport.

It will be necessary to discover the cutover requirements and any other business requirements that will drive the scope of the transition project from the constituent groups that rely on the 7-Mode environment. Cutover requirements play a major role in the scope of the transition and in transition planning activities.

7-Mode Transition Tool: Collect and Assess

Although a detailed analysis of the 7-Mode environment is not necessary initially, you can use the 7-Mode Transition Tool's (7MTT's) Collect and Assess functionality to quickly discover key information about the 7-Mode storage systems and clients (including enterprise applications) currently deployed. Doing requires that you identify the storage systems and clients by IP address or fully qualified domain name and have validated credentials the 7MTT can use to log into the systems. Additional work will have to be done to identify applications and workloads. However, the 7MTT will make clear recommendations as to which storage systems are ready to move first and which systems need additional consideration. Also, the 7MTT collects detailed technical information that will be needed for proper transition planning and cluster design. Whether or not you plan to use the 7MTT for data migration, NetApp recommends leveraging 7MTT's the Collect and Assess functionality to reduce the level of effort needed for environment discovery.

Additional tools are available that can assist with environment discovery and assessment. See the [Discovery and Assessment Tools](#) section of this document for more information.

Define Scope

Defining the scope of the transition requires that you first determine which 7MTT method should be used based on the 7-Mode environment to be transitioned. 7MTT provides two methods for transitioning your data: copy-based and copy-free. This section provides guidance about which tool to use based on your environment.

Copy-Based Transition (CBT)

There are two common approaches to dividing an environment. The first approach is to divide it by 7-Mode storage controllers and the workloads they host. In the case of controller-level granularity, each workload is normally contained within a single controller and not spread across multiple controllers (or HA pairs). The second approach is to divide the environment by the applications, workloads, or constituent groups that rely on the storage systems. As a result, only a subset of storage containers on a 7-Mode system can move with any particular grouping.

Once you establish the groupings, it will be necessary to identify the order in which the groupings will move. Considerations for grouping systems together should include:

- Cutover requirements of the applications, workloads, or storage systems
- Placement of dependent storage containers on the 7-Mode storage systems
- System warranties and tech refresh schedules
- Readiness of the storage system to move to clustered Data ONTAP; for example, if any feature or functional difference needs to be assessed before being moved.

The objectives of defining the scope are:

- Dividing the 7-Mode environment into one or more groupings that should move to clustered Data ONTAP together
- For all defined groups, a priority order for each group, identifying which group will transition first

2.2 Design

The purpose of the fundamentals contained within Design is to complete all necessary planning for the transition activity and to determine your destination cluster requirements. Although migration planning is a large portion of a transition plan, it is not the only consideration. Migration planning and more are discussed in the transition planning section below. Cluster design focuses on the requirements for your cluster based on the 7-Mode environment being migrated. Design also focuses on the general best practice decisions to consider for proper cluster operation (regardless of the 7-Mode environment).

The fundamentals and activities covered in Design are the most important activities to do correctly. That is because they determine the outcomes not only for the short-term transition activity, but also for the longer-term operational health of your cluster.

The fundamentals in the Design category are:

- **Transition planning:** This is the task of planning data migration, environment updates (offbox dependencies), and training.
- **Cluster design:** This is the task of determining what the target clustered Data ONTAP system architecture will be, based on the 7-Mode environment being moved and the desired forward-looking requirements for new or preexisting clusters.

For a list of NetApp service offerings that can assist with the activities of the Design fundamentals, see the [NetApp Service Offerings](#) section of this document.

Objective Summary

- Document the following key transition elements:
 - Detailed technical discovery that provides:
 - A mapping of 7-Mode storage containers to their clustered Data ONTAP equivalents
 - A 7-Mode feature assessment
 - Identification of existing disaster recovery and backup relationships
 - The migration method that will be used for each storage container being moved
 - A list of management applications and other operational dependencies requiring updates
 - A plan to provide training updates on clustered Data ONTAP for 7-Mode administrators
 - A list of scripts leveraging Data ONTAP APIs (ZAPIs) or commands that need to be updated
- Determine configuration elements that must be applied to the destination cluster
- 7-Mode workload characterization and utilization assessment
- Cluster sizing or headroom assessment
- Cluster architecture or configuration changes

Transition Planning

Transition planning activities focus on documenting how you intend to move your 7-Mode environment to clustered Data ONTAP. Whether using CBT or CFT, the size of your 7-Mode data footprint might be straightforward (small footprint; for example, only a couple of HA pairs) or require additional effort for a large environment (for example, tens or hundreds of systems with complex DR relationships). Although the majority of transition planning is about data migration, there are considerations beyond data migration that are addressed in this section.

The objective of Transition Planning is to document the following key transition elements:

- Detailed technical discovery that provides:
 - A mapping of 7-Mode storage environment to their clustered Data ONTAP equivalents
 - A 7-Mode feature assessment
 - Identification of existing disaster recovery and backup relationships
 - The migration methods that will be used
- A list of management applications and other operational dependencies requiring updates
- A plan to provide training updates on clustered Data ONTAP for 7-Mode administrators
- A list of scripts leveraging Data ONTAP APIs (ZAPIs) or commands that need to be updated

Whether or not you plan to migrate data by controller using CFT or by volume when using CBT, it is necessary to document all source data on the systems being moved. For some migrations, data will be migrated at volume granularity. Mapping these storage containers to their final destination in clustered Data ONTAP is necessary for these reasons:

- To assist with selecting an appropriate migration method
- To provide essential information for cluster design and configuration

In addition to listing source controller storage containers, it will be necessary to list which Data ONTAP features are in use on the source systems. Keep in mind that some features might be licensed or enabled on the 7-Mode system, but not in use. For each feature discovered, NetApp recommends verifying that the feature is in use. The feature assessment is essential to determine if the same features will be used in clustered Data ONTAP or if some features currently in use were deprecated and new functionality should be used instead. For additional information on feature assessment for 7-Mode, see the [Product and Feature Considerations](#) section of this document. Special consideration is given to the migration of qtrees from 7-Mode to clustered Data ONTAP. Additional information on the migration of qtrees is available in the [Qtree Migration Scenarios](#) section of this document.

7-Mode Transition Tool: Collect and Assess

It is possible to complete a detailed technical discovery manually; however, NetApp has provided an automated capability to conduct the detailed technical discovery through the 7MTT's Collect and Assess functionality. If the 7MTT was used when conducting environmental discovery, you can leverage the additional information that was collected at that time for transition planning. If you started with planning you can run the 7MTT to collect this detailed information. For additional information on the 7MTT and other tools that are available, see the [Discovery and Assessment Tools](#) section of this document.

NetApp has defined 10 common scenarios for disaster recovery (DR) and backup-related configurations (when SnapMirror or NetApp SnapVault® technology is being used). The DR and backup scenarios address the recommended approach for data migration when source storage containers are in a DR or backup relationship. When you conduct detailed discovery and planning, it is necessary to include this type of information for any source storage containers or systems. For additional information on the DR and backup scenarios, see the [Disaster Recovery and Backup Migration Scenarios](#) section of this document.

NetApp has developed the migration tools decision tree to assist with selecting the right migration method for each storage container, workload or controller migrated. The migration tools decision tree covers the 80/20 rule for determining which migration method should be considered first for any given storage container. Given the number of options available to move data, there can be multiple options to use. However, some options are better than others and this is reflected in the migration tools decision tree. See the [Migration Tools Decision Tree](#) section of this document for more information.

If you use NetApp OnCommand® management applications to manage your 7-Mode environment, it will be necessary to determine which versions support clustered Data ONTAP. See the [OnCommand® Management Products](#) section for additional information.

Specific versions of NetApp SnapDrive and NetApp SnapManager® products that support clustered Data ONTAP are required. You will need to review which products are in use today and their current versions to identify which products need to be updated. For information on the specific versions required for clustered Data ONTAP, see the [SnapDrive and SnapManager Products](#) section of this document.

Existing 7-Mode knowledge will form the foundation of knowledge for clustered Data ONTAP. Because of the new functionality and scale-out architecture of clustered Data ONTAP, additional training is necessary to build operational knowledge. NetApp University developed courses for clustered Data ONTAP that are specifically targeted at quickly ramping up 7-Mode administrators. For additional information on the training options available, contact your NetApp (direct or partner) representative.

Any scripts or other automation that use 7-Mode commands or Manage ONTAP® APIs will need to be reviewed for updates. Some commands are similar in updates whereas other commands are completely new or changed. The command changes are a result of the clustered Data ONTAP command interface becoming highly structured. If you use regular expressions (or something similar) to mine values from command output, it is important to note that command output in clustered Data ONTAP is generally much more structured than in 7-Mode. An excellent resource for 7-Mode administrators looking for command equivalency in clustered Data ONTAP is the [Clustered Data ONTAP 8.2 Command Map for 7-Mode Administrators](#). Additionally, you can use the [Cfor7 community forum](#) to ask questions about and get answers to any 7-Mode to clustered Data ONTAP equivalency questions.

When planning how storage containers and workloads will migrate into a destination cluster, there are two common options for achieving the desired data layout within the cluster:

- **Direct-to-node:** Migrate the source storage container directly to the destination cluster node that will host the workload.
- **Gateway node:** Specify a node in the destination cluster that will be the target of all data migrations. Once a workload is migrated, use the nondisruptive volume movement capability of clustered Data ONTAP to relocate the storage container to another node in the cluster (that is, the node intended to host the migrated workload).

The direct-to-node option is the most common approach. The gateway node approach can be considered by customers who have a large NetApp footprint with multiple workloads migrating to clustered Data ONTAP. The benefit of the second approach is that all data migration activity is isolated to a single node in the cluster.

It is important to verify the host operating system versions that are deployed in the 7-Mode environment. Consult the [Interoperability Matrix Tool](#) (IMT) for additional information on supported host operating system versions.

Cluster Design

The approach to cluster design will vary slightly based on whether a clustered Data ONTAP system already exists (a preexisting cluster that hosts production workloads) or a new cluster will be deployed as part of the transition activity.

The objectives of cluster design are to address the following:

- Configuration elements that must be applied to the destination cluster
- 7-Mode workload characterization and utilization assessment
- Cluster sizing or headroom assessment
- Cluster architecture or configuration changes

Transition planning will yield a list of the source 7-Mode storage containers that are moving and how they map into the destination cluster. The mapping will dictate which configuration elements must be created on the cluster. The mapping must specify which destination SVM and aggregate the storage container will be moved to. In the case of subvolume storage containers, it is necessary to specify additional granularity (for example, directories or qtrees). Unless you use a tool such as the 7MTT that creates the destination volume for you (it is still necessary to create higher-level containers), it will be necessary to plan to create these configuration elements when the destination cluster is configured.

In order to properly size the destination cluster or to determine if the migrating workloads require expansion of a preexisting cluster, it is necessary to complete a workload characterization and utilization assessment of the source 7-Mode storage systems. The [Workload Characterization](#) and [Utilization Assessment](#) sections of this document discuss in more detail how to accomplish these activities for 7-Mode.

The workload characterization and utilization assessment for the 7-Mode storage systems is used to either size a new cluster or determine if sufficient headroom exists within a preexisting cluster. The same methods used for utilization assessment and workload characterization on the 7-Mode storage systems can also be applied to clustered Data ONTAP (with some slight modifications). The [Workload Characterization](#) and [Utilization Assessment](#) sections of this document cover accomplishing these activities for clustered Data ONTAP in more detail.

If the transition activity is targeted to a preexisting cluster, it is likely that key design, architecture, and best practice decisions were already made. However, if a new cluster is being deployed, the key design, architecture, and best practice decisions made at that point will determine the operational health of the cluster for both the transition activity and the posttransition period. The [Cluster Design Considerations](#) and the [Product and Feature Considerations](#) sections of this document address the key design, architecture, and best practice decisions to consider when determining the design and architecture or required configuration changes for your cluster.

2.3 Implement

The fundamentals of Implement address putting your plans into action (executing). Implement includes deploying a new cluster or configuring a preexisting cluster that will be the target of the planned data migrations. Additionally, Implement covers acquiring and deploying transition tools and completing the movement of data and workloads to clustered Data ONTAP. Implement also covers completing the environmental updates that are necessary to complete the transition activity.

The fundamentals in the Implement category are:

- Deploy and Configure
 - The task of deploying and preparing a new cluster (from staging into production) or executing configuration changes on a preexisting cluster (already in production)
- Data Migration
 - The task of completing the movement of data and workloads to the destination clustered Data ONTAP system
- Environment Updates
 - The task of completing all clustered Data ONTAP environment operating dependency updates and staff training activities

For a list of NetApp service offerings that can assist with the activities of the design fundamentals, see the [NetApp Service Offerings](#) section of this document.

Objective Summary

- Deploy and configure the new cluster or execute configuration changes on the pre-existing cluster.
- Test and verify cluster configuration and offbox dependencies.
- Execute data migrations and cutover activity as planned.
- Verify data access and configuration postmigration.
- Update management applications as needed.
- Update script dependencies.
- Complete outstanding staff training requirements.
- Create new or update existing environmental procedures as needed.

Deploy and Configure

The activities Deploy and Configure will vary depending on whether there is a preexisting destination cluster for the transition activity or a new cluster is being deployed. Preexisting clusters will be a factor for customers who choose to deploy a new cluster with a new workload or application not already being serviced by the 7-Mode environment. They also will be a factor for a new application that might have been the target of a non NetApp to NetApp data migration. That being said, currently the more common case is that a new cluster is being deployed as part of the transition activity. This section addresses each objective based on whether or not there is a preexisting cluster or a new cluster.

The objective of deploy and configure is to address:

- Deploying and configuring the new cluster or executing configuration changes on the preexisting cluster
- Testing and verifying cluster configuration and offbox dependencies

The activities that occur when deploying a new cluster will be primarily dictated by the results of the cluster design activities. The following activities will be executed:

- Hardware procurement, qualification, staging, and installation
- Primary configuration of clustered Data ONTAP protocols, services, and features
- Creation and configuration of storage virtual machines, aggregates, volumes, and more, depending on which 7-Mode storage containers are being moved to clustered Data ONTAP
- Any other activities that are necessary to put the cluster into production

If the cluster preexists and CBT is the migration tool of choice, the activities being executed will change slightly:

- The introduction of new hardware might not be necessary unless the performance assessment yielded the need to nondisruptively expand the existing cluster capabilities.
- Primary configuration of clustered Data ONTAP protocols, services, and features should already be completed because the cluster is already in production and servicing workloads. That being said, it is possible that new protocol requirements are being added to the cluster (based on the workloads being moved) that will necessitate configuration of the protocol and any related services or new features.
- Some level of creation and configuration of storage virtual machines, aggregates, volumes, and more will still need to be completed based on what is being moved into the cluster.

Regardless of whether the cluster is new or preexisting, it will be necessary to thoroughly test and validate that data access and cluster operation are healthy. As with the deployment of any technology (NetApp or not), human error (typos, null entries, and the like) or misconfiguration might be overlooked and result in data access or service disruption.

For example, if the DNS server IPs changed (the servers were upgraded or replaced) but those changes did not get communicated to the people responsible for configuring the cluster CIFS servers, the change would result in the incorrect entry of IP addresses for the DNS configuration. Again, this could happen for any technology, but mentioning is intended to remind you to check your work.

Data Migration

The act of data migration is one that is familiar to many customers. Depending on the migration method being used, the data migration might be disruptive or nondisruptive. Transition planning activities provide a migration method that will be used for each storage container moved to clustered Data ONTAP, along with the cutover considerations and timing associated with that method.

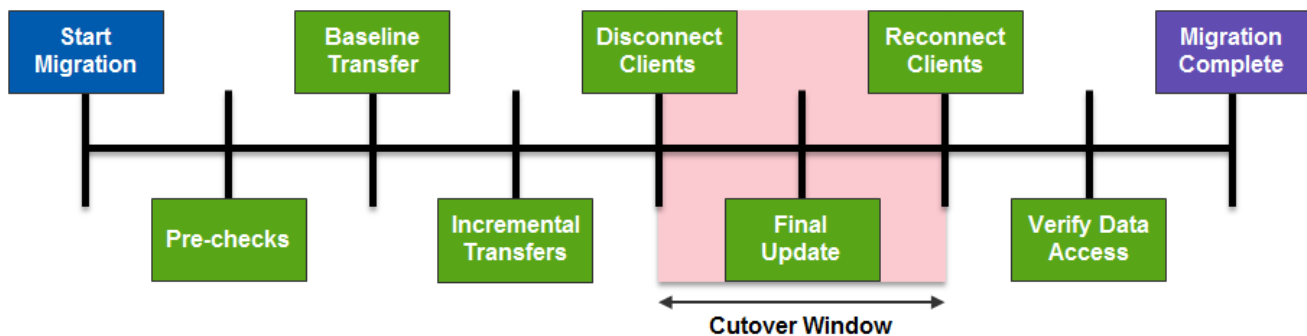
The objective of data migration is to address:

- Executing data migrations and cutover activity as planned
- Verifying data access and configuration post-migration

Certain application-based migration methods (if available) offer nondisruptive data migration capabilities. [TR-4336: Enterprise Application Transition to Clustered Data ONTAP](#) covers the available options in more detail. NetApp assumes that customers are familiar with the various steps or process involved in executing data migration using native application capabilities.

Replication-based, appliance-based, and host-based migration methods are disruptive (some might be minimally disruptive). Additional information on the available migration methods can be found in the [Data Migration Tools](#) section of this document. As a general guide, Figure 3 shows the standard steps involved in the majority of disruptive data migrations.

Figure 3) Standard disruptive data migration steps.



Various data migration tools modify the standard data migration steps that are outlined in Figure 3. However, Figure 3 is included here to set a baseline of expectations for the steps involved in disruptive data migrations.

Fundamentally, the importance of cutover is measured by customer service continuity; the customer in this case is someone who relies on the availability of the storage and/or applications. If service continuity is disrupted, there is criticality to the cutover activities and cutover window that is not present with nondisruptive cutover. If the cutover is nondisruptive, the step becomes just another transparent step (as far as service continuity is concerned) in the data migration process.

The majority of cutover activity is disruptive in some way. The following key tasks contribute to the timing of the cutover window:

- Disconnecting clients and shutdown of application services. These task are performed so that no additional writes occur on the source storage container. Additionally, shares or mounts are likely to change as a result of the data migration (but not in all cases).

- A final update (incremental transfer/data copy) that brings the destination file system in sync with the source file system. The timing of this step depends heavily on the change rate in the source file system. Even a small rate of change (measured as a percentage) if applied to a high-file-count environment can result in needing additional transfer time for the final update. That being said, the purpose of the incremental updates before cutover is to minimize the delta between the source and destination file systems.
- Any final configuration changes needed on the destination cluster. The majority of configuration should be done before cutover; however, certain configuration elements will need to be set during cutover. For example, if a storage system IP address is being moved, this process must occur after the final update and before reconnecting clients.

Note: Specifically for application cutover, it is necessary to update any SnapDrive and SnapManager products to versions that support clustered Data ONTAP during this time. For information on which versions are required, see the [SnapDrive and SnapManager Products](#) section of this document. Additional premigration, migration, and postmigration considerations for applications and SnapDrive and SnapManager products can be found in [TR-4336: Enterprise Application Transition to Clustered Data ONTAP](#).

- Reconnecting clients and restarting application services.

Once the cutover is complete, NetApp recommends that data access be verified before declaring the data migration complete. Verification of data access and service levels is recommended regardless of whether the cutover was disruptive or nondisruptive.

In addition to data access verification, some post-migration configuration might be required to address:

- SnapMirror or SnapVault relationships (this can also be done as part of the cutover activities)
- Schedules for maintenance activities such as scrubs, NetApp Snapshot[®] copies, deduplication, and more

Environment Updates

It is essential to verify that the off box environment dependencies are working once the transitioned workloads are operating (in production) on clustered Data ONTAP. It is reasonable to expect that, for large environments, a priority order might be needed as to which offbox dependencies are more critical to address first. For example, updating OnCommand management products that are used to manage your cluster would be more critical than updating a simple maintenance script that needs to use new clustered Data ONTAP commands.

The objective of environment updates is to address:

- Updating management applications as needed
- Updating script dependencies
- Completing outstanding staff training requirements
- Creating new or updating existing environmental procedures as needed

If OnCommand management products are used to manage your NetApp environment or if you are introducing them along with clustered Data ONTAP, it is necessary to determine that the proper versions that support clustered Data ONTAP (or in some cases, both 7-Mode and clustered Data ONTAP) are being deployed. Determining which versions are necessary should be done as part of transition planning, meaning that at this stage it is just necessary to deploy the correct versions. For additional information on which versions are required for clustered Data ONTAP, see the [OnCommand Management Products](#) section of this document.

Any scripts or other automation that depends on 7-Mode commands or Manage ONTAP APIs must be updated. For command updates, remember that you can use the [Clustered Data ONTAP 8.2 Command Map for 7-Mode Administrators](#) to determine command equivalencies. Additionally, you can use the [Cfor7 community forum](#) to ask questions about and get answers to any 7-Mode to clustered Data ONTAP

equivalency issues. For customers who leverage Manage ONTAP APIs, NetApp recommends working with your NetApp or partner representatives to address any required updates.

NetApp recommends closing out any outstanding training plans once the 7-Mode workloads have been cut over to clustered Data ONTAP. Training is not just for updating 7-Mode administration skills to clustered Data ONTAP, but also to learn how the operating environment has changed. Regardless of the process used for your transition activity, training can be conducted at any point before or during the transition. It is possible to conduct training afterward, although that can result in unintentional changes or issues within the environment if administrators are unaware of how the environment changed.

Although most existing operational procedures might remain unchanged as a result of transition, new procedures are warranted (if this is the first cluster being deployed) based on the new functionality of clustered Data ONTAP. Key considerations for new operational procedures are:

- **NetApp DataMotion™ for Volumes:** Enables the ability to nondisruptively move volumes between nodes in a cluster. It is necessary to define how to select a destination node based on the reasons for moving the volume. Two common scenarios for volume movement include capacity rebalancing and performance rebalancing.
- **Storage quality of service (QoS):** If you intend to use storage QoS, it will be necessary to define the operational considerations for applying QoS policies to an SVM, a volume, a LUN, or a file.
- **Flash technologies:** Having the ability to nondisruptively move volumes between aggregates and nodes in a cluster means that storage service levels (based on drive type) can be changed. As a result, volumes can be moved from and to NetApp Flash Pool™ aggregates as well as to nodes with or without NetApp Flash Cache™ intelligent caching installed.
- **Logical interfaces (LIFs):** Nondisruptive movement of LIFs between nodes in the cluster is possible. Two operational considerations given this capability are:
 - Movement of a LIF involving indirect or direct data access. Node utilization changes occur when the data access type changes in the cluster, which is a consideration for per-node utilization.
 - In the event of node disruption in the cluster, clearly define where you intend the LIFs on the disrupted node to be temporarily relocated.

3 Transition Tools

NetApp has made several tools available to assist with transition activities, from discovery and assessment through to planning and migration. This section provides an introduction to the tools that are available, where to find more information, and guidance on which tool to start looking at first for migration.

3.1 Discovery and Assessment Tools

This section provides information on the tools available that can assist with Discovery and Assessment activities. In some cases, the information provided, collected or processed by these tools directly helps with transition planning as well.

7-Mode Transition Tool: Collect and Assess

The Collect and Assess capability of the 7MTT was introduced with 7MTT version 1.4 and is included in the copy-based transition (CBT) installation software. Starting with 7MTT version 2.2, the Collect and Assess feature is not included in the Copy-Free Transition installation package. NetApp recommends that Collect and Assess be installed from the 7MTT CBT installation package, and requires installation on a separate system from the 7MTT CFT application.

The Collect and Assess feature accesses clients (hosts) and 7-Mode storage systems to collect (discover) key configuration information and assess the considerations for transition to clustered Data ONTAP.

Note: NetApp recommends using the 7MTT Collect and Assess feature for all discovery and transition planning activities regardless of transition method.

The collection component discovers information from Windows, Linux®, and VMware® ESX® hosts in addition to your 7-Mode storage systems. Key information discovered includes storage configuration (Data ONTAP 7-Mode specific features, functionality, and configuration), common enterprise applications, and host specific information. The assessment component determines if there are feature or functional differences between the existing 7-Mode feature set in use and clustered Data ONTAP 8.2.x and 8.3.

Multiple storage systems and clients can be processed in parallel by the 7MTT Collect and Assess feature. The output of the 7MTT Collect and Assess feature is an executive summary and workbook with details of all configuration elements discovered. For additional details on the 7MTT Collect and Assess feature, see the [Data and Configuration Transition Guide for Transitioning to Clustered Data ONTAP](#).

Note: Collect and Assess is not included with the 7MTT version 2.2 CFT installation package. The 7MTT version 2.2 CBT installation package does include the Collect and Assess application.





Inventory Collect Tool

If security restrictions or other considerations prevent the installation of the 7MTT in your operating environment, you can use the inventory collect tool (ICT) to execute the collect component of the 7MTT Collect and Assess feature (without installing the 7MTT). The output of the ICT execution is an XML file that can be imported into the 7MTT offline. Once the file is imported, the 7MTT will execute the assess component of the 7MTT Collect and Assess feature against the XML file. The executive summary and the detailed workbook output will be provided just as if the 7MTT was executed in the operating environment. For additional details on the ICT see the [Data and Configuration Transition Guide for Transitioning to Clustered Data ONTAP](#).

Transition Advisor

Transition Advisor (TA) is available as part of the My AutoSupport web interface. Transition Advisor allows you to execute a feature and functional assessment of 7-Mode storage systems that weekly report AutoSupport information to NetApp. When an assessment is completed using TA, a report is generated that indicates if there are functionality changes, provides recommendations for key differences that were identified, and provides as a list of any manual checks that need to be completed. TA has the capability to assess all systems reporting in AutoSupport at once (bulk assessment), individually, or by site. Assessment results can be saved in Microsoft Excel or PDF formats.

Transition Advisor color codes the items it reports on as green, yellow, red, or orange.

-  **Green:** Feature or functionality for this item is the same in clustered Data ONTAP.
-  **Yellow:** Feature or functionality for this item is different. Recommendations exist for how the feature or functionality translates to clustered Data ONTAP. NetApp recommends working with your NetApp representatives to determine the best practices for cluster design and transition planning.
-  **Orange:** Items have to be manually checked because insufficient data exists in AutoSupport to determine if the feature or functionality is the same or different in clustered Data ONTAP.
-  **Red:** Feature or functionality is very different or is not available in clustered Data ONTAP. Some items that are reported as red can be obsolete or deprecated features or functionality. This fact warrants a discussion with your NetApp representative to determine if the feature is still needed or if an alternate feature or function is appropriate.

3.2 Data Migration Tools Overview

This section provides information about which data migration tools are available from NetApp, common non NetApp tools, how data migration tools are categorized and which tools to consider first based on your migration requirements.

Data Migration Methods

Several different types of migration tools are available today, each with its own benefits and considerations. NetApp categorized these tools based on their benefits and capabilities.

- **7MTT Migration**

- **Copy-Free Transition (CFT)**

- When considering using CFT, CFT moves 7-Mode controller HA pair configurations and their attached disk storage shelves in a single instance.

- CFT does not require setting up multiple projects in order to move data, as entire data sets are moved at the same time. It is still necessary to identify all the workloads and clients that will be affected as disruption impacts entire segments of business operations.

- **Copy based transition (CBT)**

- CBT uses NetApp's proven TDP SnapMirror technology to copy data from one location to another. This proven technology preserves all data protection and storage efficiency rules throughout the data move process,

- **Application-based migration:** This migration method uses native application replication, mirroring, or copy features to migrate data to clustered Data ONTAP. Two key examples of application-based migration methods are Oracle Automatic Storage Management (ASM) and Microsoft Exchange Database Availability Group (DAG). Virtualized environment capabilities for the movement of data are also grouped in this category for the purposes for transition. For more information on application-based migration methods, see the [Non NetApp Tools](#) section of this document.
 - **Host-based migration:** This migration method uses non NetApp tools such as Rsync, NDMP copy, logical volume managers (LVMs), and more to migrate data to clustered Data ONTAP. For additional information about host-based migration tools, see the [Non NetApp Tools](#) section of this document.

Another consideration concerning the available migration tools is whether they are copy based or replication based. Replication based, as the name suggests, is based on using SnapMirror to replicate data from 7-Mode to clustered Data ONTAP. All other methods (application-, appliance-, and host-based migration methods) are copy based. Product documentation for copy-based methods may use the term replication. However, within the context of migration from 7-Mode to clustered Data ONTAP, the term is applicable only to replication-based migration methods as outlined in this section.

Migration Tools Decision Tree

NetApp developed the 7MTT as the premier transition tool for transitioning 7-Mode storage systems to clustered Data ONTAP. The 7MTT simplifies the data migration process and is able to preserve Snapshot copies, deduplication and compression savings, schedules, system configurations and SnapMirror relationships. Although the 7MTT addresses a significant number of transition use cases, it does not address all scenarios. As a result, NetApp created the migration tools decision tree which provides a starting point for which migration methods to consider first when moving data to clustered Data ONTAP.

To classify the various tools that are available, the following migration methods have been defined:

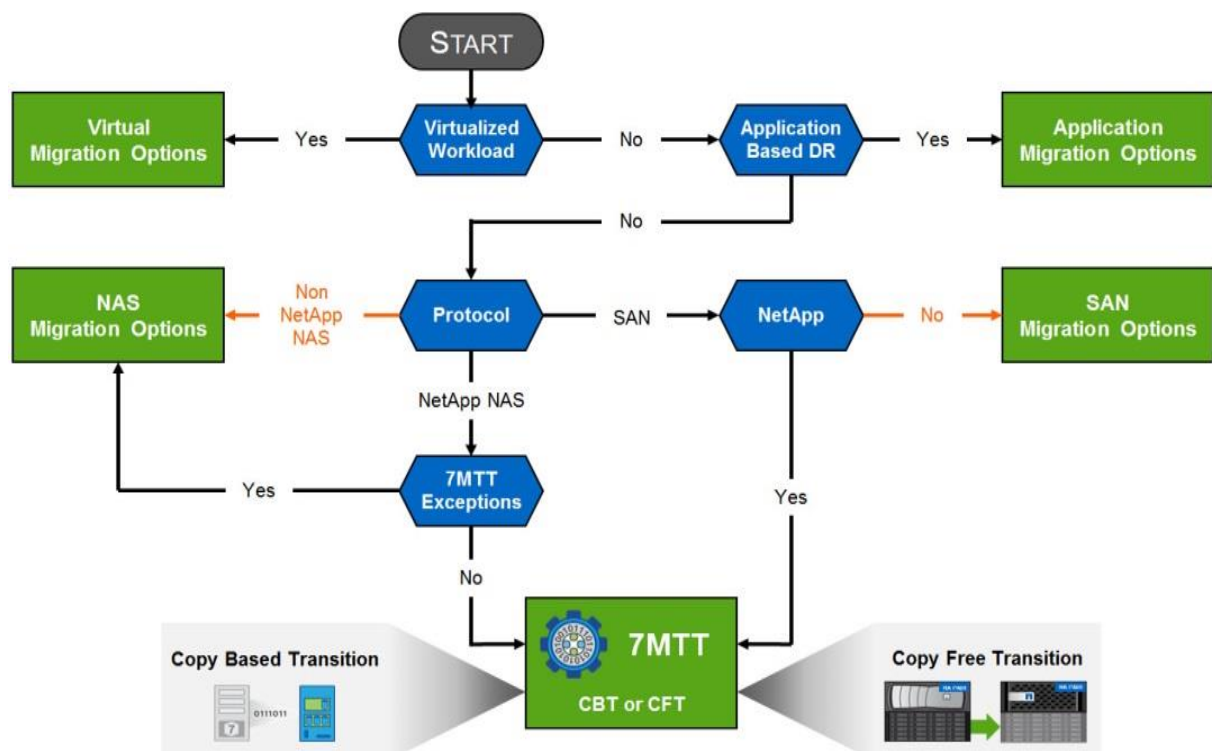
- **Virtual migration options.** This migration method uses virtualized environment capabilities to move data for the purposes for transition. Examples of virtual migration methods are vSphere Storage vMotion and Hyper-V Storage Replication.
 - **Application migration options.** This migration method uses native application replication, mirroring, or copy features to migrate data to clustered Data ONTAP. Two key examples of application-based migration methods are Oracle Automatic Storage Management (ASM) and Microsoft Exchange Database Availability Group (DAG).

- **Replication migration options.** The 7-Mode Transition Tool (7MTT) is the primary tool for replication-based data migrations. This migration method works with a Data ONTAP 7-Mode source only and uses NetApp SnapMirror technology. The 7MTT is provided at no cost to customers for replication-based migration. For additional information about the 7MTT, refer to the [7-Mode Transition Tool Library](#).
- **Copy-free migration options.** The 7-Mode Transition Tool (7MTT) is the primary tool for copy-free transition data migrations. This migration method works with Data ONTAP 7-Mode source only systems and uses NetApp copy-free transition technology. The 7MTT is provided at no cost to customers for copy-free transition migration. For additional information about the 7MTT, refer to the [7-Mode Transition Tool Library](#).
- **NAS and SAN migration options.** These migration methods use copy-based migration tools such as rsync, NDMP copy, Dell Secure Copy, logical volume managers (LVMs), SAN Foreign LUN Import (FLI), and others to move data to clustered Data ONTAP. Recommendations for host-based migration methods are beyond the scope of this document.

Before you migrate data to clustered Data ONTAP, we need to identify the recommended migration method based on the application type, the application environment, and other factors. Replication-based solutions generally provide a much higher degree of storage-level management and organization. CFT moves all the data at once with no granular organization – in a single outage or cutover window. This process reduces disruption to business operations to a single incident.

The migration tools decision tree (shown in Figure 4) was assembled to provide a starting point for which migration methods to consider first when moving data to clustered Data ONTAP. Let's take a look at this decision tree.

Figure 4) 7MTT decision tree overview.



There are a number of important points to consider when determining the migration method to use.

Virtual Migration Options

- If no virtualization environment is present, move to the next decision point because no virtual migration methods will be available.
- If a virtualized environment is present, verify that the configuration is set up to move the data within the existing operational environment. [TR-4336: Enterprise Application Transition to Clustered Data ONTAP](#) contains specific guidance on when virtualized and application data movement capabilities can be used.
- Does the virtualized infrastructure provide a nondisruptive data movement capability? If nondisruptive data movement is not possible, consider using 7MTT to preserve Snapshot copies, deduplication, compression, and SnapMirror relationships.
- Snapshot copies will not be retained through the migration activity. Snapshot copies will have to be retained on 7-Mode until the retention period expires.
- Storage efficiency savings will not be retained through the migration activity. Storage efficiency savings will have to be reapplied once the data has been moved to clustered Data ONTAP.
- SnapMirror relationships will not be retained through the migration activity. As a result, new SnapMirror relationships will have to be created in clustered Data ONTAP (a baseline transfer will be necessary to establish the new relationship).

- **Application Migration Options**

- If no application is present, move to the next decision point because no application migration methods will be available.
- If an application is present, verify that the configuration is set up to move the data within the existing operational environment. [TR-4336: Enterprise Application Transition to Clustered Data ONTAP](#) contains specific guidance on when application and virtualized data movement capabilities can be used.
- Does the application provide a nondisruptive data movement capability? If nondisruptive data movement is not possible, consider using 7MTT to preserve Snapshot copies, deduplication, compression, and SnapMirror relationships.
- Snapshot copies will not be retained through the migration activity. Snapshot copies will have to be retained on 7-Mode until the retention period expires.
- Storage efficiency savings will not be retained through the migration activity. Storage efficiency savings will have to be reapplied once the data has been moved to clustered Data ONTAP.
- SnapMirror relationships will not be retained through the migration activity. As a result, new SnapMirror relationships will have to be created in clustered Data ONTAP (a baseline transfer will be necessary to establish the new relationship).

- **Replication Migration Options**

- Replicate your data to clustered Data ONTAP either using the 7MTT or manually with TDP SnapMirror.
- The option available in 7MTT to move your data to clustered Data ONTAP:
Copy-Based Transition (CBT). CBT uses TDP SnapMirror to perform a baseline transfer and continues with incremental updates until you are ready to complete the migration.
- Snapshot copies are retained through the migration activity.
- Storage efficiency savings are retained through the migration activity.
- SnapMirror relationships are transferred through the migration activity. As a result, SnapMirror relationships are reestablished after migration in clustered Data ONTAP (a new baseline transfer is not necessary).

- **Copy-Free Migration Options**

- Starting in 7MTT 2.2, CFT provides the ability to move 7-Mode disk shelves connected to controllers running Data ONTAP 7-Mode and plug them directly into controllers running clustered Data ONTAP.
- Snapshot copies are retained through the migration activity.
- Storage efficiency savings are retained through the migration activity.
- SnapMirror relationships are transferred through the migration activity. As a result, SnapMirror relationships are reestablished after migration in clustered Data ONTAP (a new baseline transfer is not necessary however a re-sync of the SnapMirror relationship is required)

Migrations occurring as part of transition activity can use multiple migration methods. The migration tools decision tree can be applied to groups of like storage containers or on an individual storage container basis. There are always “exceptions to the rules.” The intention of the migration tools decision tree is to address the 80/20 rule and provide an initial recommendation for which migration methods will work best for the migration activity.

3.3 7-Mode Transition Tool

The 7MTT consists of components: Collect and Assess feature (as outlined in the [7-Mode Transition Tool: Collect and Assess](#) section) and two fully featured data migration capabilities. The 7MTT supports migration of data from 7-Mode storage systems to clustered Data ONTAP 8.2.x and 8.3.x.

Overview

The Key benefits of using the 7MTT for data migration are:

- 7MTT version 2.2 introduces support for:
 - The ability to migrate Data ONTAP 7-Mode disk shelves to a clustered Data ONTAP environment running version 8.3.2 and above.
 - The ability to move data that was transitioned from a 7-Mode system to a clustered DATA ONTAP SVM, to another SVM within the same cluster - using the volume rehost command.
 - The ability to migrate 7-Mode systems (HA pair) data in a single maintenance window.
- 7MTT version 2.1 introduces support for:
 - The ability to migrate Data ONTAP 7-Mode systems to clustered Data ONTAP 8.2.4 and 8.3.1.
 - SAN transition improvements: including SAN precutover testing abilities on par with NAS, staggered primary/secondary disaster recovery (DR) relationships, Data ONTAP 7-Mode backup verification (Snapshot copies) posttransition with SnapManager for Data ONTAP 8.3.1. Provides additional information to automate host-side LUN changes and auto-collection of zoning information from the FC switches.
 - Additional CIFS protocol precheck messages.
 - Fifty active projects in the 7MTT 2.1 GUI.
 - New login capability for users who are not members of the Administrator group.
 - The latest supported Data ONTAP versions are available on the NetApp [IMT](#).
- 7MTT version 2.0 includes support for:
 - SAN migration from 7-Mode to clustered Data ONTAP 8.3.
 - NetApp MetroCluster™ migration from 7-Mode to clustered Data ONTAP 8.3.
- Supports the migration of both CIFS and NFS volumes. Most data migration tools support CIFS or NFS, but not both.
- Conducts prechecks to determine that the source 7-Mode storage system and volumes are ready for the data migration activity.
- Migrates key 7-Mode storage system configuration automatically to clustered Data ONTAP.

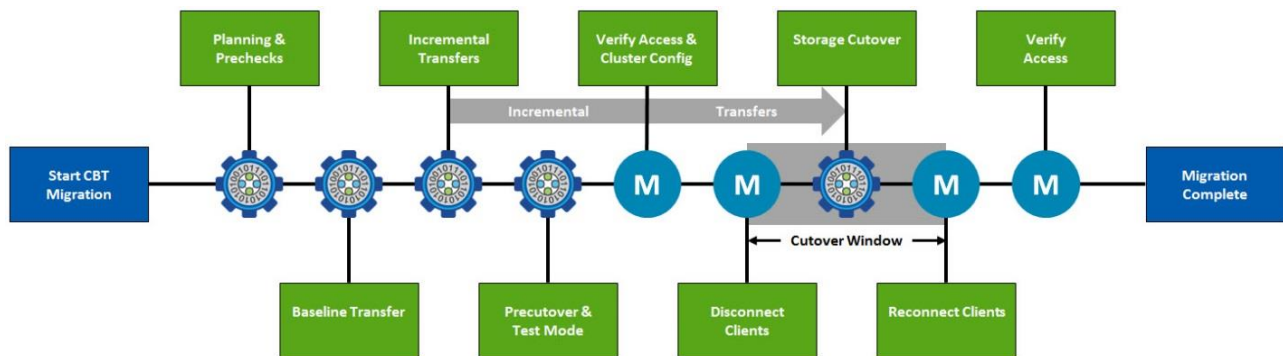
- Migrates 7-Mode volume SnapMirror relationships to clustered Data ONTAP.
- Allows the coordination of the cutover phase by automating the cutover of storage while signaling for the disconnection and reconnection of clients.
- Guides the user through the end-to-end data migration process.

The default user interface for the 7MTT is a web interface (GUI). A CLI is available for CBT in addition to the web interface; however a CLI is not currently available for CFT. The CLI enables scripts to be written that can control the 7MTT. Additionally, the 7MTT contains the NetApp Manage ONTAP kit, which can be automate the functionality of the 7MTT. For example, WFA can call the 7MTT Manage ONTAP interface to automate workflows that leverage the functionality of the 7MTT.

Copy-Base Transition Using The 7MTT

Here is an overview of the 7MTT copy-based data migration process.

Figure 5) 7MTT CBT data migration process overview.



1. **Planning and Prechecks:** 7MTT runs a series of prechecks against the source 7-Mode storage systems. Prechecks verify that the volumes and configuration are ready to be migrated to clustered Data ONTAP. Additionally, prechecks look at the cluster to verify that it is configured properly and is able to support the data migration activity. For example, verify that volumes are online, that there is network access between the systems, and so on.
2. **Baseline Transfer:** New volumes are created and a SnapMirror baseline transfer is executed between the 7-Mode source volumes and the clustered Data ONTAP destination volumes.
3. **Incremental Transfers:** After the baseline transfer is complete, incremental transfers are initiated at regular intervals (configurable by the user)
4. **Precutover and Test Mode:** This phase includes final SnapMirror incremental transfers; configuration information is applied to the clustered Data ONTAP system, storage virtual machine, and volumes; including testing of the clustered Data ONTAP volumes that are transitioned – prior to storage cutover.

Note: Although a majority of the configuration is applied, some of it is deferred to storage cutover (for example, applying quotas).

5. **Verify Access and Cluster Configuration:** Confirms cluster is accessible and ready for final storage cutover.,
6. **Disconnect Clients:** Clients are manually disconnected and access to storage suspended
7. **Storage cutover:** A final incremental transfer is completed and the SnapMirror relationship is broken
8. **Reconnect Clients:** Clients and hosts are connected to clustered Data ONTAP storage volumes and shares

9. **Verify Access:** Confirm client and application access to clustered Data ONTAP storage system

Copy-Free Transition Using The 7MTT

Copy-free transition significantly reduces the migration cost by enabling the reuse of 7-Mode disk shelves.

The overall duration for performing the transition is faster because no data copy is required.

In copy-free transition, all disk shelves are disconnected from the 7-Mode HA pair and reconnect to the an HA pair in the target cluster.

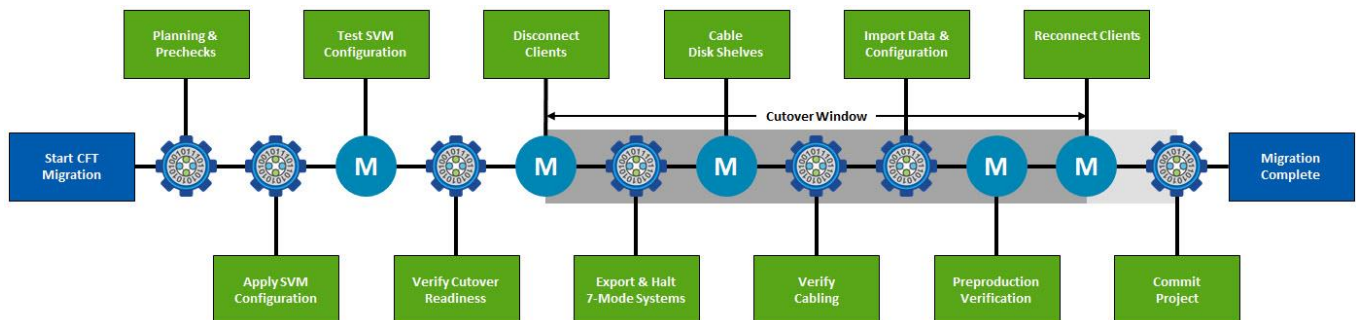
The unit of a copy-free transition is an HA pair. You must move all the disk shelves from the 7-Mode HA pair to the target cluster nodes.

The metadata of the 7-Mode aggregates and volumes is converted to the clustered Data ONTAP format by the 7-Mode Transition Tool. The time taken for this conversion is not dependent on the size of the aggregates and volumes. For example, the time take for converting a 10 GB aggregate to clustered Data ONTAP format is the same as the time required for converting a 100 TB aggregate.

Copy-free transition requires a disruption to data access. However, the total time taken to perform the data migration is faster because no data copy is required.

Here is an overview of the 7MTT copy-free data migration process.

Figure 6) 7MTT CFT data migration process overview



1. **Planning and Prechecks:** Prechecks are run against the source 7-Mode storage systems and the target clustered Data ONTAP controllers. Prechecks verify that the 7-Mode storage controllers and disk shelves are compatible, configured and ready to be migrated to clustered Data ONTAP. Additionally, prechecks look at the cluster to verify that it is configured properly and able to support the data migration activity; for example, verifying that volumes are online, that there is network access between the systems, and so on.
2. **Apply 7-mode configuration:** The 7-Mode configuration is automatically applied to the destination clustered Data ONTAP storage virtual machine.
3. **Test SVM Configuration:** Another series of pre-checks are executed to verify configuration and cutover readiness.
4. **Verify Cutover Readiness:** A final series of prechecks ensure alignment before cutover and an aggregate level snapshot is performed for rollback purposes.
5. **Disconnect Clients:** This is a manual phase which involves disconnecting all clients from the source 7-Mode storage system(s) to be transitioned.
6. **Export and Halt 7-Mode Systems:** Final configurations are exported, an aggregate level snapshot is performed for rollback purposes and 7-Mode controller services are halted.

7. **Cable Disk Shelves:** This is a manual process of disconnecting the disk shelf cables from the 7-Mode controllers and reconnecting them to the clustered Data ONTAP controllers
 8. **Verify Cabling:** This automated process verifies disk shelves are connected and communication with clustered Data ONTAP controllers.
 9. **Import Data and Configuration:** This automated process imports 7-Mode configurations and converts the 7-Mode metadata structure to clustered Data ONTAP architecture.
 10. **Preproduction Verification:** This manual process involves client application and host connectivity access validation.
- Note:** During this validation phase, any changes to configurations or data will be lost should a rollback be necessary.
11. **Reconnect Clients:** Clients and hosts are connected to clustered Data ONTAP storage volumes and shares
 12. **Commit Project:** This final phase commits the migrated data to clustered Data ONTAP and deletes the aggregate snapshot from the system.

Note: Rollback is not possible after this phase.

For more information on the 7MTT see the [Data and Configuration Transition Guide for Transitioning to Clustered Data ONTAP](#).

CFT Rollback

It is possible to roll back a transition and revert to 7-Mode at any stage of the copy-free transition--before the transitioned aggregates are committed. Rollback is a manual operation. The 7-Mode Transition Tool is used to generate the manual steps required to perform for a rollback.

The CFT process starts by taking an aggregate level snapshot of the 7-Mode storage configurations. This snapshot is preserved throughout the migration project up to the 'Commit Project' step. Any changes made to configurations or data during validation and testing will not be preserved should a rollback be necessary.

Note: Any data written to the storage prior to completing the 'Commit Project' step is considered temporary until the entire migration project is complete.

If a rollback is initiated, the 7MTT CFT interface provides instructions for completing the rollback process using the CLI (command line interface).

3.4 Host-Based And Third-Party Tools

Host- and application-based migration methods use tools that are not directly provided or supported by NetApp (because they are not NetApp products). Application-based migration methods and recommendations as to when they should be used are discussed in more detail in [TR-4336: Enterprise Application Transition to Clustered Data ONTAP](#).

Host-based tools that are commonly used for data migration are as follows:

- Logical volume managers (LVMs) from various vendors
- ScriptLogic Secure Copy
- Rsync
- Robocopy/Richcopy
- PEER Software PeerSync
- Data Dynamics StorageX
- XCP – NFSv3 migration – Linux CLI only

Some of these tools are general data migration tools; others are offerings from NetApp partners that built specific capabilities into their products to address transition from 7-Mode to clustered Data ONTAP. Note that both application- and host-based migration methods are copy based and not replication based. As a result, Snapshot copies and storage efficiency savings are not retained through the data migration activity.

3.5 Foreign LUN Import (FLI)

Foreign LUN Import (FLI) offline is an onboard feature added to the data migration tool portfolio in the clustered Data ONTAP storage operating system, beginning with version 8.3. Clustered Data ONTAP 8.3.1 includes enhancements to FLI, and now incorporates both online and Data ONTAP 7-Mode-to-clustered Data ONTAP transition workflows. The FLI Data ONTAP 7-Mode-to-clustered Data ONTAP transition can be used either online or offline. Essentially the transition is the same except that the source foreign array is a Data ONTAP 7-Mode controller.

[TR-4442: Foreign LUN Import: 7-Mode to Clustered Data ONTAP](#) covers FLI's Data ONTAP 7-Mode to clustered Data ONTAP workflow. The technical report covers using FLI as a transition solution and provides examples.

4 Cluster Design Considerations

This section addresses some of the key design points to consider when deploying a new clustered Data ONTAP system. If a preexisting cluster is the target of the transition activity, some of these design considerations might already have been determined. However, any transition activity warrants a review of the existing design and architecture to understand how the transitioned 7-Mode workloads and data will align within the existing cluster configuration. Additional design (and migration) considerations for features and products are covered in the [Product and Feature Considerations](#) section of this document.

4.1 Cluster Creation and Expansion

Certain circumstances might warrant creating a new cluster versus expanding an existing cluster. Data and workloads moving to clustered Data ONTAP from 7-Mode might not require any consideration of expansion in a preexisting cluster (depending on headroom). However, even with headroom or expansion, the configuration might warrant creating a new (separate) cluster. Key considerations for new cluster creation are:

- **Disaster recovery or backup relationships:** Given the needs of disaster recovery or backup configurations, if a cluster located in a site is interrupted because of a natural disaster or other serious event, it is likely that a separate cluster located off site will be the target of SnapMirror or SnapVault relationships.
- **Cluster maximums:** Systems do not scale infinitely in terms of the resources they can support. If the destination cluster (new or preexisting) cannot support the resources needed based on the transitioned workloads and data, it will be necessary to create a new cluster.
- **Physical location:** The physical distance between two clusters will drive the most obvious decision to deploy multiple clusters. However, even if physical distance is not a factor, other physical requirements (security, space, and more) might drive the need to deploy multiple clusters.

The considerations might seem obvious, but they are key considerations for cluster design and architecture that cannot be overlooked.

4.2 Namespaces and SVMs

In 7-Mode, individual controllers, NetApp vFiler[®] instances, and HA pairs result in clear boundaries for how workloads can be divided between systems or consolidated within the same HA pair. In short, a silo-based approach dictates how much can be hosted on any given storage system.

Clustered Data ONTAP delivers a single namespace that enables a workload to be scaled across one, some, or all of the nodes in a cluster. This process is possible because a storage virtual machine can contain a collection of flexible volumes that can reside on any node in the cluster.

The primary consideration for clustered Data ONTAP and namespace design is how the existing systems will be migrated to storage virtual machines. If possible, it is best to align with the forward-looking architecture of the clustered Data ONTAP system that will service the workloads moving forward.

In some cases, however, it might be easier to migrate each 7-Mode controller into a separate storage virtual machine. Although this might seem easier, it is not optimal because volumes cannot be nondisruptively moved between SVMs in clustered Data ONTAP. As a result, it is best to verify that volumes are migrated to clustered Data ONTAP into the appropriate SVM based on the new scale-out architecture.

If a preexisting cluster is available, it will be necessary to determine if the storage containers being moved from 7-Mode should be added to the existing SVMs or if new SVMs should be created. Additional information about namespaces can be found in [TR-4129: Namespaces in Clustered Data ONTAP](#).

4.3 Nondisruptive Movement

In clustered Data ONTAP, both volumes and logical interfaces (LIFs) can move nondisruptively between any nodes in the cluster. In the case of volumes being moved, it is necessary to plan headroom that allows space to be available on nodes in the cluster. In the case of LIFs, in addition to being able to move between nodes through planned operations, failover groups are created that dictate where LIFs will automatically move if a node is temporarily interrupted within the cluster. The movement of both volumes and LIFs in a cluster changes how data is accessed and how individual nodes are used.

The above considerations apply to cluster design and architecture regardless of whether the transition activity is targeting a new cluster or a preexisting cluster. Preexisting clusters will need to reevaluate headroom that was planned to support nondisruptive movement that might get consumed by transitioned workloads and data (that were not part of the original cluster design and architecture activities).

4.4 Cluster Lifecycle

As nodes within a cluster are refreshed and clustered Data ONTAP versions are upgraded, it is important to understand how to maintain the ability to conduct a nondisruptive platform refresh. It is also important to make sure that older platforms within the cluster don't delay your ability to upgrade to the latest Data ONTAP release. This section addresses key considerations for platform and software support as well as platform refresh.

Platform and Software Support

If you plan to redeploy 7-Mode storage systems into the destination cluster after the data is migrated to clustered Data ONTAP, it is important to consider the following:

- **Current platform support:** When planning to repurpose an existing 7-Mode storage system, it is necessary to verify that the destination clustered Data ONTAP version supports the platform that will be repurposed.
- **Future software support:** Although it might be possible now to redeploy an existing 7-Mode storage system to clustered Data ONTAP, you need to understand how older platforms might influence the capability of the cluster to upgrade Data ONTAP releases in the future. This consideration primarily applies to platforms that are two or more generations older than the existing platform family.

Platform Refresh

Although many clusters start small and grow over time (from a node-count perspective), it is important to plan ahead for the ability to conduct platform refreshes nondisruptively within the cluster. Depending on the type of platforms deployed in the cluster, clusters can support up to 4, 8, or 24 nodes. To

nondisruptively refresh a complete HA pair within a cluster, it is necessary to keep headroom within the maximum cluster size of two nodes. Maintaining the ability to add two nodes to a cluster allows an HA pair to be nondisruptively added to a cluster, an existing HA pair evacuated to the new HA pair, and the old HA pair to be removed. This process achieves a nondisruptive platform refresh.

Aggregate relocate is a feature you can use for platform refresh activities in clustered Data ONTAP. When using aggregate relocate, tech refresh activity can be accomplished within a maximum size cluster with no need to leave headroom for temporary expansion of the cluster.

5 Performance Assessment

Performance assessment of 7-Mode and clustered Data ONTAP storage systems is critical for a successful transition. The methods described in this section can be applied for any performance assessment that is required, even outside of transition. The information that is collected during a performance assessment provides specific inputs to the task of performance sizing.

Performance assessment is broken into two activities:

- **Workload characterization:** This is the task of collecting the characteristics of the workload currently applied to the controller or node. Workload characteristics include throughput, latency, operation size, concurrency, randomness, and caching.
- **Utilization assessment:** This is the task of determining CPU and disk utilization for a controller or node.

The results of the workload characterization and utilization assessment provide the key information needed for system performance sizing. Performance sizing must take into account the existing workload and utilization footprint, as well as the intended future growth of the system or cluster.

5.1 Workload Characterization

To complete a workload characterization, you need to collect the following information about the workload on the system:

- **Throughput:** How many operations occur
- **Latency:** How fast operations are completed
- **Operation size:** The size of the operations
- **Randomness:** How random or sequential are the read operations
- **Concurrency:** How many operations occur within the latency window
- **Caching:** How much caching occurs on the system

With the exception of concurrency, the workload characteristics can be easily gathered using counter manager statistics. Collecting counter manager statistics is different between 7-Mode and clustered Data ONTAP. Conducting a workload characterization for a clustered Data ONTAP node is necessary only if you plan to move a 7-Mode workload onto a node that already hosts a preexisting workload.

7-Mode

The `stats show <object>:<instance>:<counter>` command is used in 7-Mode to access counter manager statistics. Counter manager statistics for throughput, latency, and operation size must be collected for each protocol in use on the system. For example, if the 7-Mode storage system uses NFS version 3 (NFSv3), the following counters would be collected:

- Throughput is measured in operations per second.
 - NFSv3 read operations, use the `stats show nfsv3:nfs:nfsv3_read_ops` command.
 - NFSv3 write operations, use the `stats show nfsv3:nfs:nfsv3_write_ops` command.

- Latency is measured by time in milliseconds.
 - NFSv3 read latency, use the stats `show nfsv3:nfs:nfsv3_read_latency` command.
 - NFSv3 write latency, use the stats `show nfsv3:nfs:nfsv3_write_latency` command.
- Operation size is measured in kilobytes.
 - NFSv3 read operation size, use the stats `show nfsv3:nfs:nfsv3_read_size_histo` command.
 - NFSv3 write operation size, use the stats `show nfsv3:nfs:nfsv3_write_size_histo` command.

Throughput, latency, and operation size can also be collected for CIFS and SAN protocols by using the appropriate objects, instances, and counters.

Readahead counters track the percentage of random and sequential read operations that occur. The `stats show readahead:readahead:seq_read_req` command displays the percentage of sequential read requests. The `stats show readahead:readahead:rand_read_req` command displays the percentage of random read requests. It is not necessary to determine how many random writes occur because of how write operations are handled by Data ONTAP.

Note: Data ONTAP uses nonvolatile memory to log write operations. Writes are acknowledged once logged (being persistent in memory at that point) and then sent to disk at a later time. This approach allows Data ONTAP to order data in memory sequentially, turning what are random writes from a client into sequential writes when sent to disk. A solid understanding of how data flows through a NetApp storage system is critical to any performance assessment and sizing activity. For additional information on NetApp storage system performance, see [TR-4211: NetApp Storage System Performance Primer](#).

Concurrency is not tracked as a count but rather is calculated. To determine concurrency, multiply throughput by latency. The resulting value is the number of operations occurring within the latency window. Concurrency is calculated separately for read and write operations.

Caching can occur in systems memory through the use of Flash Cache or by using a Flash Pool aggregate. Use the `stats show waf1:waf1:read_io_type` command to quickly determine where data is being cached within the system. For each technology component in the system that can service reads, a percentage will be displayed that tells you which technologies are servicing the read operations on the system. The following technologies are reported:

- `waf1:waf1:read_io_type.cache`: Read operations serviced by system memory
- `waf1:waf1:read_io_type.ext_cache`: Read operations serviced by Flash Cache
- `waf1:waf1:read_io_type.disk`: Read operations serviced by HDD
- `waf1:waf1:read_io_type.bamboo_ssd`: Read operations serviced by SSD (all SSD aggregates)
- `waf1:waf1:read_io_type.hya_hdd`: Read operations serviced by Flash Pool HDD
- `waf1:waf1:read_io_type.hya_cache`: Read operations serviced by Flash Pool SSD cache

Clustered Data ONTAP

In clustered Data ONTAP it is necessary to start the collection of statistics using the `statistics start -node <node_name> -object <object>` command to start the collection of statistics. Specifying the command argument `-node` allows you to collect statistics for a specific node in the cluster. The “system” object will allow a wide range of counters to be collected for the node. Once started, statistics will be collected until the `statistics stop` command is executed.

The same application of object, instance, and counter apply to clustered Data ONTAP (just as in 7-Mode). The range of statistics available in clustered Data ONTAP is much more robust than the statistics available in 7-Mode. As a result, the objects, instances, and counters specified will be different from those used for 7-Mode. Clustered Data ONTAP provides the ability to see which objects, instances, and counters are available:

- **Objects:** The command `statistics show -object [TAB]` will list all available objects. Pressing the tab key after specifying the `-object` command argument triggers the listing.
- **Instances:** The command `statistics show -object <object> -instance [TAB]` will list all available instances for the specified object. Pressing the tab key after specifying the `-object <object> -instance` command arguments triggers the listing.
- **Counters:** The command `statistics show -object <object> -instance <instance> -counter [TAB]` will list all available counters for the specified instance and object. Pressing the tab key after specifying the `-object <object> -instance <instance> -counter` command arguments triggers the listing.

The calculation of concurrency for clustered Data ONTAP is the same as for 7-Mode.

Readahead counters are available in clustered Data ONTAP; however, they cannot be accessed with the basic administrative-level privileges. Contact NetApp support for direction on accessing the correct privilege level for the readahead counters. At the correct privilege level, the `statistics show -object readahead -instance readahead -counter rand_read_reqs` command can be used to access the random read counter information. Additionally, the `statistics show -object readahead -instance readahead -counter seq_read_reqs` command can be used to access the sequential read counter information.

NetApp WAFL® (Write Anywhere File Layout) counters are available in clustered Data ONTAP; however, they cannot be accessed with the basic administrative-level privileges. Contact NetApp support for direction on accessing the correct privilege level for the WAFL counters. At the correct privilege level, the `statistics show -object waf1 -instance waf1 -counter read_io_type` command can be used to access the caching counter information.

5.2 Utilization Assessment

A utilization assessment is meant to provide a view into the current CPU and disk utilization occurring on the system. Utilization should be measured during regular working hours/peak workload times because it is during these periods that CPU and disk utilization are most representative of the workload.

The `sysstat` command quickly measures both CPU and disk utilization. The `-x` command argument must be used to see disk utilization information in the command output. The command output will contain several columns; however, the CPU and Disk `util` columns are the specific columns to observe.

Utilization Reporting

When observing CPU utilization in Data ONTAP, it is critical to understand two key concepts:

- NetApp storage systems contain processors with multiple cores.
- Data ONTAP uses free CPU cycles to conduct background operations.

NetApp storage systems all contain processors with multiple cores. Processes occurring in Data ONTAP are not evenly distributed across all available cores, which means that some cores may be busier than others. When Data ONTAP reports CPU utilization with the `sysstat` command, it takes the highest observed CPU utilization percentage observed across all cores.

Data ONTAP uses idle system resources (including free CPU cycles) to conduct background operations. Scrubs and block reclamation are examples of background activity that use the free cycles. If an increase in workload demands access to cycles used by background operations, the cycles will be released to service the workload. As a result of background operations, reported CPU utilization can be misleading because it does not reflect just the utilization for the workload.

Disk utilization is similar to CPU utilization in that the percentage reported is based on the busiest disk in the system at the time of the observation. In most cases the busiest disk is associated to the workload on the system; however, it is possible that other system activities can result in high disk utilization being

reported. The most obvious case for high disk utilization unrelated to the workload on the system is a RAID reconstruction as the result of a drive failure.

7-Mode

In 7-Mode, use the `sysstat -x` command to observe CPU and disk utilization. Optionally you can provide a number after the command arguments to be used as the reporting interval (time in seconds).

Clustered Data ONTAP

In clustered Data ONTAP, use the `system node run -node <nodename> -command sysstat -x` command (or run the command from the node shell) to observe CPU and disk utilization. Optionally you can provide a number after the command arguments to use as the reporting interval (time in seconds).

5.3 Sizing Considerations

NetApp recommends that customers work with their NetApp or partner representatives to properly size storage systems and clusters. The information collected during the utilization assessment will be key in understanding the current workload, system utilization, and, if a preexisting cluster is in place, the headroom available on the destination node and cluster.

A few key considerations when sizing clusters are as follows.

- An accurate accounting of the existing 7-Mode workloads will establish a minimum sizing requirement for the destination cluster.
- If the existing 7-Mode workload has outgrown the 7-Mode storage system, it is necessary to plan additional resources for the destination cluster nodes to account for the increased workload requirements. General sizing considerations for this scenario are:
 - If high disk utilization is observed in 7-Mode:
 - Verify that it is not related to background operations and is a result of the workload.
 - Use a system with a larger amount of available memory.
 - Use a greater number of drives or use faster drives.
 - If high CPU utilization is observed in 7-Mode:
 - Use a system with a larger amount of available memory.
 - Use a system with larger or more advanced CPUs.
- Future cluster growth must be planned to cover new workloads that might be added and existing workloads that will likely grow over time.
- Workloads consisting of many small block random reads are excellent candidates for caching technologies.
- Cache hit rates observed in 7-Mode will be roughly equivalent to those in clustered Data ONTAP.
- Workloads moving to indirect data access might not yield equivalent throughput and latency as is observed on 7-Mode using direct data access. Any workload in clustered Data ONTAP can use direct data access depending on how clients are connected and the location of LIFs. Some workloads will only ever use direct data access such as SAN protocols using Asymmetric Logical Unit Access (ALUA).
- Background workloads such as SnapMirror transfer rates and performance, deduplication schedules, and more must be considered.


6 Disaster Recovery and Backup Migration Scenarios

This section addresses the recommended migration approaches for the most common disaster recovery (DR) and backup scenarios deployed with NetApp storage systems today. The 10 most common DR and backup scenarios are basic disaster recovery, bidirectional disaster recovery, basic backup, disaster

recovery and backup cascade, disaster recovery and backup fan-out, disaster recovery fan-out, SnapVault to disaster recovery cascade, SnapVault fan-in, disaster recovery cascade, and qtree replication.

Smaller NetApp environments might contain only one or two of the scenarios addressed in this section, whereas larger NetApp environments will likely involve several of the scenarios outlined in this section. A common disaster recovery and backup topology for medium and larger NetApp environments is the hub-spoke topology that is shown in Figure 7.

Figure 7) Hub-spoke topology for disaster recovery and backup relationships.




When migrating NetApp environments that use the hub-spoke topology, NetApp recommends migrating spoke sites and the destination volumes at the hub site in pairs. Additionally, cascaded volumes between hub sites will need to be included in the spoke and hub pair migrations. In short, divide and conquer.

6.1 Scenario 1: Basic Disaster Recovery


The basic disaster recovery scenario addresses the most common case of a single source and destination volume that are in a volume SnapMirror relationship. Figure 8 displays an overview of the 7-Mode configuration.

Figure 8) Basic disaster recovery in 7-Mode.



Migration of the volumes, associated Snapshot copies, and the SnapMirror relationship is easily accomplished using the 7MTT CBT (recommended) or by using a manual TDP SnapMirror. In this approach, the secondary volume moves to clustered Data ONTAP first, followed by the primary volume. This approach allows the one-way volume SnapMirror relationship to stay in place before migrating and

Figure 9) Basic disaster recovery in clustered Data ONTAP.



cutting over the primary volume. Figure 9 shows the final clustered Data ONTAP configuration.


The largest benefit of this approach is that once the primary and secondary volumes move to clustered Data ONTAP, the SnapMirror relationship can be resynchronized. The relationship can be resynchronized because both the source and destination volumes share a common Snapshot copy.

Note: 7MTT CFT also preserves SnapShot copies and SnapMirror relationships. Once data migration project completes, the SnapMirror relationship can be resynchronized

6.2 Scenario 2: Bidirectional Disaster Recovery

The bidirectional disaster recovery scenario sees two controllers, each hosting a secondary volume that is the destination of a primary volume located on the other controller. Figure 10 displays an overview of the 7-Mode configuration.


Figure 10) Bidirectional disaster recovery in 7-Mode.



Migration of the volumes, associated Snapshot copies, and the SnapMirror relationship is easily accomplished using the 7MTT (recommended) or by using a manual TDP SnapMirror. Although the relationships between the controllers are bidirectional, from a volume-level perspective the relationships are independent (and can move independently).

In this approach, the secondary volume moves to clustered Data ONTAP first, followed by the primary

Figure 11) Bidirectional disaster recovery in clustered Data ONTAP.




volume. This allows the one-way volume SnapMirror relationship to stay in place before migrating and cutting over the primary volume. A single project in the 7MTT can accomplish the migration of both volumes, SnapMirror relationships, and their associated Snapshot copies. Figure 12 displays an overview of the final clustered Data ONTAP configuration.

The largest benefit of this approach is that once the primary and secondary volumes are moved to clustered Data ONTAP, the SnapMirror relationship can be resynchronized. The relationship can be resynchronized because both the source and destination volumes share a common Snapshot copy, meaning a re-baseline of the relationship is not needed.

6.3 Scenario 3: Basic Backup

The basic backup scenario addresses the case of a single source and destination volume that are in a SnapVault backup relationship. Figure 12 displays an overview of the 7-Mode configuration.

Figure 12) Basic backup in 7-Mode.



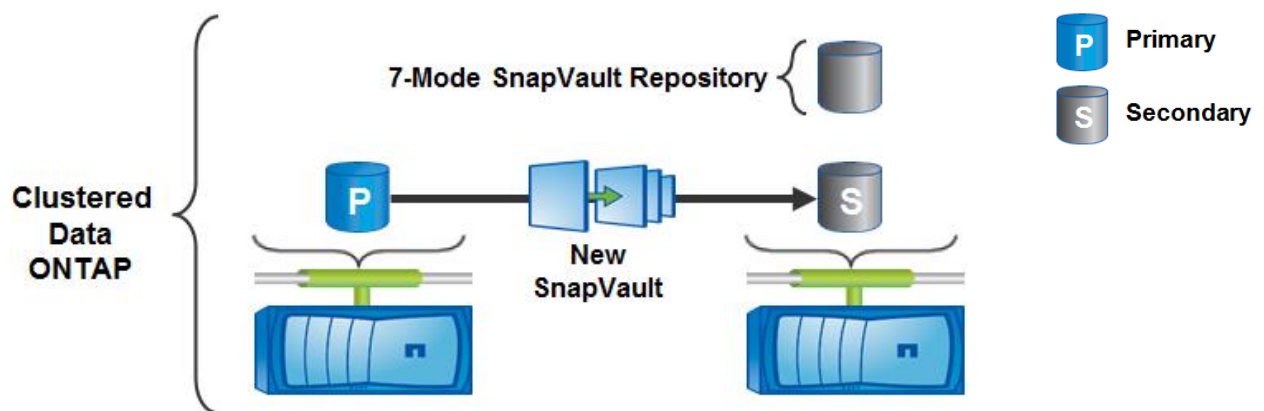
SnapVault in 7-Mode is qtree based whereas in clustered Data ONTAP SnapVault is volume based. As a result, it is necessary to create a new SnapVault relationship in clustered Data ONTAP and determine the best course of action for the 7-Mode SnapVault repository.

The primary volume can be migrated normally using the 7MTT or a manual TDP SnapMirror relationship. Movement of the secondary volume depends on the retention period for the repository.

- If the retention period is greater than three months, the repository should be migrated to a clustered Data ONTAP volume (not the secondary volume for the new SnapVault relationship in clustered Data ONTAP).
- If the retention period is less than or equal to three months, maintain the 7-Mode repository for the retention period, then retire the 7-Mode repository.

Regardless of the option used, if data must be restored from the 7-Mode repository, it will be necessary to mount the Snapshot copy containing the data to restore. Then it will be necessary to copy the data from the mounted Snapshot copy to the restore location. Figure 13 displays an overview of the final clustered Data ONTAP configuration.

Figure 13) Basic backup in clustered Data ONTAP.




Movement of the primary volume to clustered Data ONTAP and the establishment of the new SnapVault relationship does not depend on the 7-Mode secondary volume (because the clustered Data ONTAP SnapVault is a brand-new relationship). The new SnapVault relationship in clustered Data ONTAP requires a baseline transfer to be completed.

6.4 Scenario 4: Disaster Recovery and Backup Cascade

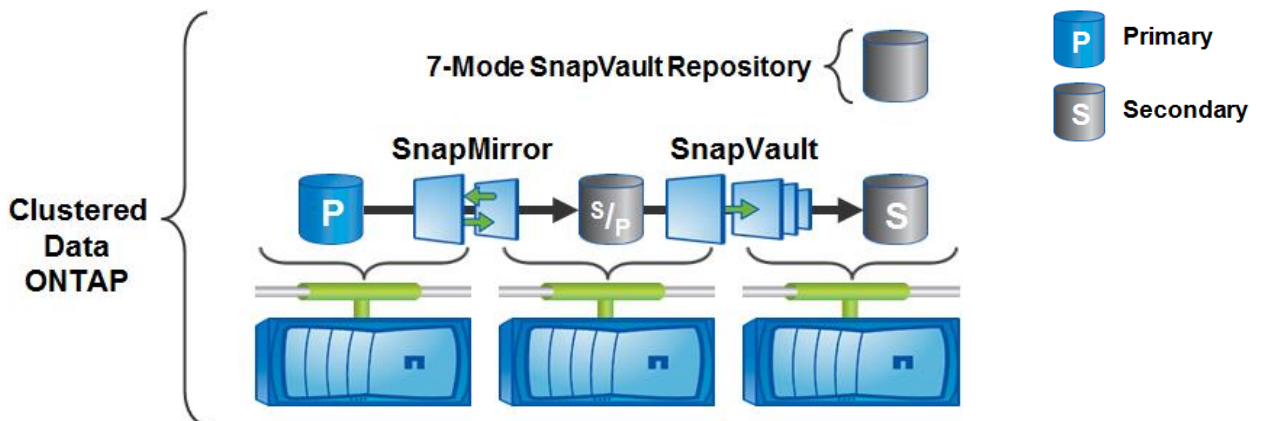
The disaster recovery and backup cascade scenario addresses the case of a primary to secondary volume SnapMirror relationship that then cascades into a SnapVault relationship. The secondary volume of the SnapMirror relationship is the primary volume for the SnapVault relationship that ends with a separate SnapVault destination volume. This scenario can span up to three 7-Mode controllers. Figure 15 displays an overview of the 7-Mode configuration.

Figure 14) Disaster recovery and backup cascade in 7-Mode.



The approach to migration in this scenario starts by breaking the SnapVault relationship. Next the 7MTT or manual TDP SnapMirror is used to migrate the primary and secondary SnapMirror volumes (the same as for [Scenario 1: Basic Disaster Recovery](#)). Once the SnapMirror volumes and relationship are resynchronized in clustered Data ONTAP, the new SnapVault relationship can be established using the SnapMirror secondary volume as the primary volume for the SnapVault relationship. The SnapVault repository (7-Mode SnapVault secondary) on 7-Mode can then be migrated or left in place in 7-Mode depending on the retention period (see [Scenario 3: Basic Backup](#)). Figure 15 displays an overview of the final clustered Data ONTAP configuration.

Figure 15) Disaster recovery and backup cascade on clustered Data ONTAP.




The new SnapVault relationship in clustered Data ONTAP requires a baseline transfer to be completed. If data must be restored from the 7-Mode SnapVault repository, it will be necessary to mount the Snapshot copy containing the data to restore and then copy the data from the mounted Snapshot copy to the restore location.

6.5 Scenario 5: Disaster Recovery and Backup Fan-Out

The disaster recovery and backup fan-out scenario addresses the case in which both a volume SnapMirror and a SnapVault relationship are in place using the same primary volume (and separate secondary volumes). Figure 16 displays an overview of the 7-Mode configuration.

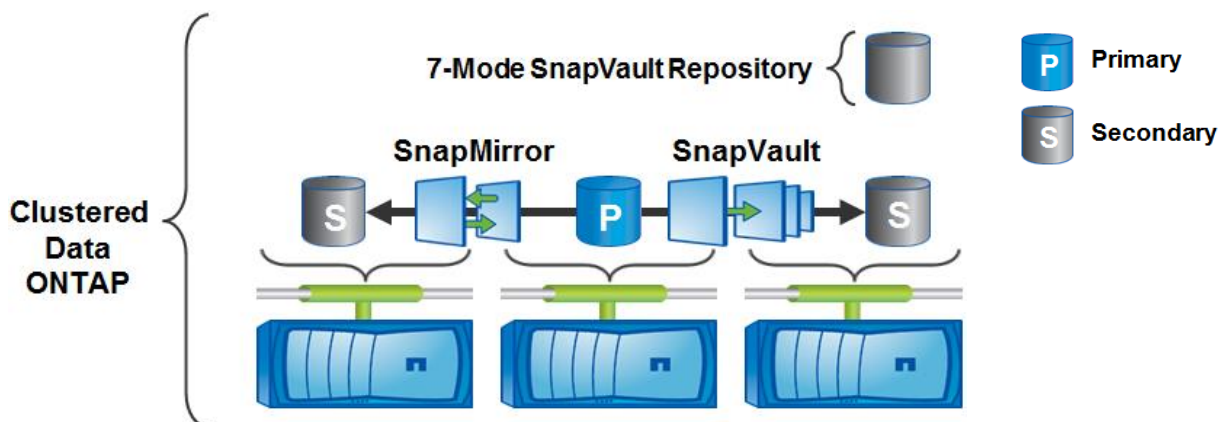
Figure 16) Disaster recovery and backup fan-out in 7-Mode.



The approach to migration in this scenario starts by breaking the SnapVault relationship. Next the 7MTT or manual TDP SnapMirror is used to migrate the primary and secondary SnapMirror volumes (the same as for [Scenario 1: Basic Disaster Recovery](#)). Once the SnapMirror volumes and relationship have been resynchronized in clustered Data ONTAP, the new SnapVault relationship can be established using the SnapMirror secondary volume as the primary volume for the SnapVault relationship. The SnapVault repository (7-Mode SnapVault secondary) on 7-Mode can then be migrated or left in place in 7-Mode depending on the retention period (see [Scenario 3: Basic Backup](#)). Figure 17 displays an overview of the final clustered Data ONTAP configuration.

The new SnapVault relationship in clustered Data ONTAP requires a baseline transfer to be completed. If data must be restored from the 7-Mode SnapVault repository, it will be necessary to mount the Snapshot copy containing the data to restore. Then it will be necessary to copy the data from the mounted Snapshot copy to the restore location.


Figure 17) Disaster recovery and backup fan-out in clustered Data ONTAP.



6.6 Scenario 6: Disaster Recovery Fan-Out

The disaster recovery fan-out scenario addresses the case in which two volume SnapMirror relationships are in place using the same primary volume (and separate secondary volumes). Figure 18 displays an overview of the 7-Mode configuration.


Figure 18) Disaster recovery fan-out in 7-Mode.



The approach to migration in this scenario starts by using manual TDP SnapMirror to migrate one of the secondary volumes to clustered Data ONTAP. After the migration of the secondary volume is complete, manually establish a TDP SnapMirror relationship between the primary volume in 7-Mode and the migrated secondary volume in clustered Data ONTAP. Use the 7MTT CBT to migrate the remaining secondary and primary volume to clustered Data ONTAP. As part of the cutover procedure for the primary volume, break the manual TDP SnapMirror relationship between the primary and manually migrated secondary volume. Once the primary volume has been moved to clustered Data ONTAP and the 7MTT has resynchronized the SnapMirror relationship for the volumes it migrated, manually resynchronize the SnapMirror relationship between the primary and manually migrated secondary volume. Figure 19 displays an overview of the final clustered Data ONTAP configuration.

Note: To resynchronize both SnapMirror relationships in clustered Data ONTAP both secondary volumes must share a common Snapshot copy with the primary volume. As a result, it is critical that the manually established TDP SnapMirror relationship not be broken until the primary volume is cut over (the 7MTT will conduct the final update).

Figure 19) Disaster recovery fan-out in clustered Data ONTAP.

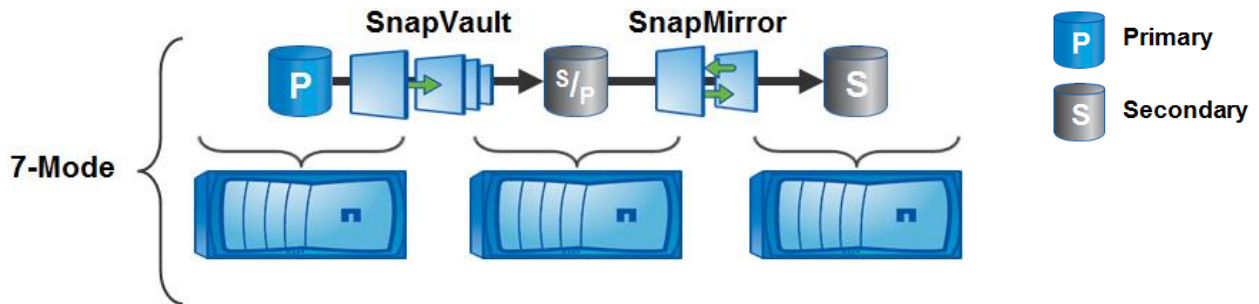


Because both SnapMirror relationships can be moved to clustered Data ONTAP while retaining Snapshot copies, there is no need to rebaseline the relationships (they are resynchronized). As a result, both SnapMirror relationships in clustered Data ONTAP can be reestablished quickly.

6.7 Scenario 7: SnapVault-to-Disaster Recovery Cascade

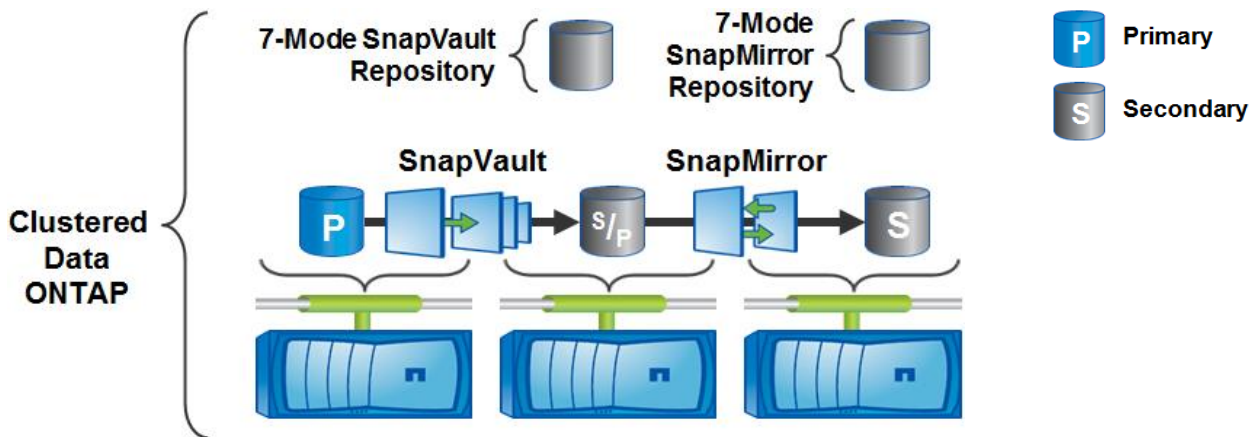
The SnapVault-to-disaster recovery cascade scenario addresses the case in which there is a SnapVault primary and secondary and the SnapVault secondary is the primary volume for the SnapMirror relationship (which uses a separate secondary volume). Figure 20 displays an overview of the 7-Mode configuration.

Figure 20) SnapVault-to-disaster recovery cascade in 7-Mode.



This approach deals directly with the fact that the SnapVault secondary as well as the SnapMirror primary and secondary volumes will contain data in Snapshot copies that cannot be directly restored in clustered Data ONTAP. The first step is to use the 7MTT or manual TDP SnapMirror to migrate the primary SnapVault volume to clustered Data ONTAP. Before cutover of the primary SnapVault volume it is necessary to break the SnapVault relationship. Once the primary SnapVault volume is on clustered Data ONTAP, create a new SnapVault relationship (using a new destination volume). Once the SnapVault relationship is established, a new SnapMirror relationship can be created between the SnapVault secondary volume and a new SnapMirror destination volume. Both the SnapVault and SnapMirror secondary volumes in 7-Mode will be subject to the retention approach outlined in [Scenario 3: Basic Backup](#) (being retained on 7-Mode or moved to a separate clustered Data ONTAP volume based on the retention period). It is common for Snapshot copy retention cycles to be short (weeks or at most a few months), which means that it is likely easier to allow the SnapMirror secondary volume to stay on 7-Mode until the Snapshot copies age out. Figure 21 displays an overview of the final clustered Data ONTAP configuration.

Figure 21) SnapVault-to-disaster recovery cascade in clustered Data ONTAP.




Both the new SnapVault relationship and the new SnapMirror relationship in clustered Data ONTAP require a baseline transfer to complete in order to establish new relationships. If data must be restored from the 7-Mode SnapVault or SnapMirror repositories, it is necessary to mount the Snapshot copy

containing the data to restore. Then it is necessary to copy the data from the mounted Snapshot copy to the restore location.

6.8 Scenario 8: SnapVault Fan-In


The SnapVault fan-in scenario addresses the case in which two SnapVault primary volumes use the same destination volume. Figure 22 displays an overview of the 7-Mode configuration.

Figure 22) SnapVault fan-in in 7-Mode.



The 7MTT or manual TDP SnapMirror will be used to migrate both primary volumes. Before cutover of the primary volumes, the SnapVault relationships must be broken. As each primary volume completes the cutover to clustered Data ONTAP, a new SnapVault relationship is created (no longer using the same destination volume). The SnapVault repository (7-Mode SnapVault secondary) on 7-Mode can then be migrated or left in place in 7-Mode depending on the retention period (see [Scenario 3: Basic Backup](#)). Figure 23 displays an overview of the final clustered Data ONTAP configuration.

Figure 23) SnapVault fan-in in clustered Data ONTAP.

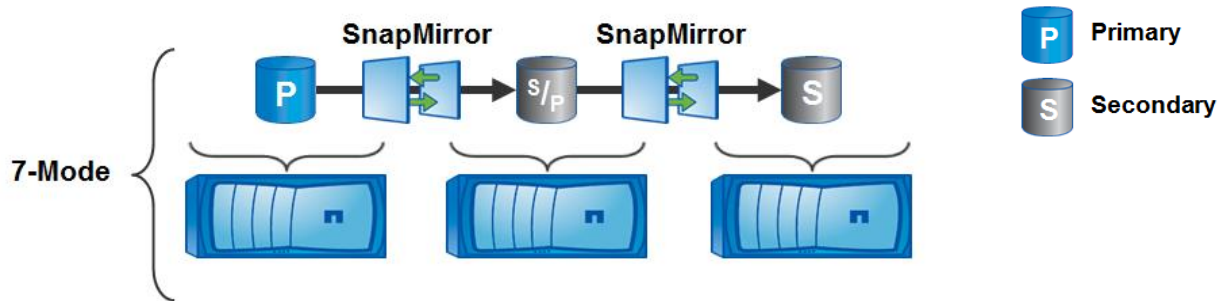


Both new SnapVault relationships in clustered Data ONTAP require a baseline transfer to complete to establish new relationships. Additionally, separate volumes must be used as the destinations (a result of SnapVault being volume based in clustered Data ONTAP versus qtree based in 7-Mode). If data must be restored from the 7-Mode SnapVault repository, it is necessary to mount the Snapshot copy containing the data to restore and then copy the data from the mounted Snapshot copy to the restore location.

6.9 Scenario 9: Disaster Recovery Cascade


The disaster recovery cascade scenario addresses the case in which there are two SnapMirror relationships and the secondary volume of one relationship is the primary volume of the other relationship. Figure 24 displays an overview of the 7-Mode configuration.

Figure 24) Disaster recovery cascade in 7-Mode.



This approach starts by using a manual TDP SnapMirror to migrate the secondary volume of the SnapMirror relationship whose primary volume is the secondary of the other SnapMirror relationship (the far-right secondary volume shown in Figure 24). Once the secondary is cut over to clustered Data ONTAP, create a manual TDP SnapMirror relationship between the primary volume on 7-Mode (which is also the secondary of the other SnapMirror relationship). Use the 7MTT to migrate the secondary and primary volumes of the SnapMirror relationship whose secondary volume is the primary volume of the other SnapMirror relationship (the left-hand primary and middle secondary/primary volumes in Figure 24). Break the manual TDP SnapMirror relationship before cutting over the secondary volume of the 7MTT migrated volumes. After the “secondary/primary” volume is cut over, immediately cut over the remaining primary volume to clustered Data ONTAP. Once all volumes have been cut over to clustered Data ONTAP, the 7MTT will reestablish the SnapMirror relationship for its volumes. It will be necessary to manually resynchronize the second SnapMirror relationship at this point. Figure 25 displays an overview of the final clustered Data ONTAP configuration.

Figure 25) Disaster recovery cascade in clustered Data ONTAP.



Because both SnapMirror relationships can be moved to clustered Data ONTAP while retaining Snapshot copies, there is no need to rebaseline the relationships (they are resynchronized). As a result, both SnapMirror relationships in clustered Data ONTAP can be reestablished quickly.

6.10 Scenario 10: Qtree Replication

Qtree replication is included as a section here for completeness. Given the considerations for qtrees during a transition from 7-Mode to clustered Data ONTAP, a dedicated section exists in this document to discuss the various qtree migration scenarios. As a result, the qtree replication scenario was broken into the three scenarios that are addressed in the [Qtree Migration Scenarios](#) section of this document.

7 Qtree Migration

In 7-Mode, qtrees can be used in one of three ways: quotas, exports, or replication. In clustered Data ONTAP, qtrees can be used for quotas and exports, but not for replication (replication is volume based in clustered Data ONTAP).

How qtrees will be moved to clustered Data ONTAP will depend not only on how you use them currently in 7-Mode, but also on what your desired organizational model is in clustered Data ONTAP. With these considerations in mind, this section addresses three scenarios for migration.

Qtree-to-qtree migration: In this scenario, volumes containing qtrees are migrated directly to clustered Data ONTAP, resulting in the same organizational structure. Each qtree in 7-Mode translates to a single qtree in clustered Data ONTAP (a 1:1 ratio of qtrees to qtrees).

Note: CFT uses this 1:1 ratio; no re-organization or consolidation of qtrees is possible when using CFT.

- **Qtree-to-volume migration:** In this scenario, each individual source qtree is migrated into a single, separate volume in clustered Data ONTAP. Each qtree in 7-Mode translates to a single volume in clustered Data ONTAP (a 1:1 ratio of qtrees to volumes).
- **Qtree consolidation:** In this scenario, qtrees with common requirements are grouped together (most likely because of replication requirements) and reorganized during the migration activity. Each group of qtrees in 7-Mode translates to a single volume in clustered Data ONTAP (a many:1 ratio of qtrees to volumes).

7.1 Should Qtrees Be Reorganized?

The most fundamental question for qtrees and transition is, “Should I reorganize my qtrees?”

Qtree reorganization refers to any movement of qtrees to clustered Data ONTAP in which the destination qtree structure will be different from the source qtree structure (qtree-to-volume or qtree consolidation scenarios, discussed later in this section). There are key considerations that will drive your decision to reorganize qtrees. Before discussing them, however, here are two common cases that result in a very specific NetApp recommendation:

- **Qtrees for quotas only:** There are no practical reasons for reorganizing qtrees if they are used only for quotas. NetApp recommends using the [qtree-to-qtree migration](#) approach in this case.
- **Qtree replication:** If qtree SnapMirror is used (required for 7-Mode SnapVault) it will most commonly be necessary to reorganize those qtrees based on the replication requirements for each qtree. A second-level consideration here is how many qtrees are in use on 7-Mode. If tens of qtrees to a few hundred qtrees are in use (the most common case), NetApp recommends the [qtree-to-volume migration](#) approach (because replication in clustered Data ONTAP is volume based). If many hundreds to thousands of qtrees are in use (certainly if maximum volume counts would be exceeded in clustered Data ONTAP), it is best to use the [qtree consolidation](#) approach.

Note: Qtrees using exports are excluded above because they can result in multiple recommendations depending on the qtree configuration being moved.

The general considerations that will drive the decision to reorganize qtrees or maintain the existing organization are as follows:


- **Data Migration:** If a copy-based migration method is used (not a SnapMirror based method), each qtree can be copied directly to the destination volume (one-step reorganization). If a replication-based migration method is used, additional steps are required to achieve qtree reorganization. For additional information on replication-based migration of qtrees, see the [qtree-to-volume migration](#) and [qtree consolidation](#) sections of this document.
- **Quality of Service:** Quality of service in clustered Data ONTAP can be applied at the SVM, volume, or file granularity level (not qtree granularity). If the qtrees in a single volume span multiple workloads, there could be difficulty when attempting to isolate a rogue workload or when setting throttles for specific storage containers.
- **DataMotion for Volumes:** Nondisruptive movement of volumes is a major feature and benefit of clustered Data ONTAP. If you encounter a situation that requires qtrees within the same volume to move to different locations within the cluster, you will be unable to nondisruptively move those qtrees separately. That is because DataMotion for Volumes moves volumes and not individual qtrees.
- **Maximum Volumes Supported:** If several hundred or several thousand qtrees are being moved to the same destination cluster it might not be possible to convert each qtree into a separate volume (because the maximum number of supported volumes in the destination cluster). If this is the case, the options are to use the [qtree-to-qtree migration](#) or [qtree consolidation](#) scenarios.
- **Qtree Dataset Requirements:** If qtrees are not required for future growth of the dataset that is currently organized into qtrees on 7-Mode, the recommended approach is to use the [qtree-to-qtree migration](#) scenario. After moving the dataset to clustered Data ONTAP, switch to a volume-provisioning approach for all new or incremental dataset requirements.

There is no single, all-encompassing correct answer to the question of whether or not qtrees should be reorganized because it will vary on a case-by-case basis. The considerations outlined in this section are provided so that an informed decision can be made during the transition planning activities and when addressing cluster design and architecture.

7.2 Scenario 1: Qtree-to-Qtree Migration Using CBT

This scenario addresses the migration of qtrees to clustered Data ONTAP in which the source 7-Mode qtree structure is retained on the destination clustered Data ONTAP system. NetApp recommends using the 7MTT to migrate the volume in this scenario. Figure 26 displays an overview of the migration scenario.

Figure 26) Qtree-to-qtree migration overview—CBT.



This is the simplest approach to qtree migration because it requires only movement of the source volume, along with its qtrees to clustered Data ONTAP. No additional planning is necessary to determine how qtrees should be reorganized on the destination clustered Data ONTAP system. Although volume SnapMirror is represented in Figure 26, it is possible to use any migration method to move the data to clustered Data ONTAP.

Note: NetApp recommends using the 7MTT as a replication-based migration method because it will retain storage efficiency savings and Snapshot copies at the destination clustered Data ONTAP system.

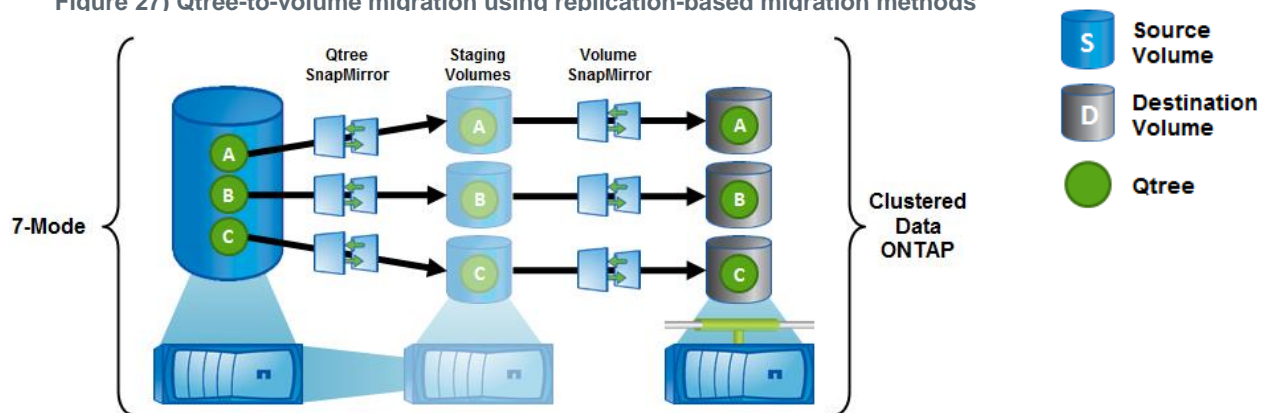
7.3 Scenario 3: Qtree-to-Volume Migration

This scenario addresses the migration of qtrees to clustered Data ONTAP in which each source 7-Mode qtree is migrated to an individual clustered Data ONTAP volume. Practically, this scenario is best used if tens of qtrees to a couple of hundred qtrees are being moved. The maximum number of supported volumes on the destination clustered Data ONTAP system can easily accommodate several hundred qtrees being moved into individual volumes. That being said, it is important to leave headroom for expansion in the destination cluster. Both replication-based and copy-based migration methods can be used in this scenario.

If a replication-based approach is used, storage efficiency savings and Snapshot copies will not be retained at the destination clustered Data ONTAP system. As a result, storage efficiency savings will have to be reapplied to the destination volumes and Snapshot copies will have to be retained on 7-Mode until the retention period expires.


Before qtrees can be migrated to clustered Data ONTAP volumes, it is necessary to break the source 7-Mode qtrees into individual volumes (while still on 7-Mode). This approach is known as the qtree staging approach. A number of staging volumes are required in order to break qtrees into separate volumes. The staging volumes can be located on the same 7-Mode controller or on a different 7-Mode controller, wherever space exists that can be used for the break-out activity. Figure 27 displays an overview of the qtree staging approach using a replication-based migration method for this scenario.

Figure 27) Qtree-to-volume migration using replication-based migration methods



If a copy-based approach is used, qtrees can be directly copied to the destination clustered Data ONTAP volumes. Although the movement of data is easier, storage efficiency savings and Snapshot copies will not be retained through the migration. Figure 28 displays an overview of the qtree staging approach using a copy-based migration method for this scenario.

Figure 28) Qtree-to-volume migration using copy-based migration methods




7.4 Scenario 4: Qtree Consolidation

This scenario addresses the migration of qtrees to clustered Data ONTAP in which source 7-Mode qtrees with common dataset requirements (for example, replication or exports) are grouped (consolidated) together and moved into volumes. Both replication-based and copy-based migration methods can be used in this scenario.


If a replication-based approach is used, storage efficiency savings and Snapshot copies will not be retained at the destination clustered Data ONTAP system. As a result, storage efficiency savings will have to be reapplied to the destination volumes and Snapshot copies will have to be retained on 7-Mode until the retention period expires. The qtree staging approach (as discussed in the [qtree-to-volume migration](#) scenario) will be required to move qtrees to clustered Data ONTAP. Figure 29 displays an overview of the qtree staging approach using a replication-based migration method for this scenario.

Figure 29) Qtree consolidation using replication based migration methods



If a copy-based approach is used, qtrees can be directly copied to the destination clustered Data ONTAP volumes. Although the movement of data is easier, storage efficiency savings and Snapshot copies will not be retained through the migration. As a result, storage efficiency savings will have to be reapplied to the destination volumes and Snapshot copies will have to be retained on 7-Mode until the retention period expires. Figure 30 displays an overview of the qtree staging approach using a copy-based migration method for this scenario.

Figure 30) Qtree consolidation using copy-based migration methods



7.5 Additional Qtree Considerations

This section addresses qtree planning considerations related to the migration of qtrees to clustered Data ONTAP. The additional considerations cover qtree SnapMirror relationships, support for qtree exports in clustered Data ONTAP, changes to export paths, and qtree 0 data.

Qtree SnapMirror Relationships

Qtree SnapMirror relationships on 7-Mode must become volume SnapMirror relationships in clustered Data ONTAP. As a result, there is no need to migrate qtree secondary volumes to clustered Data ONTAP because a new volume SnapMirror relationship will be created.

Qtree Export Support

Support for qtree-level exports was introduced in clustered Data ONTAP 8.2.1. If qtrees using exports are migrated to clustered Data ONTAP 8.2.0 (or earlier), exports will be inherited from the containing volume. If qtree exports are used on 7-Mode and qtrees will be retained in clustered Data ONTAP, NetApp recommends using clustered Data ONTAP 8.2.1 or later.

Note: NFSv4 qtree exports are supported with clustered Data ONTAP 8.3 and later.

NFS Export Paths

When qtrees are reorganized using replication-based migration methods, a single qtree still resides in the volume. The export path will change after the volume is migrated to clustered Data ONTAP. This means that either the new export path is used or a symlink can be created to preserve the original export path.

Qtree 0 (Zero)

Data that resides in a volume that uses qtrees but is not contained within a qtree itself is referred to as being in qtree 0 (zero). Qtree 0 is a reserved container in each volume that is used to contain the volume data not contained within existing user-defined qtrees. When migrating volumes that contain qtrees to clustered Data ONTAP using the qtree-to-volume or qtree consolidation scenario, make sure that data contained within qtree 0 is not overlooked during the migration.

8 Product and Feature Considerations

This section addresses the feature and functional differences between 7-Mode and clustered Data ONTAP. Given the architectural differences between them, it is expected that certain features might be implemented differently in clustered Data ONTAP. As a result, some features are configured or implemented differently, other features are replaced by improved features, and in many cases the features are the same. In rare cases, the feature in 7-Mode is not represented in clustered Data ONTAP (which might be intentional).

8.1 WAFL

The storage subsystem features of Data ONTAP function nearly the same way in clustered Data ONTAP as they do in 7-Mode. That being said, there are a few differences, such as volume name lengths, Snapshot copies, fractional reserve, traditional volumes, and 32-bit aggregates. These differences are discussed in more detail in the following sections.

Volume Name Length

In 7-Mode, the maximum volume name length is 255 characters. In clustered Data ONTAP the maximum volume name length is 203 characters. This difference is a consideration for volumes that are transitioning that use more than 203 characters.

Snapshot Copies

Snapshot copies that have been replicated (using SnapMirror) to a clustered Data ONTAP volume from a 7-Mode volume are considered migrated Snapshot copies. Starting in clustered Data ONTAP 8.2, migrated Snapshot copies support NetApp SnapRestore® technology (file and volume).

Using SnapRestore to restore migrated Snapshot copies is not supported if the Snapshot copy contains qtree SnapMirror or LUN data. Additionally, NetApp FlexClone® volumes cannot be created from Snapshot copies.

Fractional Reserve

Before Data ONTAP 8.2, the fractional reserve setting for a flexible volume could be set to any value from 0 through 100. Starting with Data ONTAP 8.2, the fractional reserve setting can be set to either 0 or 100.

- For VOLUME guaranteed volumes, the default value is 100.
- For NONE guaranteed volumes, the default value is 0.
- For FILE guaranteed volumes, the default value is 100 (cannot be modified).

If you upgrade a system with VOLUME or NONE guaranteed volumes using fractional reserve settings from 1 through 99 to Data ONTAP 8.2, the fractional reserve setting for those volumes is set to 0. For more information about fractional reserve, see the [Clustered Data ONTAP 8.2 Logical Storage Management Guide](#).

Traditional Volumes

Traditional volumes are not supported in clustered Data ONTAP. Traditional volumes are an older volume type only present in Data ONTAP 7G. Any Data ONTAP 7G systems using traditional volumes must be migrated to flexible volumes in clustered Data ONTAP.

32-Bit Aggregates

All Data ONTAP 7G storage systems and some Data ONTAP 8.x 7-Mode storage systems will use 32-bit aggregates. In clustered Data ONTAP, NetApp recommends using 64-bit aggregates. It is possible to copy data directly to a 64-bit aggregate in clustered Data ONTAP; however, if SnapMirror is used to migrate data, 32-bit Snapshot copies will also be retained. If a 32-bit Snapshot copy is restored, the aggregate will revert to a 32-bit aggregate. In Data ONTAP 8.1.x 7-Mode and later, it is possible to convert a 32-bit aggregate into a 64-bit aggregate before data movement (either through expansion, or in Data ONTAP 8.2.1 7-Mode or later an in-place conversion).

It's important to note that 32-bit aggregates are not available in clustered Data ONTAP 8.3. Copy-based migration methods can be used to migrate data regardless of the source and destination aggregate types. However, replication-based or copy-free-transition--based migration methods cannot migrate a 32-bit aggregate from a source 7-Mode storage system to a 64-bit aggregate in clustered Data ONTAP 8.3. If the destination cluster runs clustered Data 8.3, it is necessary to convert the source aggregates to 64-bit and age out (or delete) 32-bit Snapshot copies before executing the data migration. If you are unsure if you have 32-bit Snapshot copies present in your 64-bit aggregate, contact NetApp support for assistance.

NetApp recommends converting 32-bit aggregates into 64-bit aggregates and allowing the 32-bit Snapshot copies to age out before moving the associated volumes to clustered Data ONTAP. For additional information on 32-bit aggregate conversion, refer to [TR-3978: In-Place Expansion of 32 Bit Aggregates to 64 Bit](#).

Aggregate Copy

Aggregate copy is not available in clustered Data ONTAP. Volume copy is the preferred method to copy data between aggregates.

Aggregate Snapshot Copies

Aggregate Snapshot copies are not used in clustered Data ONTAP.

8.2 Nondisruptive Operations and High Availability

The term nondisruptive operations (NDO) is used to describe a wide range of features and capabilities that are available in clustered Data ONTAP. NDO features become available to a transitioned storage container after the movement to clustered Data ONTAP is completed.

Nondisruptive upgrade (NDU) refers to the upgrade of Data ONTAP software. Transition from 7-Mode to clustered Data ONTAP cannot be done using the NDU procedure. For example, if you have a system running Data ONTAP 7.3.5 and want to upgrade to clustered Data ONTAP 8.1, the system has to go through a transition from 7-Mode to clustered Data ONTAP. Then the system is able to be nondisruptively upgraded to any subsequent release of clustered Data ONTAP.

8.3 Deduplication and Compression

When transitioning to clustered Data ONTAP, there are a few things of which to be aware if the volumes on the 7-Mode storage system have deduplication or compression enabled. Additionally, note the following changes and improvements available in clustered Data ONTAP specific to deduplication and compression:

- Deduplication and compression move from the `sis` command structure to the `volume efficiency` command set. The 7-Mode `sis` commands still work in clustered Data ONTAP but do not offer as many options.
- Existing scripts that use `sis` commands will continue to work with the migrated volume. However, the scripts should be updated at a convenient time to use the more robust clustered Data ONTAP `volume efficiency` commands.
- Consider using volume efficiency policies instead of deduplication and compression schedules to make use of the duration and added flexibility in start times.

Additional information is available in TR-3966: [NetApp Data Compression and Deduplication Deployment and Implementation Guide: Clustered Data ONTAP](#).

Appliance-Based, Application-Based, or Host-Based Migration

If the source 7-Mode volume has compression and/or deduplication enabled:

- Compression and deduplication space savings are not retained during the copy operation.
- Compression and deduplication have to be reapplied on the destination in order to regain space savings.
- Both source and destination storage systems run compression and deduplication independently of each other.
- The compression and deduplication schedule must be configured on the destination (the 7-Mode schedule is not migrated to the destination).

Because savings are not preserved on the destination and it is possible to have more logical data than physical space on disk, it is important to understand the transfer size. It is also important to verify that your destination is large enough to support the copy operations.

For example, a source volume that is 100TB in size that has 150TB of logical data and 75TB of physical data (35TB, or 20% savings, from compression and 45TB, or 30% savings, from deduplication) is migrated to a destination volume that is 100TB in size. The migration would be unable to proceed because the destination volume needs to be able to accept 150TB of data before volume efficiency operations can run to regain the savings. If you had inline compression enabled, you would still need to be able to write 120TB of data to the destination volume before running volume efficiency operations to regain the additional deduplication savings.

Best Practice

NetApp recommends stopping any active `sis` operations as well as disabling the `sis` schedule on the source volume before initiating an appliance-based, application-based, or host-based migration. When the `sis` schedule is disabled and active `sis` operations are stopped, changes to blocks in the source volume are avoided. Changes to the source volume increase the amount of data that needs to be recopied to the destination volume during incremental updates.

Because the deduplication and compressing schedule is not retained during the migration, a new schedule or policy should be set up on the destination volume after the migration is complete.

Replication-Based Migration

Replication-based migration operates at a physical block level, which results in all compression and deduplication savings on the source volume being retained over the network and on the destination volume. This can reduce the amount of network bandwidth required during a replication-based migration. More information on the replication-based migration method can be found in the [7-Mode Transition Tool](#) section of this document.

8.4 Flash Cache and Flash Pool

Flash Cache and Flash Pool intelligent caching are both supported in clustered Data ONTAP with the same commands and options that are available in 7-Mode. NetApp FlexShare® software is not a feature of clustered Data ONTAP. FlexShare settings that affect Flash Cache used in 7-Mode (volume prioritization) are not available in clustered Data ONTAP.

For additional information on Flash Pool and Flash Cache, see [TR-4070: Flash Pool Design and Implementation Guide](#) and [TR-3832: Flash Cache Best Practice Guide](#).

8.5 FlexClone

In clustered Data ONTAP, FlexClone functionality has improved as compared to 7-Mode FlexClone functionality. For example, when creating FlexClone LUN clones in clustered Data ONTAP, it is not necessary to maintain Snapshot copies in the parent volume. As a result, the number of clones possible on a single volume is increased. Another advantage of FlexClone LUN clones is instant availability for read and write operations regardless of the size of the LUN.

Additional information on FlexClone volumes in clustered Data ONTAP can be found in [TR-3742: Using FlexClone to Clone Files and LUNs](#).

Non-SIS LUN Clones

In clustered Data ONTAP, non-SIS LUN clones are not available. Non-SIS LUN clones must be split before moving to clustered Data ONTAP. When the LUN clone is split, the clone will no longer share blocks with its parent and will exist as a separate volume. Subsequently, the amount of space used in the volume (previously the clone) will be greater than the space used before the split operation.

Note: All LUN clones in clustered Data ONTAP are SIS LUN clones. As a result, a FlexClone license is required for clustered Data ONTAP in order to use FlexClone functionality.

FlexClone LUN Migration

Volumes containing LUNs and LUN clones in 7-Mode have to be migrated using an appliance-based, application-based, or host-based migration method.

FlexClone Volume Types

Beginning with clustered Data ONTAP 8.2, two types of FlexClone volumes are supported. Only volumes that are a destination for a SnapVault relationship can use the data protection FlexClone type. For volumes that are not a SnapVault destination, the FlexClone type is read-write. Although not a consideration for transition specifically, this point should be considered when designing the clustered Data ONTAP system.

8.6 Data ONTAP Networking

The following 7-Mode configuration elements translate in clustered Data ONTAP in these ways:

- VIFs become interface groups.
- Single IPs or hostnames become logical interfaces (LIFs).

- Routing tables become routing groups.
- Failover of IPs becomes migration of LIFs.

7-Mode storage systems using dynamic DNS, NIS for hostname lookup, GARP VLAN Registration Protocol, or the Interdomain Routing Protocol must use alternative methods of achieving the same capabilities with clustered Data ONTAP (if these features are still required).

Note: Configurable host lookup order (file, NIS, or DNS) is available in clustered Data ONTAP 8.3 (and later).

Failover groups in clustered Data ONTAP offer the capability to configure a LIF to move around to different nodes in the cluster. Interface groups give you the ability to aggregate outgoing I/O and provide failover capability at the node level. For maximum aggregated throughput and redundancy, NetApp recommends using a combination of failover groups and interface groups.

Second-level VIFs are not available in clustered Data ONTAP. If a similar type of functionality is required, clustered Data ONTAP supports three different types of failover groups: single-mode, static multimode, and dynamic multimode. The dynamic multimode interface group satisfies the redundant and aggregate requirements that a multilayered VIF provides in 7-Mode.

Additional information about Data ONTAP networking in clustered Data ONTAP can be found in [TR-4182: Ethernet Storage Best Practices for Clustered Data ONTAP Configurations](#).

8.7 SnapMirror

In clustered Data ONTAP, SnapMirror relationships can be intercluster (using cluster peering between nodes in separate clusters) or intracluster (between nodes within the same cluster). For any SnapMirror relationship that is migrated, the decision must be made as to whether the new SnapMirror relationship will be intercluster or intracluster. Key SnapMirror design considerations are as follows.

- The physical distance between the source and destination 7-Mode controllers dictates whether an intracluster SnapMirror relationship can be used in clustered Data ONTAP.
- Migrating to an intercluster SnapMirror relationship requires that a second (destination) cluster is available that is peered to the destination (of the migration) cluster. If a peer cluster is not available, it will not be possible to set up an intercluster SnapMirror relationship until a peer cluster is available. Before a peer cluster is made available, an intracluster SnapMirror relationship can be used to provide (localized) disaster recovery.
- NetApp recommends intercluster peer relationships because they provide the capability to locate the destination volumes off site, providing improved site-level resiliency over an intracluster SnapMirror relationship.

Note: There are a maximum number of cluster peer relationships that can be in place for any single cluster. Each cluster peer relationship supports multiple SnapMirror relationships between the peered clusters.

The key transition considerations for SnapMirror are as follows:

- Clustered Data ONTAP supports asynchronous SnapMirror. Both synchronous SnapMirror and semisynchronous SnapMirror are not available in clustered Data ONTAP.

Note: In clustered Data ONTAP 8.3 and later, the synchronous mirror solution is [MetroCluster](#).

- The `snapmirror.access` command is not used in clustered Data ONTAP. Cluster peering is used in clustered Data ONTAP to provide security for intercluster SnapMirror. Cluster peering defines which clusters are allowed to replicate data bidirectionally.
- Qtree SnapMirror (QSM) is not available in clustered Data ONTAP. Qtrees that reside in volumes in clustered Data ONTAP will be replicated by volume SnapMirror. For more information on migrating qtrees to clustered Data ONTAP, see the [Qtree Migration](#) section of this document.
- Replication of NAS or LUN (SAN) configuration information, SnapMirror over FC, and volume-level fan-in are not available in clustered Data ONTAP.

- Clustered Data ONTAP 8.3 and later support:
 - Network compression
 - FlexClone volumes as destination storage containers for SnapMirror relationships

Additional information about SnapMirror in clustered Data ONTAP can be found in [TR-4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP](#).

8.8 SnapVault

In 7-Mode, SnapVault is qtree based; in clustered Data ONTAP, SnapVault is volume based. As a result, 7-Mode SnapVault relationships do not translate directly to clustered Data ONTAP SnapVault relationships from a migration perspective. See the [Disaster Recovery and Backup Migration Scenarios](#) section of this document for guidance on how to move a 7-Mode SnapVault relationship to clustered Data ONTAP. Additional information on SnapVault in clustered Data ONTAP can be found in [TR-4183: SnapVault Best Practices Guide](#).

8.9 MetroCluster

NetApp MetroCluster technology is supported in clustered Data ONTAP 8.3 and later. The clustered Data ONTAP implementation of MetroCluster is different than that of 7-Mode.

Starting in Data ONTAP 8.3.1, MetroCluster provides two-node stretch and fabric configurations and a four-node fabric configuration. The MetroCluster two-node configuration has a single node Data ONTAP cluster at each site. Although the MetroCluster four-node configuration consists of a two-node cluster at each site, all aggregates in both sites (clusters) must be mirrored when using MetroCluster, including the root aggregates. In Data ONTAP 8.3, only the four-node fabric configuration is available.

The 7MTT supports migration of MetroCluster configurations using Copy-based migration from 7-Mode to clustered Data ONTAP.

8.10 NDMP

Customers using NDMP for backup in 7-Mode should transition directly to clustered Data ONTAP 8.2. Starting in clustered Data ONTAP 8.2, NDMP is cluster aware (through CAB extensions). This means that backups continue to be functional even if volumes are moved (to another node in the cluster) after the backup is completed.

Although it is possible to migrate data to clustered Data ONTAP using NDMPcopy, NetApp does not recommend doing this for volumes with high file counts (millions of files) and/or a high rate of change in the file system. A high-file-count environment with a low rate of change or a file system with a high rate of change will very likely result in NDMPcopy incremental copies being unable to complete in time. If NDMPcopy incremental copies are unable to complete, it will not be possible to complete the cutover. See the [Data Migration Tools](#) section of this document for migration methods that should be considered before deciding to use NDMPcopy for data migration.

8.11 SMTape

Transition considerations for SMTape follow.

- Snapshot copies that have been migrated from 7-Mode to clustered Data ONTAP cannot be used to initiate a SMTape backup operation.
- 7-Mode SMTape backups cannot be restored in clustered Data ONTAP and vice versa. SMTape is not a method for movement of data from 7-Mode to clustered Data ONTAP.
- SMTape backups of clustered Data ONTAP volumes that contain Snapshot copies that have been migrated from 7-Mode are supported.
- In clustered Data ONTAP 8.2.x and before, the following SMTape functionality is available:
 - Baseline backups

- In clustered Data ONTAP 8.3 and later, the following SMTape functionality is available:
 - Baseline and incremental backups

8.12 CIFS

The key transition considerations for CIFS are:

- File folding in 7-Mode is superseded by the use of deduplication in clustered Data ONTAP.
- Kerberos authentication using DES and RC4-HMAC encryption is available in clustered Data ONTAP.
- AD-based LDAP authentication is available in clustered Data ONTAP. Non-AD LDAP and workgroup authentication is not available in clustered Data ONTAP.
- In clustered Data ONTAP 8.2.1 (and later), the following is available:
 - LDAP over SSL (using Start-TLS)
 - Off-box antivirus scanning
 - Multiple domain search for user mapping
- In clustered Data ONTAP 8.2.2 (and later), `cifs.restrict_anonymous` is available. The default for clustered Data ONTAP is `no_restriction`.
- In clustered Data ONTAP 8.3 and later, the following is available:
 - NetBIOS aliases
 - Storage-level access guard (SLAG)
 - Character mapping
 - Auditing logon/logoff events
 - Home directory visibility for administrators
 - Viewing and management of shares, open files, and sessions in the Microsoft Management Console (MMC)

Note: Windows PowerShell scripts allow clusterwide enumeration of CIFS objects and provide additional functionality for clustered Data ONTAP systems.

- Live View Audit, Perfmon, LDAP (secure) using signing and sealing, and display CIFS client activities are not available in clustered Data ONTAP.
- Microsoft GPOs are supported for Kerberos settings, refresh time interval, and refresh time interval offset in clustered Data ONTAP.

Additional information about CIFS in clustered Data ONTAP can be found in [TR-4191: Best Practices Guide for Clustered Data ONTAP Windows File Services](#).

8.13 NFS

The key transition considerations for NFS are:

- NFS exports in clustered Data ONTAP differ from those in 7-Mode because there are no junction paths in 7-Mode. In clustered Data ONTAP, to mount a volume or a qtree, the NFS client must have read-only permissions at all the parent junction paths up to the SVM's root volume junction path (/).
- In Data ONTAP 8.2 (and later), NFSv2 is no longer available. NetApp recommends NFSv3 or above for any NFS configurations in clustered Data ONTAP.
- In clustered Data ONTAP 8.2.1 (and later), NFSv3 exports at the qtree level are available; before this version NFSv3 exports were supported only at the volume level.
- In clustered Data ONTAP 8.3 and later, the following is available:
 - NFSv4 qtree exports
 - Both `krb5` and `krb5i` (`kr5p` is not available)
- NFSv4 or the Windows-free NFS client can be used as an alternate to WebNFS (which is not available in clustered Data ONTAP).

- CIFS can be used as an alternative to PCNFS (which is not available in clustered Data ONTAP).

In addition to the above considerations, it is important to understand changes to the `-actual` option (from 7-Mode), exports, and junction paths, as well as how user access has changed for NFS in clustered Data ONTAP. These changes are discussed in more detail below.

The `-actual` option that was used when mounting NFS in 7-Mode is no longer needed with clustered Data ONTAP. Clustered Data ONTAP mounts each volume at a directory junction path (a client-visible directory), which does not have to be the same as the volume name. NFS clients mount the junction path or a directory below the junction path and not the volume name (removing the need to use the `-actual` option). Doing so eliminates the requirement to have `/vol` in the export path. More details on transitioning when the 7-Mode volume uses the `-actual` option are available in [TR-4067: Clustered Data ONTAP NFS Best Practice and Implementation Guide](#).

Volumes are automatically exported when created in clustered Data ONTAP if the `-junction-path` option is specified. If the `-junction-path` option is not specified, the volume can be exported at a later time using the `vol mount` command. If an export policy is not specified when the volume is mounted, the default export policy is used.

The `anon` user ID specifies a UNIX user ID or user name that is mapped to client requests that arrive without valid NFS credentials. This user ID can include the `root` user. Clustered Data ONTAP determines a user's file access permissions by checking the user's effective UID against the storage virtual machine's specified name-mapping and name-switch methods. After the effective UID is determined, the export policy rule is leveraged to determine the access allowed for the UID. The `anon` option in the export policy rule determines which UID will map to an anonymous user. The `superuser` option in the export policy rule controls whether `root` is interpreted as `root` or is diverted to `anon`. To keep `root` as `root`, set the `superuser` to `sys` or set `anon` to `0`. Set `superuser` to `none` to divert `root` to the `anon` UID.

Additional information about NFS in clustered Data ONTAP can be found in [TR-4067: Clustered Data ONTAP NFS Best Practice and Implementation Guide](#).

8.14 SAN

This section addresses the key considerations to take into account when planning a transition to a scalable SAN solution in clustered Data ONTAP. Additional information on scalable SAN in clustered Data ONTAP can be found in [TR-4080: Best Practices for Scalable SAN in Data ONTAP 8.3](#). Table 1 outlines the key differences between SAN in 7-Mode and SAN in clustered Data ONTAP.

Table 1) Key SAN differences between 7-Mode and clustered Data ONTAP.

Configuration Detail	7-Mode	Clustered Data ONTAP
FC and iSCSI service scope	Per node/per HA pair	Per storage virtual machine
iSCSI secondary paths	Failover to partner interface	Based on Asymmetric Logical Unit Access (ALUA)
Fibre Channel indirect paths	Over NVRAM interconnect	Over cluster network
Fibre Channel ports	Physical ports	Virtual ports (NPIVs)
Interface with NetApp SnapDrive software	Any active Ethernet port	Storage virtual machine management LIF

FC and iSCSI Service Scope

In 7-Mode, a single FC World Wide Node Name (WWNN) or iSCSI Qualified Name (IQN) was used per storage controller or HA pair. Clustered Data ONTAP creates a unique WWNN or IQN for every storage virtual machine running an FC or iSCSI service and makes them available from any cluster node.

iSCSI Secondary Path and Partner Interface

7-Mode uses a partner interface to provide redundancy for iSCSI connections. During an HA failover event, the HA partner of the affected storage controller has a predesignated partner interface that is brought online with the IP address of the taken-over HA partner's iSCSI target interface and I/O resumes.

In clustered Data ONTAP, iSCSI partner interfaces are not assigned when configuring iSCSI. Instead, iSCSI connections are made to the node where the data resides, its HA partner for redundancy purposes, and, optionally, additional nodes in the same cluster. Instead of the path reappearing after an HA takeover on the HA partner, MPIO on the host determines that the primary path is no longer available and resumes I/O over the already-established connection to the taken-over node's HA partner. This action means that, unlike the 7-Mode model, iSCSI connections to a cluster will have both direct and indirect paths to data managed by host MPIO software

FC Indirect Paths

Both clustered Data ONTAP and 7-Mode use the NVRAM or HA interconnect to mirror NVRAM data onto both storage controllers in an HA pair. In 7-Mode, this HA interconnect is also used to provide a redundant FC path using the HA partner of the storage controller's FC target ports. This interconnect is commonly referred to as the virtual interconnect, or "vtic," path. In clustered Data ONTAP, the HA interconnect is used only for NVRAM mirroring and indirect I/O is forwarded over the cluster network common to all cluster nodes.

FC Ports

Clustered Data ONTAP uses N_Port ID virtualization (NPIV) to permit every LIF to log into an FC fabric with its own World Wide Port Name (WWPN). It does so rather than use a single WWNN and associated WWPNs based on the address of the HA pair's physical FC target adapter(s), as with 7-Mode. This change permits a host connected to the same FC fabric to communicate with the same SCSI target regardless of which physical node owns which LIF. NPIV is required for FC LIFs to operate correctly. Any FC fabric to which a clustered Data ONTAP system is attached must have NPIV enabled.

Management Interfaces

LIFs belonging to SVMs that serve data using block protocols cannot also be used for administration purposes. Because the logical unit of management in clustered Data ONTAP is the storage virtual machine, NetApp recommends that every storage virtual machine be provided with a management interface in addition to interfaces that serve data using block protocols. A management interface has the following attributes:

- Uses "data" as its LIF type
- Has a firewall policy that permits management access
- Has no data protocols assigned
- Is accessible to hosts that need to connect to it for administrative purposes (such as coordination with NetApp SnapDrive)

Additionally, a storage virtual machine-level administrative account should be available.

When using the Create Storage Virtual Machine (SVM) Setup Wizard in System Manager, the wizard's last step presents an option to create a dedicated LIF for SVM management. There is an option to

delegate the administration of the SVM to the SVM administrator by setting up a password for the vsadmin user account.

LUN Mapping

In 7-Mode, portsets are sometimes used to permit administrators to mask an interface group (igroup) so that the LUNs that are mapped are available on a subset of the total number of available target ports. This capability is available in both clustered Data ONTAP and 7-Mode.

In clustered Data ONTAP 8.3 and later, selective LUN mapping is available that allows storage administrators to select which storage controllers will allow access to a given LUN. For example, a storage administrator could configure LUNs on every node in a cluster for access by a single host without requiring that every LUN has a path associated with every node.

By default, new LUN mappings created will be configured to provide data access to LUNs on the node that is local to the LUN and on its HA pair partner node. LUN mappings that were originally created in clustered Data ONTAP 8.2 retain the preclustered Data ONTAP 8.3 functionality (allowing the LUN to be accessed by any LIF configured and online in the LUN's SVM). For more information on selective LUN mapping functionality, see [TR-4080: Best Practices for Scalable SAN in Data ONTAP 8.3](#).

8.15 OnCommand Management Products

This section covers the key considerations for OnCommand management products. The primary consideration for OnCommand management products for transition to clustered Data ONTAP is to understand which version of each product should be used. The version to use depends upon whether the NetApp storage systems you have within your NetApp environment run 7-Mode, clustered Data ONTAP, or a mix of the two. Table 3 covers which OnCommand management product versions to use.

Table 2) OnCommand management product versions.

Product	Description	Product Version	
OnCommand System Manager	Provides element management of NetApp storage systems. Ideal for one-off and nonrepeatable management and configuration tasks.	7-Mode: Clustered Data ONTAP: Both:	3.1.1 3.1.1 3.1.1
OnCommand Performance Manager	Provides performance management and root cause analysis of clustered Data ONTAP.	7-Mode: Clustered Data ONTAP: Both:	N/A 1.x N/A
OnCommand Unified Manager	Enables integrated policy-based data and storage management for virtual and physical environments. Used for monitoring, alerting, and automation tasks.	7-Mode: Clustered Data ONTAP: Both:	5.x 6.x See note
OnCommand Workflow Automation	Enables automation and delegation of all repeatable storage management and storage service tasks. Facilitates the customer's exact storage service.	7-Mode: Clustered Data ONTAP: Both:	2.2 2.2 2.2
OnCommand Insight	Enables storage resource management, including configuration and performance management, as well as capacity planning and advanced reporting for heterogeneous environments.	7-Mode: Clustered Data ONTAP: Both:	6.x 6.4 6.4

Note: Separate installations of OnCommand Unified Manager are required to manage both 7-Mode and clustered Data ONTAP.

OnCommand System Manager

OnCommand System Manager can be used to manage both Data ONTAP 7-Mode and clustered Data ONTAP. Information about OnCommand System Manager workflows for Data ONTAP 7-Mode and clustered Data ONTAP can be found in the “OnCommand System Manager [2.2 Workflow Guide](#)” and the [3.1 Workflow Guide](#).

Best Practice

NetApp recommends using the latest version of OnCommand System Manager for the Data ONTAP release being managed. For Data ONTAP 8.2 operating in 7-Mode and clustered Data ONTAP 8.2, OnCommand System Manager version 3.1.1 should be used.

OnCommand Unified Manager

OnCommand Unified Manager 5.x manages Data ONTAP 7-Mode and OnCommand Unified Manager 6.x manages clustered Data ONTAP. Separate installations of OnCommand Unified Manager are needed when managing environments that contain both Data ONTAP 7-Mode and clustered Data ONTAP.

Depending on the capability being used in OnCommand Unified Manager version 5.x, NetApp recommends getting the equivalent capability for managing clustered Data ONTAP as follows.

- **Operational alerting, monitoring, and reporting.** OnCommand Unified Manager version 6.x can be used for monitoring and operational alerting, protection management, and performance incident consolidation; version 5.2 can be used for reporting.
- **NetApp SnapProtect® integration.** OnCommand Unified Manager 6.x provides SnapProtect integration to mirror and vault application-consistent Snapshot copies in clustered Data ONTAP.

Host package integration, SnapManager integration and Open Systems SnapVault are features available only in 7-Mode for OnCommand Unified Manager. Data protection and provisioning automation can be performed with OnCommand WFA for clustered Data ONTAP.

Best Practice

NetApp recommends that a database backup of the current installation of OnCommand Unified Manager be created using the `dfm backup create` command. The backup file should be archived. The backup will contain all of the historical usage and performance information for historical reference.

OnCommand Workflow Automation

OnCommand WFA is used to automate storage-centric tasks in both 7-Mode and clustered Data ONTAP. If you previously used OnCommand Provisioning and Protection Manager to automate provisioning and protection tasks, in clustered Data ONTAP you will use OnCommand WFA.

Support for clustered Data ONTAP with OnCommand WFA was introduced with version 1.1. After transitioning to clustered Data ONTAP, customers should use WFA version 2.2. Depending on the capability used in OnCommand Provisioning and Protection Manager, the recommendation to get the equivalent capability for managing clustered Data ONTAP using WFA 2.2 is as follows.

- **NAS provisioning workflows.** Sample workflows for NAS provisioning use cases are provided with WFA 2.2 along with full support for NAS provisioning tasks.
- **SAN provisioning workflows.** Sample workflows for SAN provisioning use cases are provided with WFA 2.2 along with full support for SAN provisioning tasks.
- **Service catalog.** In WFA 2.2, the template feature provides similar functionality for clustered Data ONTAP.

- **vFile provisioning.** In clustered Data ONTAP, support for secure multitenant environments is achieved using storage virtual machines. Full support and sample workflows for storage virtual machines are available starting with WFA version 2.2.
- **SnapMirror and SnapVault.** Both SnapMirror and SnapVault are fully supported by WFA 2.2.
- **Open Systems SnapVault.** Open Systems SnapVault is a 7-Mode-only feature.

OnCommand Insight

Support for clustered Data ONTAP with OnCommand Insight was introduced with version 6.4 (supports clustered Data ONTAP as a data source). Environments that use both Data ONTAP 7-Mode and clustered Data ONTAP can use OnCommand Insight 6.4 or 7.0 for both deployments of Data ONTAP.

OnCommand Performance Manager

OnCommand Performance Manager is built specifically for clustered Data ONTAP 8.2 and later. OnCommand Performance Manager delivers comprehensive data storage performance troubleshooting, identifies potential bottlenecks within the cluster, and offers suggested actions to resolve performance issues based on its automated analysis.

8.16 NetApp Manage ONTAP API Suite

The NetApp Manage ONTAP API suite supports APIs for both 7-Mode and clustered Data ONTAP. The APIs used for clustered Data ONTAP are different than those used for 7-Mode. As a result, it is very likely that tasks that have been written using the APIs for 7-Mode will need to be updated to use the clustered Data ONTAP APIs.

There are two sets of APIs for clustered Data ONTAP. Cluster APIs are used for the entire cluster and its resources, whereas the storage virtual machine APIs are used for specific SVM storage and network resources.

Three sets of iterative APIs (*-iter-start, *-iter-next, and *-iter-end) are used with 7-Mode that allow operations to execute on a large set of objects. For clustered Data ONTAP, there is a single API to use for iteration (*-get-iter), which replaces the need to use the three APIs mentioned above.

Additional information on the Manage ONTAP API suite can be found in the developers guide for clustered Data ONTAP APIs, which is available on the [NetApp Developer Community](#) site.

8.17 SnapDrive and SnapManager Products

This section addresses the recommended operating versions for NetApp's suite of SnapDrive and SnapManager products. The [Interoperability Matrix Tool \(IMT\)](#) can be consulted in order to see the latest information on product interoperability. Specific versions of both SnapDrive and SnapManager products are required to support clustered Data ONTAP. Table 2 covers the minimum product versions that support clustered Data ONTAP.

Table 3) SnapDrive and SnapManager minimum product versions for clustered Data ONTAP.

Product	Description	Clustered Data ONTAP	
		8.2.x	8.3.x
SnapManager for Exchange	NetApp SnapManager for Microsoft Exchange Server lets you streamline and simplify Exchange data management.	6.1, 7.0	6.1, 7.0, 7.1
Single Mailbox Recovery	NetApp Single Mailbox Recovery for Microsoft Exchange Server enables fast retrieval of individual messages or mailboxes.	7.0	7.0, 7.1

SnapManager for Microsoft SharePoint	NetApp SnapManager for Microsoft SharePoint Server lets you improve SharePoint scalability and reduce storage costs.	8.1	8.1, 8.2
SnapManager for SAP	NetApp SnapManager for SAP data management software automates such complex tasks as backup, restore, and cloning.	3.3.x	3.3.1
SnapManager for Oracle	NetApp SnapManager for Oracle automates complex, time-consuming database backup, restore, recovery, and cloning.	3.3.x	3.3.1
SnapManager for Microsoft SQL Server	Automate and simplify backup, recovery, and database cloning with NetApp SnapManager for Microsoft SQL Server.	7.1	7.1
SnapManager for Hyper-V	NetApp SnapManager for Hyper-V speeds backup, restore, and disaster recovery in a virtualized Microsoft environment.	2.0.3	2.0.3, 2.1
SnapDrive for UNIX	Use NetApp SnapDrive to simplify and automate storage provisioning and data management in UNIX and Linux environments.	5.2.x	5.2.1P1, 5.2.2
SnapDrive for Windows	NetApp SnapDrive for Windows automates storage and data management for physical and virtual Windows environments.	7.0.3	7.0.3, 7.1

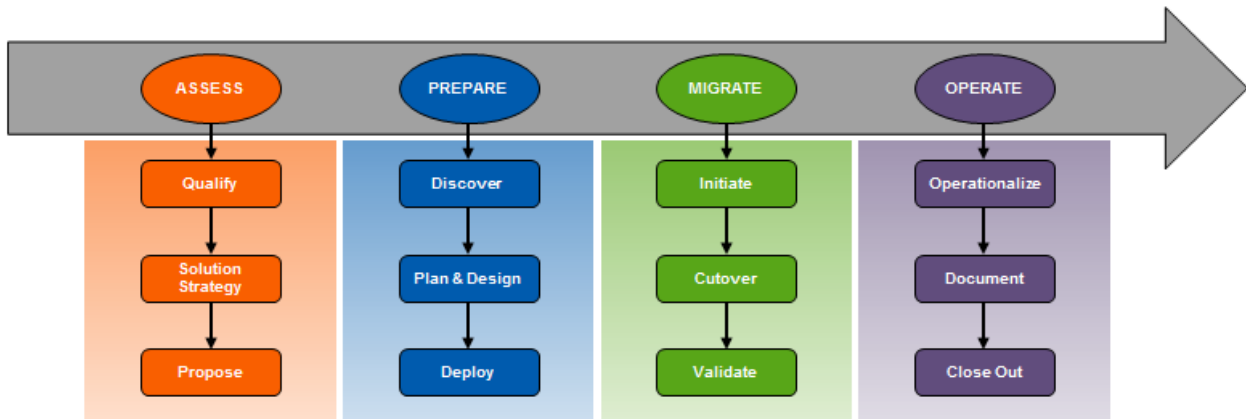
Additional information on SnapDrive and SnapManager products, as well as the applications they support, can be found in [TR-4336: Enterprise Application Transition to Clustered Data ONTAP](#).

9 NetApp Service Offerings

There is a full range of professional and support services available for those who are transitioning to clustered Data ONTAP. The service offerings are structured to help from the first transition step to business continuance with clustered Data ONTAP.

NetApp has developed a process that is used to conduct transitions. The Unified Transition Methodology (UTM) is the process NetApp uses to put the knowledge of the transition fundamentals into practice. Figure 33 displays an overview of the transition fundamentals.

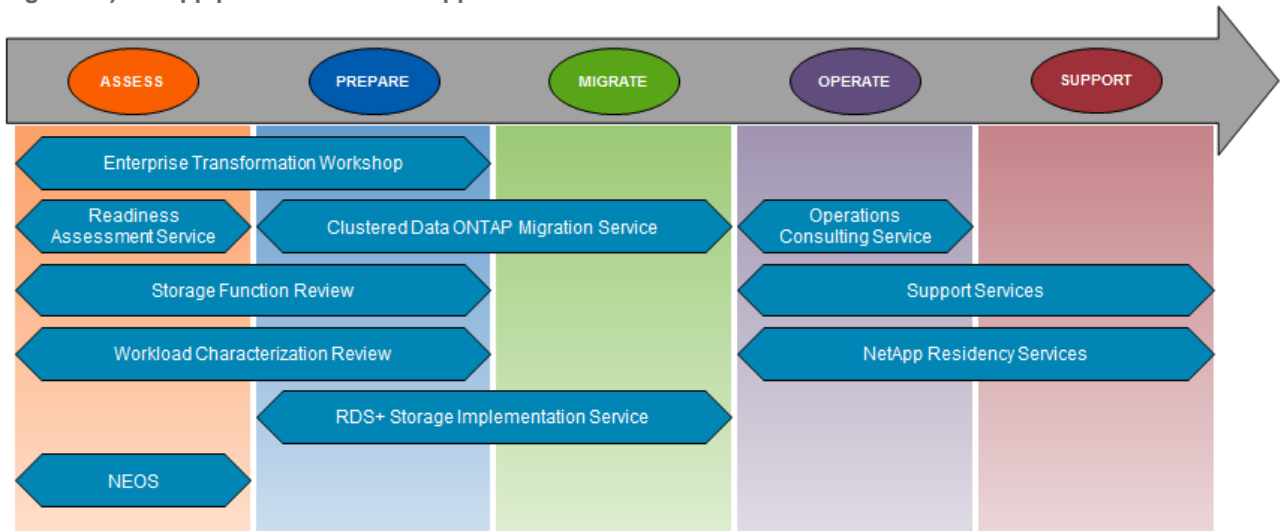
Figure 33) Unified Transition Methodology overview.



The UTM phases and steps can be used from start to finish or they can supplement an existing process. For more information on the Unified Transition Methodology, contact your NetApp representative.

NetApp professional and support services offerings are aligned to the four main phases of the UTM, with an added phase to address continued support of the NetApp environment. All of the services mentioned below can help with transition; however, several services also apply in general (not just for transition to clustered Data ONTAP). Figure 34 displays an overview of the professional and support services available from NetApp today.

Figure 34) NetApp professional and support services overview.



The following is a description of the service offerings outlined in Figure 32.

- **Enterprise Transformation Workshop:** Working with executives, services experts help customers create an IT roadmap and build an executable cloud strategy for aligning their IT capabilities and cloud services to their organization's goals. Prioritized recommendations provide business justification for future projects.
- **Clustered Data ONTAP Readiness Assessment Service:** Provides customized assessment of the NetApp environment for transition. Through identification of feature and functional differences, and by creating actionable recommendations, this service prepares a NetApp environment for transition to clustered Data ONTAP.

- **Storage Function Review:** Assess storage and management capabilities, including inputs, procedures, alerts, and reports, to gain increased visibility into business and IT processes.
- **Workload Characterization Review:** Experts confirm that the recommended clustered Data ONTAP design and architecture will support the workloads being migrated, reducing risk as the workloads transition to clustered Data ONTAP. Additionally, this service provides actionable recommendations for alignment of storage workload demand and capability.
- **NetApp Efficiency and Optimization Service:** A storage-level health check that provides expert analysis of the current state of the NetApp environment. This service includes assessment of clustered Data ONTAP design and architecture. As a result, potential problems and issues can be preempted and resolved through customized recommendations that enable increased efficiency and performance from your current investments.
- **RDS+ Storage Implementation Services:** Experts rapidly deploy and configure the new NetApp FAS series storage systems running Data ONTAP with minimal impact on end users and the bottom line.
- **Clustered Data ONTAP Migration Service:** Migrate NetApp and non NetApp (third-party) data to a fully configured, tested, and verified clustered Data ONTAP system while maintaining existing efficiencies such as Snapshot copies, deduplication, and compression.
- **Operations Consulting Services:** Based on the findings of the Storage Function Review and Workload Characterization, this service implements the recommendations resulting from those services. Services experts help to achieve operational excellence by developing and documenting procedures to boost storage management capabilities and increase workload performance.
- **NetApp Support Services:** This service provides worldwide, cost-effective support that is scaled and priced to meet your needs, allowing you to achieve greater stability, reliability, and availability of critical business data.
- **NetApp Residency Services:** On-site skills, knowledge, and expertise meet specific objectives to maximize your NetApp investment. This service enables you to manage a steady-state environment to reduce operational risk, improve utilization, achieve storage efficiency, and ultimately control costs.

10 Additional Resources

The following references were used in this technical report or are otherwise useful for making the transition to clustered Data ONTAP from 7-Mode.

- Interoperability Matrix Tool
<http://support.netapp.com/matrix/mtx/login.do>
- Using the 7-Mode Transition Tool (7MTT) to migrate your data to clustered Data ONTAP
<https://www.youtube.com/watch?v=icSPJ4Phlvk>
- 7-Mode Transition Tool 2.2 – Copy-Free Transition Guide
<https://library.netapp.com/ecmdocs/ECMLP2346396/html/index.html>
- 7MTT Copy-Free Transition Cabling Guide
http://www.netapp.com/us/media/7MTT_CFT_CablingGuide.pdf
- TR-3982: Clustered Data ONTAP 8.2: An Introduction
<http://www.netapp.com/us/media/tr-3982.pdf>
- 7-Mode Transition Tool Data and Configuration Transition Guide and Release Notes
<http://mysupport.netapp.com/documentation/docweb/index.html?productID=61851>
- DTA2800 Product Documentation
<http://mysupport.netapp.com/documentation/docweb/index.html?productID=61597>
- TR-3978: In-Place Expansion of 32-Bit Aggregates to 64-Bit
<http://www.netapp.com/us/media/tr-3978.pdf>

- TR-3966: NetApp Data Compression and Deduplication Deployment and Implementation Guide: Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-3966.pdf>
- TR-3450: HA Pair Controller Configuration Overview and Best Practice Guide
<http://www.netapp.com/us/media/tr-3450.pdf>
- TR-4129: Namespaces in Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4129.pdf>
- TR-3742: Using FlexClone to Clone Files and LUNs
<http://www.netapp.com/us/media/tr-3742.pdf>
- TR-4160: Secure Multi-Tenancy in Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4160.pdf>
- TR-4182: Ethernet Storage Best Practices for Clustered Data ONTAP Configurations
<http://www.netapp.com/us/media/tr-4182.pdf>
- TR-4067: Clustered Data ONTAP NFS Best Practice and Implementation Guide
<http://www.netapp.com/us/media/tr-4067.pdf>
- TR-3967: Deployment and Best Practices Guide for Clustered Data ONTAP 8.1 Windows File Services
<http://www.netapp.com/us/media/tr-3967.pdf>
- TR-4070: NetApp Flash Pool Design and Implementation Guide
<http://www.netapp.com/us/media/tr-4070.pdf>
- TR-3832: Flash Cache Best Practices Guide
<http://www.netapp.com/us/media/tr-3832.pdf>
- TR-4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4015.pdf>
- TR-4031: OnCommand System Manager 2.2 Guide to Common Workflows
<http://www.netapp.com/us/media/tr-4031.pdf>
- Clustered Data ONTAP 8.2 File Access and Protocols Management Guide
https://library.netapp.com/ecm/ecm_get_file/ECMP1366835
- Clustered Data ONTAP 8.2 Logical Storage Management Guide
https://library.netapp.com/ecm/ecm_get_file/ECMP1368017
- Clustered Data ONTAP 8.2 Command Map for 7-Mode Administrators
https://library.netapp.com/ecm/ecm_download_file/ECMP1366830
- Clustered Data ONTAP for the 7-Mode Administrator Portal
<https://forums.netapp.com/community/support/cfor7>
- TR-4336: Enterprise Application Transition to Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4336.pdf>

Version history

Version	Date	Document Version History
Version 4.0	January 2016	Jay Bounds: Updated for clustered Data ONTAP 8.3.2 and 7MTT 2.2, added CFT content.
Version 3.0	October 2015	Jay Bounds: Added FLI, updated decision tree and data migration diagrams, added CFT, refreshed existing content.
Version 2.0	August 2015	Roy Scaife: Added 7MTT version 2.1 feature updates.
Version 1.0	October 2014	Jay White: This is the first publically available version of this technical report.

11 Contact Us

Let us know how we can improve this technical report.

Contact us at docfeedback@netapp.com.

Include TECHNICAL REPORT 4052 in the subject line.

12 Disclaimer

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2015 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4052-1014

