

COHESITY

Version 1.1

September 2020

Use Cohesity as a Data Migration Engine

Simplify Enterprise Data Protection & Data Center Migration

Table of Contents

- Introduction4**
 - Goal 4
 - Infrastructure Assumptions 4
- Traditional Data Center Migration Methods5**
 - Power Down and Physically Move All Primary Storage and Compute 5
 - Storage Array Replication 5
- Leverage Cohesity for Migration7**
 - Cohesity Clusters Leverage Replication 7
- Common Types of Migrations8**
 - Relocation (One Data Center to One Data Center) 8
 - Consolidation (Many Data Centers to One Data Center) 8
- Data Migration Customer Requirements9**
- Cohesity Solution for Data Center Migration10**
- Design11**
 - New Data Center Infrastructure Requirements 11
 - Planning for Migration 11
- Execution12**
 - Replicate the Existing Environment to the New Data Center..... 12
 - Test Restoration Before to Cutover 13
 - Cutover: Restore VMs to the New Data Center..... 13
 - Complete the Transfer of Operations and Migration..... 14
- Conclusion.....15**
- Your Feedback.....16**
- About the Authors.....16**
- Document Version History16**

Figures

- Figure 1: Relocation — from One Data Center to Another 8
- Figure 2: Consolidation — from Many Data Centers to One..... 8
- Figure 3: Before Data Center Migration 9
- Figure 4: Enable Replication Between Cohesity Clusters 13
- Figure 5: Cutover of VMs 14
- Figure 6: Transfer of Operations and Migration Completion 14

Tables

- Table 1: Customer Business Requirements 9
- Table 2: Benefits of Cohesity Features for Data Center Migrations 10

Introduction

Cohesity is ideally suited for much more than just data protection and file services. Because of its efficiencies with backup, replication, and restore, it is also well positioned to complement data center migrations and consolidations.

This chapter covers the goal of this white paper and the assumptions it relies on for the procedures it describes.

Goal

Migrating large amounts of data from one data center to another can be a daunting task. The goal of this document is to outline different data center migration options and how you can use Cohesity to simplify and reduce the downtime window involved in the process.

Infrastructure Assumptions

This document assumes there is network connectivity between data centers and that sufficient compute and storage exists at the target or destination data center. If there is insufficient compute and storage for migrating all servers and applications at once, splitting the migration process into phases might be advantageous.

While we focus here primarily on migrating vSphere VMs from data center to data center, you can apply many of the same concepts and mechanisms for other workloads, such as file services, databases, NAS, cloud migrations, and more.

Traditional Data Center Migration Methods

Before digging too deep into the details of a data center migration, it is important to note there are many different ways you can handle such a migration. Each approach has its own pros and cons. Below, we explain a few of the typical migration methods but note that each migration is different and there might be other methods or variations that are not specially called out here.

Power Down and Physically Move All Primary Storage and Compute

Historically, data center moves involve simply powering down all primary storage along with compute hardware and then physically moving the hardware from the old data center to the new data center by vehicle. This type of migration might make sense for less critical environments, but for larger or complex and crucial data centers, such a downtime or multiple downtimes might be unacceptable.

Pros:

- Little to no additional hardware to purchase.

Cons:

- Longest downtime, which equates to the highest cost of the different methods because of lost revenue.
- All VMs/servers using SAN arrays being moved must be migrated at one time.
- Long downtime, depending on how long it takes to un-rack, physically move, re-rack, power on, and verify there are no SAN or network issues after the migration.
- If there are issues with the migration and the issues can't be resolved, one might have to move the equipment back to the original data center, extending the downtime and introducing other risks inherent with physical moving equipment.

Storage Array Replication

Migrating VMs via SAN Array-based replication involves pairing two compatible arrays and enabling replication from one to the other. This often involves needing to replicate at the LUN level or a group of LUNs. After replication/synchronization, the LUNs at the new data center will contain the same data as the old data center. At this point, the LUNs must be presented to the vSphere compute hosts and the VMs imported, modified if necessary, and then powered up.

Using Storage Array-based replication for migration is a good option for simple migrations or those where almost no changes are needed and the VMs can be easily imported from the replicated LUNs at the new data center.

Pros:

- Shorter downtime.
- In the case of unforeseen issues, it takes less time to fail-back to the original data center.

Cons:

- Longer downtime, which equates to higher costs because of lost revenue.
- Requires the same or similar arrays at each site that are capable of replicating between each other.
- Might require physically moving compute hosts, unless you purchase additional hosts and install them in the destination data center.
- Breaking up a migration into phases can reduce the need to purchase all new compute hosts, but this can get tricky if many reside on just a few large LUNs.
- Requires a re-scan of VMWare datastores and re-import of VMs from a SAN Array snapshot. This is often a manual process and can be time-consuming for a large number of VMs, snapshots, and/or datastores.
- Importing VMs that span multiple datastores can also be problematic and time-consuming.

Leverage Cohesity for Migration

Leveraging Cohesity to migrate between data centers has a number of advantages, outlined below, over legacy and more traditional migration methods. No one solution can handle all aspects of a data center migration, but that said, Cohesity can enable a more seamless migration with greater flexibility than other migration methods.

Cohesity Clusters Leverage Replication

Organizations can achieve enterprise-level resiliency with site-to-site replication between Cohesity clusters. You can configure replication at multiple granularity levels for maximum flexibility, including cluster-wide — that is, all data on a single Cohesity cluster is replicated to one or more other clusters).

After VM snapshots are replicated (copied) to a remote Cohesity cluster, you can use them to:

- **Recover VMs** from the replicated snapshots to a new location.
- **Clone VMs** from the replicated snapshots.
- **Recover files** from the replicated snapshots.
- **Run the Analytics Workbench** on the replicated snapshots.
- **Download VMX files** from the replicated snapshots.

Pros:

- Ability to leverage Cohesity for more than just data protection.
- Efficient site-to-site replication.
- The ability to do one or more test migrations without consuming SAN Array storage before actual migration.
- Granular recovery and migration of VMs, databases, and other workloads via Cohesity's intuitive UI, REST API calls, or PowerShell scripts.
- Ability to choose to restore to new or different vSphere datastores, reconfigure the VMs to use different NICs, or even switch to a different version of vSphere altogether if desired.

Cons:

- Might require physically moving compute hosts, unless you purchase additional hosts and install them in the destination data center.
- Breaking up a migration into phases can reduce the need to purchase all new compute hosts.

Common Types of Migrations

There are several different migration types, but two common ones are *relocation* and *consolidation*. Regardless of the reason or type of migration, Cohesity can play a vital role in the migration process.

Relocation (One Data Center to One Data Center)

In this type of migration, only two data centers are involved. Typical use cases are:

- Moving from one colocation facility to another.
- Moving from an existing company-owned data center to a colocation facility.
- Moving to a newly built or purchased data center.

Figure 1: Relocation — from One Data Center to Another



Consolidation (Many Data Centers to One Data Center)

In this type of migration, typically more than two data centers are involved. For example, a consolidation of three data centers into one, be it a physical data center or colocation facility. Consolidation might be the result of reducing the number of data centers to achieve greater efficiencies or the result of a merger or acquisition. In addition to simply migrating data, there might be the requirement to retain backups from merger or acquisition as well, of which is no problem with Cohesity.

Consolidation can be complicated but with Cohesity's ability to replicate backups from multiple clusters to a single cluster, helps remove and reduce that complexity.

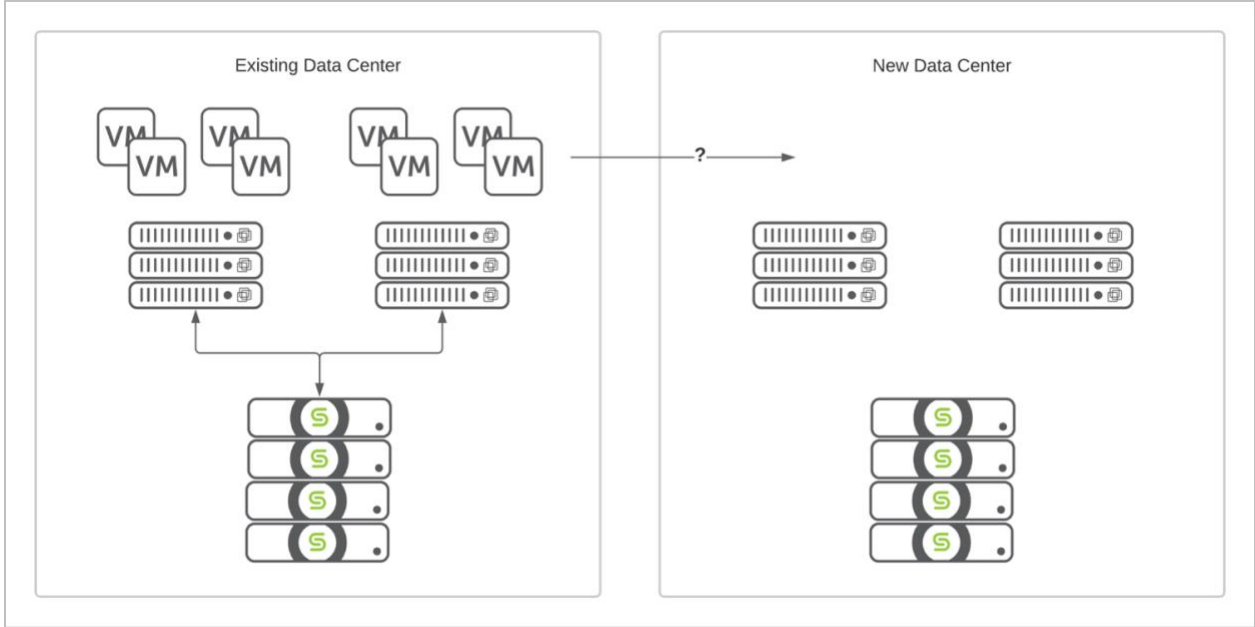
Figure 2: Consolidation — from Many Data Centers to One



Data Migration Customer Requirements

Think of the data migration framework as our starting point, in which the old data center exists with all existing workloads while the new data center exists but does not yet have workloads.

Figure 3: Before Data Center Migration



For a migration to meet various important business requirements, as outlined in Table 1.

Table 1: Customer Business Requirements

SOLUTION REQUIREMENT	STRATEGY
Minimum-to-Zero Downtime on Primary Operations	Stakeholders want a seamless transfer of <i>primary operations</i> from one data center to another without affecting ongoing operations.
Minimum-to-Zero Downtime on Backups	Stakeholders want a seamless transfer of <i>backup operations</i> from one data center to another without downtime or suspension of operations.
Cost-effective	Utilize any data optimization techniques to ensure only the minimum utilization of the data pipe between data centers.
Seamless Switchover	Ensure the disaster recovery (DR) solution provides turn-key capability when switching from the old data center to the new data center.
Data Integrity	The solution should not just include encryption, but also ensure <i>end-to-end</i> data security and integrity.

Cohesity Solution for Data Center Migration

Cohesity's web-scale distributed file system that provides unlimited scale across any number of industry-standard nodes — using inherent data optimization and reduction techniques — while providing a cost-effective mechanism for data mobility tasks like replication or disaster recovery. No matter the size of the data center migration, be it the number of ESXi hosts, vCenters, or VMs, Cohesity can handle it.

Table 2: Benefits of Cohesity Features for Data Center Migrations

COHESITY FEATURE	HOW IT HELPS	BENEFITS
Global Deduplication and Compression	Reduces the amount of data that needs to be stored on Cohesity.	Reduced data transfer and better cross-site network pipe usability.
Optimized Data Migration through Cohesity Replication 2.0	Parallel data transfer, MegaFile, metadata-aware synchronization between Cohesity clusters.	Reduced data transfer and better cross-site network pipe usability.
Test/Dev and Instant Cloning	Cohesity has the capability to provide instant availability of VMs and applications.	Flexibility to use Cohesity as primary storage temporarily. Bring up VMs, MS-SQL Databases, or other applications near-instantaneously.
Data and Permission Integrity	Cohesity has inherent immutable capability. Ability to restore VMs to the latest backup (most common) or any point in time.	Data integrity is built into the Cohesity platform.
Advanced VM Restore Options	Ability to choose any suitable vCenter, Resource Pool, Datastore, VM Folder, and/or Network.	Flexibility to choose where to restore, all from within Cohesity.
Intuitive UI or Scripting/Automation	Choose to use Cohesity's intuitive UI for restores or leverage the Cohesity REST APIs or PowerShell Cmdlets.	For single restores of one or more VMs, Cohesity's UI is intuitive and efficient. For repeat restores, in the case of multiple mock restores or migration, leveraging Cohesity REST APIs or PowerShell Cmdlets might be desired.

Design

Instead of a lift-and-shift model, you can use Cohesity's replication model that allows the data centers to be in sync *before* cutting over from one to the other.

New Data Center Infrastructure Requirements

This document assumes some basic infrastructure is already in place at the destination data center, such as Primary SAN storage, networking, compute hosts, Active Directory, DNS, DHCP, and NTP.

It is possible during the migration to restore the previous Active Directory, DNS, and NTP servers or services, for this document, we assume that some of these core services are already in place, either through a prior restore or via some other method these are already setup and functioning.

Planning for Migration

Planning how things will look at the new data center requires some thought. With the ability to choose any suitable vCenter, Resource Pool, Datastore, VM Folder, and/or Network, you can choose to restore VMs originally residing on different vCenters to a single vCenter at the new data center or split them up on multiple vCenters, depending on the final design and objectives. It might make sense to deploy vSphere 7.0 at the new data center and restore VMs to *that* even though they were running on vSphere 6.7 in the old data center. Perhaps restoring to a VSAN datastore instead of a traditional Datastore makes sense. A huge amount of flexibility exists to either keep things as close as possible to the old data center or take advantage of the downtime and move to reorganize and re-architect.

Execution

Once you have completed the required planning outlined above, you're ready to begin executing the data center migration.

Replicate the Existing Environment to the New Data Center

Setting up replication is the first step in executing the migration. Following the high-level steps below allows operations to be available from the existing data center while seamlessly replicating data into the new data center *prior to the cut-over date*.

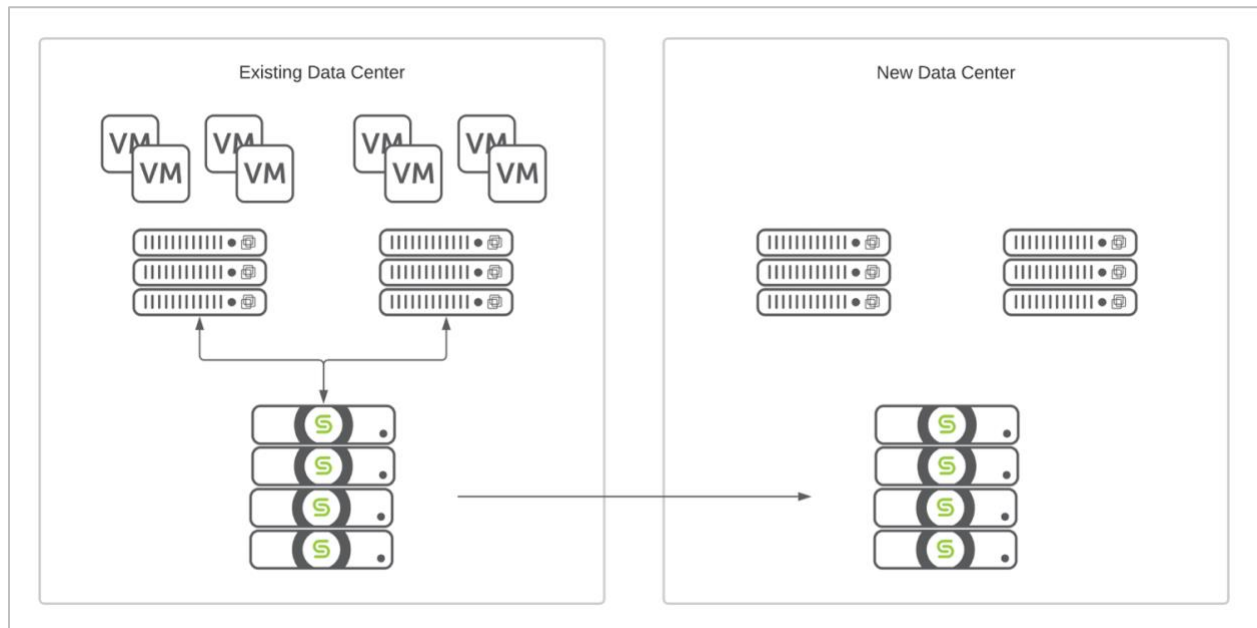
1. Ensure your Cohesity clusters at both data centers can communicate with each other.
2. Create or identify which Cohesity Storage Domains you will use at the new data center.
3. Pair Cohesity Storage Domains.
4. Modify or create a new Protection Policy that includes Replication between the Cohesity cluster at the old data center with the Cohesity cluster at the new data center.
5. Modify or create new Protection Groups to use the Protection Policy you created with Replication.
6. After the Protection Run, verify replication has completed.
7. If necessary, make adjustments to the Protection Group schedule to reduce the time between Protection Runs.

TIP: We recommend you do some testing prior to the actual cutover.

8. Before you cut over, power off the VMs to be migrated and manually run one last/final backup.

For details, see [About Replication](#) in the online Help.

Figure 4: Enable Replication Between Cohesity Clusters



Test Restoration Before to Cutover

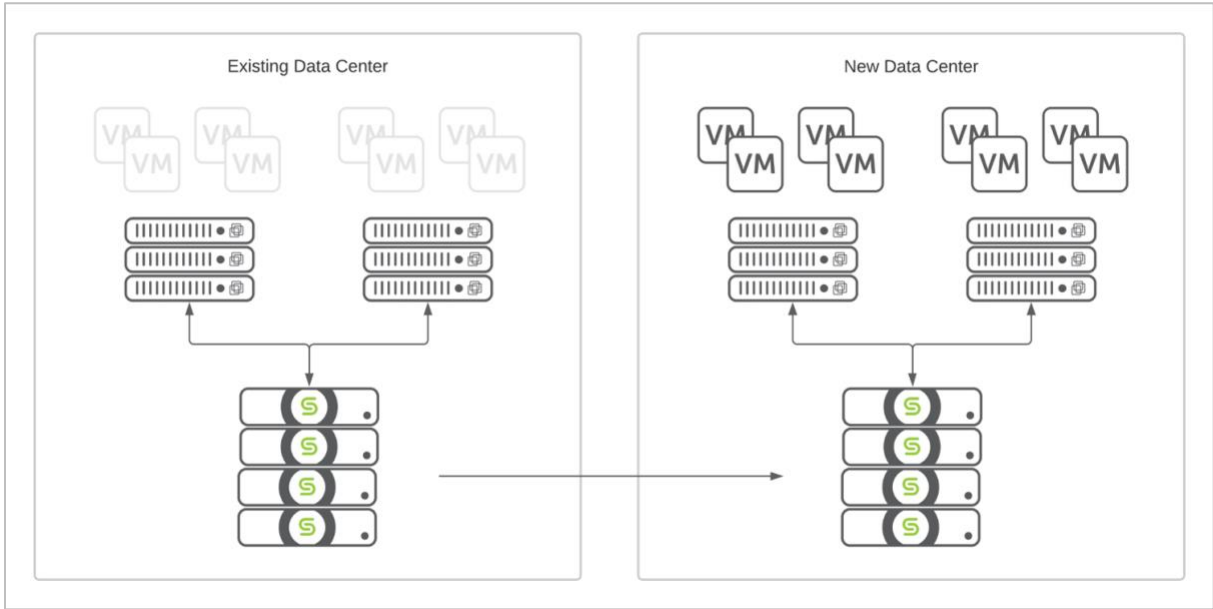
With Cohesity's Test/Dev Clone workflow, you can bring up VMs nearly instantly at the new data center to verify the new infrastructure or to test data integrity. This can be a great way to quickly verify how things will work during the actual cutover without using space on the primary storage.

Cloning and tearing down VMs is a quick and easy process that you can perform as often as needed. To prevent issues with duplicate IPs interfering with the production network, you can create a separate subnet for testing. You can then assign this separate testing network to the VMs being tested. If the VMs don't require network access for testing, you can simply restore or clone VMs with their vNICs disconnected. This is also an option during restore or clone workflows.

Cutover: Restore VMs to the New Data Center

As soon as replication completes and you're done with pre-cutover testing, you can begin restoring VMs. You can leverage [Cohesity's Instant Mass Restore](#) to bring up the VMs as soon as possible, without having to wait for them to fully restore to the primary storage array in the new data center.

Figure 5: Cutover of VMs

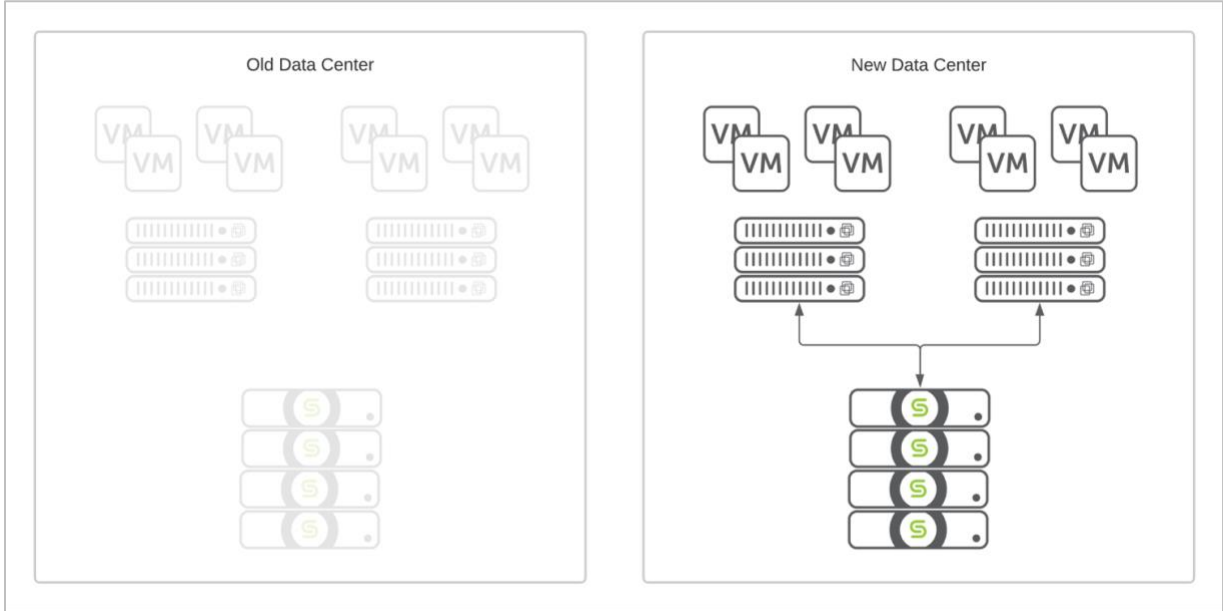


Complete the Transfer of Operations and Migration

The next steps to do a full transfer of operations from the existing data center to the new data center. Follow these high-level steps:

- 1. Once the restores have completed, perform integrity testing.
- 2. Set up new Protection Groups to protect the migrated VMs.
- 3. Decommission the old data center.

Figure 6: Transfer of Operations and Migration Completion



Conclusion

Cohesity is much more than just a backup solution or insurance policy. Because of its efficiencies with backup, replication, and restore, it can also complement data center migrations and consolidations. Data center migrations frequently involve much more than just moving data and VMs from one location to another. Often, the new data center will have different networking or storage requirements, different vCenters, or the need to organize VMs in a completely different way. On many occasions, it's also a good idea to test *before* the migration. Cohesity provides the workflows to accomplish all of this and more.

Your Feedback

Was this document helpful? [Send us your feedback!](#)

About the Authors

Justin Willoughby is Solutions Engineer at Cohesity. In his role, Justin architects, builds, tests, and validates business-critical applications, databases, and virtualization solutions with Cohesity's DataProtect platform.

Contributor:

- Bharath Nagaraj, Sr. Systems Engineer

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.1	Sep 2020	First public release

ABOUT COHESITY

[Cohesity](#) ushers in a new era in data management that solves a critical challenge facing businesses today: [mass data fragmentation](#). The vast majority of enterprise data — backups, archives, file shares, object stores, and data used for dev/test and analytics — sits in fragmented infrastructure silos that makes it hard to protect, expensive to manage, and difficult to analyze. Cohesity consolidates silos onto one web-scale [platform](#), spanning on-premises, cloud, and the edge, and uniquely empowers organizations to run apps on that platform — making it easier than ever to back up and extract insights from data. Cohesity is a [2019 CNBC Disruptor](#) and was named a [Technology Pioneer by the World Economic Forum](#).

Visit our [website](#) and [blog](#), follow us on [Twitter](#) and [LinkedIn](#) and like us on [Facebook](#).

© 2020. Cohesity, Inc. Confidential & Proprietary.

Cohesity, the Cohesity logo, SnapFS, SnapTree, SpanFS, and SpanOS, are registered trademarks, and DataPlatform, DataProtect, and Helios are trademarks of Cohesity, Inc. All rights reserved.

The information in this document is provided “as is,” for information purposes only, and is non-binding. Cohesity uses good faith efforts to ensure the information in this document is accurate at the time of publication but makes no representations or warranties of any kind with respect to it, and specifically disclaims implied warranties. Cohesity reserves the right to change its policies and practices without notice.

2000039-001-EN