

# DELL EMC DATA DOMAIN REPLICATOR

A Detailed Review

## ABSTRACT

Increasing frequency of catastrophic events like hurricanes, floods, fire, etc. have raised the urgency to have disaster recovery (DR) procedures. One of the most crucial steps for DR is to have a copy of the data at a remote site. To improve reliability of disaster recovery and meet stringent recovery time objectives (RTO) imposed by the business, organizations are increasingly replicating backups to create this offsite copy of their critical data. Reducing the amount of backup and archive data replicated through deduplication and compression reduces the network bandwidth required, and makes replication over existing networks economically viable. Dell EMC Data Domain Replicator software offers broad scalability for throughput and plug-in and provides the industry's most flexible and robust disaster recovery solution for the enterprise.

August, 2017

# Table of Contents

<b>EXECUTIVE SUMMARY</b> .....	<b>3</b>
INTRODUCTION .....	3
AUDIENCE .....	4
<b>DATA DOMAIN REPLICATOR OVERVIEW</b> .....	<b>4</b>
LEVERAGING LOGICAL STORAGE LAYERS TO MEET DIFFERENT REPLICATION REQUIREMENTS .....	5
Directory replication .....	5
Managed file replication .....	8
MTree replication .....	8
Collection replication .....	9
<b>CAPABILITIES OF DATA DOMAIN REPLICATOR</b> .....	<b>10</b>
ONLY DEDUPLICATED DATA .....	10
INDEPENDENT RETENTION POLICIES AT SOURCE AND DESTINATION .....	10
COMPRESSION .....	11
ENCRYPTION OF DATA-AT-REST .....	11
ENCRYPTION OF DATA-IN-FLIGHT .....	11
DATA DOMAIN RETENTION LOCK .....	12
DATA DOMAIN EXTENDED RETENTION .....	12
DATA DOMAIN CLOUD TIER .....	12
SECURE MULTI-TENANCY .....	13
FLEXIBLE REPLICATION TOPOLOGIES .....	13
NETWORK MANAGEMENT .....	14
STREAM MANAGEMENT .....	15
CONTENT AWARE REPLICATION .....	15
<b>CHOOSING BETWEEN REPLICATION APPROACHES</b> .....	<b>16</b>
<b>COMPARING DEDUPLICATION STORAGE: RPO, RTO, AND TIME-TO-DR</b> .....	<b>16</b>
RECOVERY POINT .....	17
RECOVERY TIME .....	17
TIME-TO-DR READINESS SUMMARY .....	18
<b>CONCLUSION</b> .....	<b>18</b>

## EXECUTIVE SUMMARY

Backups have long been used to provide operational recovery for business critical data. In addition to operational recovery, IT administrators also need to provide disaster recovery (DR) capabilities to protect against catastrophic natural events like hurricanes and floods or man-made disasters like chemical spills and terrorist attacks. Recent disasters like “Superstorm Sandy” have also increased the urgency and importance of providing DR capabilities for backup and archive data. While backups create a copy of the data, archiving moves the original version of the data off primary storage to secondary storage. Therefore, it is even more important to provide DR capabilities for archive data.

To enable disaster recovery, a copy of the data must be sent to an offsite location. Historically, backup tapes have been the primary method used to transport data to the DR site. However, in today’s age of rapid data growth, handling and tracking tapes introduces a significant management cost and complexity. Further, lost and misplaced tapes have forced IT to consider other approaches for disaster recovery.

In comparison, replication uses the wide area network (WAN) as the transport mechanism for data instead of tapes and trucks, which significantly reduces the cost, complexity and risk. Hence, replicating backups has become the preferred approach for enabling disaster recovery. Replication can also reduce the security risk by encrypting data in-flight over the WAN. However, not all replication is created equal and traditional replication is not suitable for moving backup and data due to the volume of data backed up and archived every day. Replicating backup and archive data requires deduplicated replication, which significantly reduces the network bandwidth required by sending only unique data over the network.

Dell EMC® Data Domain® leverages dynamically variable-length deduplication coupled with local compression and can eliminate up to 99 percent of the bandwidth used by normal replication methods. By lowering the cost floor for replication deployments, it encourages much broader use of this simplifying technology. Feedback from many customers applying deduplication to enable WAN replication has led Dell EMC to three critical conclusions about what matters in its deployment.

- **Speed:** Time-to-DR readiness is critical. While backup and archive applications do not require the synchronous behavior of transactional replication, they still have to be designed to meet or exceed the recovery requirements of the tape-centric solutions they replace. Data Domain design elements, such as true inline deduplication, continuous data consistency at the replica, and fast restore streams from replicas, all contribute to easy, fast recovery at the replica site and as quickly as possible after data is initially stored on the source system.
- **Flexibility:** Network characteristics like latency, bandwidth, and packet loss differ from one deployment to the other. It is important to offer choices in setting policies to achieve a balance in speed versus efficiency over these different network types. Further, modern enterprises have multiple remote offices and data centers and require a wide variety of replication topologies to support their DR needs. Data Domain Replicator offers numerous replication types and policies and also supports a wide variety of topologies to meet the needs of various deployments.
- **Simplicity:** The amount of data to be protected is growing rapidly. However, the number of employees responsible for handling these growing volumes of data is not increasing at the same rate. This means that replication needs to be simple to configure and manage, and not require any complex maneuvering. Data Domain systems offer policy-driven management of deduplicated replication and a graphical user interface to configure and monitor replication, thereby simplifying the life of the administrator.

## INTRODUCTION

This white paper introduces Dell EMC Data Domain Replicator software and explains how it delivers flexible replication topologies for enhanced disaster recovery in various enterprise environments. Read this white paper to find out how DD Replicator addresses DR needs for backup and archive data in centralized and distributed enterprises.

In the following sections, we will describe the unique characteristics of DD Replicator, including cross-site deduplication, as well as interactions with different capabilities of the Data Domain system.

## AUDIENCE

This white paper is intended for Dell EMC customers, system engineers, partners, and members of the Dell EMC and partner professional services community who are interested in learning more about Data Domain Replicator software.

## DATA DOMAIN REPLICATOR OVERVIEW

Within a Data Domain system, there are several levels of logical data abstraction above the physical disk storage, as illustrated in Figure 1. In the DD OS file system, protocol-specific namespaces are presented to clients/applications for accessing the logical file system layer. The files and directories within MTrees as well as MTree snapshots, all reference the same pool of unique segments, called a collection, which is made up of log-structured containers that organize the segments on disk to optimize throughput and deduplication effectiveness.




Figure 1: Data Domain system levels of Data Abstraction

These layers are described below:

1. **Protocol-specific namespaces:** As an external interface to applications, there are protocol namespaces, such as Data Domain Virtual Tape library (over Fibre Channel), Data Domain Boost storage units (for use with Dell EMC NetWorker, Dell EMC Avamar, Pivotal Greenplum, Symantec OpenStorage, NetVault, vRanger and Oracle RMAN), and CIFS/NFS fileshares (over Ethernet). A Data Domain deployment may use any combination of these simultaneously to store and access data.
2. **Data Domain file system: MTrees, Directories, and snapshots:** Files and directories for each namespace are stored in an MTree within the Data Domain file system. With DD VTL, the virtual tape cartridges are stored as files under special directories. MTree snapshots in Data Domain Operating System (DD OS) are logical; they share the same underlying data segments in the collection, and are very space-efficient.

3. **Unique segment collection:** A 'collection' is the set of files (or virtual tapes) and logical MTree snapshots. The system identifies and eliminates duplicates within each container and then writes compressed deduplicated segments to physical disk. Segments are unique within the collection (not including specific duplicates maintained in DD OS to enable self-correction or recovery). Each Data Domain system has a single collection that is stored in a log of segment locality containers. For more about segment localities, see the white paper [Dell EMC Data Domain SISL Scaling Architecture](#).
4. **RAID-enabled disk storage:** These collection containers layer over RAID-enabled disk drive blocks. Data Domain deduplication storage systems use Data Domain RAID 6 internal disk and storage expansion shelves to protect against dual disk failures.

## LEVERAGING LOGICAL STORAGE LAYERS TO MEET DIFFERENT REPLICATION REQUIREMENTS

Data Domain Replicator software offers four replication types that leverage these different logical levels of the system for different effects. All four replication types are designed to deal with network interruptions that are common in the WAN and recover gracefully with very high data integrity and resilience. This ensures that the data on the replica is in an application usable state. This is critically important for optimizing utility of the replica for DR purposes.

At a high level, the four replication types are:

- **Directory replication** transfers deduplicated changes of any file or subdirectory within a Data Domain file system directory that has been configured as a replication source to a directory configured as a replication target on a different system. Directory replication offers flexible replication topologies including system mirroring, bi-directional, many-to-one, one-to-many, and cascaded, enabling efficient cross-site deduplication but does not support more advanced features such as Secure Multi-tenancy, Data Domain Virtual Edition (DD VE), and Data Domain Cloud Tier (DD CT) for long term retention.
- **Managed file replication** is used by the DD Boost software option, for optimized levels of performance and integration with Dell EMC NetWorker, Dell EMC Avamar, Symantec OpenStorage and Oracle RMAN. Managed file replication directly transfers a backup image from one Data Domain system to another, one at a time on request from the backup software. The backup software keeps track of all copies, allowing easy monitoring of replication status and recovery from multiple copies. This form of replication provides the same cross-site deduplication effects and flexible network deployment topologies as directory replication.
- **MTree replication** is used to replicate MTrees between Data Domain systems. Periodic snapshots are created on the source and the differences between them are transferred to the destination by leveraging the same cross-site deduplication mechanism used for directory replication. This ensures that the data on the destination is always a point-in-time copy of the source with file-consistency. This also reduces replication of churn in the data leading to more efficient utilization of the WAN. MTree replication supports all the replication topologies supported by directory replication plus newer, more advanced features such as Secure Multi-tenancy, DD VE, and DD CT. MTree replication is also generally faster than directory replication and is recommended by Dell EMC in order for customers to take advantage of increased replication performance and all the advanced features that are available today and new MTree replication features that are planned for the future.
- **Collection replication** performs whole-system mirroring in a one-to-one topology, continuously transferring changes in the underlying collection, including all of the logical directories and files of the Data Domain filesystem. While collection replication does not support the flexibility of the other three types, it is very simple and lightweight, so it can provide higher throughput and support more objects with less overhead, which is ideal in high-scale enterprise cases.

A detailed examination of each replication type follows.

### DIRECTORY REPLICATION

With directory replication, a replication context pairs a directory (and all files and directories below it) on a source system with a destination directory on a different system, as seen in Figure 2. During replication, deduplication is preserved since

data segments that already reside on the destination system will not be resent across the WAN. The destination directory will be read-only as long as the replication context is configured.




Figure 2: Data Domain Directory Replication

The replication destination can contain other replication destination directories, replication source directories, and other local directories, all of which will share deduplication in that system's collection. As a result, directory replication offers a wide variety of topologies: simple system mirroring, bi-directional, many-to-one, one-to-many, and cascading. As illustrated in Figure 2 above, with directory replication, the source and destination can have independent collections.

In directory replication, the file transfer is triggered by a file closing, and the order of the closes is preserved. In cases where closes are infrequent, DD Replicator will force the data transfer periodically. As metadata and corresponding unique data segments are transferred, the files are separately created and maintained on the remote system—that is, the collection of the destination is independent of the source. Figure 2 shows that System A replicates to System B and each has its own separate collection. Once the destination system receives the complete file, it is immediately made visible to the namespace (CIFS, NFS, or VTL) at the destination and can be used for recovery purposes, writing to tape, etc.

In Figure 3 below, Metadata exchange between the source and destination ensures that a data segment only needs to be sent to the destination once, irrespective of where the data comes from. This provides significant efficiencies over the WAN in many-to-one deployments since common segments on different sources only need to be sent once.




Figure 3: Cross-Site Deduplication

The effect of cross-site deduplication provides WAN replication efficiencies comparable to the deduplication effect on storage, and the benefits aggregate in a multi-site topology. As an illustration, imagine that there are three sites, S1, S2, and R, in a two-to-one topology as seen in Figure 3. R is the destination for source directories replicating from sources S1 and S2. Assume that the replicating directories in S1 and S2 have identical data, but only S1 has replicated already, and S2 is just getting started replicating to R. The time and bandwidth required for S2's data to replicate to R are very small; the data is already there so just the metadata needs to transfer. The effect on bandwidth between S1's initial data being sent versus S2's redundant data is similar to the difference between the first full and the second full in local deduplication storage capacity used. S1 would typically have had to send data about a third the size of a full backup to synchronize to the replica system R; S2, with the same data coming later, would send about 1/60th the size of that same full backup.

While a secondary copy of data is sufficient for many organizations, some require a tertiary (or even greater number) copy, particularly in highly distributed enterprises. Creating an additional copy provides increased data protection and the ability to distribute data for multi-site usage. For example, Q/A testing content or training material can be reliably and efficiently replicated to different remote sites. DD Replicator supports two powerful replication topologies, one-to-many and cascaded, that enable the creation of multiple copies of data. One-to-many replication creates multiple copies from the source system, and cascaded replication creates copies from successive replication of the source system data. Combining these two provides the greatest flexibility in leveraging existing networks with complex topologies and varying bandwidths.

As shown below in Figure 4, one-to-many replication allows the same source directory to be replicated in parallel to multiple remote sites. Setting up one-to-many replication is similar to creating multiple independent replication contexts, one at a time, all with the same source directory. One-to-many replication allows a directory to be replicated concurrently to multiple remote systems. A replication context is created for each destination from the same source directory.




Figure 4: One-to-Many Directory Replication

With cascaded replication, a directory on a Data Domain system can be configured to be both the destination of one replication context and the source of another. This allows the replication of a directory (and all files and sub-directories) from Data Domain system A or Data Domain system B to Data Domain system C, and a subsequent replication of the same directory to Data Domain system D, as shown in Figure 5. This enables datasets to be replicated to two sites in sequential hops.

Cascaded replication allows directory replication contexts to be configured in a chain of sequential replication hops across multiple systems. Bi-directional replication is supported as shown in Figure 5 below.




Figure 5: Cascaded Directory Replication

Depending on requirements either one-to-many or cascaded replication may be preferable. For example, since one-to-many creates copies from the same source directory, data arrives at the destination sooner than cascaded replication, which requires the data to arrive at the intermediate system first before it is replicated to its final destination. Therefore, if the fastest speed to multiple copies is a priority, one-to-many may be preferable. However, if network bandwidth is limited at the source and/or because getting one copy offsite first is more critical for DR readiness than any additional copies, then cascaded replication may be preferable.

### MANAGED FILE REPLICATION

Managed file replication using DD Boost allows the backup software to control the replication on a per-file basis. When integrated with DD Boost, the backup software's users can configure policies to selectively replicate the individual backup image or dataset to another system after completion of the backup. Unlike traditional vaulting or cloning to tape, the data is not read by the backup server to be written elsewhere. Instead, the backup software delegates the data movement to the DD system; thereby leveraging the most efficient method available to create a DR copy of the data.

The backup software decides when to get started, and knows when it is finished, based on interactive signaling between DD Boost and the Data Domain system. Using this approach, the backup software knows that the destination holds a copy of the file that is separate and different from the source's file, and retention periods for the two can be managed independently, for example, to keep full backups longer on the DR site. Furthermore, the backup operator has the flexibility to decide which backup images need to be replicated, and which ones do not require DR protection; e.g. user may decide that daily incremental backups do not need to be replicated, but weekly full backups should be replicated offsite.

### MTREE REPLICATION

MTrees are user-defined logical partitions of the Data Domain file system that enable more granular management of the Data Domain filesystem. MTree replication enables the creation of disaster recovery copies of MTrees at a secondary location. In addition, one can also enable DD Retention Lock on an MTree-level at the source, which will get replicated to the destination.

MTree replication creates periodic snapshots at the source and transmits the differences between two consecutive snapshots to the destination. At the destination Data Domain system, the latest snapshot is not exposed until all the data for that snapshot is received. This ensures the destination will always be a point-in-time image of the source Data Domain system. In addition, files will not up show out-of-order at the destination and provides file-level consistency that

simplifies recovery procedures and reduces RTOs. Users are also able to create a snapshot at the source Data Domain system to capture a consistent point-in-time image (for example, after archiving a user's emails), which gets replicated to the destination where the data can be used for recovery in a disaster scenario.

MTree replication groups all changes between snapshots and replicates them. Consequently, any churn in the data between snapshots will not be replicated. This makes MTree replication suitable for applications that make frequent updates to the file system, for example filesharing and archiving workloads or certain backup applications that create, modify and delete temporary lock files within a short interval.

MTree replication has a lot of commonality with directory replication. It uses the same WAN deduplication mechanism used by directory replication to avoid sending redundant data over the WAN. It also supports all the topologies supported by directory replication (one-to-one, bi-directional, one-to-many, many-to-one, cascaded). In addition, one can configure MTree replication to replicate MTree data on a system that already leverages directory replication and/or managed file replication.

As stated earlier, MTree replication supports all the replication topologies supported by directory replication plus newer and more advanced features such as Secure Multi-tenancy, DD VE, and DD CT. MTree replication is also generally faster than directory replication and is recommended by Dell EMC in order for customers to take advantage of increased replication performance and all the advanced features that are available today and new MTree replication features that are planned for the future.

Dell EMC has added Directory Replication (DREPL) to MTree replication (MREPL) migration CLI commands to help customers migrate their older Directory Replication contexts to MTree replication so they can take advantage of all the advanced MTree replication features and improved replication performance. These commands include the ability to monitor the status of the migrations as they are being completed.

## COLLECTION REPLICATION

The fastest and lightest impact replication type is at the collection level. Unlike the prior three, there is no on-going negotiation between the systems regarding what to send. Collection replication is mostly unaware of the boundaries between files. Replication operates on segment locality containers that are sent once they are closed. By leveraging the log structure of the collection, collection replication tracks the delta between the head of the source and destination collections, and transfers each container containing unique segments, in order, until it catches up. This is illustrated below in Figure 6.




Figure 6: Collection Replication

Because there is only one collection per Data Domain system, this is specifically an approach to system mirroring. The destination system cannot be shared for other roles. It is read-only and shows only data from one source. Once the data is on the destination, files (and virtual cartridges) become immediately visible for recovery.

As previously described, the collection's container set is a log structure. Therefore, transferring data in this way means simply comparing the heads of the source and destination logs, and catching up one container at a time, as shown in Figure 6. If it gets behind, it will catch up later. This approach is very well adapted to enterprise deployments wishing to minimize the resource overhead of the selectivity and cross-site filtering overheads of directory or MTree replication (for

example for very large DR deployments using high-bandwidth WANs), or systems containing millions of files in an archiving deployment. Due to this lightweight approach, collection replication can provide a logical throughput of up to 52 Tb/hr on a 10 Gb network.

As shown in Figure 7, the final hop in a cascaded chain of replication contexts can also be configured to use collection replication when the entire contents of the intermediate system need to be replicated to a secondary DR site.




Figure 7: Collection Replication in a Cascaded Topology

## CAPABILITIES OF DATA DOMAIN REPLICATOR

The following capabilities of Data Domain Replicator software are important to be aware of when designing the DR solution.

### ONLY DEDUPLICATED DATA

In DD OS, data is deduplicated as it is written to the source system and replication preserves deduplication. This ensures that the network is efficiently utilized for creating a DR copy of backup and archive data.

### INDEPENDENT RETENTION POLICIES AT SOURCE AND DESTINATION

With Data Domain replication, it is possible to retain data for different periods on the source and destination; e.g. data may be retained for 30 days on the source Data Domain system and for 6 months on the destination Data Domain system. With DD Boost managed file replication, users can configure separate retention periods within the backup application for copies on the source and destination Data Domain systems. With directory replication, users can create independent snapshots on the source and destination systems after the backup has been replicated and retain these snapshots for the desired durations. With MTree replication, users can create a snapshot with the appropriate expiration time on the source system after completing the backup, and this snapshot gets replicated to the destination. Once the snapshot has been replicated to the destination, the user can modify the snapshot expiration time on this system.

## COMPRESSION

DD OS offers a choice of local compression types: LZ-style (default), GZ-fast, GZ, or no compression. Data transferred over the network using DD replicator is always compressed using the same algorithm as that of the destination. If source and destination have different compression types, then the data is first uncompressed at the source and then recompressed using the destination's algorithm before being sent across the network.

## ENCRYPTION OF DATA-AT-REST

Dell EMC Data Domain Encryption software allows the user to encrypt data at rest by using RSA BSAFE FIPS 140-2 compliant libraries with standard 128-bit or 256-bit Advanced Encryption Standard (AES) algorithms. Depending on IT security policies, the block cipher modes for the AES algorithm can be selected either as Cipher Block Chaining (CBC) or Galois Counter Mode (GCM). DD Replicator is compatible with DD Encryption and any data transferred over the network is always encrypted using the encryption key of the destination.

With collection replication, both the source and destination systems share the same encryption key. If source and destination have different encryption keys, then, as shown in Figure 8, data at the source is first decrypted, the source system will obtain the encryption key of the destination and data is re-encrypted using the destination's encryption key before sending the data across the network.




Figure 8: Replication source uses the destination's encryption key

## ENCRYPTION OF DATA-IN-FLIGHT

DD Replicator supports encryption of data-in-flight by using standard Secure Socket Layer (SSL) protocol version 3, which uses the ADH-AES256-SHA cipher suite to establish secure replication connections. As shown in Figure 9, encryption for data-in-flight can be configured for each individual context for directory and MTree replication. For managed file replication and collection replication, the setting to enable encryption of data-in-flight is set at the system level, and either all or none of the replicated traffic is encrypted.




Figure 9: Data being replicated over the WAN can be encrypted using SSL

Encryption over-the-wire has no impact on replication throughput for directory, MTree and managed file replication. With collection replication, there can be up to 50% performance degradation when encryption over-the-wire is turned on. This capability can be used in conjunction with data-at-rest encryption; the only caveat being that the data payload will be encrypted twice when this configuration is used.

## DATA DOMAIN RETENTION LOCK

Dell EMC Data Domain Retention Lock software provides immutable file locking and secure data retention capabilities for customers to meet both corporate governance and compliance standards (such as SEC 17a-4(f)). DD Retention Lock comes in two editions – Dell EMC Data Domain Retention Lock Governance edition and Dell EMC Data Domain Retention Lock Compliance edition. Both editions provide the capability for IT administrators to configure minimum and maximum retention periods at the MTree level and apply retention policies at an individual file level.

For DD Retention Lock Governance, collection replication, MTree replication, and directory replication can be used to replicate the retention attributes (locked or unlocked, retention period) of the files. If directory replication is used (for archive data within “backup” MTree), then the Min and Max retention periods configured on the “backup” MTree are not replicated.

When DD Retention Lock Compliance is used to store archive data, then MTree replication or collection replication can be used to replicate the retention attributes of the individual files and the associated MTree-level retention settings.

## DATA DOMAIN EXTENDED RETENTION

Dell EMC Data Domain Extended Retention software increases the storage scalability of a Data Domain system to enable cost-effective long-term retention of backup data on deduplicated disk. When a system is enabled with the DD Extended Retention software, data can be moved from the active tier to a cost-effective retention tier for long-term retention. DD Replicator is compatible with DD Extended Retention and is agnostic to the location of data; i.e. DD Replicator does not care if the data is in the active tier or the retention tier.

For a system configured with DD Extended Retention software, collection replication, MTree replication or managed file replication (for DD Boost) can be used to create a DR copy of the backup data. A system with DD Extended Retention software can be used as the destination for directory replication, MTree replication or managed file replication (for DD Boost). In such a deployment, the source system does not need to have the DD Extended Retention software (shown in Table 1). When directory replication is used to replicate data into a system with DD Extended Retention, then it is not possible to configure data migration policies to move the data from the active tier to the retention tier on the destination Data Domain system.

Source Data Domain System	Destination Data Domain System	
	Without DD Extended Retention	With DD Extended Retention
Without DD Extended Retention	Directory, managed file, MTree, collection	Directory (no data migration), MTree, collection, managed file
With DD Extended Retention	MTree, managed file	Collection, MTree, managed file

Table 1: Replication support with Extended Retention

## DATA DOMAIN CLOUD TIER

Data Domain Cloud Tier is a native feature of DD OS for moving data from the Active Tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention.

Directory replication only works on the /backup MTree, and this MTree cannot be assigned to the Cloud Tier. So, directory replication is not affected by Cloud Tier. Collection replication is not supported on Cloud Tier enabled Data Domain systems.

Managed file replication and MTree replication are supported on Cloud Tier enabled Data Domain systems. One or both systems can have Cloud Tier enabled. If the source system is Cloud Tier enabled, during replication, data may need to be read from the cloud if the file was already migrated to the Cloud Tier. A replicated file is always placed first in the Active Tier on the destination system even when Cloud Tier is enabled.

## SECURE MULTI-TENANCY

Secure Multi-tenancy (SMT) is a feature set that enables large enterprises and service providers to offer data protection as a service with Data Domain systems in a private, hybrid, or public cloud.

With secure Multi-tenancy, a Data Domain system can logically isolate data and restrict each tenant's visibility and read/write access to only their data. In addition, SMT provides management and monitoring by tenant to enable chargeback, trending and other reporting.

Only Mtree replication and Managed File replication are supported in a SMT environment.

MTree replication is supported on MTrees assigned to Tenant Units. During MTree replication, an MTree assigned to a Tenant Unit on one system can be replicated to an MTree assigned to a Tenant Unit on another system. When setting up SMT-aware MTree replication, security mode defines how much checking is done on the tenant. The default mode checks that the source and destination do not belong to different tenants (e.g., you do not need to assign a tenant to the MTree). The strict mode makes sure the source and destination belong to the same Tenant.

DD Boost managed file replication is supported between Storage Units, regardless of whether one Storage Unit, or both, are assigned to Tenant Units.

During DD Boost Managed File replication, Storage Units are not replicated in total. Instead, certain files within a Storage Unit are selected by the backup application for replication. The files selected in a Storage Unit and assigned to a Tenant Unit on one system can be replicated to a Storage Unit assigned to a Tenant Unit on another system. DD Boost Managed File replication can also be used in DD Boost AIR deployments.

To protect against man-in-the-middle (MITM) attacks when replicating over a public Internet connection, authentication can validate SSL certificate-related information at the replication destination and, optionally, at the replication source, providing secure replication over public Internet.

## FLEXIBLE REPLICATION TOPOLOGIES

To enable enterprise-wide data protection, DD Replicator provides multiple replication topologies – system mirroring, selective data replication, bi-directional replication, many-to-one replication, one-to-many replication and cascaded replication (see Figure 10). With many-to-one replication, up to 540 Data Domain systems in geographically distributed locations can replicate into a single DD9800 at the central data center.




Figure 10: Replication supports a wide variety of topology DR needs

## NETWORK MANAGEMENT

There are many network management capabilities that benefit all of the Data Domain replication approaches.

- Identifying status.** Dell EMC Data Domain System Manager provides a GUI for setting up replication and managing all replication choices. Figure 11 is an example pane that highlights current status and current replication completeness on a many-to-one directory replication destination node.



Figure 11: Data Domain System Manager GUI

To further monitor replication, additional views are available on a per-context basis. In these views, shown in Figure 12, additional focus is provided for the most common administrative questions regarding DR-readiness, such as synchronization status, along with a graph to quickly view the ongoing delay curves between data storage and data transmission rates.




Figure 12: Detailed views of replication configuration and status

- **Throttling.** As a basic form of quality of service (QoS), administrators may establish times of day during which data may or may not be sent, along with limits to the amount of bandwidth that can be used for replication. For more advanced QoS functionality, use of a more sophisticated system in the network itself is recommended.
- **Multi-streaming of a replication context for high-bandwidth networks.** To maximize the use of high-bandwidth WAN links, Data Domain systems can allocate multiple streams for each replication context. This improves the replication throughput and subsequently the time-to-DR readiness.
- **Optimization of a replication context for low-bandwidth networks.** For Enterprises with small datasets and 6 Mb/s or less bandwidth networks, DD Replicator can further reduce the amount of data to be sent using the low-bandwidth optimization mode. This enables remote sites with limited bandwidth to use less bandwidth or to replicate and protect more of their data over existing networks.
- **Secure connection authentication.** All connections use Diffie-Hellman key exchange between source and destination systems.
- **Replication-level data verification.** Data Domain Replicator uses its own large checksum to verify correctness of all sent data, above and beyond what TCP provides. The TCP embedded checksum is not strong enough for deduplication, and can have thousands of errors per week on many WANs. Additional verification is also provided at the storage level, as discussed in the [Dell EMC Data Domain Data Invulnerability Architecture](#) white paper.
- **Use of alternative ports.** By default, DD Replicator will use TCP port 2051 for replicating data between two Data Domain systems. However, administrators can configure any suitable TCP port to be used by replication.
- **Scripting and reporting tool integration.** Data Domain systems have a rich command line interface that includes full support for replication management. All event and status information is stored in ASCII logs for easy adoption by third-party reporting tools. Warnings and summary status are provided in email reports distributed by the Data Domain autosupport process. These reports are also sent to Dell EMC Support to provide history for optimal customer service on demand.

## STREAM MANAGEMENT

Data Domain Replicator provides comprehensive stream management on both source and target systems, which improves the quality of service. It can prevent poor performance when a system exhausts all of its available streams. Internally, there is an automated mechanism that can manage the number of replication streams used per context. DD replication source stream management, together with target stream management ensures that replication streams will be fairly distributed among replication contexts on both source and target systems.

Replication also supports stream quotas for an MTree replication context. Users can set the maximum number of streams an MTree replication context can use either on source, destination or both. Users can use stream quota to control how much resource (CPU/Memory/Network Bandwidth, etc.) an MTree replication context can use.

## CONTENT AWARE REPLICATION

Backup applications can write virtual synthetic full backups to Data Domain systems using the DD Boost protocol. The virtual full backups are synthesized from existing backups on the DD system, and provide significant performance improvements and network utilization reduction when writing backups.

DD Replicator applies the same synthesis optimizations for synthetic full backups to deliver similar performance improvements and network utilization reduction. The synthetic replication optimization is applicable with both Managed File replication and MTree replication.

Applications such as Dell EMC Avamar, Dell EMC Networker, and Dell EMC ProtectPoint benefit from the synthetic replication optimizations.

## CHOOSING BETWEEN REPLICATION APPROACHES

To determine the best replication approach for a given Data Domain deployment refer to Table 2 to figure out the supported replication types based on the protocol being used.

Protocol	Directory	Managed File	MTree	Collection
NFS/CIFS	✓		✓	✓
DD Boost		✓	✓	✓
VTL/NDMP	✓		✓	✓

Table 2: Replication protocol support

Then use the following guidelines to select the appropriate replication type.

- If all that is needed is one-way system mirroring between two sites, then use collection replication. It is the fastest and lowest impact method for DR readiness. In some large enterprise data centers, based on the size of the systems and link speeds involved, collection replication may be the most appropriate means of creating a DR copy of the data.
- If you are using applications with DD Boost support (e.g. Dell EMC NetWorker, Dell EMC Avamar, Pivotal Greenplum, Symantec OpenStorage, vRanger and NetVault Backup, and Oracle RMAN), then use managed file replication when available. It offers simple methods for more advanced policies, such as selectively replicating specific backup images and maintaining separate retention periods for original and replica images.
- If you are creating MTrees for logically partitioning the Data Domain file system, then use MTree replication to get the benefits of flexibility in topologies and WAN efficiency as well as better performance compared to directory replication.
- If you want to create snapshots on a source MTree (to get a consistent point in time copy) and have those snapshots available at the destination as well, use MTree replication.
- For a system with DD Extended Retention license, use MTree replication for creating DR copy of data written via CIFS, NFS, DD VTL or NDMP and managed file replication for replicating the data written via DD Boost.
- If the system contains compliant archive data (whether using DD Retention Lock Compliance edition or DD Retention Lock Governance edition), then use MTree replication for creating a DR copy of this data.
- If you want to leverage newer Data Domain features such as SMT, DD VE, or DD CT, then use MTree replication.
- For archived data stored on Data Domain systems, use MTree replication for creating the DR copy of this data.

## COMPARING DEDUPLICATION STORAGE: RPO, RTO, AND TIME-TO-DR

Two well-known metrics, recovery point objective (RPO, how old is the recoverable data at the replica?) and recovery time objective (RTO, how long does it take to recover the data to usability at the replica site?), are useful when considering replication techniques. To compare deduplication storage systems, consider one additional, composite metric. Starting with a suitably large full backup definition such as 20 TB, time how long it will take to:

1. Back up and deduplication at the original system
2. Replicate across an IP network
3. Restore the data from the replica to a different set of servers at the DR site

For the sake of brevity, call this composite metric time-to-DR. This is the end-to-end time from the beginning of a backup to completion of data recovery at the replica site. In proof-of-concept tests comparing new deduplication storage systems intended for DR use, this is the most telling indicator.

## RECOVERY POINT

The recovery point of the data at the replica will be older if it takes longer to replicate. Whatever most recent complete replica exists already, it will still be the best restorable copy until the new one gets there. For example, assume a new backup is starting, and the recovery point at that time is from yesterday's backup. In deduplicated replication, users typically only want to replicate the deduplicated (smaller bandwidth) data. In a Data Domain system, deduplication is fast and inline, and replication can be simultaneous with backup, so it can finish shortly thereafter. The restore image is available immediately on arrival at the replica.

In a slower deduplication-rate system, especially in one that delays the beginning of deduplication through being a "post-process" deduplication system, replication takes much longer. In a post-process system, the deduplication rate is typically less than half the ingest rate to non-deduplication storage (otherwise, no one would bother with the two-step process and its complexity, boundary conditions, and extra disk provisioning requirements). Since replication can only complete when deduplication is finished, this typically compromises the real arrival rate at the replica site. If half the deduplication rate of a Data Domain system, that means data can get to the replica site at no more than half the speed. So the restore point would be from yesterday plus two times the backup window, or worse.

Some systems compromise even further:

- Some deduplication systems do not compress on the WAN, resulting in either twice the bandwidth cost or twice the time for data to arrive at the replica.
- Some deduplication systems are not continuously consistent at the replica. They have to do a batch or manual process to enable the newly replicated data to be readable. The timing and effort required for this need to be taken into account in planning for recovery. For example, in a system that synchronizes in batch on a daily schedule, after deduplicated data is presumed to have crossed the wire, there could be an additional window where even though data has arrived, it is not restorable. Therefore, yesterday's backup remains the recovery point.

## RECOVERY TIME

On the replica, there is only deduplicated data. The recovery time is the same as the restore rate from the deduplication pool in the replica. This should be measured carefully with a large dataset to ensure sustained performance characteristics. Because of the SISL architecture, data domain deduplication storage provides fast restores at both the originator and replica. The only performance rates published by data domain are from deduplicated storage. Figure 13 below illustrates how various approaches to deduplication can drastically impact time-to-DR readiness.

Post-process systems, at the time of publication, do not provide their restore rates from their deduplication pools, especially on replicas. There seems to be a significant drop-off in performance from their rated specifications (which are all benchmarked against their non-deduplicated cache storage). Recovery time is somewhere between slower and infinite.




Figure 13: Deduplication approaches impact time-to-DR readiness

## TIME-TO-DR READINESS SUMMARY

When considering disaster recovery solutions for backup and archive data, it is essential to consider time-to-DR readiness in addition to the traditional RPO and RTO metrics. Time-to-DR readiness measures the time required to get data offsite and recoverable and helps to determine the potential data loss in a disaster recovery scenario. As shown in Figure 13, different approaches to deduplication dramatically impact how quickly data is replicated offsite and is therefore how quickly the system is DR ready. Data Domain systems deduplicate data inline and begin replication immediately, enabling the remote system to be DR-ready faster. Post-process approaches first ingest backup data to a disk cache and then perform deduplication, typically at < 50% of ingest speed. Consequently, the point at which these systems are DR-ready is delayed.

## CONCLUSION

Most large enterprise users require a global disaster recovery (DR) strategy that protects the entire organization by having one of more copies of their data at offsite locations. Dell EMC Data Domain Replicator software asynchronously transfers only the compressed, deduplicated data over the WAN, making network-based replication cost-effective, fast and reliable without requiring manual intervention.

The key features of Data Domain Replicator software are as follows:

- Safe and network-efficient replication
  - Cross site deduplication
  - Low-bandwidth optimization
  - Encrypted replication
  - Up to 99 percent bandwidth reduction
- Scalable replication throughput
  - Up to 52 TB/hr logical throughput

- Multi-stream optimization
- Enterprise deployment flexibility
  - Flexible replication topologies
  - Consolidate data from up to 540 remote Data Domain systems
  - Policy-based data management
  - Significant advantages in RPO, RTO, and fasted “time-to-DR”
- Easy integration
  - Supports leading enterprise applications for database, email, content management, and virtual environments
  - Compatible with all DD OS functionality – Compression, Encryption, Extended Retention, Retention Lock
  - Replicates CIFS, NFS, DD VTL, NDMP, and DD Boost data
- Advanced MTree replication features
  - Support for Secure Multi-tenancy for large Enterprises and Service Providers
  - Support for Data Domain Virtual Edition (DD VE)
  - Support for Data Domain Cloud Tier (DD CT) for long term retention to private or public clouds

**For additional information on Data Domain systems, please refer to:**

[Data Domain SMT Integrating Replication as a Service \(RaaS\) - Solution Brief](#) - White Paper

[Data Domain Encryption - A Detailed Review](#) - White Paper

[Data Domain Extended Retention - A Detailed Review](#) - White Paper

[Why DD Secure Multi-Tenancy](#) – Business Value Paper

[Secure Multi-Tenancy with Data Domain](#) – Lightboard Video


[Data Domain Virtual Edition Overview](#) – Lightboard Video

[Data Domain Virtual Edition for ROBO Environments](#) – Narrated Demo

[Data Domain Cloud Tier](#) – Narrated Demo

[Back to Basics: Data Domain Systems overview](#) – Video w/slides

[Why Data Domain](#) – Business Value Paper



[Learn more](#) about Dell EMC Data Domain solutions



[Contact](#) a Dell EMC Expert