

Replicator Software

Network-efficient Replication of Backup and Archive Data

Key Benefits

Enterprise Solution

- > Offsite vaulting for disaster recovery
- > Branch office data protection
- > Multi-site tape consolidation

Network-efficient Replication

- > 99% bandwidth reduction
- > Cross-site deduplication
- > Cost-efficient disaster recovery

Scalable Throughput

- > Fastest replication throughput
- > Up to 21 TB/hr logical throughput

Enterprise Deployment Flexibility

- > Flexible replication topologies
- > Policy-based data management
- > Multi-site disaster recovery

Improved DR Readiness

- > Immediate availability of replicated data
- > Fastest 'time-to-DR'

Easy Integration

- > Supports leading backup and archive applications from:

Symantec	EMC
HP	IBM
Microsoft	CommVault
BakBone	CA
Atempo	
- > Supports leading enterprise applications including:
 - > Database: Oracle, SAP, DB2, SQL
 - > Email: Microsoft Exchange
 - > Virtual environments: VMware
 - > Content management: Microsoft SharePoint
- > Simultaneous use of VTL, NAS and Symantec OpenStorage (OST)

Most large enterprise users require a global disaster recovery (DR) strategy that protects the entire organization by having one or more copies of their data at offsite locations. Traditionally, backup applications are used to copy data to tape, which are then shipped offsite. This labor-intensive process is error prone, introduces security risks and is extremely slow for recovery. Network-based alternatives are limited because of the large size of the data, whether geographically centralized or distributed. There is just not enough affordable bandwidth or time in the day to move backup and archive data over a conventional Wide Area Network (WAN).

Data Domain Replicator software provides network-efficient, high-throughput, automated, ultra-safe replication for DR, remote office data protection, multi-site tape consolidation and long-term, offsite retention. Replicator asynchronously transfers only the compressed, deduplicated data over the WAN, making network-based replication economically and operationally feasible.

Safe and Network-efficient

With Replicator, all or selected business-critical backup and archive data is safely and efficiently replicated from one system over the WAN to another system at a secure offsite location. Data Domain deduplication storage systems first deduplicate data at the originating Data Domain system, reducing the volume of data stored by 10-30x, on average. Replicator sends only these new and unique data segments, significantly reducing the amount of data sent to or from remote locations.

Reduce daily network bandwidth requirements by up to 99%

Data Domain Replicator performs two levels of bandwidth reduction: local and cross-site deduplication.

Data Domain deduplication massively reduces data volume stored locally, thereby reducing the amount of data that needs to be replicated. Typically less than 1% of a full backup, for example, is actually new, unique compressed sequences to be replicated over a WAN.

Cross-site deduplication takes place when multiple sites replicate to the same destination system. Any redundant segment previously transferred by any remote site or as a result of a local backup is not replicated again by any other remote site, saving critical network bandwidth. Cross-site deduplication improves network efficiency across all sites and reduces the storage needed in the remote disaster recovery site, further contributing to the deduplication efficiency at the destination.

Scalable Replication Throughput

Data Domain Replicator software is an ideal DR solution for enterprise data centers that have large amounts of backup and archive data and which require high-throughput replication bandwidth to get the data offsite.

Replicator provides significant throughput scalability to meet these demanding requirements. Logical throughput performance can scale up to 21 TB/hr over a 1 Gb network in replication deployments where one Data Domain system is mirroring its data to another.

High-speed replication is also obtained by multi-streaming data further maximizing the utilization of available network bandwidth and reducing the impact of network packet loss on replication throughput.

Enterprise Deployment Flexibility

Replicator provides a robust and flexible enterprise-ready replication solution allowing a variety of optimizations to focus on different topologies, data management policies and multi-site recovery options.

Flexible replication topologies

Replicator allows enterprises to deploy multiple replication topologies, such as full system mirroring, selective, bi-directional, many-to-one and cascaded.

Flexible deployment options such as site-to-site, bi-directional replication provide added disaster protection, allowing each site to keep local stores while serving as the secure recovery site for the other location. All Data Domain systems can simultaneously host local stores and host replicated images from other sites.

Replicator Software

Network-efficient Replication of Backup and Archive Data

Using Replicator software, up to 180 geographically distributed locations can simultaneously replicate selected backup and archive data to a single large, high-end Data Domain system at a central hub. This high fan-in scalability enables a flexible data retention model and enterprise-wide site recovery. For increased enterprise-wide data protection, more copies of this data can be replicated to additional disaster recovery sites in a cascaded fashion with the added flexibility to replicate the entire system to the final location. Enterprise-wide backup and recovery can now be achieved orders of magnitude faster, more cheaply, more securely and more reliably than with tape, globally across all sites.

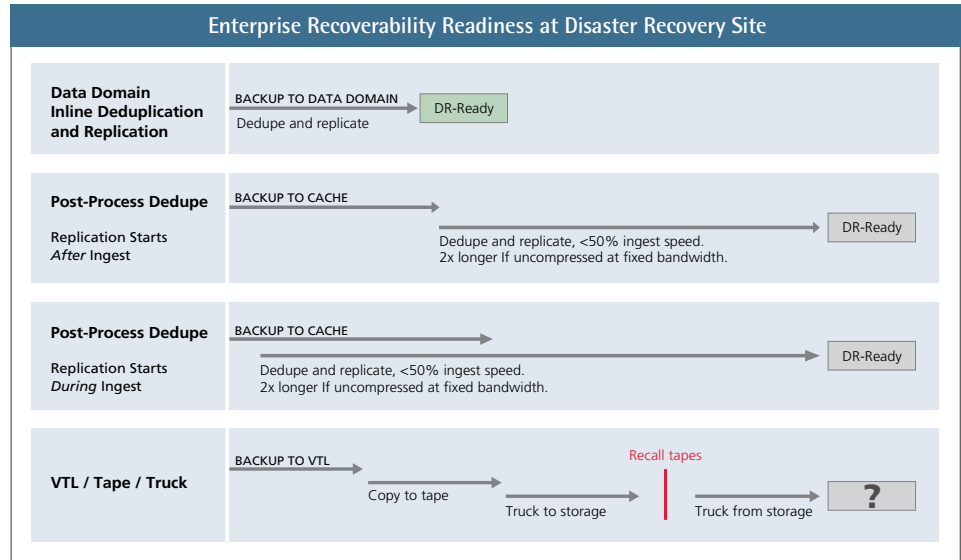
Policy-based data management

Replicator software allows users to choose all or a subset of data to be replicated along with the type of compression used. Policies range from complete system mirroring to selectively choosing one or more directories or virtual tapes to individual backup policies using Symantec NetBackup OpenStorage optimized duplication.

With schedule-based network bandwidth throttling, data can be replicated only when the network is the least used by other applications. It provides the ability to control when the data can be replicated and how much network bandwidth it should use at different times of the day. WAN links used for mission-critical applications can now be safely shared, removing the need to dedicate costly WAN links just for data protection and archive traffic.

Multi-site disaster recovery

Once replicated across the WAN, data can be recovered or copied to tape from the Data Domain storage system at either location. Fast online recovery is possible using the local system or, if there is a problem with the onsite originating system, a server can access the replicated data over the WAN to get key information back onsite quickly. This flexibility



Today, enterprises backup data to enable both onsite and offsite recovery in the event of a disaster. Data Domain systems deduplicate this 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. Regardless of whether the deduplication and replication begins during or after ingest, the point at which these systems are DR-ready is delayed.

also enables organizations to consolidate the creation of tape to a single location to further improve their operational effectiveness.

Improved DR Readiness

The end-to-end time from beginning of backup to completion of data recovery at the DR site is an important metric in defining an organization's readiness after a disaster. The shorter the time, the faster the enterprise can resume operations. The sooner an organization is prepared, the less exposure in the event of the disaster. There can be a high degree of variability in this 'time-to-DR' metric between deduplication products so it is important to understand implementation details that will ultimately dictate the outcome of your DR readiness.

Fast, inline deduplication combined with the high-throughput replication inherent in Replicator software enables data to quickly and automatically be replicated once the new deduplicated data is stored on the originating system. This replicated data is then immediately usable at the disaster recovery site.

Next, the faster the restore time from the replicated data, the sooner the disaster recovery process can continue. The larger the data set, the more critical the restore time. Because of the Data Domain SISL™ (Stream-Informed Segment Layout) scalability architecture, users can realize high-throughput restore performance on large data sets for both the DR and the originating Data Domain system.

Using Data Domain as a DR solution gives an organization the greatest flexibility in defining their disaster recovery strategy.

Easy Integration

Data Domain Replicator is qualified with leading enterprise backup and archive software and easily integrates into existing enterprise infrastructures. Additional deployment flexibility exists with support for multiple simultaneous access protocols including using Data Domain VTL over Fibre Channel, through NFS and CIFS file service protocols over Ethernet, or as a disk-based target using application-specific interfaces such as Symantec NetBackup OpenStorage.