

EMC DATA DOMAIN REPLICATOR SOFTWARE

Network-efficient replication for backup and archive data

ESSENTIALS

Next-generation Disaster Recovery

- Offsite vaulting for disaster recovery
- Branch office data protection
- Multisite tape consolidation

Safe and Network-Efficient Replication

- 99 percent bandwidth reduction
- Cross-site deduplication
- Cost-efficient disaster recovery
- Low bandwidth optimization

Scalable Replication Throughput

- Up to 52 TB/hr logical throughput
- Multi-stream optimization

Enterprise Deployment Flexibility

- Flexible replication topologies
- Point-in-time replication
- Policy-based data management
- Multisite disaster recovery
- Encrypted replication

Improved DR Readiness

- Immediate availability of replicated data
- Fastest “time-to-DR”

Easy Integration

- Supports leading backup and archive applications
- Supports leading enterprise applications for database, email, content management, and virtual environments
- Simultaneous use of EMC Data Domain Boost, VTL, CIFS, NFS, and NDMP

NEXT-GENERATION BACKUP AND RECOVERY

Most large IT organizations 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 customers write their backup and archive data to tapes and ship them offsite for DR. However, 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 volume of 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).

EMC® Data Domain® Replicator software provides fast, network-efficient and encrypted replication for DR, remote office data protection, multisite tape consolidation, and long-term offsite retention of both backup and archive data. DD Replicator asynchronously transfers only the compressed, deduplicated data over the WAN, making network-based replication cost-effective, fast, and reliable.

SAFE AND NETWORK-EFFICIENT REPLICATION

With DD 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. EMC Data Domain deduplication storage systems first deduplicate data at the originating Data Domain system, reducing the volume of backup data stored by 10 to 30 times, on average. DD Replicator sends only these new and unique data segments, significantly reducing the amount of data sent to or from remote locations.

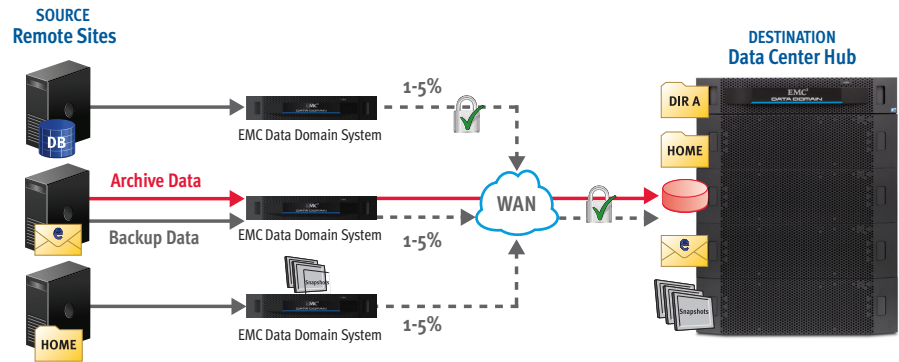
REDUCE DAILY NETWORK BANDWIDTH REQUIREMENTS BY UP TO 99 PERCENT

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 one percent 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 or archive will not be replicated again, 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.

95-99 Percent Cross-Site Bandwidth Reduction

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 or archive will not be replicated again, saving critical network bandwidth. Cross-site deduplication improves network efficiency across all sites and reduces the storage needed at the destination, further contributing to deduplication efficiency.



For enterprises with small datasets and 6 Mbps or less bandwidth networks, DD Replicator can further reduce the amount of data to be sent by using a low bandwidth optimization mode. This mode uses a delta compression technique to further compress the deduplicated data and eliminate redundant data within segments before sending it over the WAN. This enables remote sites with limited bandwidth to replicate and protect more of their data over existing networks.

SCALABLE REPLICATION THROUGHPUT

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

This replication software provides significant throughput scalability to meet these demanding requirements. Logical throughput performance can scale up to 52 TB/hr over a 10 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

Data Domain Replicator is a robust enterprise-ready replication solution that provides the flexibility to optimize a variety of topologies, data management policies, and multisite recovery options.

FLEXIBLE REPLICATION TOPOLOGIES

Data Domain Replicator provides multiple replication topologies, such as full system mirroring, selective, bi-directional, many-to-one, one-to-many, 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 replicated images from other sites.

Up to 270 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 one-to-many or cascaded fashion with the added flexibility to replicate the entire system to the final location. Secure, reliable, enterprise-wide disaster recovery needs for both backup and archive data can now be achieved globally across all sites with less time and cost than with tape.

Additionally, DD Replicator enables data distribution for multisite usage using one-to-many replication. Development or QA/testing can reliably and efficiently replicate the same content to different remote sites.

POINT-IN-TIME REPLICATION

DD Replicator offers a replication type that leverages snapshots to ensure that the destination Data Domain system is always a point-in-time image of the source Data Domain system.

With this replication type, snapshots created on the source Data Domain system are automatically replicated to the destination. To simplify DR readiness verification, users can create snapshots on the source Data Domain system at defined intervals (for example, upon completion of a backup) to easily verify that a data set has been completely replicated to the destination Data Domain system.

POLICY-BASED DATA MANAGEMENT

DD Replicator software lets you 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 Data Domain Boost backup application-managed replication.

With schedule-based network bandwidth throttling, data can be replicated only when the network is the least used by other applications. This allows you 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.

MULTISITE 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 also means you can consolidate the creation of tape to a single location to further improve operational effectiveness.

ENCRYPTED REPLICATION

Organizations with stringent security compliance requirements may be mandated to securely transfer data over public or even private networks. For the highest level of security, DD Replicator can encrypt data being replicated between Data Domain systems using the standard Secure Socket Layer (SSL) protocol, independent of the replication type and topology used.

While the EMC Data Domain Encryption software option can be used to encrypt data when stored at-rest on the Data Domain storage system, DD Replicator encrypts and de-encrypts data in-flight when being replicated between Data Domain systems. Encryption of data-at-rest and data-in-flight can be used concurrently to achieve different security goals.

For secure data transfer, encrypted replication uses standard 256-bit Advanced Encryption Standard (AES) encryption key strength with CBC Cipher Block Mode, which can be enforced on all or only a selected portion of the replicated data set. Data Domain systems' point-to-point encrypted replication is also a simple and highly secure alternative for remote sites, which cannot rely on or justify deploying dedicated VPN network devices.

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. There can be a high degree of variability in this "time-to-DR" metric between deduplication products, so it is important to

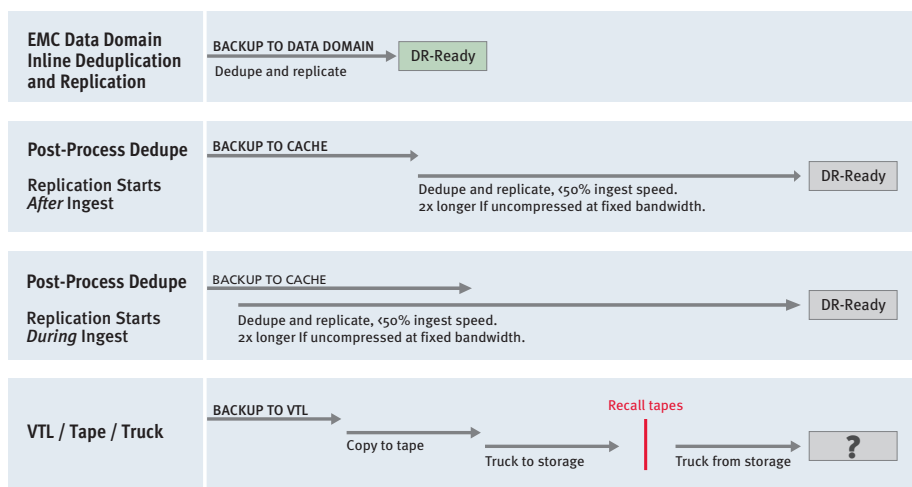
understand implementation details that will ultimately dictate the outcome of your DR readiness.

Fast, inline deduplication, combined with the high-throughput replication inherent in DD 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.

Time-to-DR Readiness

Today, enterprises back up and archive data to enable both onsite and offsite recovery in the event of a disaster. Data Domain systems deduplicate this data inline—before storing to disk—and can begin replication quickly, enabling the remote system to be DR-ready faster.

Post-process approaches first ingest data to a disk cache and then perform deduplication, typically at less than 50 percent 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.



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

Using Data Domain deduplication storage as a DR solution provides the greatest flexibility in defining a disaster recovery strategy for the user's specific requirements and service-level agreements.

EASY INTEGRATION

Data Domain Replicator is qualified with leading enterprise backup and archive software applications and easily integrates into existing enterprise infrastructures. Additional deployment flexibility exists with support for multiple simultaneous data access methods including the use of EMC Data Domain Virtual Tape Library software over Fibre Channel, or NFS and CIFS file service protocols over Ethernet, or as a disk-based target using application-specific interfaces such as EMC Data Domain Boost.

CONTACT US

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, contact your local representative or authorized reseller—or visit us at www.EMC.com.

EMC², EMC, Data Domain, SISL, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.
© Copyright 2011, 2012 EMC Corporation. All rights reserved. Published in the USA. Data Sheet 05/12 H6803.4