

EMC DATA DOMAIN OPERATING SYSTEM

Scalable, fast, and ultra-safe operating environment

ESSENTIALS

Scalable Deduplication

- Fast, inline deduplication with up to 26.3 TB/hour of throughput
- Extended retention providing up to 28.5 PB of logical storage
- 10 to 30x average data reduction

Easy Integration

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

Ultra-Safe Storage for Reliable Recovery

- Continuous recovery verification, fault detection, and healing
- Dual disk parity RAID 6

Operational Simplicity

- Lower administrative costs
- Reduced footprint; power, cooling, and space efficiencies for green operation
- Multisite disaster recovery
- Supports any combination of backup and archive applications in a single system

Conventional disk backup storage systems are more reliable than tape, and help backups run faster, but because they don't reduce data size, they are generally much too expensive to replace the retention of a tape library. They are usually just a front-end cache to a tape library infrastructure, temporarily alleviating backup window problems. By not reducing data size, they leave data volumes too large to replicate affordably. They also lack robust data invulnerability because they are based on primary storage designs, which favor random block update speed over data invulnerability. This means recovery and offsite disaster recovery (DR) protection still depend on tape. How did your last tape-based DR audit go?

EMC® Data Domain® Operating System (DD OS) software has revolutionized backup, archiving, network-efficient DR, and remote office data protection with high-speed, inline deduplication and consistent, integrated system management tools. DD OS enables the simplest approach to going tapeless for operational recovery and longer term retention.

SCALABLE DEDUPLICATION STORAGE

DD OS offers the most aggressive inline data deduplication available for backup/recovery data, averaging 10 to 30x data reduction. As data is written to an EMC Data Domain deduplication storage system, it is quickly scanned for patterns that have been stored before. Large patterns across the history of all data stored are identified in the Data Domain filesystem regardless of application or workload. DD OS stores each unique data sequence only once and saves significant physical storage capacity by substituting small references for each identical redundant sequence.

Local compression scans the unique data sequence for small strings across a local small window of comparison, like a tape drive. The combination of both deduplication and compression provides the resulting optimized data reduction.

SCALABLE DATA PROTECTION

DD OS supports petabytes of storage for a typical enterprise data set and backup policy. Up to 28.5 PB of logical storage is available in the EMC Data Domain Global Deduplication Array for long term disk-based retention. Multiple months of retention on disk is now possible using the same number of "floor tiles" that would normally provide only a couple days of disk staging. Snapshot technology further enables extended local and offsite retention on disk.

HIGHEST THROUGHPUT DATA REDUCTION

DD OS provides modular, scalable throughput to minimize backup and recovery times. The Global Deduplication Array offers up to 26.3 TB/hour of throughput, and the EMC Data Domain Appliance Series offers up to 14.7 TB/hour of throughput, significantly exceeding LTO-5 performance. Modular scalability enables simple expandability as your business challenges grow.

ULTRA-SAFE STORAGE FOR RELIABLE RECOVERY

CONTINUOUS PROTECTION

The Data Domain Data Invulnerability Architecture provides the industry's best defense against data integrity issues. Continuous recovery verification, along with extra levels of data protection, continuously detect and protect against data integrity issues during the initial backup and throughout the data life cycle. Unlike any other enterprise array or filesystem, each appliance ensures recoverability is verified and then continuously re-verified.

DUAL DISK PARITY RAID

DD OS includes extra levels of data protection to protect itself against storage related faults that threaten data recoverability. Dual disk parity RAID 6 is part of the foundation for DD OS continuous fault detection and healing. RAID 6 protects against two simultaneous disk faults, can rebuild a failed disk even if there are read errors on other sectors, and can detect and correct errors on-the-fly during reads. This added protection ensures the highest levels of data availability.

UNIQUENESS VERIFICATION

In determining global uniqueness, DD OS leverages very strong cryptographic hashing for speed and security, but it does not stop there. A universal hash ensures against random and malicious hash collisions. A separate independent algorithm for verification ensures early warning in the event of a challenge and adds additional data protection.

EASY INTEGRATION INTO EXISTING INFRASTRUCTURES

DD OS simultaneously supports multiple access methods. It presents itself as a file server, offering NFS or CIFS access over Ethernet, as a virtual tape library (VTL) over Fibre Channel, as an NDMP tape server over Ethernet, or as a disk-based target using application specific interfaces like EMC Data Domain Boost (for use with EMC Avamar®, EMC NetWorker® and Symantec OpenStorage). Increasingly, storage software is taking advantage of file servers or VTLs as optimized disk stores because they are so simple to manage and access. This deployment flexibility enables IT organizations to rapidly adjust to changing business requirements. DD OS is qualified with all leading enterprise backup and archiving software and other business critical applications.

OPERATIONAL SIMPLICITY

LOWER ADMINISTRATIVE AND OPERATIONAL COSTS

Because of DD OS, Data Domain systems are very simple to install and manage. Administrators can access DD OS through command line over SSH or through EMC Data Domain Enterprise Manager, a browser-based graphical user interface. Initial configuration and configuration updates can easily be made for multiple systems along with the monitoring of system states and the state of system operations. Simple scriptability along with SNMP monitoring provides additional management flexibility.

All Data Domain systems have an automatic call-home system reporting feature, called auto-support, which provides email notification of complete system status to Data Domain and to a selected list of administrators. This non-intrusive alerting and data collection capability enables proactive support and service without administrator intervention, further simplifying ongoing management.

Because of the massive data reduction, less physical equipment is required. This makes the physical footprint significantly smaller, and consequently the systems are energy efficient because they require less power and cooling.

MULTISITE DISASTER RECOVERY

Once you back up to a Data Domain system running DD OS, it is easy and cost-effective to use the system for local retention and to limit the use of tape to provide offsite and archive support. DD OS protects your data with network-efficient and encrypted replication, which provides remote office data protection and enables multi-site tape consolidation.

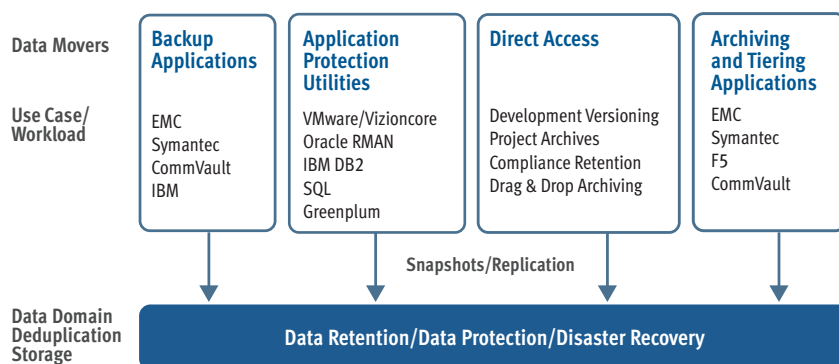
Data Domain systems replicate only the deduplicated and compressed unique changes across any IP network, using up to 99 percent less bandwidth. A Data Domain system facilitates a copy of the entire retained data set, online and disaster protected. If multiple systems replicate to the same destination, the destination will only store each segment uniquely across all inbound replication streams, further minimizing bandwidth and storage. If confidentiality is required, deduplicated and compressed data can be encrypted in-flight when being replicated between Data Domain systems, independent of the replication topology used.

EFFICIENT DATA ARCHIVING

Data Domain deduplication storage can be used to efficiently archive data. Data Domain systems reduce the overall cost of storage by utilizing less space and power for improved data archiving and retention. In addition, you can leverage Data Domain storage infrastructure for both backup and archive applications in the same system. This improves the efficiency across backup and archive applications and data types, as well as reduces management overhead by combining multiple applications' storage on a single system.

BACKUP AND ARCHIVE APPLICATION SUPPORT

Data can be moved onto an EMC Data Domain system in a variety of ways. For data protection, all leading backup applications are supported, as are applications for archiving data. Users can also make Data Domain systems the target for application protection utilities like Vizioncore for virtual environments or Oracle RMAN for databases. Systems can even be accessed directly to support additional workloads. All of these data movers and workloads can be supported in the same Data Domain system, at the same time:



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, Avamar, Data Domain, NetWorker, 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 EMC Corporation. All rights reserved. Published in the USA. Data Sheet 08/11 H6808.2

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000
In North America 1-866-464-7381
www.EMC.com

EMC Backup Recovery Systems
Santa Clara, California 95054
1-408-980-4800
In North America 1-866-933-3873

EMC²