

EMC DATA DOMAIN ARCHIVER

Long-term retention system for backup and archive

ESSENTIALS

Cost-Optimized Scalability

- Provides up to 570 TB usable capacity in two storage tiers
- Fast, inline deduplication with up to 9.8 TB/hour of throughput
- Enables tape reduction

Long-Term Data Retention Architecture

- Transparent separation of data in tiers by incorporating a large archive tier
- Archive units sealed for fault isolation

Ultra-Safe Storage for Reliable Recovery

- Continuous recovery verification, fault detection, and healing
- Dual disk parity RAID 6
- Storage of last resort with fault isolation

Operational Simplicity

- Replication and retention managed from backup application console with EMC Data Domain Boost
- Simple system administration

Storage and Network Efficiency

- Power, cooling, and space efficiencies for green operation
- Supports any combination of backup and archive applications in a single system
- Up to 99 percent bandwidth reduction

Easy Integration

- Supports leading backup and archive applications
- Simultaneous use of CIFS, NFS and EMC Data Domain Boost

Enforced Retention

- Satisfy IT governance policies
- Retain files on disk in a non-rewriteable and non-erasable format

EMC® Data Domain® deduplication storage systems have revolutionized disk backup, disaster recovery, and remote office data protection with high-speed, inline deduplication. Backup data can be reduced in size by an average of 10-30x making disk backup storage cost-effective for short-term retention to meet the operational recovery needs. However, many still rely on tape technology for long-term data retention due to perceived advantages in total cost of ownership, energy savings, and media portability. In reality, these perceived advantages do not materialize due to tape infrastructure costs, personnel costs incurred to manage the infrastructure, and media migration costs required every 5 to 7 years to upgrade to newer tape drive technology. In addition, from an operational point of view, accessing a file stored on tape takes a significant amount of time and the measured failure rate for restoring data from tapes is relatively high.

The best alternative to storing long-term retention data on tape has been using specialized archiving processes or applications to move the data from primary storage to an archive platform, but this has not yet become as universal a practice as backup. As a consequence, when customers are faced with a long-term retention requirement, keeping copies of backups longer on tape is the dominant approach despite the known operational challenges.

EMC Data Domain Archiver, the industry's first system for long-term retention of backup and archive data, extends the Data Domain architecture with an internal tiering approach designed to enable cost-effective long-term retention of data on deduplicated disk. Unlike other archive storage platforms, it offers all the advantages of a traditional Data Domain system in backup throughput and resilience to address long-term backup retention. Unlike other backup platforms, it offers significant optimization for long-term cost efficiency and granular fault-containment and recovery. DD Archiver enables customers to transition from using backup processes for long-term data retention through the evolving adoption of archiving best practices, while minimizing reliance on tape infrastructure in the data center.

COST-OPTIMIZED SCALABILITY

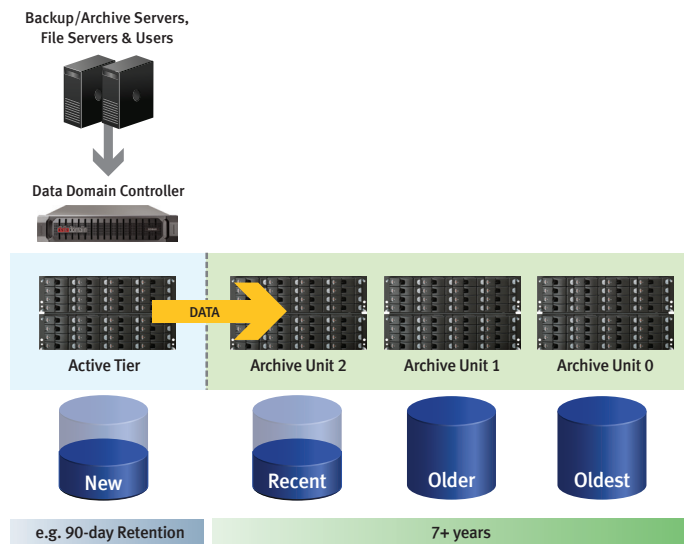
Data Domain Archiver transparently incorporates two tiers of storage in order to achieve cost-effective scalability while delivering the high throughput required to ingest hundreds of terabytes of backup data. DD Archiver provides up to 9.8 TB/hour throughput and scales up to 570 TB of usable capacity and up to 28.5 PB of logical capacity. This combination of high throughput and cost-optimized storage built on proven Data Domain system technology positions DD Archiver as the ideal tape reduction solution for long-term retention of backup and archive data.

LONG-TERM DATA RETENTION ARCHITECTURE

Data Domain Archiver provides a transparent separation of short-term and long-term data by storing it in different tiers. Data is initially stored in the active tier, which operates similarly to a standard Data Domain system for backup and operational recovery. Based on user-defined policies, the aged data is then moved to a very large archive storage tier that is optimized for long-term data retention—usually measured in years.

Long-Term Retention Architecture

Backup or archive data initially lands in the active tier, which operates similarly to a standard Data Domain system for backup and operational recovery. Based on user-defined policies, the aged data is then moved to a large archive storage tier that is optimized for long-term data retention—usually measured in years.



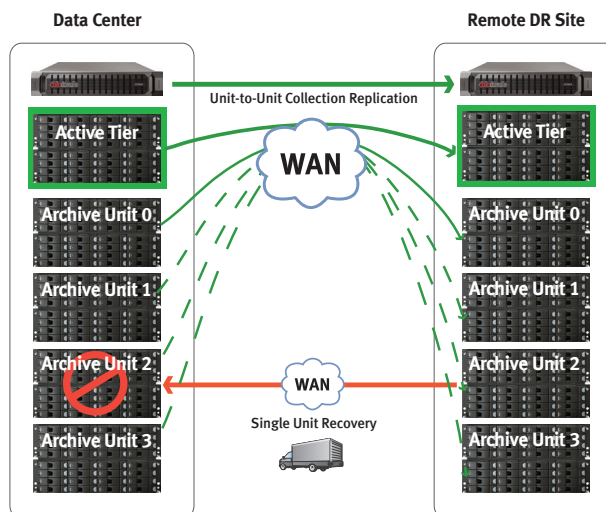
DD Archiver is designed as the storage of last resort—a petabyte-scale system that is the persistent location for all data when it comes to protection and retention. The system can grow to a large scale and operate for a much longer lifespan than traditional storage systems. To ensure long-term data access and recoverability, DD Archiver offers a new and unique “fault isolation” feature. Specifically, the archive tier consists of multiple archive units. Each archive unit is sealed when it reaches full capacity and becomes self-contained to ensure long-term data preservation. As a result, even if a failure or a catastrophe affects major components of the DD Archiver, the system continues to operate with all unaffected components.

Leveraging this new architecture, DD Archiver also provides granular replication configurations for disaster recovery. In the event of a connectivity issue or a disaster affecting replication for example, smaller and more manageable data units can be replicated or recovered instead of replicating or recovering petabytes of data.

Disaster Recovery

Instead of having a system-to-system replication relationship, the DD Archiver provides a more granular unit-to-unit replication relationship between the source system in the data center and the replica system in the remote site.

Once archive units are sealed for fault isolation, no additional data is transmitted between the source and destination archive units. In the event of a connectivity issue affecting the replication process, DD Archiver only needs to replicate the impacted unit to resynchronize. In addition, if a disaster impacts the availability of the primary archive unit, only the affected unit has to be recovered.



ULTRA-SAFE STORAGE FOR RELIABLE RECOVERY

DD Archiver takes advantage of existing Data Domain technology including the EMC Data Domain Data Invulnerability Architecture, which provides the industry’s best defense against data integrity issues. Continuous recovery verification, fault detection, and self-healing protect data during the initial backup and throughout the data lifecycle. DD Archiver is configured with dual disk parity RAID 6, enabling the system to withstand dual disk failures without interruption. Fans and power supplies are redundant and easy to replace for added system resilience. New fault isolation capabilities in DD Archiver ensure long-term data access and recoverability.

OPERATIONAL SIMPLICITY

DD Archiver is simple to manage. From the EMC Data Domain Enterprise Manager graphical user interface (GUI), administrators can centrally manage the file system, replication, and DD Boost interface, simplifying deployment. It also provides administrators with a consolidated view of capacity and performance utilization. Additionally, autosupport proactively reports on complete system status by e-mail. Administrators can also manage the DD Archiver through a command line interface (CLI) over SSH. SNMP monitoring allows administrators to easily integrate the DD Archiver with existing heterogeneous SNMP monitoring tools. Simple scriptability provides additional management flexibility.

STORAGE AND NETWORK EFFICIENCY

Industry-leading Data Domain inline deduplication technology offers an average of 10-30x data reduction for long-term retention of backups, so enterprises can dramatically reduce the amount of required storage capacity, data center space, and power and cooling, lowering total cost of ownership over the lifecycle of the data.

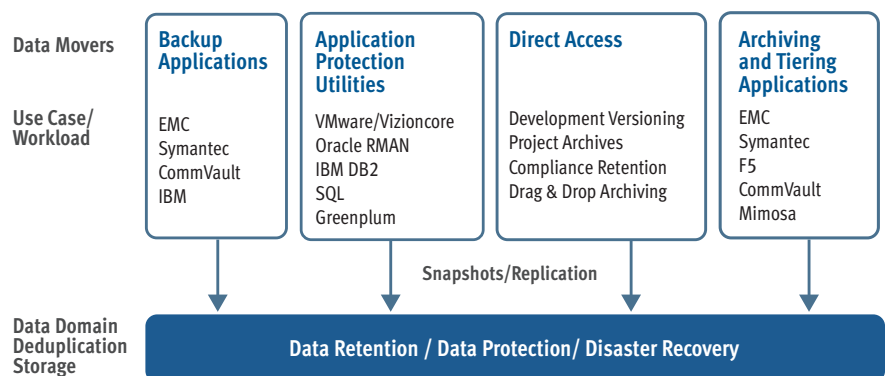
Leveraging EMC Data Domain Replicator with DD Archiver reduces bandwidth required for replication by up to 99 percent. In turn, this dramatically reduces the time needed to copy backups and archives over the WAN for consolidation or disaster recovery purposes. This network-efficient replication, in addition to the Data Domain system's ability to start replication while backing up, enables faster "time-to-DR" readiness.

EASY INTEGRATION

Unlike tape or traditional disk systems, DD Archiver provides deduplication storage for consolidated protection across a broad spectrum of applications. DD Archiver is qualified with leading enterprise backup and archive software applications and easily integrates into existing enterprise infrastructures. The system supports simultaneous data access methods through NFS and CIFS file service protocols over Ethernet, or as a disk-based target using application-specific interfaces such as EMC Data Domain Boost. DD Boost enables advanced integration for environments with Symantec NetBackup and Backup Exec.

Users can leverage the same DD Archiver system for both backup and archive workloads. 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.

All of these data movers and workloads can be supported in the same Data Domain system, at the same time.



ENFORCED RETENTION

Taking advantage of EMC Data Domain Retention Lock software, Data Domain Archiver helps organizations comply with internal IT governance for information retention by making files non-rewriteable and non-erasable. This ensures that critical business information is available for a specified retention period, at which time the information can be deleted as required. The retention parameters can be set on a file-by-file basis, and minimum and maximum retention periods can be set globally.

DD Retention Lock leverages industry-standard management interfaces supported by other leading storage vendors. As a result, it can be used seamlessly with most archive and data management products, providing an end-to-end solution.

SPECIFICATIONS

DD Archiver Specifications

Capacity, Raw ¹	Up to 768 TB
Logical Capacity, Standard ^{1,2,3}	5.7 PB
Logical Capacity, Redundant ^{1,3,4}	28.5 PB
Max. Throughput (Other) ⁵	4.3 TB/hr
Max. Throughput (DD Boost) ⁶	9.8 TB/hr
Power Dissipation ⁷	523 W
Cooling Requirement ⁷	1,783 BTU/hr

1. All capacity values are calculated using Base 10 (i.e., 1 TB = 1,000,000,000,000 bytes) and the maximum raw capacity configuration
2. Mix of typical enterprise backup data (filesystems, databases, e-mail, developer files). The low end of capacity range represents a full backup weekly or monthly, incremental backup daily or weekly, to system capacity. The top end of the range represents full backup daily, to system capacity.
3. Mix of typical enterprise data (filesystems, databases, e-mail, developer files), full backup daily, to system capacity.
4. Includes support for add-on shelves.
5. Maximum throughput achieved using CIFS and 10 Gb Ethernet.
6. Maximum throughput achieved using DD Boost and 10 Gb Ethernet.
7. Controller only.

SOFTWARE

Software Features

Global Compression™, Data Invulnerability Architecture including end-to-end verification (ongoing) and integrated dual disk parity RAID 6, snapshots, telnet, FTP, SSH, e-mail alerts, scheduled capacity reclamation, Ethernet failover and aggregation, Link Aggregation Control Protocol (LACP), VLAN Tagging, IP Aliasing; EMC Data Domain Boost, EMC Data Domain Replicator, and EMC Data Domain Retention Lock optional software

Management

EMC Data Domain Enterprise Manager, SNMP, and command line interface

Data Access

NFS v3 over TCP, CIFS, DD Boost (for use with Symantec OpenStorage)

SYSTEM EXPANSION

Up to 768 TB raw capacity

- Up to twenty-four 32 TB expansion shelves
- Up to twenty-four 16 TB expansion shelves
- Support for a mix of 32TB and 16TB expansion shelves up to raw capacity of 768 TB

REGULATORY APPROVALS

Safety: UL 60950-1, CSA 60950-1, EN 60950-1, IEC 60950-1, GS, SABS, GOST, IRAM Emissions: FCC Class A, EN 55022, CISPR 22, VCCI, BSMI, MIC, ICES-003

Immunity: EN 55024, CISPR 24 Power Line Harmonics: EN 61000-3-2

HARDWARE PLATFORM

2U 19-inch, rack mountable, use in 4-post rack, hotplug disks, redundant fans, redundant power supplies, serial port, 2 copper 10/100/1000 Ethernet ports and optional dual-port copper or optical 1 Gb Ethernet, quad-port copper 1 Gb Ethernet, and dual-port copper or optical 10 Gb Ethernet

System Weight

52 lbs (23.6 kg)

System Dimensions (WxDxH)

19" x 29.5" x 3.5" (48.3 cm x 74.9 cm x 8.9 cm)

2 EIA units

Minimum Clearance

Front, with bezel closed: 1.56" (4.0 cm)

Rear: 5" (12.7 cm)

Power (VA)

100-120 / 200-240 V~, 50/60 Hz; 550 VA

System Thermal Rating

1,783 BTU/hr

Operating Temperature / Altitude

10°C to 35°C (50°F to 95°F), derate 1.1°C/1000 feet above 7500 feet to 10,000 feet

Operating Humidity

20% to 80% non-condensing

Non-operating (Transportation) Temperature

-40°C to +65°C (-40°F to +149°F)

Operating Acoustic Noise

Declared noise emission values per ISO 9296:

Sound power, LWAd: 7.52 bels

Sound pressure, LpAM: 56.4 db

CONTACT US

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

EMC², EMC, where information lives, Data Domain, Global Compression, and SISL 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 01/11 H7509

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
where information lives®