

DD OS Essentials

High-speed, scalable deduplication

- Up to 68 TB/hr performance
- Reduces protection storage requirements by 10 to 55x
- CPU-centric scalability

Data invulnerability architecture

- Inline write/read verification, continuous fault detection
- Dual disk parity RAID 6

Seamless integration

- Supports leading enterprise backup and archive applications
- Tightest integration with Dell EMC data protection suite family

Backup and archive consolidation

- Efficiently protect backup and/or archive data
- Meets governance policies and compliance regulations for archive data

Data Domain Virtual Edition

- Software-defined protection storage

Cloud-enabled

- Simple and efficient long-term retention to a public, private or hybrid cloud
- Secure multi-tenancy provides logical isolation of user data
- Low-cost disaster recovery to the cloud

Fast and efficient disaster recovery

- Reduce bandwidth requirements by up to 98%

Instant access

- Boot VMs from protection copies

Operational simplicity

- Lower administrative costs
- Reduced footprint

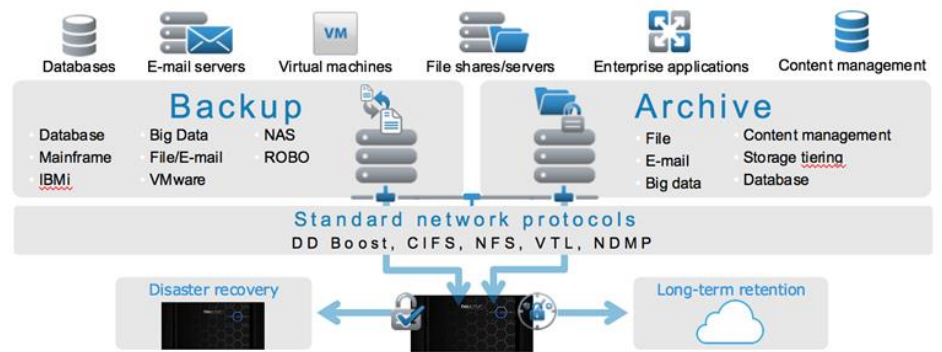
DELL EMC DATA DOMAIN OPERATING SYSTEM

Powering Dell EMC protection storage

The Data Domain Operating System (DD OS) is the intelligence that powers Dell EMC Data Domain. It provides the agility, security and reliability that enables the Data Domain platform to deliver scalable, high-speed, and cloud-enabled protection storage for backup, archive and disaster recovery.


High-speed scalable deduplication

The Data Domain Operating System powers Data Domain to deliver industry-leading speed and efficiency with throughput up to 68 TB/hour, enabling more backups to complete sooner and reducing pressure on backup windows. DD OS employs variable-length deduplication to minimize disk requirements and ensures data lands on disk already deduplicated. This reduces backup and archive storage requirements by an average of 10 to 55x, making disk a cost-effective alternative to tape. These ratios reach an even greater average of 36x when using the Dell EMC Data Protection Suite Family to backup to Data Domain. Data on disk is available online and onsite for longer retention periods and restores and retrievals become fast and reliable. This efficiency enables Data Domain to protect up to 50 PB logical capacity on a single system, with the ability to natively-tier another 100 PB of deduplicated data to the cloud for long-term retention thanks to support for Data Domain Cloud Tier. With deduplication, months of retention on disk is possible using the same number of “floor tiles” that traditionally provided only a couple of days of disk staging.



DD OS powers Dell EMC Data Domain to deduplicate data during either the backup process or archive process, which maximizes performance while minimizing disk storage requirements. Deduplicated data can be stored onsite for immediate restores and longer-term retention on disk. The deduplicated data can also be replicated over the WAN to a remote site or a service provider site in the cloud for disaster recovery operations, eliminating the need for tape-based backups, or for consolidating tape backups to a central location. Data Domain provides the capability to consolidate both backup and archive data on the same infrastructure allowing greater consolidation by eliminating silos of backup and archive storage and associated overhead.

The key that enables DD OS to provide industry-leading performance while minimizing disk requirements is the Dell EMC Data Domain Stream-Informed Segment Layout (SISLTM) scaling architecture. Specifically, SISL leverages the continued advancement of CPU performance to continuously increase Data Domain system performance by minimizing disk accesses required to deduplicate data. SISL deduplicates data by identifying duplicate data segments in memory, which minimizes disk usage. This enables Data Domain throughput to be CPU-centric, not “spindle bound.



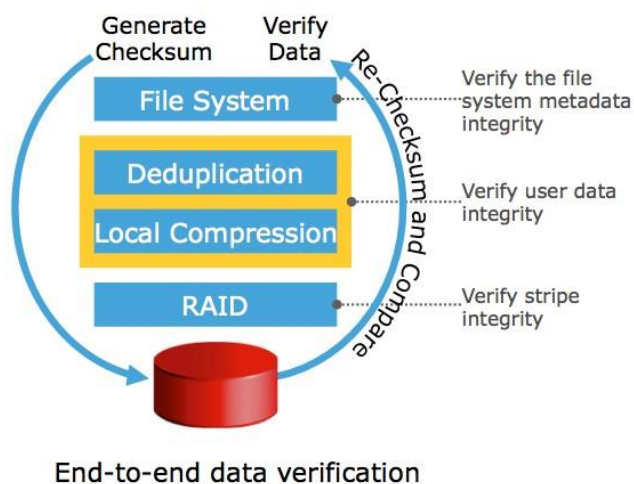
CPU-Centric Storage

Data Domain Stream Informed Segment Layout (SISL) scaling architecture takes the pressure off of disk I/O as a bottleneck, so the remaining system design is CPU-Centric. Other deduplication methods require more disks to increase their throughput speeds.

Data invulnerability architecture

Data Domain is designed as the storage of last resort – built to ensure you can reliably recover your data with confidence. The Dell EMC Data Domain Data Invulnerability Architecture is built into DD OS to provide the industry’s best defense against data integrity issues. Inline write and read verification protects against and automatically recovers from data integrity issues during data ingest and retrieval.

Capturing and correcting I/O errors inline during the backup process eliminates the need to repeat backup jobs, ensuring backups complete on time and satisfy service-level agreements. In addition, unlike other enterprise arrays or file systems, continuous fault detection and self-healing ensures data remains recoverable throughout its lifecycle on Data Domain.



End-to-end data verification

End-to-end data verifications reads data after it is written and compares it to what was sent to disk, proving that it is reachable through the file system to disk and that the data is not corrupted. Specifically, when the Data Domain Operating System receives a write request from backup software, it computes a checksum over the data. After analyzing the data for redundancy, it stores the new data segments and all of the checksums. After all the data is written to disk, the Data Domain Operating System verifies that it can read the entire file from the disk platter and through the Data Domain, and that the checksums of the data read back match the checksums of the written data. This confirms the data is correct and recoverable from every level of the system.

Backup and archive data stored on a Data Domain system can be protected with encryption performed inline—before the data is written to disk. For advanced security of data-at-rest, Data Domain’s inline encryption uses industry-standard RSA® BSAFE FIPS 140-2 validated cryptographic libraries. Data Domain supports internal encryption key management on each system or external key management through RSA Data Protection Manager, to deliver a robust encryption key lifecycle management solution for the entire enterprise.

Performance and capacity						
	DDVE¹	DD3300	DD6300	DD6800	DD9300	DD9800
MAXIMUM THROUGHPUT	4 TB/hr	7.0 TB/hr	8.5 TB/hr	14 TB/hr	20 TB/hr	31 TB/hr
MAXIMUM THROUGHPUT (DD BOOST)	11.2 TB/hr	4.2 TB/hr	24 TB/hr	32 TB/hr	41 TB/hr	68 TB/hr
LOGICAL CAPACITY	Up to 4.8 PB	200 TB – 1.6 PB	1.4 - 7 PB	2.8 - 14.4 PB	7.2 -36 PB	10-50 PB
W/ DD EXTENDED RETENTION²	-	-	-	5.6 - 28.8 PB	14.4 - 72 PB	20 - 100 PB
W/ DD CLOUD TIER²	Up 14.8 PB	600 TB – 4.8 PB	-	8.4 – 43.2 PB	21.6 - 108 PB	30 - 150 PB

1. Throughput drawn running DD VE in the following environment with 16 TB & 96 TB instance: Host server: 2x Intel Xeon CPU (6 Cores each) @ 2GHz, 128GB memory, 2x10GbE NIC; Storage: DAS with 3TB 7200RPM SAS Drives, RAID6, Battery Powered HBA Cache Enabled, Disk Cache Disabled.

2. DD Cloud Tier and DD Extended Retention are mutually exclusive long-term retention solutions.

Capacity is based on a mix of typical enterprise backup data (file systems, databases, email, 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. All capacity values are calculated using Base10 (i.e., 1TB = 1,000,000,000,000 bytes).

Seamless integration

Data Domain integrates easily with existing infrastructures and can be used seamlessly with leading backup and archiving applications. Integrating a Data Domain system into your environment does not require any change in process or infrastructure, so you can realize the value of deduplication quickly and efficiently. In addition, Data Domain can integrate directly with leading enterprise applications such as Oracle RMAN or write directly over CIFS or NFS to support a variety of workloads.

For a complete data protection solution, leverage the powerful combination of Dell EMC data protection software with protection storage. Experience reduced risk of data loss, industry-leading deduplication, and superior performance when using the Data Protection Suite Family with Data Domain.

Since Data Domain can simultaneously support multiple access methods including NFS, CIFS, VTL, NDMP and Data Domain Boost™ all applications and utilities can be supported in the same Data Domain system at the same time to enable greater protection storage consolidation. A system can present itself as a file server, offering NFS, CIFS access over Ethernet; as a virtual tape library (VTL) over Fibre Channel; as an NDMP tape server over Ethernet; or as a disk target using application specific interfaces like Data Domain Boost. Data Domain VTL is qualified with leading open systems and IBM i enterprise backup applications.

Backup and archive consolidation

The flexibility of DD OS enables Data Domain to be the only protection storage to simultaneously support backup and archive data. This enables Data Domain to reduce overall total cost of ownership (TCO) by sharing resources across backup and archive data. Specifically, a single Data Domain can be used for backup and recovery of the entire enterprise (including Oracle, SAP, Microsoft, and VMware®, as well as IBM i and mainframe environments) as well as protecting archive data (including file, e-mail, enterprise content management, database and Virtual Machine archiving). Data Domain cost-effectively protects archive data through integration with industry-leading archiving applications including Data Protection Suite for Archive and Veritas Enterprise Vault. By consolidating to a common protection storage platform, you can eliminate silos of storage and associated overhead - including management, floor space, power and cooling. In addition, with Data Domain Retention Lock software, Data Domain systems can meet internal governance policies or compliance regulations for archive data including SEC 17a-4(f). Data Domain Retention Lock allows storage administrators, backup administrators, and compliance officers to meet data retention requirements for data stored on a Data Domain system. DD Retention Lock prevents files from being modified or deleted for a user-defined retention period. This feature enables secure file locking of archive data at an individual file level; enabling these files to be intermixed with unlocked files on the same Data Domain system – thereby allowing for a broader consolidation of backup and retention based archive data.

Data Domain Virtual Edition

Data Domain Virtual Edition (DD VE) leverages the power of DD OS to deliver software-defined protection storage. DD VE is fast and simple to download, deploy and configure - and can be up and running in minutes. DD VE can be deployed on standard hardware, and runs in VMware vSphere, Microsoft Hyper-V, as well as in the cloud with AWS and Azure. DD VE is also certified with VxRail and Dell PowerEdge servers. With DD VE existing infrastructure can be utilized to deploy virtual protection storage. An assessment tool can be run during deployment to check the underlying infrastructure and ensure it meets recommended requirements. A single DD VE instance can scale from .5 TB to 96 TBs or up to 16 TB if deployed in the cloud. Capacity can be easily be moved between virtual systems and/or locations and can be purchased in 1 TB increments allowing you to grow capacity as the business demands it. DD VE maintains the core DD OS features and includes DD Boost, DD Encryption and DD Replicator. Configure and manage a DD VE instance using DD System Manager and centrally manage multiple DD VE instances through DD Management Center.

Cloud enabled

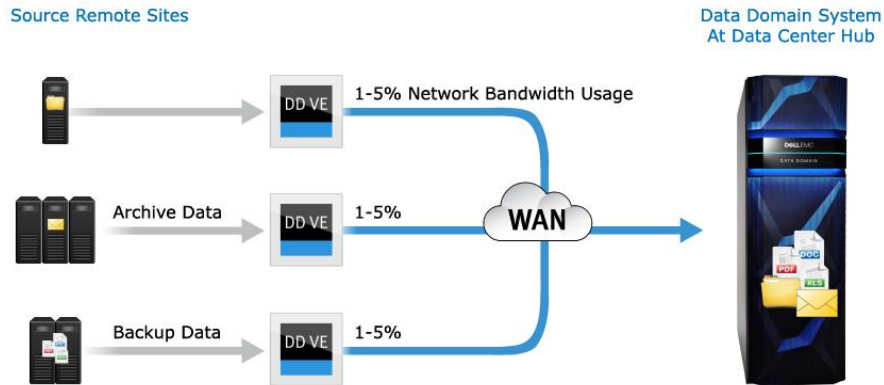
With Data Domain Cloud Tier, DD OS can natively tier data to a public, private or hybrid cloud for long-term retention. Only unique data is sent directly from Data Domain to the cloud and data lands on the cloud object storage already deduplicated. With deduplication ratios of 10 – 30x, storage footprint is greatly reduced lowering overall TCO. DD Cloud Tier can scale up to 2x the max capacity of the active tier. With DD Encryption, data in the cloud remains secure. A broad ecosystem of backup and enterprise applications and a variety of public and private clouds are supported with DD Cloud Tier including EMC Elastic Cloud Storage (ECS) and Virtustream Storage Cloud.

DD OS also provides secure multi-tenancy, which enables large enterprises and service providers to deliver data protection as a service with Data Domain in a private or hybrid cloud. With secure multi-tenancy, Data Domain will logically isolate tenant data and network isolation provides strict data access isolation between tenants ensuring that each tenant's data is only visible and accessible to them and allowing for tenant self-service. SSL Certificate authentication ensures secure replication. Physical capacity measurement allows you to capture how much physical capacity is being consumed at a file, directory, MTree, tenant, or tenant-unit level. This serves as an effective mechanism for managing shared Data Domain protection storage capacity between individual departments or tenants.

Data Domain Cloud DR (DD Cloud DR) allows enterprises to copy backed-up VMs from their on-premises Data Domain environments to the public cloud (AWS) and to orchestrate DR testing and failover of workloads to the cloud in a disaster scenario.

Fast, efficient and scalable disaster recovery

As data lands on Data Domain, it can immediately begin replicating it to a disaster recovery site. To meet strict DR requirements, Data Domain Replicator software can replicate at 52 TB/hr over a 10 Gb network connection. DD OS replicates only unique compressed data across the network, requiring a fraction of the time, bandwidth and cost of traditional replication methods. With cross-site deduplication only unique data is transferred across any of the WAN segments. This can reduce WAN bandwidth requirements up to 99%, making network-based replication fast, reliable and cost-effective. For the highest level of security, data being replicated between Data Domain instances can be encrypted using the standard Secure Socket Layer (SSL) protocol. Data Domain provides flexible replication topologies including full system mirroring, bi-directional, many-to-one, one-to-many, and cascaded. In a many-to-one deployment, data from up to 540 remote offices can be replicated to a single DD9800.



Operational simplicity

Data Domain is very simple to install and manage resulting in lower administrative and operational costs. Administrators can access the Data Domain Operating System through command line over SSH or through Data Domain System Manager, a browser-based graphical user interface. Multiple Data Domain systems can be managed and monitored through a single interface through Data Domain Management Center. Customizable dashboards provide visibility into aggregate status and the ability to drill-down to system-level details. Role based access allows different levels of access via assigned user roles for various levels of expertise within the organization. Simple scriptability as well as SNMP monitoring provides additional management flexibility.

In addition, Data Domain has an automatic call-home system reporting called autosupports, which provides email notification of complete system status to Dell EMC support and 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.



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



[Contact](#) a Dell EMC Expert