

# EMC DATA DOMAIN BOOST FOR SYMANTEC NETBACKUP

## Increase backup performance with advanced OpenStorage integration

### ESSENTIALS

#### Faster Backups and Greater Resource Utilization

- Distributed deduplication process dramatically increases throughput
- Efficient CPU utilization on NetBackup media servers
- Reduced network bandwidth utilization
- Enables faster “time-to-DR” readiness

#### Seamless Integration with NetBackup

- Centralized management and tighter integration
- Heterogeneous backup server environment support
- Support for Data Domain Appliance Series, Data Domain Global Deduplication Array, and Data Domain Archiver

#### NetBackup Managed Replication

- 99 percent bandwidth reduction
- Cost-efficient disaster recovery
- Multi-site tape and disk consolidation
- Encrypted replication

#### Advanced Load Balancing and Link Failover

- Scalable link aggregation at the application layer
- Simplified NetBackup configuration
- Seamless load balancing of jobs among available ports
- Link failover keeps backups operational

#### Ultra-Safe Storage for Reliable Recovery

- Continuous recovery verification, fault detection, and healing
- End-to-end data integrity

### NEXT-GENERATION DATA PROTECTION

Symantec™ NetBackup™ (NBU) is one of the most pervasive data management tools used in medium to large size data centers around the world. It is a critical component of backup, recovery, and disaster preparedness strategies. NBU directs the operation of disk and tape media that stores typically five times more data than what can be found on the primary storage in the same data center. As data continues to grow unabatedly, and as more attention is focused on efficient data protection and disaster recovery (DR), traditional approaches to backup are no longer acceptable. Disk-based backup and deduplication solutions are becoming mandatory in many data management approaches.

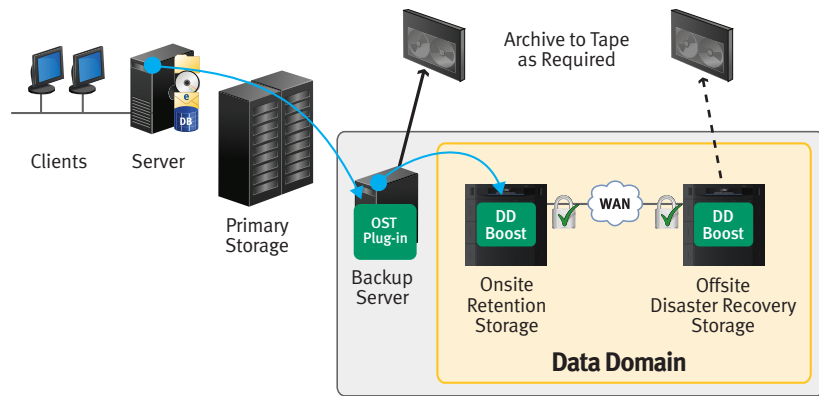
These transformational backup and DR solutions must have a foundation of simplicity and flexibility that enable consolidated management from a single management pane within NetBackup. Unfortunately, many previous disk-based backup technologies have only focused on improving the speed and reliability of backups, not advancing DR and ease-of-management. For example, some virtual tape libraries (VTLs) actually worsened the complexities and liabilities of tape-based data protection. With many traditional VTLs, users have to provision imaginary tape drives, slots, and cartridges which could increase management complexity. The inability of data management tools to handle multiple instances of the same tape barcode or image makes management of DR copies a breaking point.

EMC® Data Domain® deduplication storage systems offer an alternative to VTLs with a simple NFS/CIFS interface. Data Domain systems are simple to integrate with NetBackup, and allow users to enjoy the retention and recovery benefits of inline deduplication as well as the off-site disaster recovery protection of replication over the wide area network (WAN).

Data Domain Boost extends the optimization capabilities of Data Domain solutions. DD Boost significantly increases performance by distributing parts of the deduplication process to the media server, simplifies disaster recovery procedures, and serves as a solid foundation for additional integration between Symantec NetBackup and Data Domain systems.

### SIGNIFICANT REDUCTION IN BACKUP TIME

Prior to DD Boost, NetBackup media servers would send all data, unique or redundant, to a Data Domain system for deduplication processing. With the distributed segment processing feature of DD Boost, parts of the deduplication process are distributed to the media server, enabling it to send only unique data segments to a Data Domain system. This dramatically increases the aggregate throughput, up to 26.3 TB/hr, and reduces the amount of data transferred over the network by 80 to 99 percent. These efficiencies can help eliminate future costs by leveraging existing backup servers and Ethernet networks.



### EMC Data Domain Boost

Data Domain Boost significantly increases performance by distributing parts of the deduplication process to the backup server, simplifies disaster recovery procedures, and serves as a solid foundation for additional integration between backup applications and Data Domain systems.

In addition to performance improvements and network bandwidth benefits, the reduction in data transferred over the network also decreases CPU utilization on the media servers since sending data is significantly more CPU intensive than the distributed deduplication process. Consequently, this improved efficiency eliminates the need for upgrading the media server infrastructure and also provides an ability to drive more backups from a single media server.

DD Boost for NetBackup also increases the speed of restart and completion of failed backups. Since only unique data is sent over the network, once a failed job restarts, the data that has already been sent to the Data Domain system for a given backup job does not need to be sent again. This not only reduces the load on the network substantially, but also improves the overall throughput for the failed backups upon retry.

Overall, DD Boost for NetBackup significantly increases aggregate throughput, substantially reduces backup windows, and improves media server efficiency. Best of all, DD Boost is transparent to the backup application because the distributed segment processing is handled by the DD Boost Library, which is embedded in the OST plug-in installed on the media server.

### SEAMLESS INTEGRATION WITH NETBACKUP

The combination of a Data Domain system and DD Boost for NetBackup creates an optimized connection to provide a tightly integrated solution. DD Boost for NetBackup offers operational simplicity by enabling the media server to manage the connection between the backup application and one or more Data Domain systems.

Unlike a VTL implementation, DD Boost for NetBackup does not require any of the artifacts of tape or tape emulation. Similar to the Data Domain NFS/CIFS implementation, an image can be simultaneously written to and read from Data Domain systems in NetBackup's native image formats. Images written to these systems can be expired one at a time, leading to improved storage efficiency over VTL-only implementations.

DD Boost for NetBackup simplifies management and allows easy sharing and simultaneous use of disk pools among multiple heterogeneous media servers. This facilitates media server load balancing and the use of the "best" available media servers for completing a particular backup job based on pre-configured policies. NetBackup's storage lifecycle policies also allow administrators to fully automate the retention and duplication of images stored on

Data Domain storage systems. Data Domain Boost software is certified by Symantec and can be found on its Hardware Compatibility List (HCL).

DD Boost works on multiple media server operating systems, providing a seamless experience to administrators in heterogeneous environments. Further, DD Boost serves as a consolidated interface to various Data Domain systems—Appliance series, Global Deduplication Array, and DD Archiver. This simplifies the management of Data Domain deduplication storage systems.

## **NETBACKUP-MANAGED REPLICATION—OPTIMIZED DUPLICATION**

Network-efficient replication technology sends backup images from one Data Domain system to one or more systems, when all systems are configured with DD Boost software. Using this approach, the bandwidth required for replication is reduced by up to 99 percent. This dramatically reduces the time and WAN bandwidth needed to create multiple copies of backups for tape consolidation or disaster recovery purposes. Resource usage on the media servers is also reduced as they are no longer in the data path when duplicate copies of the backup are being created and transferred.

With DD Boost for NetBackup, the media server controls the replication of backup images between multiple Data Domain systems, and provides backup administrators with a single point of management for tracking all backups and duplicate copies. This paradigm is designed to be very familiar to NetBackup users because replication over the WAN using EMC Data Domain Replicator software feels similar to vaulting a copy of the data to tape. All copies of the backup images are present and consistent in the NetBackup catalog, which provides a single management console view to administrators for easy disaster recovery. This also enables administrators to manage retention periods for various copies individually, resulting in a more flexible DR deployment. 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.

Leveraging WAN-efficient replication, DD Boost for NetBackup eliminates the need for tape-based backup and recovery procedures at remote sites. Recovery from backup copies at the central site also becomes simplified because all copies are tracked in the NetBackup catalog. Since tape is no longer required at the remote sites, it is entirely practical to deploy Data Domain systems with DD Boost for NetBackup software as key foundational elements of a multisite tape consolidation strategy. However, if tape is still a requirement, NetBackup's storage lifecycle policies allow administrators to automate centralized tape management.

## **ULTRA-SAFE STORAGE FOR RELIABLE RECOVERY**

Systems with DD Boost for NetBackup benefit from the EMC Data Domain Data Invulnerability Architecture—continuous verification, continuous fault detection and healing, and other resiliency features, all transparent to NetBackup.

DD Boost for NetBackup extends this protection to the NetBackup media server by generating checksums on the data that it sends to the Data Domain system. The checksums are then transferred along with the data. The Data Domain system receiving the data computes new checksums on the incoming data and then compares them to the computed values from DD Boost for NetBackup for verification purposes. This ensures end-to-end verification of data.

Features	Benefits
Distributed segment processing	<ul style="list-style-type: none"> <li>Significant reduction in backup time</li> <li>Up to 26.3 TB/hr aggregate throughput performance</li> <li>Improved network bandwidth utilization</li> <li>Efficient CPU utilization on media servers</li> <li>Less time required to restart failed backup jobs</li> </ul>
Seamless integration	<ul style="list-style-type: none"> <li>Centralized management and operational simplicity</li> <li>Media server load balancing</li> <li>Heterogeneous media server environment support</li> <li>Facilitates hierarchical duplication and centralized tape out operations</li> </ul>
NetBackup managed replication	<ul style="list-style-type: none"> <li>Optimized duplication using WAN-efficient DD Replicator</li> <li>99 percent bandwidth reduction</li> <li>Cost-efficient disaster recovery</li> <li>Multisite tape consolidation</li> <li>Encryption of replication session between Data Domain systems</li> </ul>
Advanced load balancing and link failover	<ul style="list-style-type: none"> <li>Scalable link aggregation at the application layer</li> <li>Simplified configuration</li> <li>Optimize throughput of multiple links</li> <li>Link failover keeps backups operational in case of temporary network glitches and failures</li> </ul>
Data Invulnerability Architecture	<ul style="list-style-type: none"> <li>Ultra-safe storage for reliable recovery</li> <li>End-to-end data integrity</li> </ul>

## SPECIFICATIONS

### SOFTWARE

EMC Data Domain Operating System 5.0 or later

EMC Data Domain OST plug-in 2.3 or later

Symantec NetBackup 6.5 or later

### HARDWARE

EMC Data Domain Appliance Series

EMC Data Domain Global Deduplication Array Series

EMC Data Domain Archiver

### MEDIA SERVER PLATFORMS

Solaris 9/10 (SPARC)

Red Hat Enterprise Linux 4/5 (x86, x86\_64)

SuSE Linux Enterprise Server 10/11 (x86, x86\_64)

Windows® Server 2003 (x86, x86\_64)

Windows Server 2008 (x86\_64)

Windows Server 2008 R2 (x86\_64)

HP-UX 11.23/11.31 (PA-RISC)

HP-UX 11.23/11.31 (IA64)

AIX 5.3/6.1 (POWER)

### NETWORK CONNECTIVITY

IP connectivity between backup servers and the Data Domain system

## 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](http://www.EMC.com).

EMC<sup>2</sup>, EMC, Data Domain, 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. Solution Overview 08/11 H7143.2