

## ESG Lab Review

# Symantec OpenStorage

**Date:** February 2010 **Author:** Tony Palmer, Senior ESG Lab Engineer

**Abstract:** This ESG Lab review documents hands-on testing of consolidated management and automated data protection services across multiple sites and diverse storage technologies using Symantec NetBackup OpenStorage.

## The Challenges

Relentless information growth is necessitating greater investments in IT infrastructure. Data protection processes, such as backup, compound capacity growth issues since multiple copies of primary data are made for operational and disaster recovery. In fact, ESG research respondents cited data protection as *the* application that will be most responsible for storage growth over the next 24 months.<sup>1</sup> Disk-based systems inherently offer faster restores and are more reliable than tape, but they can make backup environments more complex and difficult to manage.

As shown in Figure 1, a typical enterprise has multiple separate and distinct backup and recovery technologies distributed throughout the environment for a variety of purposes. Entry level backup software and tape autoloaders protect remote offices. Disk-based systems and tape libraries in the data center protect mission-critical applications and provide deep archive and offsite capabilities. In this scenario, each copy of data is treated as a separate entity and managed using device and application specific tools.

Figure 1. Distributed Data Protection Challenges



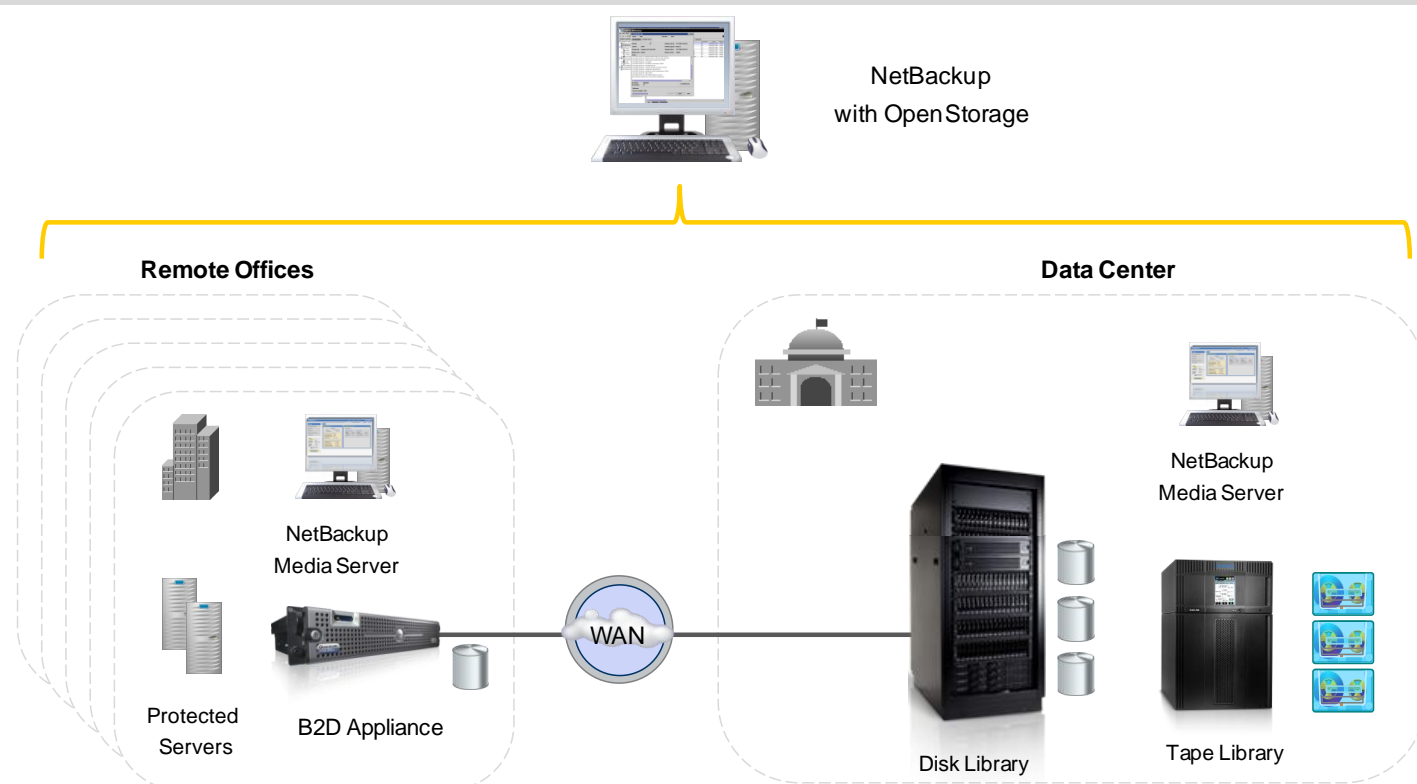
<sup>1</sup> Source: ESG Research Report, *Enterprise Storage Systems Survey*, November 2008.

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was commissioned by Symantec.

## The Solution: Edge-to-Core Data Protection with Symantec OpenStorage

Symantec NetBackup software and its OpenStorage API option have been designed to provide centrally managed, edge-to-core data protection that spans multiple sites and provides disk to disk to tape (D2D2T) functionality as shown in Figure 2. The OpenStorage API introduced in NetBackup version 6.5 provides automated movement of data between sites and storage tiers and a single point of management and catalog for backup data—regardless of where it resides (remote office or corporate data center), what type of media it is stored on (disk or tape), or its age (recent backup or long term archive). Symantec NetBackup provides a single point of backup data management, including “direct path to tape” backup images that can be created without having to pass through a media server when supported by the backup target.

Figure 2. Symantec OpenStorage Edge-to-Core Data Protection



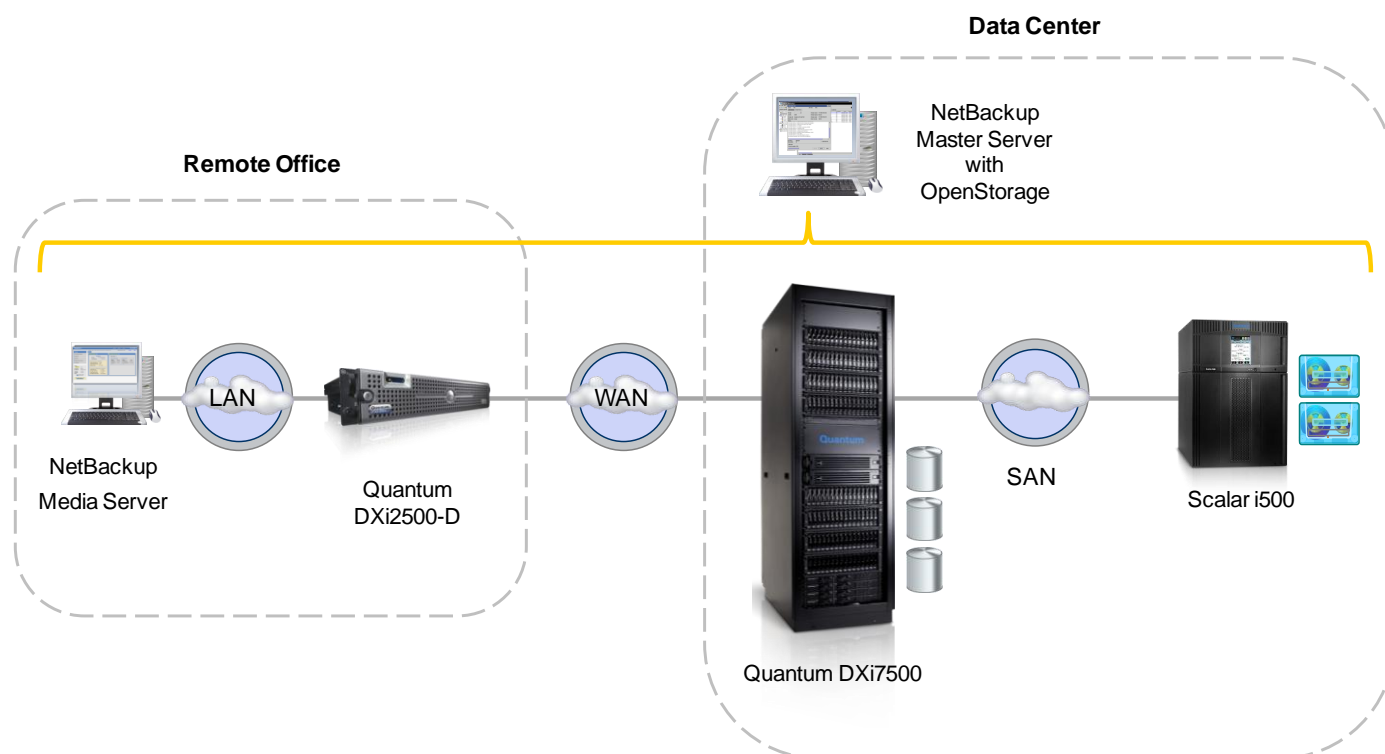
## ESG Lab Tested

ESG Lab used the Symantec OpenStorage API in Symantec NetBackup version 6.5 to configure, automate, and track the migration of backup data residing on Quantum disk and tape products as shown in Figure 3.<sup>2</sup> An edge-to-core D2D2T data protection strategy was implemented using a Quantum DXi2500-D appliance located in a simulated remote office. Remote office backup data was replicated over a simulated WAN to a Quantum DXi7500 Enterprise disk library in a corporate data center with data movement carried out by the DXi systems. The direct path to tape capabilities of the DXi7500 and the Symantec OpenStorage API were used to automatically write a copy of the data to removable media in a Fibre Channel SAN attached Quantum Scalar i500 tape library.

In our tests, the Quantum DXi systems moved all data directly, while the Symantec NetBackup master server in the corporate data center was used to create a single point of control and catalog for all the versions of the backup data residing anywhere within the D2D2T infrastructure.

<sup>2</sup> Quantum is one of several qualified NetBackup OpenStorage Partners. For more information on OpenStorage and the Symantec Technology Enabled Program (STEP), visit: <http://www.symantec.com/business/theme.jsp?themeid=openstorage>

Figure 3. Symantec OpenStorage Test Bed



At a high level, configuring Symantec NetBackup software and DXi-Series hardware to take advantage of OpenStorage-enabled replication and tape creation was a four-step process:

1. Configure backup to disk appliances to take advantage of the OpenStorage API
2. Install the storage-specific OpenStorage plug-in on NetBackup servers
3. Configure NetBackup to use OpenStorage-enabled disk pools
4. Create NetBackup storage lifecycle policies for backup, replication, and tape creation

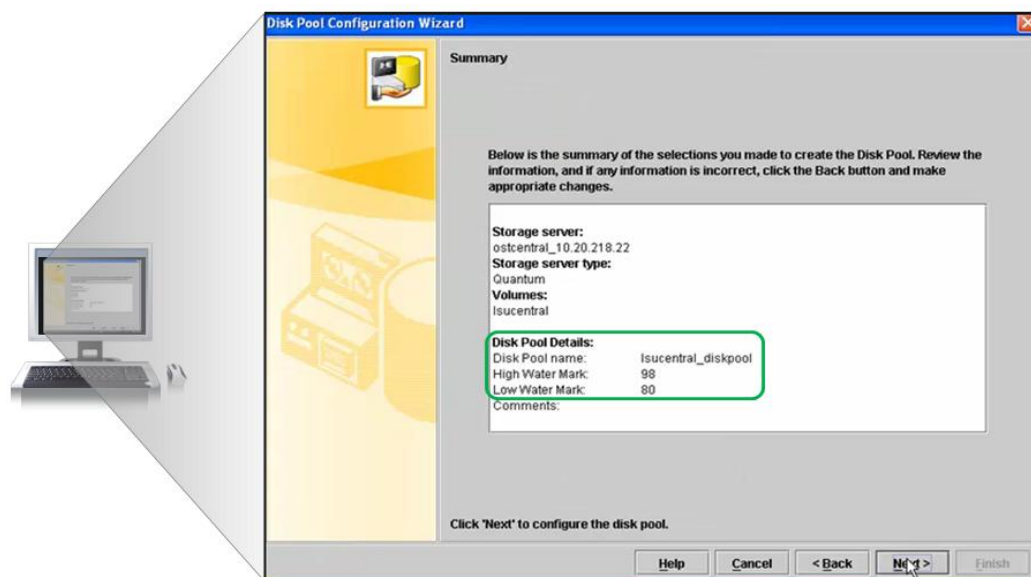
**Step 1:** The process began with the configuration of an OpenStorage Storage Server and a Logical Storage Unit (LSU) on the Quantum DXi7500. An LSU named *lsucentral* was defined to use the *ostcentral* storage server. This process was repeated on the DXi2500-D in the remote office (*ostremote/lsuremote*).

**Step 2:** Next, the OpenStorage plug-in for the DXi series was downloaded from Quantum and installed on the NetBackup servers in the corporate data center and the remote site. A Quantum configuration file was copied to the */usr/quantum* directory on each of the Linux NetBackup servers. This phase of the installation ended with a restart of the NetBackup daemon on each of the Linux NetBackup servers.

**Step 3:** Symantec NetBackup commands were used to create two storage servers: one for the central site and another for the remote site. The names of the DXi-Series storage servers created earlier (*ostcentral* and *ostremote*) and the IP addresses of each of the DXi systems was provided. The NetBackup GUI was used to start a NetBackup path to tape wizard which scanned and discovered each of the DXi-Series systems.

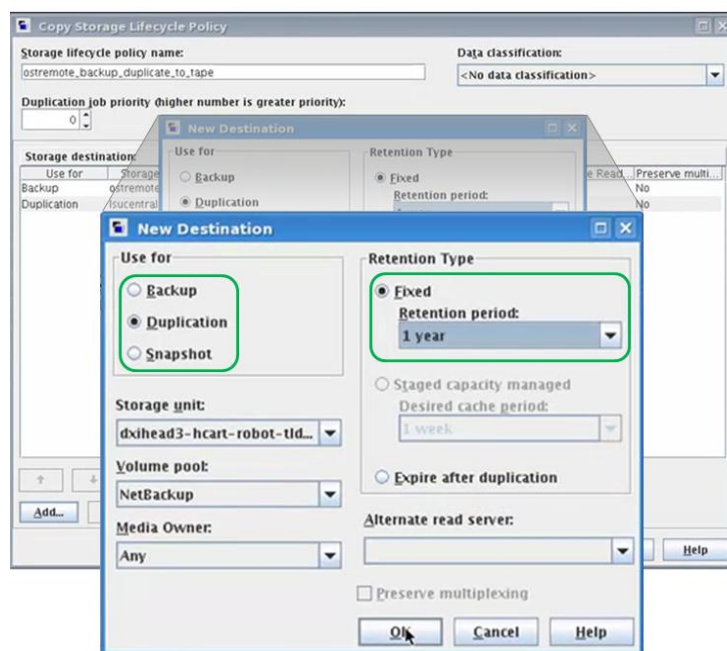
The NetBackup wizard shown in Figure 4 was used to define OpenStorage enabled disk pools. In this example, the disk pool for the disk pool on the DXi7500 in the corporate data center has been defined (*Isucentral*). The process was repeated for the DXi2500-D in the remote office (*Isuremote*).

Figure 4. Configuring a Disk Pool in NetBackup



**Step 4:** As shown in Figure 5, a NetBackup storage lifecycle policy was configured to automatically perform remote replication and direct path to tape duplication.

Figure 5. Configuring a Storage Lifecycle Policy

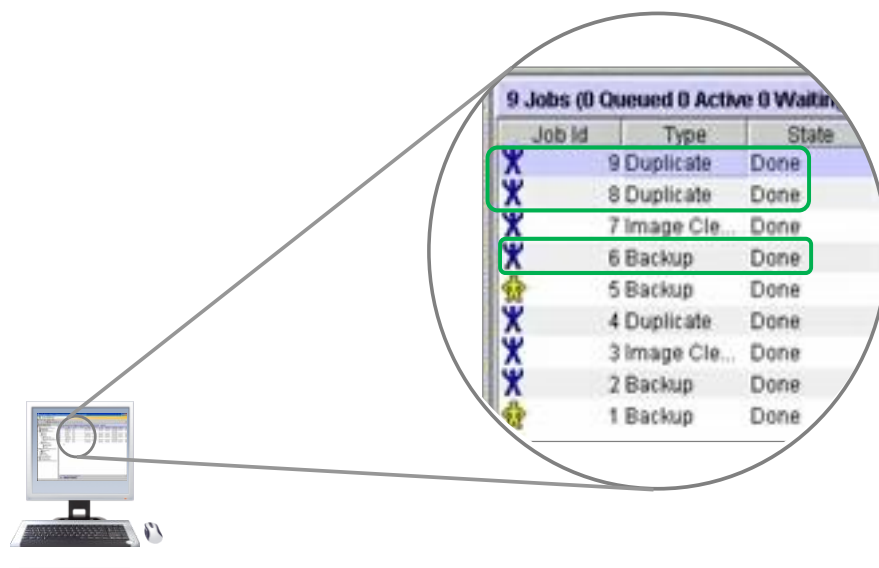


After completing the four step configuration process, an existing backup policy was modified to take advantage of the new storage lifecycle policy. Twenty minutes after getting started with the configuration on a pre-wired test bed, ESG

Lab was monitoring the automated execution of backup, replication, and tape creation. A backup was started manually and monitored from the NetBackup console.

As shown in Figure 6, NetBackup automatically started the DXi-enabled replication after the backup job finished. When that completed, a third job was started automatically to create a removable copy of the remote office backup data in the corporate data center.

Figure 6. Config The NetBackup Activity Monitor – After D2D2T



ESG Lab finished the evaluation with a quick tour of the NetBackup catalog. As expected, all three copies of the remote office backup job were found and could be used for recovery from a single console:

1. A disk-based copy at the remote site, which is ideally suited for quick and reliable recovery of lost or corrupt files
2. A disk-based copy in the corporate data center, which is ideally suited for recovery after a disaster in a remote office (e.g., server failure, fire, flood).
3. A removable tape-based copy, which is ideally suited for legal discovery, compliance, and disaster recovery at corporate headquarters.

## Why This Matters

Until recently, extending the benefits of a D2D2T protection strategy to remote and branch offices has been expensive and hard to manage. If disk-based storage systems were used for replication and tape creation, the backup software couldn't keep track of where the copies resided. If the backup software was making tape copies, media servers' resources were consumed. If full copies of data were stored on disk and replicated, disk capacity and WAN bandwidth was cost prohibitive.

ESG Lab has confirmed that the Symantec OpenStorage initiative solves these problems and provides the benefits of a centrally managed D2D2T approach that can be extended across sites and across storage technologies. Deduplication-enabled backup to disk appliances reduce the cost of disk capacity and WAN bandwidth and provide direct tape creation that reduces the load on backup servers. NetBackup initiates all the data movement, provides the ability to define separate retention policies for disk and tape copies in multiple locations, and maintains a single catalogue to manage it all. For users with multiple sites and multiple tiers of backup data, NetBackup OpenStorage can be used to dramatically reduce complexity, saving time and money.

## The Bigger Truth

Organizations are challenged to keep pace with the volume of data maintained for recovery purposes. The capacity of data to be backed up makes it difficult to complete backup processing within prescribed windows. To accelerate backup and recovery, IT organizations have deployed disk in the backup data path—improving both performance and reliability of backup and recovery processes at the cost of increased complexity of the backup environment.

Symantec OpenStorage enabled ESG Lab to manage backups for a remote office from creation, through replication, and finally to tape—all using NetBackup storage policies and native tools, enabling an edge-to-core data protection strategy covering remote and branch offices as well as multiple data centers. Path-to-tape integration with backup software enables enterprises to meet offsite and deep archive requirements using the familiar tools and techniques of their backup application without burdening media servers or consuming SAN bandwidth with the task of tape creation.

As with all disk-based backup systems today, when a retained backup image is deleted or expired by the backup application, space on the disk backup appliance is not reclaimed until the image is deleted, expired, or overwritten via the backup appliance's management application and a reclamation process is run. Tighter integration with NetBackup to automatically trigger a policy-based reclamation on the backup appliance when images have been expired by NetBackup would be a useful enhancement.

ESG Lab believes that tight integration between centrally managed backup software and intelligent storage devices allows customers to increase the efficiency and performance of their data protection environment. The combination of comprehensive data protection management software and storage-based data deduplication and replication provides such an integrated approach, optimizing data protection and recovery in the enterprise. Using the Symantec NetBackup OpenStorage API can increase the role of disk in the protection of critical data and dramatically improve recovery times.