

High-Availability Tape Backup

in a Cluster Environment

Many vendors offer tape backup software to support clustering, but not all such software maximizes the benefits of a cluster environment. A true cluster-aware tape backup configuration has the ability to integrate tape software and hardware with cluster services, and to recover from backup operations interrupted during a failover. VERITAS Backup Exec™ 9.0 *for Windows Servers* software can be used in combination with the Dell™ PowerVault™ 132T Fibre Channel tape library or the PowerVault 136T Fibre Channel tape library, or both, to help provide the benefits of cluster-aware tape backups.

BY RICHARD GOLASKY AND NAM NGUYEN

System administrators increasingly use server clusters in network environments to achieve the high availability required for 24x7 business operations. In a cluster configuration, failover nodes help prevent downtime by taking control of cluster resources from the failed cluster node and then restarting the applications. Some examples of cluster resources are file shares, network names, databases, and tape backup software.

Implementing a tape backup system in a cluster environment allows administrators to back up data even when the tape backup server fails. A cluster-aware tape backup solution should perform the following duties:

- Maintain all configuration settings for tape backup
- Use a shared database, or catalog, for backing up and restoring data
- Automatically assume the role of a tape backup server
- Automatically restart any backup job(s) interrupted during the failover
- Automatically remove media left in tape drives during the failover

- Back up and restore the system state data, including the cluster quorum
- Support disaster recovery of a cluster

In addition, in a storage area network (SAN) configuration, a cluster failover ought not to affect tape backup operations on other servers or clusters connected to the same SAN.

Cluster-aware tape backups using Dell and VERITAS products

Tape backup software to support clustering has existed for years, but each vendor has its own interpretation of how tape backup systems should operate in a cluster environment. VERITAS Backup Exec™ 9.0 *for Windows Servers* software can be combined with the Dell™ PowerVault™ 132T Fibre Channel tape library or the PowerVault 136T Fibre Channel tape library, or both, to enable cluster-aware tape backups. This combination of Dell hardware and VERITAS® software allows a failover node of a tape backup server to perform the cluster-aware tape backup solution functions

listed in the previous section. VERITAS Backup Exec cluster components support the Microsoft® Windows® 2000, Windows NT®, and Windows .NET operating systems.

VERITAS Backup Exec 9.0 *for Windows Servers* provides enhanced tape backup capabilities in cluster environments for Dell and Dell|EMC storage. Older versions of Dell tape libraries and Backup Exec software required human intervention to remove cartridges from the tape drives after a failover. In addition, if backup jobs were in progress during a node failover, backup operations would restart on the failover node at the beginning of the data set rather than from the point of failure. Even if a backup operation were 99 percent complete when a failover occurred, the backup would start over again, extending the backup window and using extra tape to back up data that had already been backed up.

Backup Exec 9.0 *for Windows Servers* includes a CheckPoint Restart feature. If a cluster failover occurs during an active backup operation, backup jobs restart from the point of failure instead of from the beginning. Once new backup tapes are allocated, Backup Exec restarts the backup operation at the last file that was being backed up. The CheckPoint Restart feature includes some restrictions: database backups and backups using the VERITAS Open File Option or Intelligent Image Option that are interrupted during failover do not start from the point of failure, but instead restart from the beginning of the last drive being backed up.

Backup Exec 9.0 *for Windows Servers* also includes a cluster installation wizard. The wizard guides operators step by step through the installation process, improving installation compared to previous Backup Exec versions.

Failover process of a tape backup server

In a multinode cluster environment in which all cluster nodes have backup server capability (such as the third configuration option described in the “Options for tape backup configurations” sidebar), each node in the cluster can take ownership of resources, including the tape backup software services. The server in control of the tape backup resources is known as the controlling node, and standby servers are known as the failover nodes. If a controlling node fails, then all tape backup resources move to a failover node, which then becomes a new controlling node. The new controlling node begins to function as the tape backup server, thus continuing the original backup operations and maintaining all previous job scheduling and media management services.

In a cluster configuration,
failover nodes help
prevent downtime by
taking control of cluster
resources from the
failed cluster node.

The Backup Exec shared SQL database is the key to accomplishing tape backup device and media management between cluster nodes. The database must reside on a shared clustered disk owned by the controlling node, so that the controlling node always has direct access to the database. When a failover occurs, the database resource must move with the other Backup Exec resources. The database contains information about all tape hardware devices, media, and backup and restore sessions.

When a tape backup server fails during backup operations, all I/O to each tape drive that was in use by the failing server will stop. During the failover process, the failover node invokes tape backup software components to automatically remove the cartridges left in the tape drive by the failing server, and return them to the original slots. Once the tape is unloaded, the tape backup software loads a different cartridge into the drive, and backup operations automatically resume. Another cartridge is used because, when a backup operation is unexpectedly interrupted, the tape software has no opportunity to write an end-of-data marker to the tape. The absence of the marker prevents Backup Exec from appending the rest of the backup job to the original cartridge.

The automatic removal of cartridges from the tape drives represents an improvement over past versions of Backup Exec and Dell PowerVault libraries. In prior releases, if a node failed during backup operations, the failover node could not unload and eject the tapes because of SCSI reservation conflicts, and orphaned cartridges remained in the drives. As a result, tape drives were not available for use until an administrator manually removed the cartridges. If a tape drive was the only one available for that backup job, the backup job might not have continued or restarted on the failover node.

Backup disk selection in a cluster environment

Administrators defining backup sets in a cluster environment must select the correct backup source node. Selecting the correct source enables the tape backup server to guarantee itself access to the resource regardless of which cluster node owns the disk resource. In a cluster environment using Microsoft Cluster Service (MSCS) and VERITAS Backup Exec software, disk resources for tape backups appear under the following node names:

- Name of the physical server
- Virtual name of the MSCS cluster
- Virtual name of the Backup Exec cluster
- Virtual server name of the disk resource

In the following example, consider a two-node cluster using Backup Exec 9.0 *for Windows Servers*:

- Server 1: Asia
- Server 2: Europe

OPTIONS FOR TAPE BACKUP CONFIGURATIONS

In a cluster environment, tape backup configurations vary depending on data center requirements. Three general configuration options exist:

- Configured with stand-alone dedicated backup server
- Configured with clustered dedicated backup server
- Integrated configuration in which all cluster nodes have backup server capability

Configuration 1: Stand-alone dedicated backup server

A stand-alone dedicated tape backup server backs up cluster nodes through a Gigabit Ethernet* network connection (see Figure A). The backup server uses the virtual server name of the cluster resources to maintain continuous access to the cluster nodes, regardless of which cluster node owns the cluster resources. Although this method helps guarantee a backup of the cluster nodes, backups are performed across the network, which can degrade performance. This configuration lacks a failover node to resume backup operations if the tape backup server fails, and thus is not a true cluster-aware tape backup solution.

Configuration 2: Clustered dedicated backup server

In this scenario, one node in the cluster is designated as a tape backup server (see Figure B). This method offers a slight performance advantage over the first configuration, because cluster resources are backed up directly, rather than over the network—provided that the tape backup server controls the resources. If

cluster resources move to a failover node, the tape backup server can conduct backups through the network using the resource's virtual server name. However, if the tape backup server fails, backup operations are completely suspended. This configuration is not a true cluster-aware tape backup solution.

Configuration 3: Backup server capability for all cluster nodes

Figure C shows the true cluster-aware tape backup solution in which all nodes in the cluster can act as a tape backup server, using the tape backup software's cluster-aware components. Tape backup operations continue on a failover node with minimal interruption when a backup server fails; backup jobs are automatically restarted and can complete in full after a failover.




Figure B. Configuration with a clustered dedicated backup server

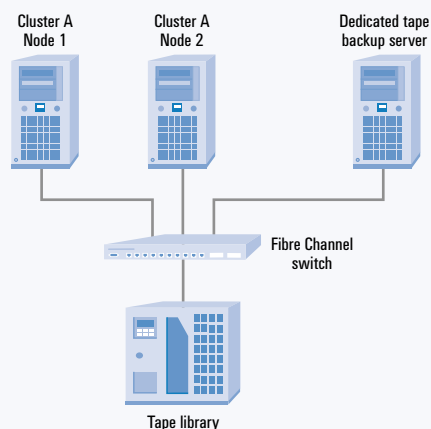


Figure A. Configuration with a stand-alone dedicated backup server

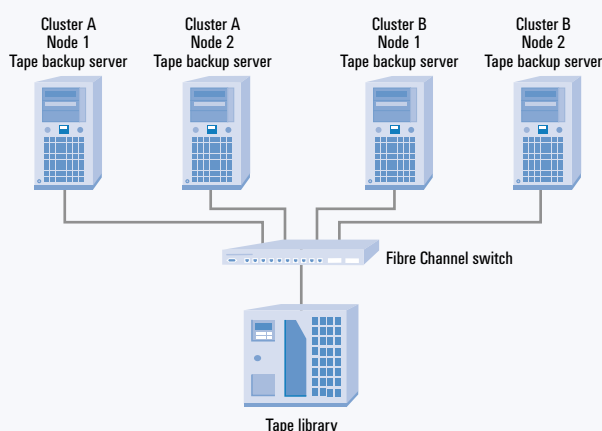


Figure C. Integrated configuration in which all cluster nodes have backup server capability

*Gigabit Ethernet indicates compliance with IEEE® 802.3ab and does not connote speeds of 1 Gbps.

- MSCS cluster virtual name: ASIA-EUROPE
- Backup Exec cluster virtual name: BE-ASIA-EUROPE
- Three virtual server names for the three fileshare groups: Drive G, Drive H, and Drive I

The cluster disk resources are drives E, F, G, H, I. Server 1 (Asia) controls resources F, G, and I. Server 2 (Europe) controls resources E and H.

The proper way to guarantee access to any logical drive of a cluster disk resource is to back up the node through its virtual server name, which is created in the MSCS Cluster Administrator. At first glance, the backup may appear as if it will be conducted through a network connection (\\UNC\share), but Backup Exec can differentiate between a cluster disk local to the tape backup server and one remotely connected through the network. In this scenario, drives E, F, G, H, and I must be backed up. Selecting the improper node name may result in the tape backup server being unable to access the disk drive.

The Backup Exec Backup Selection window displays all server names under the Microsoft Windows Network heading (see Figure 1). The selection process for backing up drives E, F, G, H, and I is as follows:

- Drive E (quorum disk) is always assigned to the MSCS virtual name. Select drive E from the ASIA-EUROPE server name list.
- Drive F (Backup Exec shared database) is always assigned to the Backup Exec virtual name. Select drive F from the BE-ASIA-EUROPE server name list.
- Drives G, H, and I each have a separate virtual server name assigned. Select each drive under its corresponding virtual server name.

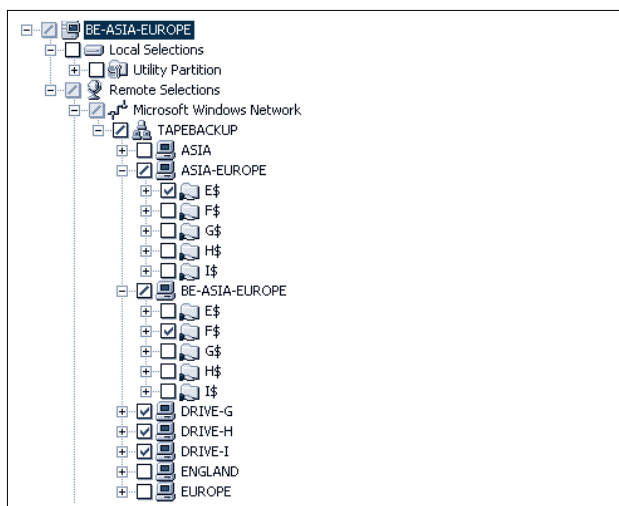


Figure 1. The Backup Exec Backup Selection window

During the failover

process, the failover node

invokes tape backup

software components to

automatically remove the

cartridges left in the tape

drive by the failing server,

and return them to the


original slots.

Once saved, the backup selections are permanently recorded and always available to any tape backup controlling node.

Tape backup that maximizes failover capabilities in cluster environments

When used in combination, Backup Exec 9.0 for Windows Servers and Dell PowerVault tape backup libraries integrate functionality to enhance the high-availability features of a PowerVault cluster backup solution based on Fibre Channel.

Administrators can use Backup

Exec to maintain configuration settings, back up and restore system data, and restart backup jobs from the point of failure. The enhanced architecture of VERITAS Backup Exec 9.0 for Windows Servers incorporates several new features for cluster environments, such as a shared database for restoring data and a cluster installation wizard for easy installation. VERITAS Backup Exec 9.0 for Windows Servers and the Dell PowerVault 132T Fibre Channel tape library or the PowerVault 136T Fibre Channel tape library, or both, can help provide the benefits of true cluster-aware tape backups. 

Acknowledgments

The authors would like to thank Pat Hanavan, director of engineering at VERITAS Software Corporation, who provided assistance with this article.

Richard Golasky (richard_golasky@dell.com) is a solution and systems engineer in the Dell Product Group—Data Protection. His responsibilities include all areas of tape backup, including Fibre Channel, SAN, network attached storage (NAS), and high-availability cluster and performance analysis. He has a B.S. in Electrical Engineering from Florida Atlantic University.

Nam Nguyen (nam_nguyen@dell.com) is a systems engineer in the Dell High Availability Cluster Development Group. Nam is the lead engineer for several Fibre Channel–based Dell PowerEdge™ cluster products. He has both a B.S. and an M.S. in Electrical Engineering from The University of Texas at Austin.

FOR MORE INFORMATION

<http://www.dell.com>

<http://www.veritas.com>