HP StorageWorks Integrated Non-disruptive Backup for Exchange 2003 using VERITAS Backup Exec

Solution overview

The protection of Microsoft Exchange databases is vitally important, and with new governmental regulations concerning the retention of e-mail, it is becoming increasingly harder to maintain reasonable backup windows. Now that e-mail in most cases must be maintained for years at a time, databases are getting larger and larger. Consequently, the time it takes to back up these databases is growing rapidly, sometimes to the point of being unmanageable. While there are many new solutions available today that specifically address e-mail backup and retention (such as ILM solutions), the vast majority of Exchange deployments are traditionally backed up over the LAN using a backup application and its agent/client package for Exchange.

The HP StorageWorks Integrated Non-disruptive Backup for Exchange 2003 using VERITAS Backup Exec provides a fast, effective Exchange backup solution without the need for unfamiliar technology. It makes use of the Advanced Disk-Based Option in VERITAS Backup Exec, Microsoft's Volume Shadow Copy Service (VSS), the HP StorageWorks VSS Hardware Provider for EVA software, and the new EVA4000/6000/8000 storage array series with the HSV200/210 controller set.

With these elements in place, along with a supported HP StorageWorks tape library, a simple backup job can be created within Backup Exec to create *snapclones* of the Exchange volumes on the EVA, present the snapclones to the Backup Exec server over a SAN fabric, and perform a local backup of the databases directly to tape. By using this method, the backup window necessary for completing Exchange backups is dramatically reduced. Since all the backup data is transferred over the SAN for vastly increased data throughput, backup resources can be freed for use elsewhere up to 75% faster than with traditional backup methods. In addition, manipulation of the backup data is easily handled via the Backup Exec interface, in traditional backup/restore fashion.

Evaluation overview

In order to fully understand the differences between this solution and other available options, three available methods of backing up an Exchange database using Backup Exec were examined:

- The first method is the traditional backup method in which a Backup Exec database agent is installed on the Exchange server. This method maintains availability of the storage groups to Exchange users while the data is being backed up over the LAN to the Backup Exec media server.
- The second method involves the Backup Exec Advanced Open File Option (AOFO) feature, which uses the Microsoft Volume Shadow Copy Service (VSS) to create snapclones of the LUNs that contain the Exchange data being backed up. This snapclone is presented to the Exchange server as a separate LUN, which is then backed up over the LAN to the Backup Exec media server.
- The third method is the focus of this solution—using the Backup Exec Advanced Disk-Based Option (ADBO) feature to create snapclones of the LUNs containing the Exchange data. The EVA presents the LUNs to the Backup Exec media server where the data can be backed up over the SAN. This option provides the best performance writing to tape, while minimizing the impact to the Exchange server during backup periods. This approach minimizes user access disruptions.

Note
The ADBO backup backs up data by storage group. Consequently, restores of individual mailboxes are not available with this solution. Because the snapclone involves the entire LUN, it is not feasible to use the off-host snapclone method to perform backups of individual mailboxes. The focus of this solution is to improve backup speed to tape and disaster recovery in the event of a failed or corrupt storage group.

Prerequisites

The requirements for this solution are as follows:

- Microsoft Exchange 2003 Server
- Veritas Backup Exec 10.0 with the following options:
  - SAN Shared Storage
  - Agent for Microsoft Exchange
  - Advanced Disk-Based Option
  - Remote Agent for Windows (installed on Exchange Virtual Server)
- HP EVA4000, 6000, and 8000 storage arrays (with the HSV200 and HSV210 controller sets)
  - HP Business Copy EVA license, installed on the EVA storage management server (this enables the snap functionality on the EVA)
- HP StorageWorks tape library
- HP EVA VSS hardware provider
- Microsoft patches 831112, 833167, 867667 (these must be installed on both the Exchange server and the Backup Exec media server)
- Multipathing software, such as Microsoft MPIO (required for connectivity to the EVA), must be installed on both the Exchange server and the Backup Exec media server

In addition, the following system modifications should be made to both the Exchange nodes and the Backup Exec server:

Note
These changes are not necessary for the solution to work. Rather, they are tuning parameters that optimize how the Exchange server and the Backup Exec server access devices on the SAN. Always use caution when editing the Windows registry.

- On servers having 1 GB or more of RAM, add the `/3GB` switch and the `/USERVA` switch to the `boot.ini` file on each cluster node, as per Microsoft recommendations for optimizing memory usage with Exchange Server 2003.
- Add the following settings to the registry entries of the HBAs in both the nodes and the Backup Exec server (all HBAs were Emulex-based). These settings are recommended to improve performance with the EVA when using Emulex-based HBAs:
  - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lpxnds\Parameters\Device\DriverParameter: QueueTarget=0`
  - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lpxnds\Parameters\Device\DriverParameter: QueueDepth=128`
  - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lpxnds\Parameters\Device\NumberOfRequests: 0x00000096 (150)`

Test configuration

Testing was done using a Microsoft Cluster environment to house the Exchange 2003 databases. The two-node cluster was created using two HP ProLiant DL580 servers. The cluster was connected via the SAN to an EVA8000 storage array containing the HSV210 controller set. The Exchange instance supported 2,000 users spread across four storage groups. Each storage group used two LUNs presented from the EVA as volumes—one LUN for the databases and one for the transaction logs. This created a total of eight LUNs dedicated to Exchange. Because this was a cluster setup, an Exchange Virtual Server was created on the cluster to actually present the databases to users.

The EVA8000 was configured with separate disk groups to house the database LUNs and the transaction log LUNs. An additional disk group was allocated to house the snapclones created during the solution backup. This configuration was done in accordance with Microsoft best practices for Exchange and HP recommendations for EVA/Exchange configurations.

A separate server was set up as the Backup Exec media server, running Backup Exec 10.0. This server was connected to the SAN and had access to both the EVA (in order to receive the snapclones for backup) and an HP MSL6030 tape library with two HP LTO-3 tape drives.

The Exchange instance was populated for testing using Microsoft's LoadSim 2003 utility. Two additional SAN servers were enlisted to act as LoadSim servers to share the load more effectively. Each LoadSim server handled 1,000 user accounts. LoadSim was used to initialize the databases and transaction logs with data. For the portions of the testing in which Exchange user activity was necessary, the LoadSim servers "logged on" all 2,000 users and performed various e-mail, schedule, and public folder functions. To measure the impact of this solution on the Exchange server as compared to the other methods, Performance Monitor was set up on an alternate Windows 2003 server. A counter log was created with counters from the active node, and performance data was gathered by the alternate server during the testing.

The HP StorageWorks VSS Hardware Provider for EVA must be installed on both the Exchange server (in this case, both cluster nodes) and the Backup Exec media server. This package provides the communication link between the Microsoft VSS service and the EVA when creating snapshots/snapclones. The provider software requires that the servers have TCP/IP connectivity to the SAN Appliance that is used to configure the EVA. During installation of the provider, a window is displayed requiring the IP address of the SAN Appliance and the login/password of the administrator account used to log in to Command View EVA. After the connection is made, choose the appropriate HSV storage system and choose Snapclone under Snapshot Type. Enter the snapclone disk group created on the EVA to separately house the snapclones during the backup job.

NOTE:
Snapclones (rather than snapshots) should be selected as the preferred method of point-in-time copy when configuring the hardware provider. Maintaining a snapshot during a backup job generates excess I/O overhead on the Exchange server because the snapshot constantly references the original LUN. This referencing, especially on the transaction log LUN, can cause enough overhead to make the VSS process fail during a backup job. Snapclones do not generate this excess I/O.

Figure 1 illustrates the test configuration.

Figure 1. Test configuration

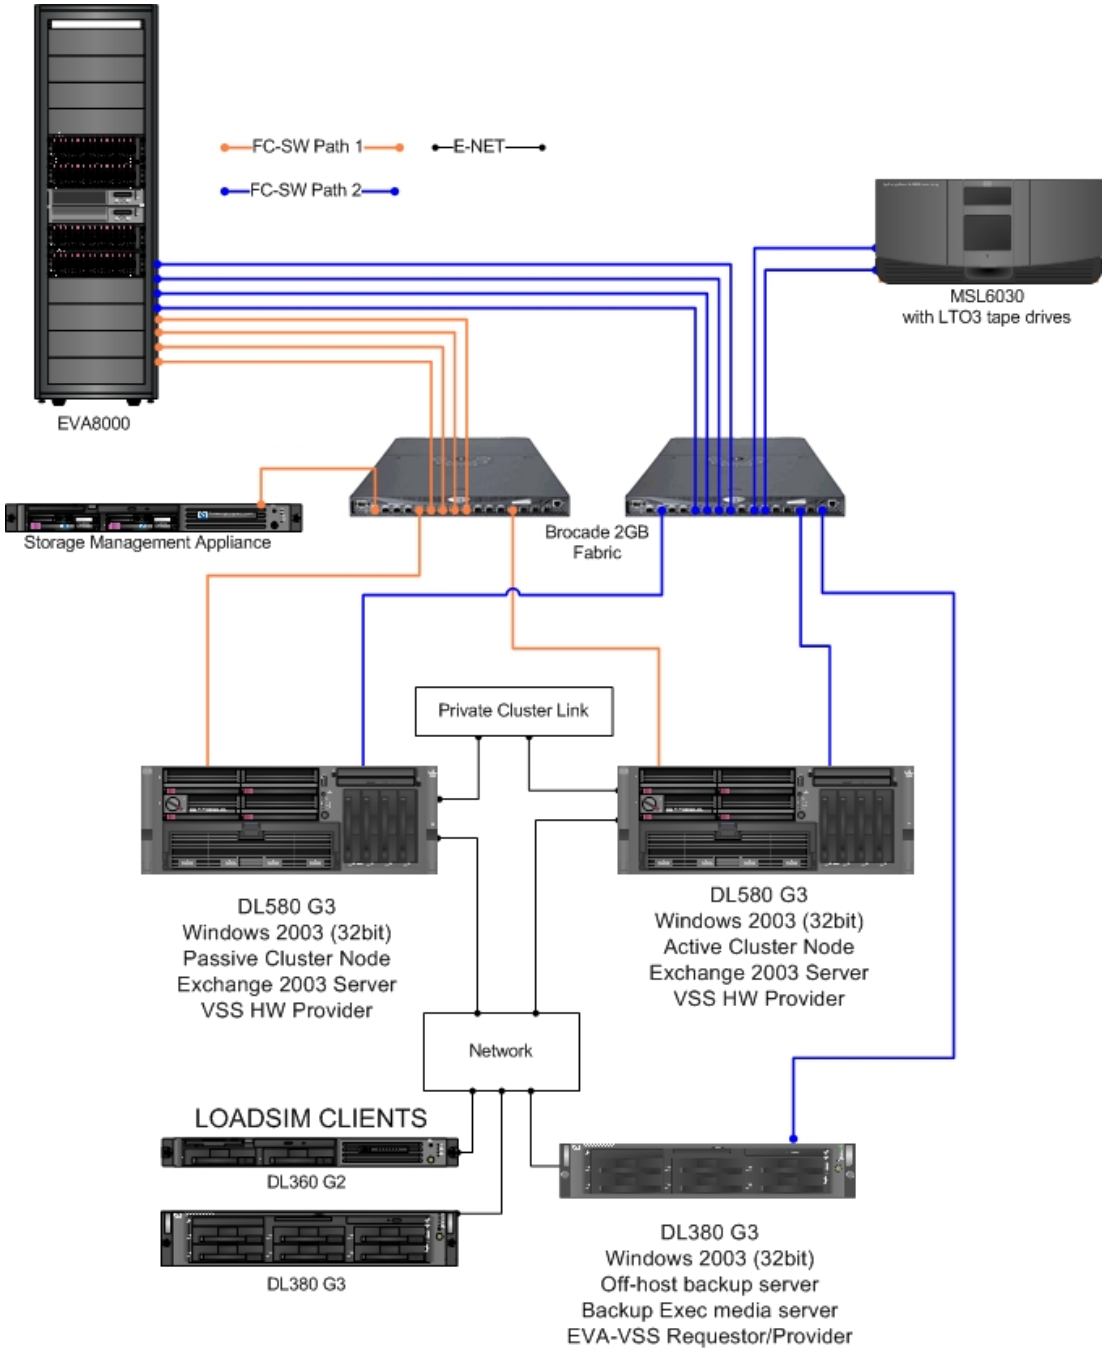

The Solution

To perform the ADBO Exchange backup:

1. Within the Backup Exec administration console, click the **Backup** tab. The Backup Job Properties screen is displayed.
2. Enter a unique name for the backup job in the Selection list name dropdown box.
3. In the Selection List window, click **Remote Selections** and search in the proper Windows domain for the Exchange Virtual Server name.

Note:
Choose the EVS name rather than the active cluster node name. Otherwise, the Backup Exec Exchange agent will not recognize the Information Store and will not handle the database quiescence properly. If a cluster setup is not being used, this does not apply.

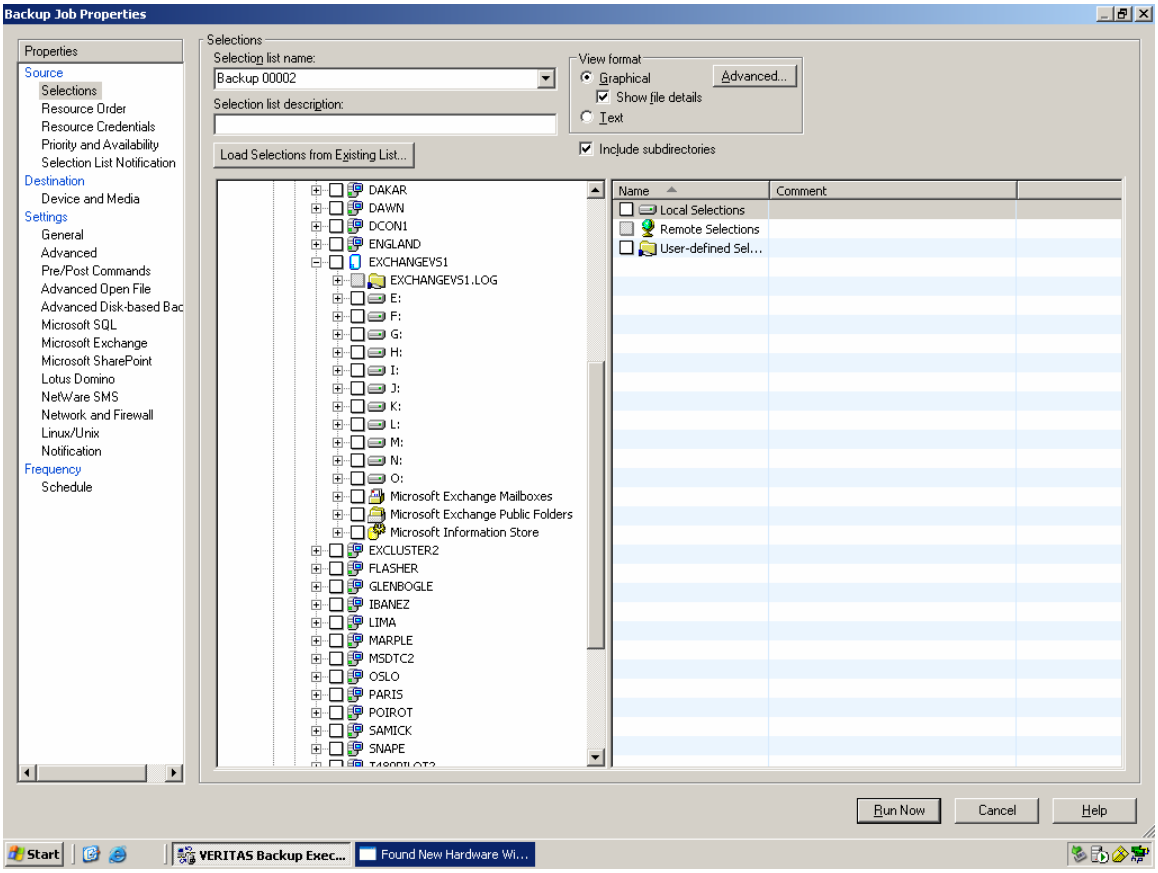

4. Locate **Microsoft Information Store** under the EVS selection. Click this check box to back up all the storage groups, or drill down and check each individual storage group as required. If the Microsoft Information Store is not listed, verify that the information store on the Exchange server is online.

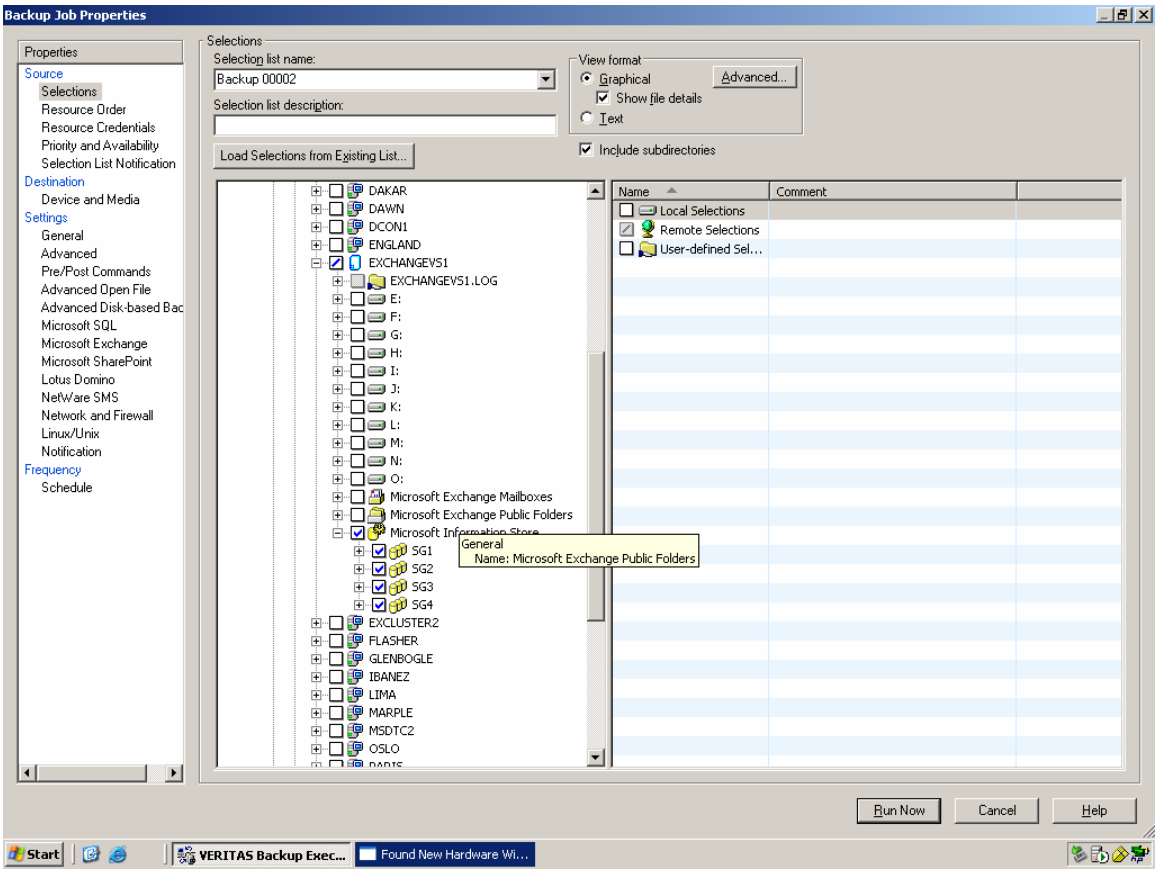

5. Choose the appropriate backup settings. Click **Microsoft Exchange** under **Settings** in the Properties window. Select the proper backup method from the **Exchange server backup method** dropdown box. Because this solution does not deal with individual mailbox backups, the mailbox backup method can be ignored. HP recommends that the last two boxes, dealing with performing a consistency check of the snapclone and proceeding with the backup if the consistency check fails, remain checked. If the check fails, the backup proceeds over the LAN as with a traditional Exchange backup. If this occurs, backup time is significantly increased, and other backup windows can be hampered.

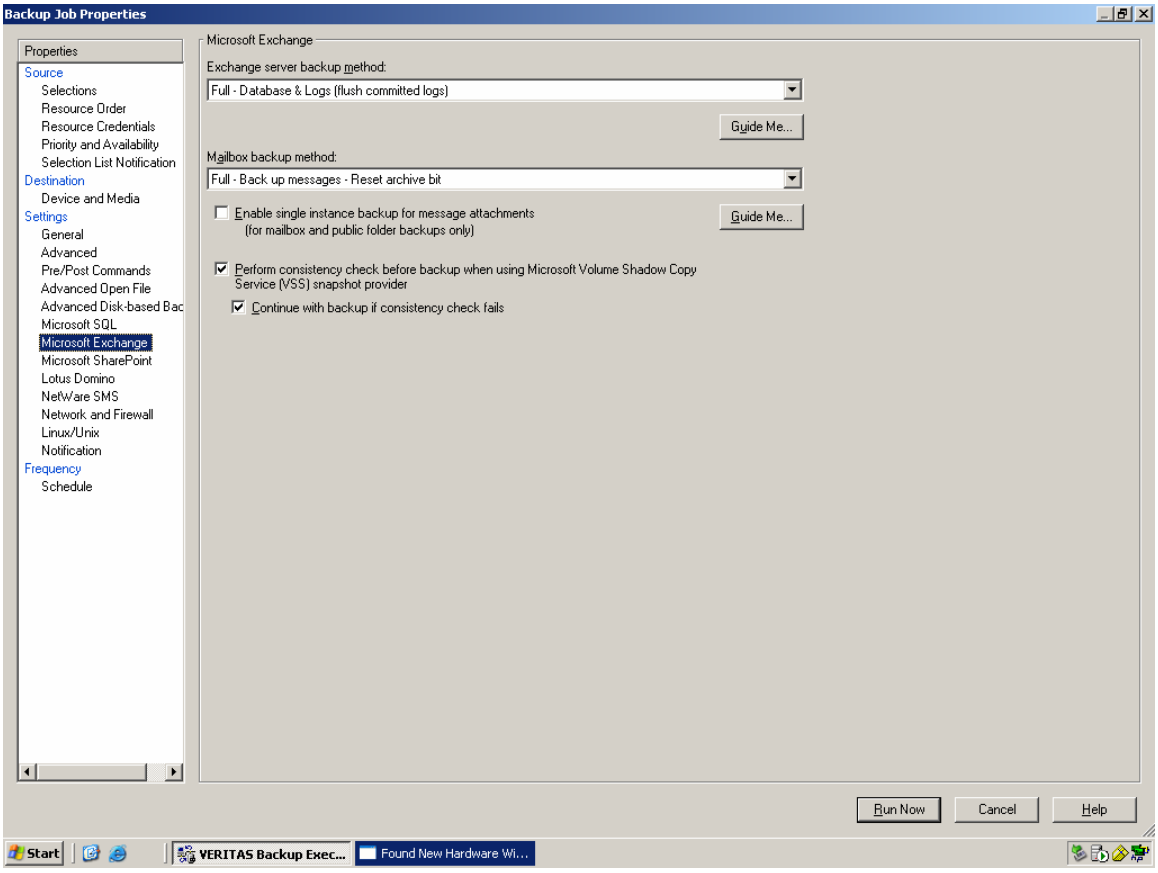

6. Select **Advanced Disk-Based Backup** under **Settings** in the Properties window. Click the **Use offhost backup…** check box. The ADBO can now be configured. The snapshot provider is the HP VSS Hardware Provider, so select the **Hardware** option in the **Snapshot provider** dropdown box.

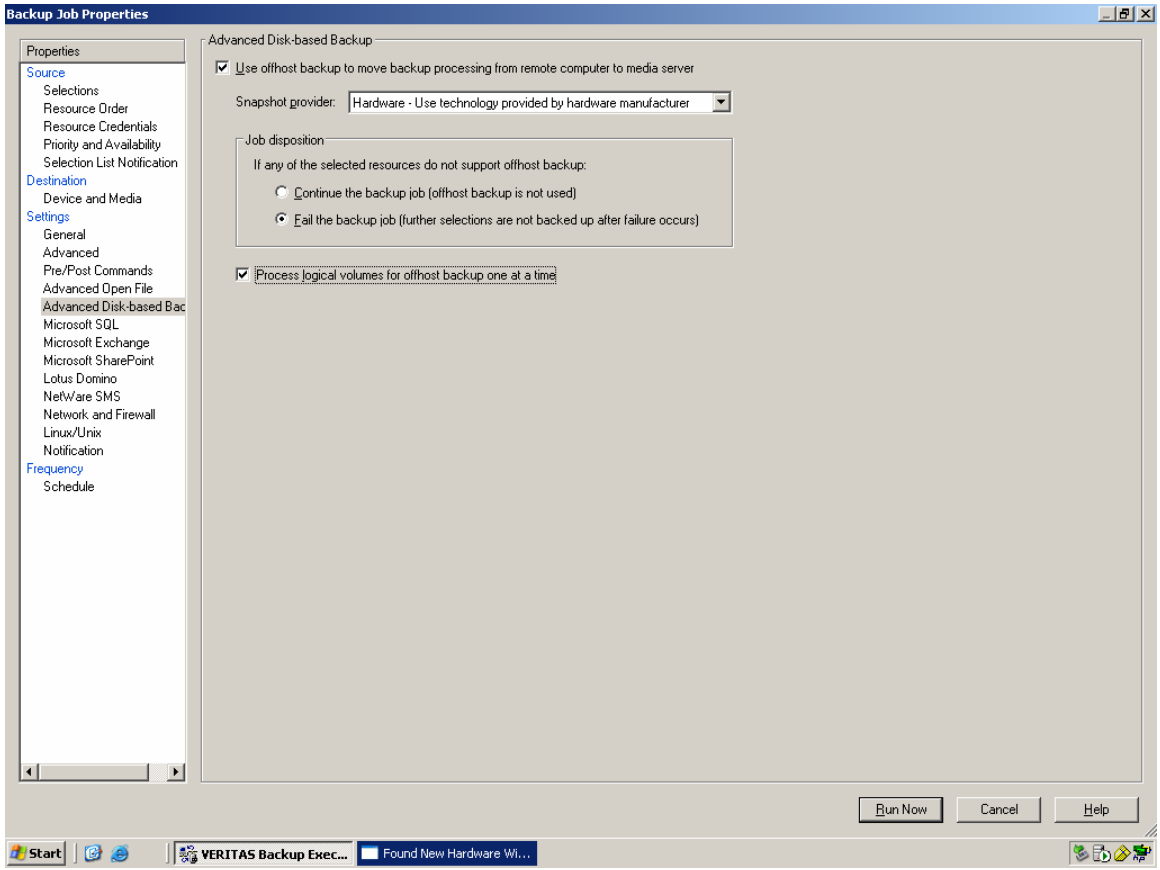

7. If more than one storage group is being backed up in a single job, check the **Process logical volumes one at a time** option. During testing, attempts to handle multiple (more than two) simultaneous snapclone requests to the VSS writer during the backup job were unsuccessful. By selecting the “one at a time” option, VSS can handle a single storage group snapclone request. Backup Exec performs the backup of the snapclone volumes for the storage group (one for the database volume and one for the log volume), and then the snapclones are discarded, making room for the next request. This pattern continues until the backup job is complete.

8. If required, you can create a schedule for this job so that the job can be repeated as necessary. You can create additional jobs to perform copy, differential, or incremental backups. These jobs can be scheduled between full backups for a more precise level of protection. To create a backup schedule, select **Schedule** under **Frequency** in the Properties window.

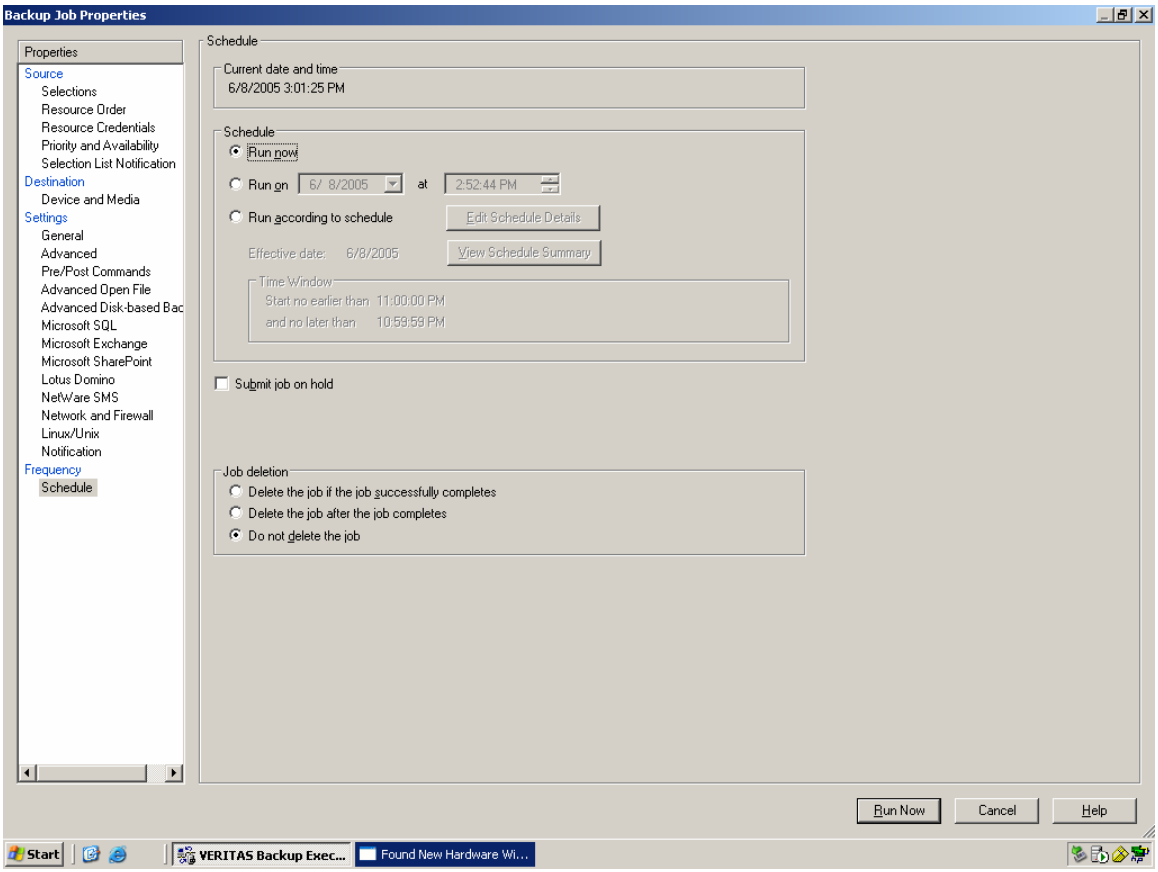

9. Click **Run Now** to submit the job. The job is run immediately, or whenever the schedule dictates. The progress of any job can be monitored in the Job Monitor tab.

   During the job, the Backup Exec media server communicates with its remote agent on the Exchange Virtual Server to initiate the backup. The remote agent then communicates the request to create a snapclone on the EVA for each of the requested Exchange LUNs via the HP VSS Hardware Provider. The Provider is the conduit between VSS and the EVA that facilitates the snapclone process. After the snapclone has been created but not yet normalized, it is presented by the EVA to the Backup Exec media server and mounted as a volume. The Exchange *eseutil* utility and appropriate DLLs are copied from the Exchange server to the Backup Exec media server so that a consistency check can be run on the Exchange data after the snapclone has been presented. When that is complete, Backup Exec begins the backup to tape.

Test Comparison Results

To establish a baseline for comparison purposes, the two LoadSim 2003 servers ran a simulation to mimic typical Exchange usage for the 2,000 users. Performance Monitor was configured on an alternate Windows 2003 server to capture counter data from the active node of the Exchange cluster. The counters chosen for measuring the performance of the Exchange cluster node are listed below:

- Processor
  - % Processor Time
- Memory
  - Available MB
  - Pool non-paged bytes
- MS Exchange Mailbox
  - Receive queue size
  - Send queue size
  - Avg. del. Time
- SMTP Server
  - Categorizer queue length
  - Local queue length
  - Remote queue length
- Logical Disk
  - Avg. queue length
  - Current queue length
  - Avg. sec. read
  - Avg. sec. write
- Physical Disk
  - Avg. disk queue length
  - Current disk queue length

Per LoadSim recommendations, the simulation was run for at least 8 hours. This allowed a two-hour "warm-up" period as well as a two-hour "warm-down" period, leaving at minimum a four-hour window when the simulation is at peak levels. During the testing of the various backup options, LoadSim was required to run much longer than 8 hours, as the backup windows for LAN-based and AOFO backups far exceeded the window. Therefore, the LoadSim configuration was adjusted for each test to ensure that all backup jobs were processed during the "steady-state" period. In addition, averages of all the performance data counters were taken to gain an overall view of the server impact for the duration of each job.

After the baseline performance of the Exchange cluster was established, measurements were done with the three listed BackupExec backup options (LAN, AOFO, and ADBO) to show the impact on the Exchange server. The comparisons of the Performance Monitor counters for each of the three backup options are listed in Table 1.

Table 1. Backup performance comparison

PERFORMANCE MONITOR COUNTERS	LOADSIM BASELINE	LAN BACKUP	AOFO BACKUP	ADBO BACKUP
Logical Disk: Avg. Disk Queue Length	12.98	40.51	51.81	119.43
Logical Disk: Avg. Disk sec/Read	0.0135	0.0276	0.0425	0.1186
Logical Disk: Avg. Disk sec/Write	0.00517	0.02145	0.03594	0.1356
Logical Disk: Current Disk Queue Length	11.97	36.36	43.54	98.10
Memory: Available MBytes	1855.23	1823.73	1757.40	1704.20
Memory: Pool Nonpaged Bytes	60628305.78	60830646.92	61125427.2	62453965.88
MSExchangeIS Mailbox: Avg. Delivery Time	26532.06	2682536.68	724869.42	6435969.57
MSExchangeIS Mailbox: Send Queue Size	10.78	569.58	62.01	227.14
Physical Disk: Avg. Disk Queue Length	13.01	40.58	100.97	119.69
Physical Disk: Current Disk Queue Length	11.96	36.39	47.72	98.21
Processor: % Processor Time	31.30	36.95	39.63	35.74
SMTP Server: Categorizer Queue Length	0.085	70.77	5.038	404.62
SMTP Server: Local Queue Length	12.31	507.30	385.29	2771.40

The counters that measure disk usage increase when performing snapclone backups. This is due to the normalization of the snapclones by the EVA during the backup job. However, the overall impact to the Exchange server during an ADBO backup is minimized due to the reduction in the backup window. This is where the benefit of this solution is really seen.

Table 2 shows an average of the times for the jobs run in the table above. The ADBO backup produces an almost 75% reduction in the backup window from traditional LAN backups, and a 70% reduction over the AOFO onhost-snapclone backup job. This is a critical enhancement that will dramatically reduce the cost in time that administrators must spend on Exchange backups. In addition, the elapsed time below brings into perspective the counters in the table above—the higher disk usage impact is only felt over the three-hour ADBO backup window, while the LAN-based backup counters are measuring impact over a nearly 12-hour backup window. In addition, the measured impact on system memory and processor usage is minimal, even when compared to the LoadSim baseline counters.

Table 2. Average backup times

Backup Option	Elapsed Time (hours:minutes)
LAN-based	11:42
Advanced Open File Option (onhost)	9:45
Advanced Disk-Based Option (offhost)	2:58

Restoring

The main limitation of this solution is the restore process. By default, restores run over the LAN. However, an administrator does have options if a LAN-based restore would be too lengthy. A combination of manual manipulation of the EVA and scripting on the Exchange server can be used to restore the data. In addition, third-party software solutions can be added to help manage the snapclones. However, the simplest method is to install the Backup Exec media server software on the Exchange server for use only in restore situations. Using the SSO option to share devices on the SAN with the Backup Exec primary media server, the secondary media server instance on the Exchange server can redirect the Exchange restore to itself. In this case, data is received from tape over the SAN and restored directly to the storage group volume(s). This method dramatically reduces the restore time from several hours to minutes, depending upon the size of the data. The downside of this solution is that you have to purchase software and licenses for a package that, hopefully, will never be used. In addition, the restore job creates a larger load on the Exchange server during the job. However, the tremendous reduction in time required to get Exchange back up and running for the users should outweigh these issues in most situations.

Conclusion

The HP StorageWorks Integrated Non-Disruptive Backup Solution for Exchange 2003 with VERITAS Backup Exec provides a significant improvement in Exchange backup performance. *Because of the dramatically reduced backup times, the overall impact on Exchange is minimized*. Shorter backup jobs free up tape and library components for use with other backup jobs, thereby having a further positive impact on backup windows as a whole. This solution can be a very cost-effective and time-effective tool for Exchange administrators.

For more information

www.hp.com/go/ebs

406689-001, September 2005

invent