BrightStor® ARCserve® Backup for Windows

Enterprise Option for Microsoft SQL Using HDS Snap-Shot Guide

r11.5

Contents

Chapter 1: Introducing the Option

BrightStor® ARCserve® Backup is a comprehensive storage solution for applications, databases, distributed servers, and file systems. It provides backup and restore capabilities for databases, business-critical applications, and network clients.

Among the options BrightStor ARCserve Backup offers is the BrightStor® ARCserve® Backup Enterprise Option for Microsoft SQL Using HDS Snap-Shot. This option combines the power of high performance storage management tools with a custom option set. It is specifically designed for environments that use Hitachi ShadowImage to back up frequently accessed, mission-critical data in a Microsoft SQL Server 2000 database. Since the option enables backups to be performed to the local S-VOLs and resynchronized to the production server, businesses that require continuous access to information can back up and restore extensive amounts of data with minimal impact on production time while preserving maximum data integrity.

The Enterprise Option for Microsoft SQL Using HDS Snap-Shot takes advantage of the Microsoft SQL Server 2000 ShadowImage replication feature, which creates a consistent image of both the database and the transaction log. The option includes certain metadata used for recovering the database as part of the snapshot image.

Using BrightStor ARCserve Backup and the Enterprise Option for Microsoft SQL Using HDS Snap-Shot you can also perform the following tasks:

- Remotely manage backups of Microsoft SQL Server 2000 databases

- Schedule your backup

- Back up to a wide range of media storage devices

Option Architecture

Two servers are involved in a snapshot backup and restore:

- **Production server:** A production server connects to the P-VOLs on the Hitachi disk array. The following software must be installed on a local disk of the production server:

  – Microsoft SQL Server 2000

  – BrightStor® ARCserve® Backup Agent for Microsoft SQL Server

  – Hitachi specific software

- **Backup server:** BrightStor ARCserve Backup uses the option to back up a Microsoft SQL Server 2000 database that resides on one or more Hitachi disk array P-VOLs. The backup server connects to the S-VOLs on the Hitachi disk array. A tape library is connected to the backup server to store the backup data. The following software must be installed on the local disk of the backup server:

  – BrightStor ARCserve Backup

  – The Enterprise Module

  – The Enterprise Option for Microsoft SQL Using HDS Snap-Shot

The following illustration displays this configuration:
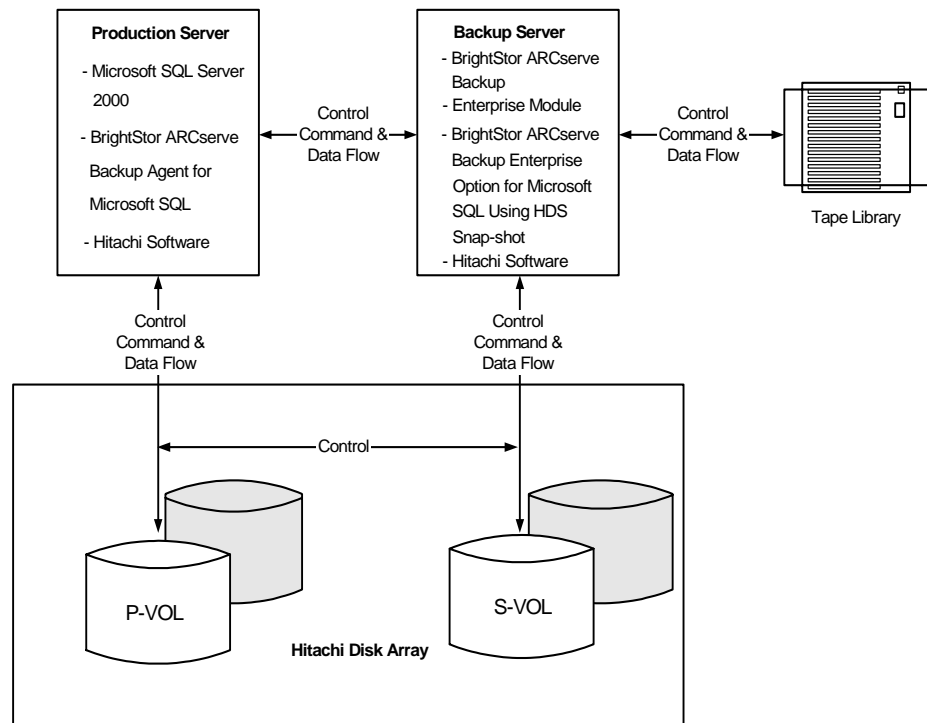
How the Option Works

The Hitachi ShadowImage is a storage manager Logical Unit Number (LUN) mirroring product that runs on the Hitachi Freedom Storage™ Lightning 9900™ and 7700E™. The Hitachi ShadowImage mirrors one or more volume images within a disk array between a primary LUN (called a P-VOL) and a secondary LUN (called an S-VOL) with no server CPU usage.

Using the Hitachi ShadowImage, you can synchronize, split, resynchronize, and reverse resynchronize the mirrored P-VOLs and S-VOLs. *Synchronization* is the creation of a mirrored copy of a P-VOL on an S-VOL. *Splitting* ends the mirroring between a P-VOL and an S-VOL and makes them independent. *Resynchronization* recreates a mirrored copy of a P-VOL on an S-VOL after splitting. *Reverse resynchronization* creates a mirrored copy of an S-VOL on a P-VOL.

The option integrates with the Hitachi ShadowImage to support backup and restore of Microsoft SQL Server 2000 databases. The option provides the ability for the Hitachi ShadowImage to perform multiple tasks without affecting the production systems. Using the option, you can perform the following:

- Backup, restore, or restore staging recovery operations

- Instantaneous recovery from a catastrophic failure

- Full data copies for testing and development

- Concurrent operations such as parallel data backups

The following steps describe the workflow when BrightStor ARCserve Backup uses the option to back up a Microsoft SQL Server database on a machine on which the BrightStor ARCserve Backup Agent for Microsoft SQL Server is installed. The Agent for Microsoft SQL Server facilitates communication between the option and the SQL Server database.

1. BrightStor ARCserve Backup sends a request to Agent for Microsoft SQL Server to freeze a particular Microsoft SQL Server 2000 database to prepare for a backup.

2. The Agent for Microsoft SQL Server initiates a snapshot of a database and creates a file in the agent home directory containing the metadata for the snapshot.

3. The Enterprise Option for Microsoft SQL Using HDS Snap-Shot instructs the Hitachi disk array to synchronize the S-VOLs with the P-VOLs, split the P-VOLs and the S-VOLs, and then mount the S-VOLs on the backup server.

4. The agent releases the snapshot of the database.

5. BrightStor ARCserve Backup sends a request to the agent to access the metadata file and back up the metadata to the storage media.

6. BrightStor ARCserve Backup backs up the Microsoft SQL Server database files and transaction log files from the local mounted S-VOLs to the storage media.

7. Depending on the backup options selected, BrightStor ARCserve Backup may then unmount the S-VOLs from the backup server and resynchronize the S-VOLs with the P-VOLs.

Chapter 2: Installing the Option

This chapter provides the information you must install and configure the Enterprise Option for Microsoft SQL Using HDS Snap-Shot. It also describes the pre-installation and post-installation tasks you must do.

System Requirements

Before installing the Enterprise Option for Microsoft SQL Using HDS Snap-Shot, verify that the following are installed and working properly:

- Hitachi Disk Array

- The production server

- The backup server

- Hitachi specific software (Hitachi DF500 Disk Array Management Program V6.00/3.50, Hitachi Command Control Interface Software)

- Agent for Microsoft SQL Server

- Microsoft SQL Server 2000

- BrightStor ARCserve Backup

Note: Hitachi DF500 Disk Array Management Program (DAMP) V6.00/3.50 and Hitachi Command Control Interface Software must be installed on both the production and backup servers. BrightStor ARCserve Backup must be installed on the backup server, and Microsoft SQL Server 2000 and Agent for Microsoft SQL Server must be installed on the production server.

If these applications are not already installed, you must install them before you install the Enterprise Option for Microsoft SQL Using HDS Snap-Shot.

Installation Prerequisites

Before you begin the installation of the Enterprise Option for Microsoft SQL Using HDS Snap-Shot, perform the following tasks:

- Create Basic disks and a command device.

- Configure the ShadowImage on the production server

- Configure the ShadowImage on the backup server

- Create mirrored volumes on the Hitachi disk array

Important! *Configuration requires complete knowledge of the Hitachi Disk Array Manager. Configuration should only be performed by the technician who configures and maintains the disk array.*

Before installing also verify that:

- Your system meets the software requirements needed to install the option. For a list of these requirements, see the readme file and the section System Requirements in this chapter.

- You have configured the hardware based on your hardware environment.

- You have made a note of any changes to the default installation path.

- You know the name and the password of the machine you are installing the option on.

- You have administrative privileges or the proper authority to install software on the computers where you will be installing the option.

Create Basic Disks and a Command Device

To create Basic disks and a command device, do the following:

1. Use the Hitachi DF500 Disk Array Management Program to create Basic disks and the command device on the production server after you install the production server.

The following example displays how the Basic disks are managed:
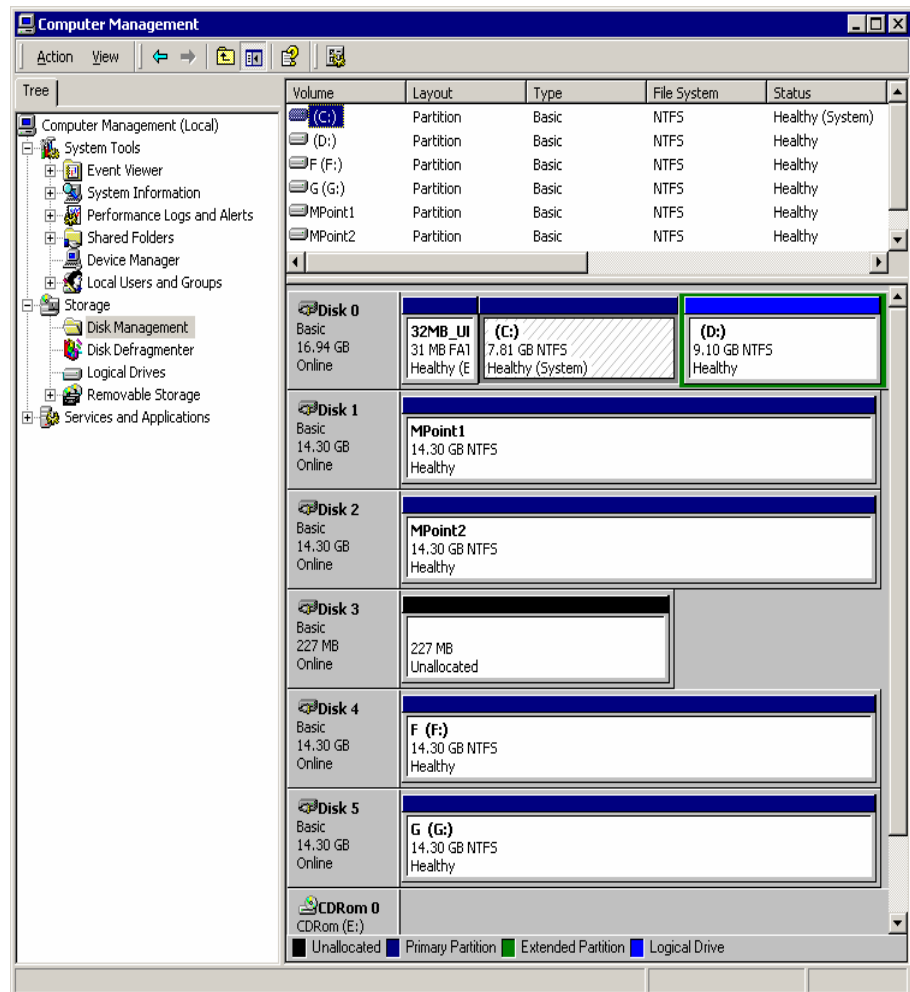
In this example, Basic disk 0 is a local disk. Basic disks 1, 2, 4, and 5 are Hitachi disk array P-VOLs. Basic disk 3 is the command device.

Note: Changing your hardware configuration may require updating the volume name. For example, pulling out physical disks, adding new physical disks, or rearranging internal physical disks may cause the volume names to change. You can verify changes using the "inqraid" utility from the command line.

2. Create mount points MPoint1 and MPoint2, Volume F: and G: using Basic disks 1, 2, 4 and 5 respectively on the production server, as displayed in the following window:
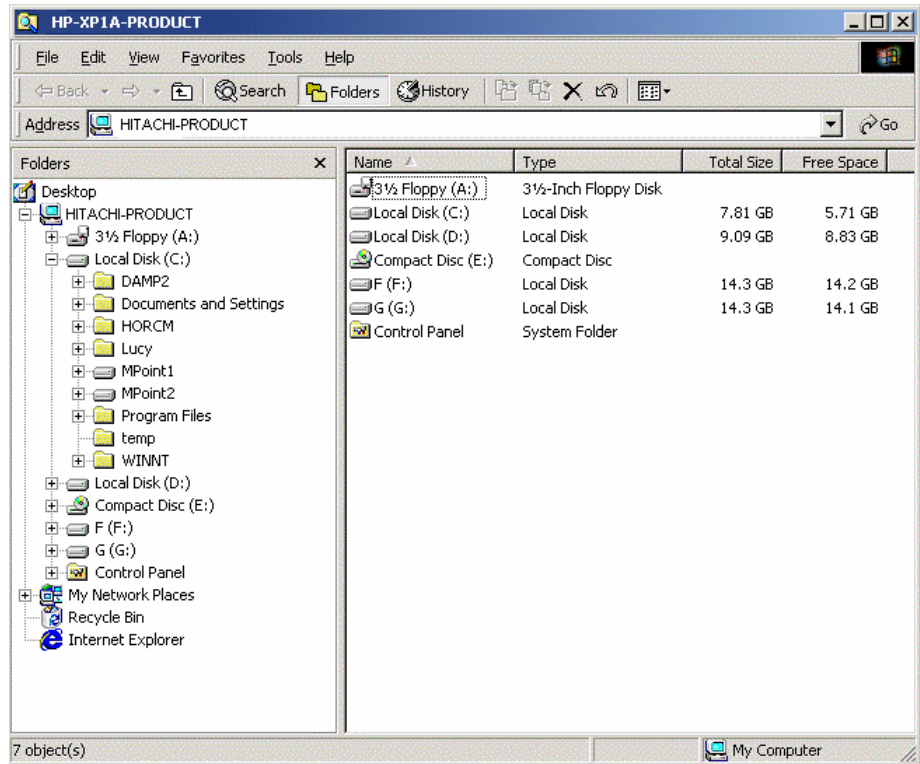
3. Create directories on the backup server to match the mount points MPoint1 and MPoint2.

Note: You must use the Basic disk type for all Hitachi disk array P-VOLs and S-VOLs.

Configure the ShadowImage on the Production Server

ShadowImage is in-system replication software that provides nondisruptive, high-speed data replication within Hitachi storage system to restore data from corruption. You must configure the ShadowImage after you install the Hitachi specific software on the production server. The following procedures provide a sample configuration of the ShadowImage. For more information about how to configure the ShadowImage, see Hitachi product documentation.

To configure the snapshot on the production server, follow these steps:

1. Use the Hitachi DF500 Disk Array Management Program to create a command device.

2. Identify the current command device by issuing the following command from a command prompt:

   ```
   horcctl -x findcmddev drive#(0-n),
   ```

   where *n* is the number of logical disks visible to your system.

3. Modify the file C:\WINNT\system32\drivers\etc\services on the production server.

   To modify the file, open the file with Windows Notepad and add the following line to the end of the file with 9900 as the port number:

   ```
   horcm              9900/udp                          #horcm
   ```

   Note: You can modify the port number value, but this value must not be used as a port number elsewhere in this file.

4. Modify the file C:\WINNT\horcm.conf.

   ```
   HORCM_MON
   #ip_address     service    poll(10ms)    timeout(10ms)
   111.222.33.11   horcm      12000         12000


   HORCM_CMD
   #dev_name                dev_name            dev_name
   \\.\PhysicalDrive3


   #specifications for this host.
   HORCM_DEV
   #dev_group    dev_name          port#        TargetID      LU#       MU#
   bab           C:\MPoint1        CL1-A              1        16        0
   bab           C:\MPoint2        CL1-A              1        17        0
   bab           F:                CL1-A              1        19        0
   bab           G:                CL1-A              1        20        0
   #specifications for the peer host.
   HORCM_INST
   #dev_group       ip_address            service
   bab              111.222.33.12         horcm
   ```

   - Replace **111.222.33.11** with the IP address of the production server.

- Replace **111.222.33.12** with the IP address of the backup server.

- Replace **\\.\PhysicalDrive3** with the command device.

- Use **bab** as the name of a group of paired logical volumes.

- Use **C:\MPoint1, C:\MPoint2, F: and G:** as the names of the paired logical volume within a group.

Configure the ShadowImage on the Backup Server

You need to configure the ShadowImage on the backup server after you install the Hitachi software. To configure the ShadowImage, follow these steps:

1. Use the Hitachi DF500 Disk Array Management Program to create a command device.

2. Identify the current command device by issuing the following command from a command prompt:

   ```
   horcctl –x findcmddev drive#(0-n),
   ```

   where *n* is the number of logical disks visible to your system.

3. Modify the file C:\WINNT\system32\drivers\etc\services on the backup server.

   To modify the file, open the file with Windows Notepad and add the following line to the end of the file with 9900 as the port number:

   ```
   horcm            9900/udp                          #horcm
   ```

   Note: You can modify the port number value, but this value must not be used as a port number elsewhere in this file.

4. Modify the file C:\WINNT\horcm.conf.

   ```
   HORCM_MON
   #ip_address      service    poll(10ms)    timeout(10ms)
   111.222.33.12    horcm      12000         12000

   HORCM_CMD
   #dev_name                 dev_name                dev_name
   \\.\PhysicalDrive5
   HORCM_DEV
   #dev_group      dev_name            port#        TargetID      LU#       MU#
   bab            C:\MPoint1          CL1-A        1             21        0
   bab            C:\MPoint2          CL1-A        1             22        0
   bab            F:                  CL1-A        1             23        0
   bab            G:                  CL1-A        1             24        0
   #specifications for the peer host.
   HORCM_INST
   #dev_group          ip_address              service
   bab                 111.222.33.11           horcm
   ```

- Replace **111.222.33.11** with the IP address of the production server.

- Replace **111.222.33.12** with the IP address of the backup server.

- Replace **\\.\PhysicalDrive5** with the command device.

- Use **bab** as the name of a group of paired logical volumes.

- Use **C:\MPoint1, C:\MPoint2, F: and G:** as the names of the paired logical volume within a group.

Create Mirrored Volumes on the Hitachi Disk Array

You must create mirrored volumes on the Hitachi disk array after you install the Hitachi disk array. Use Hitachi HORCM commands to create mirrored volumes. Every HORCM instance must be up and running on both the production and backup servers for the HORCM commands to work; otherwise, you will receive an error message that you cannot access HORCM.

To start the HORCM instance, open a command prompt and enter the following HORCM command:

```
horcmstart
```

Note: Run this command on both production and backup servers.

After you have modified the configuration files as shown in the sections, Configure the ShadowImage on the Production Server and Configure the ShadowImage on the Backup Server in this chapter, you need to do the following:

- Open the command prompt on either the production or backup server, and issue the following Hitachi command to create the logical device group bab:

  ```
  paircreate –g bab –vl –c 15
  ```

- Split the volume and issue the command:

  ```
  pairsplit –g bab
  ```

- Check the status for the Hitachi disk array and issue the command:

  ```
  pairevtwait
  ```

Once you have completed the creating process, you are ready to install the software *components.*

Installation of the Option

The Enterprise Option for Microsoft SQL Using HDS Snap-Shot follows the standard installation procedure for the system components, agents, and options of BrightStor ARCserve Backup. For more information about this procedure, see the *Getting Started*.

After you complete the installation procedure, restart your computer when prompted.

Post-Installation Procedures

After installing the option, two default configuration scripts are created:

- HSSSQL_PHOST.ini—Located on the production server.
- HSSSQL.ini—Located on the backup server.

Modify these default configuration scripts to match your environment. The following sections explain this process.

Modification of the Production Server Script File

The script file HSSSQL_PHOST.ini contains all the information that the Enterprise Option for Microsoft SQL Using HDS Snap-Shot needs to configure and use Hitachi ShadowImage. This script file is located on the production server in the Agent for Microsoft SQL Server home directory.

You must update or modify the following fields and values in the script to match your production server environment, but you should not modify other fields:

- **PAIRSPLIT_PATH**—Specifies the path for Hitachi pairsplit.exe utility.
- **PAIRRESYNC_PATH**—Specifies the path for Hitachi pairresync.exe utility.
- **PAIREVTWAIT_PATH**—Specifies the path for Hitachi pairevtwait.exe utility.
- **HSSCHECKVOL_PATH**—Specifies the Agent for Microsoft SQL Server home directory.
- **bab**—Specifies the device group name. For more information about how to create a device group, see Configure the ShadowImage on the Production Server in this chapter.
- **C:\MPoint1, C:\MPoint2, F: and G:**—Specifies the mount point and volume names for the Hitachi disk array P-VOLs.

  To mount a volume or mount point, issue the following Hitachi command:

PAIRSPLIT -g bab -x mount C:\Mpoint1 \Vol4

For more information about the mount points and volumes created on the production server, see Configure the ShadowImage on the Production Server in this chapter.

- **Vol4, Vol5, Vol6 and Vol7**—Contains the Hitachi disk array logical volume names.

- **NUM_OF_CMDS**—Updates this parameter in each section to match the number of commands in that section.

To find the matching volume parameter, open a Windows command prompt and issue the Hitachi command:

inqraid -CLI $Vol –fv

Note: The device group name, mount point name and volume name are case-sensitive.

Modification of the Backup Server Script File

The script file HSSSQL.ini contains all the information that the Enterprise Option for Microsoft SQL Using HDS Snap-Shot needs to configure Hitachi ShadowImage. This script is located on the backup server in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory.

You must update or modify the following fields and values to match your backup server environment, but you should not modify other fields:

- **PAIRSPLIT_PATH**—Specifies the path for Hitachi pairsplit.exe utility.

- **PAIRRESYNC_PATH**—Specifies the path for Hitachi pairresync.exe utility.

- **PAIREVTWAIT_PATH**—Specifies the path for Hitachi pairevtwait.exe utility.

- **HSSCHECKVOL_PATH**—Specifies the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory.

- **RESTORE_OPTION**—Sets the restore option. If no value is specified, the restore job will fail.

- **Set RECOVERY_MODE to RECOVER**—Recovers the Microsoft SQL Server database after the restore.

- **Set RECOVERY_MODE to NO_RECOVER**—Does not recover the Microsoft SQL Server database after the restore.

- **bab**—Specifies the device group name. For more information about how to create the device group, see Configure the ShadowImage on the Production Server in this chapter.

- **C:\MPoint1, C:\MPoint2, F: and G:**—Specifies the mount point and volume names for the Hitachi disk array P-VOLs.

To mount a volume or mount point, issue the following Hitachi command:

PAIRSPLIT –g bab -x mount C:\MPoint1 \Vol4

For more information about the mount points created on the production server, see Configure the ShadowImage on the Backup Server in this chapter.

- **Vol4, Vol5, Vol6 and Vol7—**Contains the Hitachi disk array logical volume names.

- **NUM_OF_CMDS—**Updates these parameters in each section to match their environment.

  To find the matching volume parameter, open a command prompt and issue the Hitachi command:

  inqraid -CLI $Vol –fv

  Note: The device group name, mount point name, and volume name are all case-sensitive.

Chapter 3: Performing Backups

This chapter provides the guidelines and procedures you need to successfully back up your databases using the option.

Practices for Successful Backups

The following guidelines can help ensure a successful backup:

- Databases residing on local disks, including the Microsoft SQL Server 2000 databases, **cannot** be backed up using the option.

- To back up a Microsoft SQL Server 2000 database with the Agent for Microsoft SQL Server, the database must be created on the production server's P-VOLs.

- During the backup, the S-VOLs containing the Microsoft SQL Server 2000 database data and transaction log files should not be accessed by other applications and services on the backup server.

- If the system has antivirus software installed, configure the antivirus software so that it will not scan the S-VOLs during the backup.

- All Windows Explorer open folders should be closed with only the desktop and the BrightStor ARCserve Backup Manager interface active.

- All terminal services or remote desktop users should be logged off with the exception of the user performing the backup. If this user is logged on remotely, no other user should be logged on through the console.

- To back up multiple databases, each database must be created on its own set of P-VOLs; that is, each P-VOL should only contain files from one database.

- To achieve the best backup performance, you should create database files and transaction log files for a database on multiple volumes or mount points, and use a tape library with multiple drives. This arrangement ensures that the backup job creates multiple child jobs, and that the Microsoft SQL data files and transaction log files on each volume or mount point are backed up to separate tape drives simultaneously.

- The backup jobs must be submitted as multiple streaming jobs.

- BrightStor ARCserve Backup does not allow certain special characters to be used in file names, file group names, and database names. Examples include slashes, asterisks, greater than or less than signs, question marks, quotes, and commas.

- BrightStor ARCserve Backup does not support database names that contain parentheses or square brackets.

- If you have named a Microsoft SQL Server database with the characters that you are not allowed to use, by default, Microsoft SQL Server names the file, file group, or database with the same characters. To perform a backup job, you must first rename the file, file group, or database so it does not include one of these characters.

Backup Options

The option provides the following backup options:

- **Resynchronize after Backup**—After the back up of the metadata and database files to the storage media is complete, the option unmounts the S-VOLs from the backup server and resynchronizes the S-VOLs with the corresponding P-VOLs.

- **Leave the Mirror split after Backup**—After the back up of the metadata and database files to the media is complete, the option leaves the S-VOLs mounted on the backup server. The P-VOLs and S-VOLs remain split in the Hitachi disk array.

- **Get Mirror copy only (no data backup to media)**—By splitting P-VOLs and S-VOLs, a mirrored copy of the database is created on the backup server's S-VOLs. No data is backed up to the storage media.

Back Up Databases

Important! *Hardware snapshot backups are not supported on local disks of the production server. Backups for these types of databases must be packaged without Hardware Snapshot selected.*

To begin a basic backup, follow these steps:

1. Select the Backup Manager using the BrightStor ARCserve Backup Manager interface on the backup server.

2. From the Source tab, expand the tree until you reach Microsoft SQL Server on the production server.

3. Right-click a database in the Microsoft SQL Server folder to display a pop-up menu containing the following options:

   – Database Agent

   – Hardware Snapshot

   Note: If a parent object is selected, you can right-click a database agent to customize local agent options. For more information about explicit job packaging, see the *Getting Started.*

4. Select Hardware Snapshot. The Agent Backup Options dialog opens.
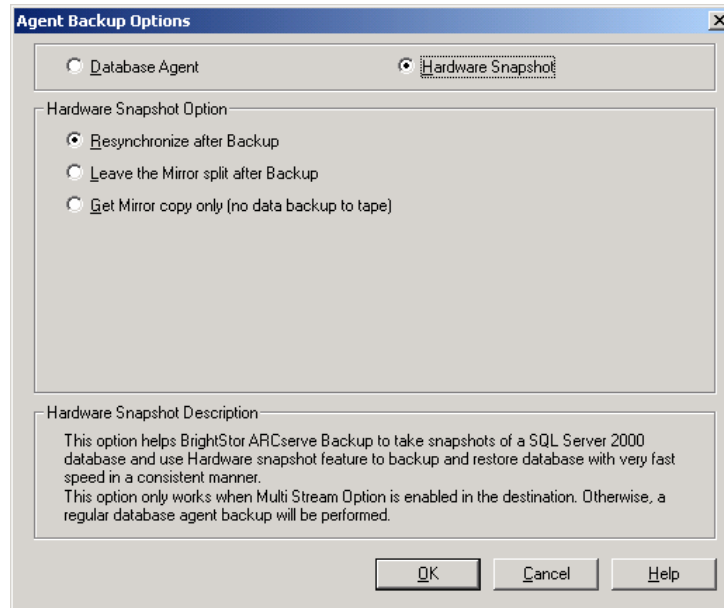
5. Select an option, and click OK.

   For more information about the three backup options, see Backup Options in this chapter.

6. Click the Destination tab of the Backup Manager.

7. Select the Multi Stream option. Hardware Snapshot backup jobs require the Multi Stream option. For more information about multiple streaming jobs, see the *Administrator Guide*.

8. Click Start on the Backup Manager to begin the backup job.

You can double-click the backup job in Job Status window to monitor the job.

Back Up a Database Instance

This section explains how to back up a Microsoft SQL Server 2000 instance. In the example used in this procedure, the Microsoft SQL Server instance contains nine databases. The databases created on the mirrored volumes are:

- HSSF on volume F
- HSSG on volume G
- HSSMP1 on mount point C:\MPoint1
- HSSMP2 on mount point C:\MPoint2

The other five databases reside on local disks of the production server and include the Microsoft SQL Server system databases and standard sample databases.

To back up the databases in an instance, follow these steps:

1. Select the Backup Manager using the BrightStor ARCserve Backup Manager interface on the backup server.

2. In the Source tab, expand the tree until you see Microsoft SQL Server on the production server called HITACHI-PRODUCT.
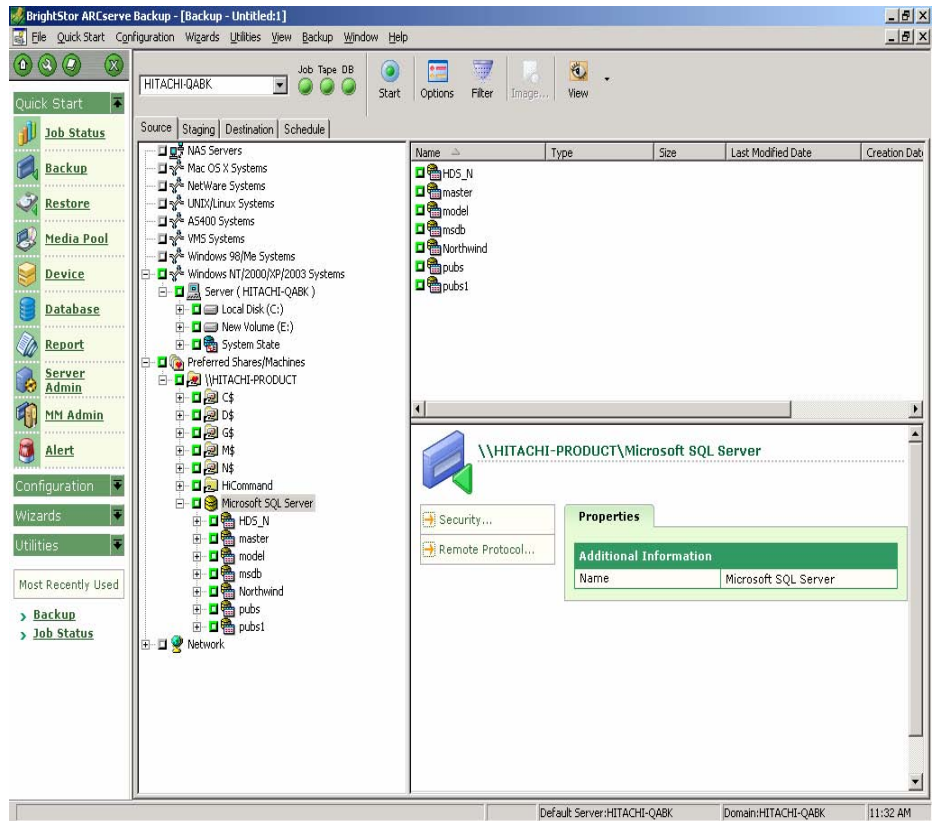
3. Select the database named HSSF that resides on a mirrored volume. Then right-click HSSF and select Backup Agent Options from the pop-up menu.

4. Select Hardware Snapshot on the Agent Backup Options dialog.

5. Repeat Steps 2 and 3 to set the Hardware Snapshot option for all the other databases residing on the mirrored volumes. These databases are HSSG, HSSMP1, and HSSMP2 in the example.

6. Verify that the Database Agent option is selected for the other databases residing on local volumes. These databases are the system databases master, model, msdb, and the standard sample databases Northwind and pubs in the example.

7. Click the Destination tab and select the Multi Stream option.
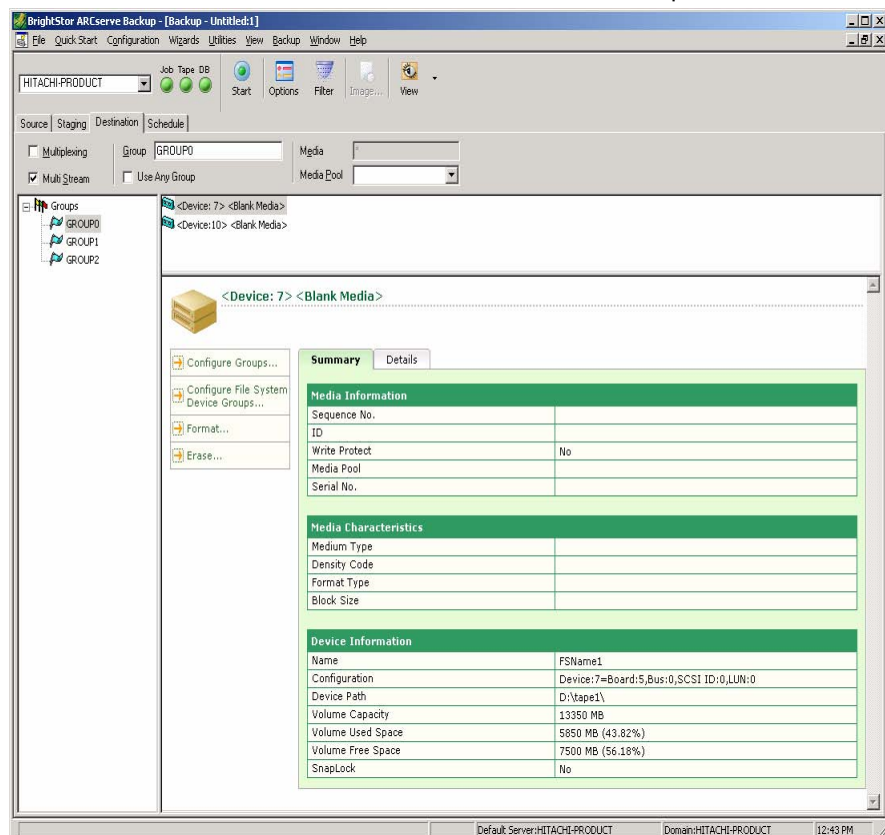
Note: Hardware Snapshot backup jobs require the Multi Stream option.

8. Click Start. A Security dialog opens. Enter the authentication information for the production server and Microsoft SQL Server database, and then click OK.

9. Schedule the job to either run at a specific time or run immediately. Click OK, and the backup job is added to the job queue.

The backup operation starts with the master job running a pre-job script on the backup server. The pre-job script synchronizes the S-VOLs with the corresponding P-VOLs, and then splits the volumes to mount the S-VOLs on the backup server.

The master job also creates a number of child jobs. The first child job backs up the metadata created by the Microsoft SQL snapshot operation. For each Microsoft SQL data file or transaction log file, one additional child backup job is created by the master job. Child jobs may run simultaneously depending on how many storage devices are available.

After the child jobs finish, the master job runs the post-job script. This will unmount the volumes and mount points, and resynchronize the S-VOLs with the corresponding P-VOLs depending on the options you have selected.

You can double-click a backup job in Job Status Manager to view details and monitor progress of that job. If you open the Job Status Manager, you will find:

- A master job to control the backup. This entry was created when you submitted the backup job. The master job finishes last and creates a management session on the storage device.

- A remote child backup job, which backs up the Microsoft SQL Server 2000 databases residing on the local volumes of the production server.

- Multiple snapshot child jobs to back up the database data files and transaction log files from the local S-VOLs of the backup server. The number of jobs equals the number of S-VOLs.

- Multiple remote child backup jobs, which back up the Microsoft SQL Server database snapshot metadata for databases residing on the mirrored volume. In this example, four such child jobs are created.

After the backup jobs are complete, check the results using the BrightStor ARCserve Backup Manager interface.

Chapter 4: Restoring Data Using the Option

You can use the Enterprise Option for Microsoft SQL Using HDS Snap-Shot to restore the data previously backed up by BrightStor ARCserve Backup. The option enables the restore to be performed from the local tape library and resynchronized to the production server. The restore capabilities allow you to set policies, schedule and manage your restore operations, and perform a complete or partial restore of data.

Restore Operation Guidelines

The following guidelines can help ensure a successful restore:

- During a restore, the P-VOL and S-VOL that contain the Microsoft SQL database data and transaction log files should not be accessed by other applications and services on the production and backup servers.

- If either system has antivirus software installed, configure the antivirus software so that it will not scan the P-VOLs and S-VOLs during the restore.

- All Windows Explorer open folders should be closed with only the desktop and the BrightStor ARCserve Backup Manager interface active.

- All terminal services or remote desktop users should be logged off with the exception of the user performing the restore. If the user is logged on remotely, no other users should be logged on through the console.

- BrightStor ARCserve Backup does not allow certain special characters to be used in file names, file group names, and database names. Examples include slashes, asterisks, greater than or less than signs, question marks, quotes, and commas.

- BrightStor ARCserve Backup does not support database names that contain parentheses or square brackets.

- If you have named a Microsoft SQL Server database with the characters that you are not allowed to use, by default, Microsoft SQL Server names the file, file group, or database with the same characters. To perform a restore job, you must first rename the file, file group, or database so it does not include one of these characters.

Restore Methods, Options, and Modes

This section describes the methods used to restore a Microsoft SQL Server 2000 database, the options to perform the restore operation, and the modes to recover the Microsoft SQL Server database.

Restore a Database with the Restore Methods

Two restore methods are used to restore a Microsoft SQL Server 2000 database that was backed up with the option:

- **Restore by Tree:** Restores a specific Microsoft SQL Server database.
- **Restore by Session:** Restores a specific management session.

Follow these steps to restore a specific Microsoft SQL Server database:

1. Select Restore Manager on the backup server.
2. Select Restore by Tree from the drop-down list in the Source tab.
3. Locate Microsoft SQL Server 2000 and select the database to restore.

Follow these steps to restore a specific management session:

1. Select Restore Manager on the backup server.
2. Select Restore by Session from the drop-down list in the Source tab.
3. Locate the management session for the Microsoft SQL database to restore.

After you select the restore methods and the database or session to restore, submit the job.

To submit a restore job, follow these steps:

1. Click Start.

2. The Session User Name and Password dialog opens. Double-click the session to open the Enter User Name and Password dialog.

3. Enter the login information of a Windows administrator account on the production server, and click OK to return to the Session User Name and Password dialog.

4. Select the DBAgent folder tab, and double-click the session.

5. The Enter User Name and Password dialog opens. Enter the same login information or the login information of a Microsoft SQL Server administrator, as appropriate.

6. Click OK to close the Enter User Name and Password dialog and return to the Session User Name and Password dialog. Click OK to submit the restore job.

The initial restore job runs pre-job and post-job scripts and creates child jobs to restore all related Microsoft SQL data files to each volume or mount point simultaneously.

Double-click a job to view the job status.

Restore Options and Modes

The restore options and restore modes allow you to manage your restore or recovery operations according to your specific restore or recovery need. You cannot select restore options or modes using the BrightStor ARCserve Backup Manager interface. You must set the restore options or modes in the backup server script file, HSSSQL.ini. For more information about the backup server script file, see Modification of the Backup Server Script File in the chapter "Installing the Option."

The following are supported restore options:

- **RESTORE_RESYNC**—Performs a complete restore operation. It first restores the database files to the backup server S-VOL, unmounts the S-VOLs from the backup server and P-VOLs from the production server, and then reverse resynchronizes the P-VOLs with the S-VOLs. Finally the P-VOLs are mounted on the production server.

- **RESTORE_ONLY**—Restores database files only to the backup server S-VOLs.

- **RESYNC_ONLY**—Unmounts the S-VOLs from the backup server, reverse-resynchronizes the P-VOLs with the S-VOLs, and then mounts the P-VOLs on the production server. No data is read from backup media.

The following are supported restore modes:

- **RECOVERY**—The Microsoft SQL database will be recovered after restore.

- **NO_RECOVERY**—The Microsoft SQL database will not be recovered after restore.

Appendix A: Troubleshooting

This appendix provides troubleshooting information to help you identify and resolve problems that you may encounter when using BrightStor ARCserve Backup and the Enterprise Option for Microsoft SQL Using HDS Snap-Shot. To help you quickly find the information you need, this appendix lists some common error messages and possible reasons and solutions for these messages.

Common Error Messages

This section explains the most common error messages for BrightStor ARCserve Backup and the Enterprise Option for Microsoft SQL Using HDS Snap-Shot.

Activity Log

Many of the actions to resolve error conditions advise you to check the BrightStor ARCserve Backup Activity log. The Activity log contains comprehensive information about the operations performed by BrightStor ARCserve Backup. It provides an audit trail of all BrightStor ARCserve Backup activity for every job that is run. You can scan this log whenever necessary to see if any errors have occurred. The log is available from the Job Status Manager. For more information about using the Activity log, see the *Administrator Guide.*

E29041

Hardware Snapshot error: Unable to retrieve the file list for database [*DBNAME*].

Reason:

An error has occurred while trying to retrieve the list of database files for backup.

Action:

When you encounter this error message, follow these steps:

1. Check network connections.

2. For Microsoft SQL Server, verify in the SQL Server Enterprise Manager that the database [DBNAME] is accessible, and not in a Suspect or Loading state.

3. If all network connections are functional and the database is accessible, stop and restart the agent services on the production server. These services are the CA BrightStor Backup Agent RPC server for all database agents, and the CA BrightStor Backup Agent Remote Service for the Agent for Microsoft SQL Server.

E29045

Hardware Snapshot error: Failed to login to BrightStor ARCserve Backup database. Please check the security information.

Reason:

The BrightStor ARCserve Backup database may not be running or the security information may be incorrect.

Action:

When you encounter this error message, follow these steps:

1. Check the authentication under which all the BrightStor ARCserve Backup services are set to run on the backup server.

2. Verify that all the BrightStor ARCserve Backup services are running on the backup server.

3. If using Microsoft SQL Server for the BrightStor ARCserve Backup database, run the Backup Server Admin utility, select MS SQL User Account from the Admin menu, and update the information in the dialog.

4. Verify the connection to the BrightStor ARCserve Backup database. If you are using Microsoft SQL Server as a BrightStor ARCserve Backup database, verify settings for the system DSN ASNT in the Windows ODBC Data Source Administrator.

E29048

Hardware Snapshot error: Failed to retrieve snapshot file name(s) for database [*DBNAME*].

Reason:

Communication failure between the BrightStor ARCserve Backup server and the Agent for Microsoft SQL Server, or a Microsoft SQL Server query failure.

Action:

When you encounter this error message, follow these steps:

1. Check network connections.

2. Verify in the SQL Server Enterprise Manager that the database [DBNAME] is accessible and not in a Suspect or Loading state.

3. If all network connections are functional and the database is accessible, stop and restart the agent services on the production server. These services are the CA BrightStor Backup Agent RPC server for all database agents, and the CA BrightStor Backup Agent Remote Service for the Agent for Microsoft SQL Server.

E29049

Hardware Snapshot error: Failed to backup the management session information for database [*DBNAME*].

Reason:

Indicates a tape drive error or BrightStor ARCserve Backup database error.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for error messages retrieved from the Tape Engine.

2. Check the Activity log for error messages retrieved from the Database Engine.

E29058

Hardware Snapshot error: Failed to begin restore of the management session for database [*DBNAME*].

Reason:

Indicates a tape drive error or BrightStor ARCserve Backup database error.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for error messages retrieved from the Tape Engine.
2. Check the Activity log for error messages retrieved from the Database Engine.

E29059

Hardware Snapshot error: Failed to restore extended session information for database *[DBNAME]*.

Reason:

Indicates a tape drive error or BrightStor ARCserve Backup database error.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for error messages retrieved from the Tape Engine.
2. Check the Activity log for error messages retrieved from the Database Engine.

E29060

Hardware Snapshot error: Failed to restore a data stream for the management session for database [*DBNAME*].

Reason:

Indicates a tape drive error or BrightStor ARCserve Backup database error.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for error messages retrieved from the Tape Engine.
2. Check the Activity log for error messages retrieved from the Database Engine.

E29061

Hardware Snapshot error: One of the stream header signatures is invalid for the management session of database [*DBNAME*].

Reason:

Indicates a tape drive error or BrightStor ARCserve Backup database error.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for error messages retrieved from the Tape Engine.
2. Check the Activity log for error messages retrieved from the Database Engine.

E29062

Hardware Snapshot error: Failed to restore the management session for database *[DBNAME]*.

Reason:

Indicates a tape drive error or BrightStor ARCserve Backup database error.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for error messages retrieved from the Tape Engine.
2. Check the Activity log for error messages retrieved from the Database Engine.

E29063

Hardware Snapshot error: Failed to restore the snapshot file for database [*DBNAME*].

Reason:

Indicates a network connection error, tape drive error, or incorrect username or password for production server login.

Action:

When you encounter this error message, follow these steps:

1. Check the Activity log for any error messages related to the Tape Engine.
2. Check network connections.
3. Verify authentication information for the production server.

E29064

Hardware Snapshot error: Failed to connect to the Hardware Snapshot SQL Agent for database [*DBNAME*].

Reason:

The agent services may be down, the security information provided is not correct, or a network failure may have occurred.

Action:

When you encounter this error message, follow these steps:

1. Check network connections.

2. Verify the Windows and database authentication information for the production server.

3. If all network connections are functional and authentication is correct, stop and restart the agent services on the production server. These services are the CA BrightStor Backup Agent RPC server for all database agents and the CA BrightStor Backup Agent Remote Service for the Agent for Microsoft SQL Server.

E29065

Hardware Snapshot error: Failed to begin restore for Hardware Snapshot SQL Agent for database [*DBNAME*].

Reason:

Indicates an agent error, database error, or network communication failure.

Action:

When you encounter this error message, follow these steps:

1. Check network connections.

2. Check the agent log and database error log for error messages, and if found, respond accordingly.

3. If network connections are functional and no agent or database errors are found, stop and restart the CA BrightStor Backup Agent RPC server service and CA BrightStor Backup Agent Remote Service on the production server.

E29072

Hardware Snapshot error: Failed to communicate with Hardware Snapshot SQL Agent for database [*DBNAME*].

Reason:

Indicates an agent error, database error, or network communication failure.

Action:

When you encounter this error message, follow these steps:

1. Check network connections.

2. Check the agent log and database error log for error messages, and if found, respond accordingly.

3. If network connections are functional and no agent or database errors are found, stop and restart the CA BrightStor Backup Agent RPC server service and CA BrightStor Backup Agent Remote Service on the production server.

E29073

Hardware Snapshot error: Failed to disconnect from the Hardware Snapshot SQL agent for database [*DBNAME*].

Reason:

Indicates an agent error, database error, or network communication failure.

Action:

When you encounter this error message, follow these steps:

1. Check network connections.

2. Check the agent log and database error log for error messages, and if found, respond accordingly.

3. If network connections are functional and no agent or database errors are found, stop and restart the CA BrightStor Backup Agent RPC server service and CA BrightStor Backup Agent Remote Service on the production server.

E29076

Hardware Snapshot error: No restore option specified for database [*DBNAME*] in the Backup Server snapshot script file.

Reason:

The backup server script file may have been modified or has incorrect information and the restore option setting is missing or invalid.

Action:

Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server. The restore option should be one of the following: RESTORE_RESYNC (the default), RESYNC_ONLY, or RESTORE_ONLY.

E29077

Hardware Snapshot error: Pre-job process failed on Backup Server: command="command".

Reason:

This error message is caused by the following reasons:

- When you encounter **command="mount"**, it indicates that the Hardware Snapshot mount script command failed on the backup server.

- When you encounter **command="umount"**, it indicates that the Hardware Snapshot umount script command failed on the backup server.

- When you encounter **command="split"**, it indicates that the Hardware Snapshot split script command failed on the backup server.

- When you encounter **command="resync"**, it indicates that the Hardware Snapshot resync script command failed on the backup server.

Action:

To correct the errors, follow this instruction:

- When you encounter **command="mount"**, follow these steps:

1. Check the script HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Check that the drive letter or mount point is not used by another volume.

3. Check that the volume is not mounted as another drive letter or mount point.

4. Verify the mount command using the disk array control tools.

- When you encounter **command="umount"**, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Identify and close all other applications that may be accessing this volume or mount point, including Windows Explorer and open folder windows, antivirus real-time scanners, and command prompts. Verify that all desktop sessions except the session from which the backup is being managed are logged out, and that the drive or mount point being unmounted is not being accessed remotely through file sharing.

3. Verify the unmount command using the disk array control tools.

- When you encounter **command="split"**, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Check that the pairs are properly created using the disk array control tools.

3. Verify the split command using the disk array control tools.

- When you encounter **command="resync"**, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Check that the pairs are properly created using the disk array control tools.

3. Verify the resync command using the disk array control tools.

E29079

Hardware Snapshot error: No recovery option specified for database [*DBNAME*] in the Backup Server snapshot script file.

Reason:

The backup server script file may have been modified or has incorrect information, and the recovery option setting is missing or invalid.

Action:

Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server. The recovery option should be one of the following: RECOVERY (the default) or NO_RECOVERY.

E29080

Hardware Snapshot error: Pre-job process failed on Production Server, command="command" for database [*DBNAME*].

Reason:

When you encounter **Command="mount"**, it indicates that the Hardware Snapshot mount script command failed on the production server.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL_PHOST.ini in the Agent for Microsoft SQL Server home directory on the production server.

2. Check that the drive letter or mount point is not used by another volume.

3. Check that the volume is not mounted as another drive letter or mount point.

4. Verify the mount command using the disk array control tools.

Reason:

When you encounter **Command="umount"**, it indicates that the Hardware Snapshot umount script command failed on the production server.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL_PHOST.ini in the Agent for Microsoft SQL Server home directory on the production server.

2. Identify and close all other applications that may be accessing this volume or mount point, including Windows Explorer and open folder windows, antivirus real-time scanners, and command prompts. Verify that all desktop sessions except the session from which the backup is being managed are logged out, and that the drive or mount point being unmounted is not being accessed remotely through file sharing.

3. Verify the umount command using the disk array control tools.

E29081

Hardware Snapshot error: Post-job process failed on Backup Server: command="command".

Reason:

When you encounter **Command="umount"**, it indicates that the Hardware Snapshot umount script command failed on the backup server.

When you encounter **Command="resync"**, it indicates that the Hardware Snapshot resync script command failed on the backup server.

Action:

To correct the errors, follow this instruction:

- When you encounter **Command="umount"**, follow these steps:

1. Check the script in the file "HSSSQL.ini" in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Identify and close all other applications that may be accessing this volume or mount point, including Windows Explorer and open folder windows, antivirus real-time scanners, and command prompts. Verify that all desktop sessions except the session from which the backup is being managed are logged out, and that the drive or mount point being unmounted is not being accessed remotely through file sharing.

3. Verify the umount command using the disk array control tools.

- When you encounter **Command="resync"**, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Check that the pairs are properly created using the disk array control tools.

3. Verify the resync command using the disk array control tools.

E29440

Hardware Snapshot error: Post-job process failed on Production Server: command="command".

Reason:

When you encounter **Command="mount"**, it indicates that the Hardware Snapshot mount script command failed on the production server.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL_PHOST.ini in the Agent for Microsoft SQL Server home directory on the production server.
2. Check that the drive letter or mount point is not used by another volume.
3. Check that the volume is not mounted as another drive letter or mount point.
4. Verify the mount command using the disk array control tools.

E29441

Hardware Snapshot error: Failed to read the Hardware Snapshot script file on Backup Server.

Reason:

The file cannot be read from the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

Action:

Make sure that the file HSSSQL.ini exists on the backup server in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory and is readable.

E29443

Hardware Snapshot error: Failed to load 'dbares.dll'.

Reason:

BrightStor ARCserve Backup server installation is damaged.

Action:

Verify that the BrightStor ARCserve Backup server software is properly installed and verify that the file dbares.dll exists. Uninstall and reinstall BrightStor ARCserve Backup server.

E29444

Mount failed.

Reason:

The script mount command may not be correct, the volume may be in use by other applications, or the volume may be already mounted to a different volume.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Check that the drive letter or mount point is not used by other disks.

3. Check that the disk is not mounted to another drive letter or mount point.

4. Verify the mount command using the disk array control tools.

E29445

Dismount failed.

Reason:

The script umount command may not be correct, or the volume may be in use by other applications.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Identify and close all other applications that may be accessing this volume or mount point, including Windows Explorer and open folders, antivirus real-time scanners, and command prompts. Verify that all desktop sessions except the session from which the backup is being managed are logged out, and that the drive or mount point being unmounted is not being accessed remotely through file sharing.

3. Verify the umount operation by the disk array control tools.

E29446

Split failed.

Reason:

The split command from the script is not correct or the pairs are not created by disk array control tools.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.

2. Check that the pairs are created by the disk array control tools.

3. Verify the split operation by the disk array control tools.

E29447

Resynchronization failed.

Reason:

The resync command from the script is not correct or the pairs are not created by disk array control tools.

Action:

When you encounter this error message, follow these steps:

1. Check the script file HSSSQL.ini in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory on the backup server.
2. Check that the pairs are created by the disk array control tools.
3. Verify the mount operation by the disk array control tools.

E29448

Pre-script failed.

Refer to the description of error E29077 for a script failure on the backup server, or to error E29080 for a script failure on the production server.

E29449

Post-script failed.

Refer to the description of error E29081 for a script failure on the backup server, or to error E29440 for a script failure on the production server.

E29456

Hardware Snapshot error: The hardware snapshot process is not started on Backup Server.

Reason:

The HORCM instance for the disk array management utilities is not started.

Action:

To start the HORCM instance, go to HORCM home directory, open a command prompt and type the following command:

```
horcmstart
```

For more information about HORCM, see Hitachi documentation packaged with this release.

Microsoft SQL Server Error Messages

Microsoft SQL Server errors can occur for a variety of reasons, and can appear within several different BrightStor ARCserve Backup error messages. For more information, see the *Agent for Microsoft SQL Server Guide*, and online Microsoft SQL Server documentation.

Appendix B: Disaster Recovery

Disaster recovery is a backup and recovery process used to protect computing environments against the loss of data. This chapter discusses disaster recovery backup strategy and provides sample scenarios with procedures to help you understand how you should perform a disaster recovery.

Backup Strategy for Disaster Recovery

Administrators should always maintain backups of the following files to expedite disaster recovery:

- **HORCM.conf**—This is the configuration file for the disk array that contains disk mirroring information. It is located in the Windows home directory (usually C:\WINNT) on both production and backup servers.

- **HSSSQL.ini**—This script file contains commands to perform the Hardware Snapshot operations for backup and restore. It is located on the backup server in the Enterprise Option for Microsoft SQL Using HDS Snap-Shot home directory.

- **HSSSQL_PHOST.ini**—This script file contains commands to perform the Hardware Snapshot operation for restore function. It is located on the production server in the Agent for Microsoft SQL Server home directory.

Disaster Recovery Procedures

This section provides the procedures of recovering from a disaster on various scenarios.

Respond to a Failed Disk Array

When the hardware Hitachi disk array fails, follow these steps:

1. Replace the failed disk array.

2. Configure your replacement disk array. For more information about creating mirrored volumes, see Create Mirrored Volumes on the Hitachi Disk Array in the chapter "Installing the Option". Ensure that you recreate all of the volumes to at least their original size.

3. Ensure that the new configuration matches the original configuration, or carefully modify the files HORCM.conf, HSSSQL.ini, and HSSSQL_PHOST.ini to incorporate the differences. On both the production and backup servers, make sure the group name, dev_name, prot#, TargetID LU# MU# are correct in HORCM.conf. For HSSSQL.ini and HSSSQL_PHOST.ini, make sure the drive letters, mount point paths, and commands are correct.

4. Start Microsoft SQL Server on the production server.

5. Perform a restore job to its original location on the backup server for databases that reside on the Hitachi disk array.

Restore a Microsoft SQL Server

When the Microsoft SQL Server is down, follow these steps:

1. If an offline backup (a backup for the Microsoft SQL directory) exists, restore it. Otherwise, reinstall Microsoft SQL Server on the production server.

2. Reinstall the Agent for Microsoft SQL Server on the production server.

3. Start Microsoft SQL Server in Single User mode on the production server and restore the master database on the backup server. After the restoration is complete, Microsoft SQL Server will automatically stop on the production server.

4. Restart Microsoft SQL Server on the production server as you would normally do.

5. From the backup server, restore the system database msdb, then system database model, and finally all of the user databases (including databases that reside on the Hitachi disk array) and transaction logs to their original location. When this step is finished, recovery is complete.

Restore the Backup and Production Servers

When the backup and production servers are down, follow these steps:

1. Replace two server machines.

2. Reinstall the operating system, if necessary, on both machines.

3. Configure the ShadowImage on your servers. For more information about configuring the ShadowImage, see Configure the ShadowImage on the Production Server and Configure the ShadowImage on the Backup Server in the chapter "Installing the Option". Make sure you recreate all of the volumes to at least their original size.

4. Install BrightStor ARCserve Backup on the backup server and reconfigure it as needed.

5. Install the Enterprise Option for Microsoft SQL Using HDS Snap-Shot on the backup server.

6. If an offline backup exists in the Microsoft SQL Server directories and system databases, restore it. Otherwise, reinstall the Microsoft SQL Server on the production server.

7. Reinstall the Agent for Microsoft SQL Server on the production server.

8. Perform a merge on the backup server of the media that contains the restore sessions for the Microsoft SQL Server databases.

9. Make sure the new configuration matches the original configuration, or carefully modify the files HORCM.conf, HSSSQL.ini, and HSSSQL_PHOST.ini to reflect the differences. On both the production and backup servers, make sure the group name, dev_name, prot#, TargetID LU# MU# are correct in HORCM.conf. For HSSSQL.ini and HSSSQL_PHOST.ini, make sure the drive letter, mount point path, and command are correct.

10. Start Microsoft SQL Server in Single User mode on the production server and restore the master database on the backup server. After the restoration is complete, the Microsoft SQL Server automatically stops on the production server.

11. Restart Microsoft SQL Server for normal operation on the production server.

12. From the backup server, restore the system database msdb, then system database model, and finally all of the user databases (including databases that reside on the Hitachi disk array) and transaction logs to their original location. When this step is finished, recovery is complete.

For more information about Microsoft SQL Server, see the *Agent for Microsoft SQL Server Guide* and Microsoft SQL Server documentation.

Appendix C: Enabling the High Performance Feature

The Enterprise Option for Microsoft SQL Using HDS Snap-Shot provides the high performance feature. The following sections present the procedures of how to activate the high performance capability.

You should see an increase in performance after activating the high performance feature. If you do not see an increase in performance, it may indicate one or both of the following:

- Low memory
- High disk or CPU activity

The use of other applications may also cancel the performance increase provided by this tuning.

Note: You need at least 512 MB of free physical memory to use the high performance feature.

Identify the SCSI Card Driver Name

To identify your SCSI card driver name, follow these steps:

1. In Windows Explorer, right-click My Computer and select Manage to open the Device Manager.

2. In the Computer Management window tree in the left pane of the Computer Management window, find and select Device Manager under the System Tools category.

3. In the Device Manager tree in the right pane, locate and expand the entry for SCSI controllers.

4. Right-click the SCSI controller to which your tape drive is attached. Select Properties from the pop-up menu. The Device Properties dialog opens:

5. Click the Driver tab, and then click Driver Details. The Driver File Details dialog opens.

6. In the Driver files list, the drive name for your SCSI card is the portion of the driver file name between the directory path and the .sys file type extension. In the following example, the driver name is AIC78XX:
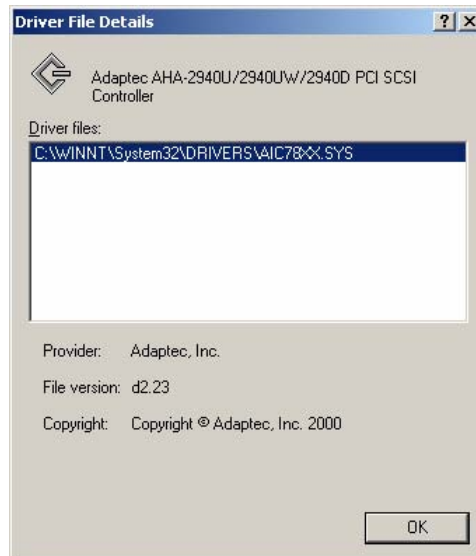
Tune the SCSI Card

You must tune the SCSI card for high throughput. Skip this section and go directly to Tune the Option if you are using a fibre channel controller instead of a SCSI card. The default settings for the SCSI card have a MaxTransferLength value of only 64 KB. You must increase the value to attain higher throughput.

To increase this value, follow these steps:

1. Run the Registry Editor (regedit or regedt32).

2. Under HKEY_LOCAL_MACHINE, follow these steps:

   - Expand the System branch

   - Expand the CurrentControlSet branch

   - Expand the Services branch

3. Under Services, find the key that matches the driver name for your SCSI adapter and expand it.

4. Under the key that matches the driver name for your SCSI adapter, find the key named Parameters. If it does not exist, create it.

5. Expand the Parameters key, and locate the key named Device. If it does not exist, create it.

6. Select the Device key. Add a new DWORD value to this key named MAXIMUMSGLIST.

7. Double-click the MAXIMUMSGLIST value you just created. In the dialog, enter the hexadecimal value FE (decimal value 254) as the data value in the edit box. Click OK to close the dialog.

8. Restart the operating system.

Tune the Option

You must tune the Enterprise Option for Microsoft SQL Using HDS Snap-Shot for high performance.

Important! *You must tune your SCSI card first if you are using a SCSI controller. If you do not, tuning the option will cause SCSI port errors.*

To tune the option, follow these steps:

1. Run the Registry Editor.

2. Under HKEY_LOCAL_MACHINE, expand the Software branch. Expand the Computer Associates branch. Expand the BrightStor ARCserve Backup branch. Expand Base, and then Tape Engine. Find CONFIG.

3. Select the CONFIG key. Add a new DWORD Value UseTaskBufferSize to the CONFIG key, and set its value to 1.

4. Close the Registry Editor and open the Server Admin.

5. From the Admin menu, select Configuration. The Configuration dialog opens:
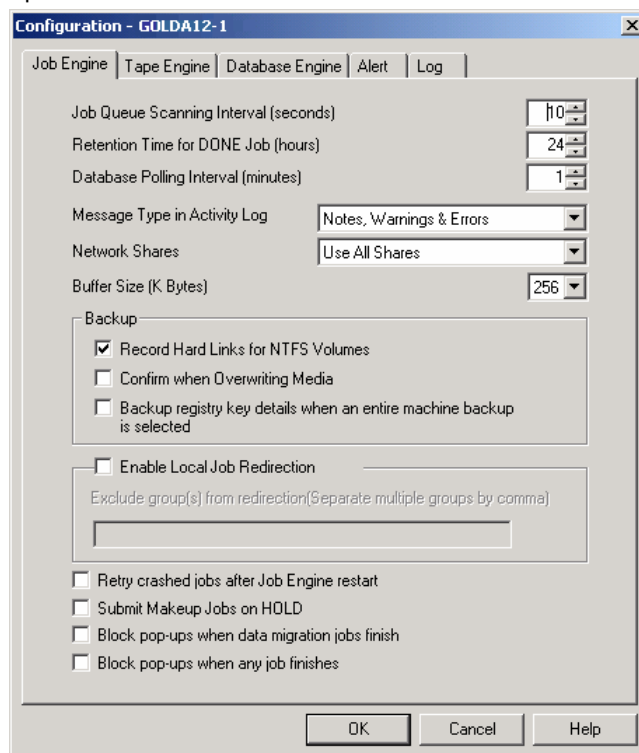
6. In the Job Engine tab, find the Buffer Size option drop-down list and select an appropriate buffer size. The default buffer size is 256 KB. Depending on system resources and hardware capabilities, a different buffer size such as 512 KB may get better throughput. If increasing this setting causes jobs to produce SCSI port errors, reduce this setting until the errors no longer occur.

Index

production server • 6
P-VOL • 6

R

Registry Editor • 57
restore methods • 28
    Restore by Session • 28
    Restore by Tree • 28
restore modes • 29
    NO_RECOVERY • 29
    RECOVERY • 29
restore options • 29
    RESTORE_ONLY • 29
    RESTORE_RESYNC • 29
    RESYNC_ONLY • 29
resynchronization • 7
reverse resynchronization • 7

S

script files • 17
    HSSSQL.ini • 19, 51
    HSSSQL_PHOST.ini • 18, 51
SCSI controller • 55
ShadowImage • 13
Source tab • 22, 24
splitting • 7
submit a restore job • 28
S-VOL • 5
synchronization • 7
system requirements • 9

T

tape library • 6
transaction log file • 7, 21, 24

U

updating the volume name • 10

V

volume name • 18

W

workflow • 7