

HP StorageWorks DButil user guide

Part number: T3680-96031
Second edition: January 2007



Legal and notice information

© Copyright 2005-2007 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Contents

Preface	7
About this guide	7
Intended audience	7
Prerequisites	7
Related documentation	7
Document conventions and symbols	8
HP technical support	8
Subscription service	9
HP websites	9
Documentation feedback	9
Product feedback	9
1 Using HP StorageWorks DButil	11
Supported software	11
Installing the utility	11
Running the utility	11
HP DButil commands	11
DButil -pre	11
DButil -post	12
DButil -mfreeze	12
DButil -mthaw	12
DButil -recover	12
DButil -restoresns	12
DButil -setcred	12
Optional parameters and other considerations	13
Named instances	13
Database names that contain spaces	13
Metatdata files	13
2 Creating snapshot backups using HP DButil	15
Procedure	15
Determining which databases require backup	15
Setting credentials	15
Creating the HP DButil batch files	16
Single SQL Server database	16
Three SQL Server databases on a common set of volumes	16
Determining which volumes need to be replicated	17
Single SQL Server database	17
Three SQL Server databases on a common set of volumes	17
Constructing HP Business Copy EVA/MA/EMA or HP Replication Solutions Manager jobs	17
Sample HP Business Copy EVA/MA/EMA job	18
Sample HP Replication Solutions Manager job	18
Testing the process	19
Backing up transaction logs	19
3 SQL Server recovery using HP DButil	21
Recovery approaches	21
Volume replacement	21
Selective file replacement	21

Applying transaction logs to a restored database	22
Example recovery scenarios	22
Recovery scenario 1	22
Recovery scenario 2	23
Recovery scenario 3	24
Recovery scenario 4	24
4 Troubleshooting HP DButil	27
Index	29

Tables

1	Document conventions	8
2	Troubleshooting HP DButil error messages	28

Preface

About this guide

This guide describes procedures for installing, configuring, and managing HP StorageWorks DButil. Topics include:

- Installing HP DButil
- Running HP DButil
- Creating snapshots using HP DButil
- Recovering SQL Server using HP DButil
- Troubleshooting HP DButil

Intended audience

This guide is intended for network and storage administrators and HP-authorized service providers with the knowledge of:

- Storage area networks
- SAN fabrics
- HP StorageWorks Enterprise Virtual Array
- Operating systems in your EVA and EVA management configuration
- HP StorageWorks Replication Solutions Manager or HP StorageWorks Business Copy EVA/MA/EMA

Prerequisites

Prerequisites for using this product include:

- Supported version of Microsoft SQL Server installed on a Microsoft Windows 2000 or 2003 host
- Supported versions of HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager (or HP StorageWorks Business Copy EVA/MA/EMA) installed on a management server
- HP Replication Solutions Manager host agent software installed on the host where SQL Server is installed if you will use HP Replication Solutions Manager for backups (optional)

For supported software versions and Microsoft operating systems, see *HP StorageWorks EVA software compatibility reference*.

Related documentation

The following documents provide related information:

- *HP StorageWorks Business Copy EVA administrator guide*
- *HP StorageWorks EVA software compatibility reference*
- *HP StorageWorks Replication Solutions Manager online help and user guide*
- *HP StorageWorks EVA replication software consolidated release notes*

You can find these documents from the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text: Table 1	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	website addresses
Bold text	<ul style="list-style-type: none">• Keys that are pressed• Text typed into a GUI element, such as a box• GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none">• File and directory names• System output• Code• Commands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none">• Code variables• Command variables
Monospace, bold text	Emphasized monospace text

 **CAUTION:**

Indicates that failure to follow directions could result in damage to equipment or data.

 **IMPORTANT:**

Provides clarifying information or specific instructions.

 **NOTE:**

Provides additional information.

 **TIP:**

Provides helpful hints and shortcuts.

HP technical support

Telephone numbers for worldwide technical support are listed on the HP support website: <http://www.hp.com/support/>.

Collect the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Error messages
- Operating system type and revision level
- Detailed questions

For continuous quality improvement, calls may be recorded or monitored.

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/e-updates>

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

HP websites

For additional information, see the following HP websites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- http://www.hp.com/service_locator
- <http://www.docs.hp.com>

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to storagedocs.feedback@hp.com. All submissions become the property of HP.

Product feedback

To make comments and suggestions about HP Business Copy EVA or HP Replication Solutions Manager, please mail your comments and suggestions to EVAREplication@hp.com.

1 Using HP StorageWorks DButil

HP StorageWorks DButil is a command-line utility to facilitate application-consistent replication of a Microsoft SQL Server database. HP DButil commands place the database in a transactionally consistent state and suspend write operations while HP StorageWorks replication software is used to copy the database. Other HP DButil commands then resume the database and normal write operations.

Users are not logged out of the database while it is suspended and the database remains online.

Supported software

HP DButil is supported with the following software:

- Microsoft SQL Server
- HP StorageWorks Business Copy EVA/MA/EMA
- HP StorageWorks Replication Solutions Manager

For supported software versions and Windows operating systems, see the *HP StorageWorks EVA software compatibility reference*.

Installing the utility

1. Locate the file `dbutil.exe`. It is available on the HP Business Copy EVA software and drivers website at <http://www.hp.com/support/downloads>. The file is also installed with the HP Replication Solutions Manager Windows Host Agent.
2. Determine the storage hosts or host clusters where Microsoft SQL Server databases are to be replicated.
3. On each identified host, copy the file `dbutil.exe` to a directory of your choice. HP recommends that you copy HP DButil to the directory of the SQL Server database.

Running the utility

HP DButil is run from a Windows command line, typically using the following methods:

- Windows batch files
- HP Business Copy EVA/MA/EMA jobs that call batch files or scripts
- HP Replication Solutions Manager jobs that call batch files or scripts

HP DButil creates a metadata file in the directory where it is run. HP DButil uses this metadata to restore the database. A metadata file is valid only for the snapshot that it is created for. If you run HP DButil from the directory of the SQL Server database, you can keep separate metadata files for restoring multiple databases.

HP DButil commands

HP DButil supports the following commands.

DButil `-pre`

The `DButil -pre` command suspends an SQL Server database and places it in a transactionally consistent state. Write operations to the database are suspended until the `DButil -post` command is issued. Use this command just before you replicate the database.

Syntax:

```
DButil -pre <sqlserver_name> <database_name> <username> <password>
```

DButil -post

The `DButil -post` command resumes a suspended SQL Server database and allows write operations to continue. Use this command just after you replicate the database.

Syntax:

```
DButil -post <sqlserver_name> <database_name> <username> <password>
```

DButil -mfreeze

The `DButil -mfreeze` command suspends multiple databases residing on a single LUN or a set of LUNs that must be replicated as a group. Write operations to specified databases are suspended until the `DButil -mthaw` command is issued. For example, if three databases reside on the same LUN, the following commands would suspend all three:

```
DButil -mfreeze 1 of 3 <sqlserver_name> <database_name> <username> <password>
```

```
DButil -mfreeze 2 of 3 <sqlserver_name> <database_name> <username> <password>
```

```
DButil -mfreeze 3 of 3 <sqlserver_name> <database_name> <username> <password>
```

Complete this procedure by issuing a matching `DButil -mthaw` command before issuing any subsequent `DButil -mfreeze` or `DButil -pre` commands.

DButil -mthaw

The `DButil -mthaw` command resumes multiple SQL Server databases that were suspended using `DButil -mfreeze`.

Syntax:

```
DButil -mthaw 3
```

DButil -recover

Use this command to restore a database from a snapshot or snapclone. You cannot apply transaction logs to a database that has been restored using this command.

Syntax:

```
DButil -recover <sqlserver_name> <database_name> <username> <password>
```

DButil -restoresns

Use this command to restore a database from a snapshot or snapclone when you want to apply transaction logs.

Syntax:

```
DButil -restoresns <sqlserver_name> <database_name> <username> <password>
```

DButil -setcred

The `DButil -setcred` command stores an encrypted copy of the user name and password for a given server-database combination in a local "lock box" file. Once a user name and password are so stored, subsequent HP `DButil` operations for the same server-database do not need a user name or password. HP `DButil` automatically extracts the required credentials from the lock box, decrypts them, and sends them to SQL Server. This allows you to create scripts to control HP `DButil` operations without storing SQL Server authentication information in plain text.

Syntax:

```
DButil -setcred <sqlserver_name> <database_name> <username> <password>
```

Optional parameters and other considerations

Named instances

HP DButil supports SQL Server named instances. Specify an instance with the server name, as shown in the following sample syntax:

```
DButil -pre <sqlserver_name\instance> <database_name> <username> <password>
```

Database names that contain spaces

Enclose the database name in quotes if it contains an embedded space. For example:

```
DButil -pre <sqlserver_name> "My Database" <username> <password>
```

Metadata files

The `DButil -pre` and `DButil -mfreeze` commands create metadata files that contain checkpoint and control information about the database(s) being replicated. This data is sent to SQL Server during DButil restore operations (`-restoresns` or `-recover`) and allows SQL Server to restore from a snapshot or snapclone.

By default, `DButil -pre` and `DButil -mfreeze` create metadata files in the same directory where HP DButil is running and use the following naming convention:

```
servername_databasename.meta
```

If a file of the same name already exists, it is overwritten. Also by default, `DButil -restoresns` and `DButil -recover` automatically look for metadata files that conform to this naming convention.

Because metadata files are valid for the specific snapshot that created them and cannot be used to restore another database or for a different backup of the same database, metadata files must be organized and identified so that they cannot become confused. In general, it is best to accept the default location and naming. However, if you must specify a name for the metadata files, use the following syntax:

```
DButil -pre <sqlserver_name> <database_name> <username>  
<password> [metafilename]
```

Where `[metafilename]` is the path and name of the file for the metadata. If you specify the metadata file when you create a snapshot, you must specify the same name in `DButil -restoresns` or `DButil -recover` to ensure that the correct metadata is used.

2 Creating snapshot backups using HP DButil

Procedure

To back up an SQL Server database using snapshots or snapclones and HP DButil:

1. Determine which databases need to be backed up.
2. Set the necessary login credentials (optional).
3. Create the HP DButil batch files to suspend and resume the desired databases.
4. Determine which volumes need to be replicated.
5. Construct the HP Business Copy EVA/MA/EMA or HP Replication Solutions Manager job to replicate the volumes identified in [Step 4](#).
6. Test the job.

Determining which databases require backup

Keep the following in mind when choosing databases to be backed up:

- Write operations to the database are suspended for the duration of the snap creation. It typically takes 10 to 20 seconds per snapshot or snapclone but can be longer.
- You can use HP DButil to prepare any SQL Server database for backup, but you can use only HP Business Copy EVA/MA/EMA and HP Replication Solutions Manager to create the snapshot or snapclone.
- Once you start a snapshot backup, you cannot start another one until the first one has completed. If you need multiple jobs to back up multiple databases, ensure HP DButil operations do not overlap.

Setting credentials

Regardless of how HP DButil is used, the recommended first step is to establish login credentials for each database. The alternative is to specify a user name and password in the HP DButil batch files.

1. Open a DOS window on the host where SQL Server is installed.
2. Navigate to the directory where HP DButil is located.
3. For each database to be replicated using HP DButil, execute the following command:

```
DButil -setcred <sqlserver_name> <database_name> <username> <password>
```

This command stores the authentication information in a lock box file so that it is not needed in each HP DButil command. If you change the user name or password after setting the credentials, use `DButil -setcred` again to update the information in the lock box file.

If you use multiple instances of HP DButil (for multiple sets of physical LUNs, for example), you must set the credentials for each HP DButil instance. For a given instance of HP DButil, you need only set the credentials for the databases that the instance will interact with.



NOTE:

HP DButil has been validated against SQL server authentication.

Creating the HP DButil batch files

Once you know which databases must be replicated, you can create the batch files to issue the HP DButil commands. Two batch files are needed: one to suspend the database and another to resume it. The actual names of the batch files are not important but `pre.bat` and `post.bat` are used in the examples in this guide. The following examples show two methods. Each example uses the default metadata file names, and assumes that the login credentials are set using the DButil `-setcred` command.

Single SQL Server database

In this example, a single SQL Server database and its log files reside on one or more volumes. The details are as follows:

- HP DButil is installed in the directory `C:\DButil`
- The server name is: `Server1`
- Named instances are not used
- The database name is: `TestDB`

An example `pre.bat` file for this configuration would look like this:

```
@ echo off
rem pre.bat file for suspending a single SQL server database
cd \dbutil
dbutil -pre Server1 TestDB
```

An example `post.bat` file for this configuration would look like this:

```
@ echo off
rem pre.bat file for suspending a single SQL server database
cd \dbutil
dbutil -post Server1 TestDB
```

Three SQL Server databases on a common set of volumes

In this example, three SQL Server databases and their log files are co-located on a set of volumes. Details for this example are as follows:

- HP DButil is installed in the directory `C:\DButil`
- The server name is: `Server1`
- Named instances are not used
- The first database name is: `TestDB1`
- The second database name is: `TestDB2`
- The third database name is: `TestDB3`

An example `pre.bat` file for this configuration would look like this:

```
@ echo off
rem pre.bat file for suspending multiple SQL Server databases
cd \dbutil
dbutil -mfreeze 1 of 3 pre Server1 TestDB1
dbutil -mfreeze 2 of 3 pre Server1 TestDB2
dbutil -mfreeze 3 of 3 pre Server1 TestDB3
```

An example `post.bat` file for this configuration would look like this:

```
@ echo off
rem post.bat file for thawing multiple suspended SQL Server databases
```



```
cd \dbutil
dbutil -mthaw 3
```

Determining which volumes need to be replicated

A full database backup includes all data and log files that are associated with that database. A simple way to determine precisely the needed files is to execute the `sp_dbhelp` stored procedure as follows:

```
exec sp_helpdb <database_name>
```

This procedure returns a list of files used by the database. The snapshots and snapclones of all the volumes identified in this step must be created at the same time and by the same job. Using multiple jobs to create the snapshots or snapclones would result in the files being captured at different times and would invalidate the backup. Repeat this procedure to determine the volumes that need to be replicated for each database.

The following examples expand the examples used in the previous step.

Single SQL Server database

To determine where the data and log files for the `TestDB` database are located, execute the `sp_helpdb` procedure:

```
exec sp_helpdb TestDB
testdb_Data          F:\SQLdata\testdb_Data.MDF
testdb_Log           E:\SQLlogs\testdb_Log.LDF
```

In this case, the database is a single data file located on the `E:` drive and a single log file located on the `F:` drive. A backup of this database requires snapshots or snapclones of both `E:` and `F:`.

Three SQL Server databases on a common set of volumes

To determine where the data and log files for the `TestDB1`, `TestDB2` and `TestDB3` databases are located, execute the `sp_helpdb` procedure for each database:

```
exec sp_helpdb TestDB1
testdb1_Data        F:\SQLdata\testdb1_Data.MDF
testdb1_Log         E:\SQLlogs\testdb1_Log.LDF
```

```
exec sp_helpdb TestDB2
testdb2_Data        G:\SQLdata\testdb2_Data.MDF
testdb2_Log         E:\SQLlogs\testdb2_Log.LDF
```

```
exec sp_helpdb TestDB3
testdb3_Data        H:\SQLdata\testdb3_Data.MDF
testdb3_Log         E:\SQLlogs\testdb3_Log.LDF
```

In this case, the logs files are all contained on the `E:` drive but the data files are spread across the `F:`, `G:`, and `H:` drives. A backup of this database requires snapshots or snapclones of `E:`, `F:`, `G:`, `H:`.

Constructing HP Business Copy EVA/MA/EMA or HP Replication Solutions Manager jobs

Once you have created the `pre.bat` and `post.bat` files and determined the volumes to be replicated, you can construct the HP Business Copy EVA/MA/EMA or HP Replication Solutions Manager jobs. The following sample jobs are based on [Three SQL Server databases on a common set of volumes](#), described in the previous steps. These samples use fully allocated snapshots. Depending on the storage array, other replication methods may be available.

These samples are for informational purposes only and are not intended to be used as is.

Sample HP Business Copy EVA/MA/EMA job

This sample job uses the `snap two volumes.suspend.mount` template in HP Business Copy EVA. Two SNAP VOLUME and MOUNT steps were added to accommodate additional volumes to be replicated in this job, and the comments were removed for clarity:

```
;Sample BC job for Example2
SUSPEND WAIT TRUE TRUE INTEGER 0 Server1 pre.bat
SNAP VOLUME Server1 E: $BCV1 FULLY_ALLOCATED SAME_AS_SOURCE
SNAP VOLUME Server1 F: $BCV2 FULLY_ALLOCATED SAME_AS_SOURCE
SNAP VOLUME Server1 G: $BCV3 FULLY_ALLOCATED SAME_AS_SOURCE
SNAP VOLUME Server1 H: $BCV4 FULLY_ALLOCATED SAME_AS_SOURCE
RESUME WAIT TRUE TRUE INTEGER 0 Server1 post.bat
MOUNT VOLUME_SINGLE SS $BCV1 Backup-Server N/A N/A 1 L:
MOUNT VOLUME_SINGLE SS $BCV2 Backup-Server N/A N/A 1 M:
MOUNT VOLUME_SINGLE SS $BCV3 Backup-Server N/A N/A 1 N:
MOUNT VOLUME_SINGLE SS $BCV4 Backup-Server N/A N/A 1 O:
LAUNCH WAIT TRUE TRUE INTEGER 0 Backup-Server backup.bat
```

Sample HP Replication Solutions Manager job

This sample job uses the Replicate host volume(s), mount to a host template. The # of volumes to replicate was set to 4, and Suspend source before replication and Launch backup after replication were selected. The comments and Validation steps normally included with this template were removed for clarity:

```
// Sample RSM job for Example 2
Launch ( "Server1", "pre.bat", "", WAIT, "0" ) onerror pauseat E1:
DO {
$Rep1 = SnapshotHostVolume ( "\\Server1\E:\", FULLY_ALLOCATED, SAME, NOWAIT ) onerror pauseat E1:
$Rep2 = SnapshotHostVolume ( "\\Server1\F:\", FULLY_ALLOCATED, SAME, NOWAIT ) onerror pauseat E2:
$Rep3 = SnapshotHostVolume ( "\\Server1\G:\", FULLY_ALLOCATED, SAME, NOWAIT ) onerror pauseat E3:
$Rep4 = SnapshotHostVolume ( "\\Server1\H:\", FULLY_ALLOCATED, SAME, NOWAIT ) onerror pauseat E4:
} ALWAYS {
Launch ( "Server1", "post.bat", "", WAIT, "0" )
}
WaitForStorageVolumesDiscovery ( $Rep1 ) onerror pauseat E5:
WaitForStorageVolumesDiscovery ( $Rep2 ) onerror pauseat E5:
WaitForStorageVolumesDiscovery ( $Rep3 ) onerror pauseat E5:
WaitForStorageVolumesDiscovery ( $Rep4 ) onerror pauseat E5:
$HV1 = CreateHostVolume ( "\\Server1\E:\", $Rep1, "Backup-Server" ) onerror pauseat E5:
$HV2 = CreateHostVolume ( "\\Server1\F:\", $Rep2, "Backup-Server" ) onerror pauseat E6:
$HV3 = CreateHostVolume ( "\\Server1\G:\", $Rep3, "Backup-Server" ) onerror pauseat E7:
$HV4 = CreateHostVolume ( "\\Server1\H:\", $Rep4, "Backup-Server" ) onerror pauseat E8:
$MP1 = MountHostVolume ( $HV1, "L:" ) onerror pauseat E9:
$MP2 = MountHostVolume ( $HV2, "M:" ) onerror pauseat E10:
$MP3 = MountHostVolume ( $HV3, "N:" ) onerror pauseat E11:
$MP4 = MountHostVolume ( $HV4, "O:" ) onerror pauseat E12:
Launch ( "Backup-Server", "backup.bat", "", WAIT, "0" ) onerror pauseat E13:
Pause ( )
//Unwind section
E13: UnmountHostVolume ( $MP4 ) onerror pauseat E12:
E12: UnmountHostVolume ( $MP3 ) onerror pauseat E11:
E11: UnmountHostVolume ( $MP2 ) onerror pauseat E10:
E10: UnmountHostVolume ( $MP1 ) onerror pauseat E9:
E9: DeleteHostVolume ( $HV4 ) onerror pauseat E9:
E8: DeleteHostVolume ( $HV3 ) onerror pauseat E8:
E7: DeleteHostVolume ( $HV2 ) onerror pauseat E7:
E6: DeleteHostVolume ( $HV1 ) onerror pauseat E6:
E5: DeleteStorageVolumes ( $Rep4 ) onerror pauseat E5:
E4: DeleteStorageVolumes ( $Rep3 ) onerror pauseat E4:
E3: DeleteStorageVolumes ( $Rep2 ) onerror pauseat E3:
E2: DeleteStorageVolumes ( $Rep1 ) onerror pauseat E2:
Exit ( SUCCESS )
```

```
E1: Exit ( FAILURE )
```

Testing the process

Thoroughly test the entire backup and restore process before implementing it in a production environment. The following chapter explores the recovery and restore process in more detail.

Backing up transaction logs

Snapshots and snapclones are point-in-time images of the database as it exists when the snapshot or snapclone is created. By themselves, these backups can be used to restore the database to the same point in time. To roll the database forward from the point of the backup to a more recent time, you need a continuous set of transaction log backups for the database.

Consult SQL Server documentation for information about backing up transaction logs. Regardless of the method used, a restored database can be rolled forward only to the point in time of the most recent transaction log. This determines the “data exposure” or the maximum amount of data that may be lost in a failure. Give this critical aspect of the backup and restore process careful consideration when planning backups.

3 SQL Server recovery using HP DButil

Recovery approaches

The appropriate recovery method depends on the situation and the type of backup. Choosing an incorrect method can prevent transaction logs from being applied and the database from being restored to the latest log backup.

Two basic recovery methods can be used with snapshots or snapclones: volume replacement and selective file replacement.

Volume replacement

Volume replacement involves completely removing the existing source volume and replacing it with a backup. This method is preferred when the source volume has failed or if the combined size of the files that need to be restored is large enough that selective file replacement (described below) would take too long. There are two ways that volume replacement can be accomplished:

Direct snapclone: With this approach, the source volume is removed from the source host and replaced with a snapclone using the same drive letter or mount point.

From tape using intermediate device: If there is no snapclone but the required fileset is available on tape, you can create a volume of the desired size and redundancy level, present it to the backup server, restore the data from tape to the new volume and use this new volume as a replacement for the original source.

△ CAUTION:

The volume replacement method replaces the entire contents of the source. If the source volume contains data or log files for databases that were not suspended when the copy was created, then replacing the entire volume with a snapclone copy will effectively corrupt these databases. When using the volume replacement method, be sure that ALL files on the replacement volume belong to databases that were backed up at the same time, by the same HP DButil operation.

NOTE:

Backups based on snapshots (as opposed to snapclones) are unsuitable for volume replacement because snapshots are virtual volumes that depend on the integrity of the source. If the source volume has physically failed, any snapshots of it will have likely failed as well.

Selective file replacement

This method involves leaving the original source volume in place and replacing some or all of its files with backup copies. This method is generally preferred over volume replacement in cases where only a single database needs to be restored and other databases also reside on the volume. There are two ways that selective file replacement can be accomplished:

Drag and drop from snap: If a snapshot or snapclone of the required data and log files is available, the snapshot or snapclone devices can be presented to the source host and manually mounted using operating system utilities. The required files can simply be dragged and dropped from the snapshot or snapclones back to the original source volumes.

From tape: If the backup exists only on tape, it may be possible to restore the required files from tape directly back to the source volume. This process can be complicated if the files being restored were originally backed up from snapshots or snapclones on a different server. To restore the source volume, the restore process must be redirected to the desired location. Most backup applications support this feature.

Applying transaction logs to a restored database

Snapshots and snapclones are point-in-time images of the database as it exists when the backup is created. This point in time may be hours or even days old, depending on when the backup was taken. In order to bring the database current, it is necessary to apply a series of transaction log backups to the restored database.

In some cases, though, you may want to recover the database to the point in time of the backup, without rolling it forward. HP DButil supports both cases but the recovery approach is different for each.

Example recovery scenarios

The following recovery scenarios illustrate different circumstances involving the database failure, the condition of the database and transaction logs, and the type of backup.

Recovery scenario 1

Situation: A source volume has failed and needs to be replaced. SQL Server is on a cluster. Multiple transaction log backups need to be applied to roll the database forward to the most recent point in time. Snapclones of the volumes that contain the database files are available and do not contain other data.

Solution: You decide to restore the database using volume replacement and to apply transaction logs.

1. Take the bad SQL database offline using Enterprise Manager or Query Analyzer.
2. Run the `DButil -restoresns` command to start restoring the database. You will receive a prompt to replace the files.
3. Stop the SQL Server cluster service.
4. Remove the database volumes as SQL Server cluster dependencies in Cluster Administrator, and delete the database cluster disks.
5. Unpresent the failed database disks from the cluster hosts. (For example, using HP Command View EVA, select the appropriate virtual disk and click the **Unpresent** tab. Unpresent the failed disk from each cluster node.)
6. Present the snapclones to the cluster hosts. (For example, using HP Command View EVA, select the snapclone virtual disk and click the **Present** tab. Present the snapclone to each cluster host.)
7. Run the Disk Management utility on the SQL Server host and scan for new disks. Repeat as needed until all snapclone disks have been found.
8. Change the drive letters on the snapclone to reflect the drive letters of the original database volumes using the Disk Management utility.
9. Add the snapclone volumes as cluster resources and SQL Server cluster dependencies.
10. Start the SQL Server cluster service.
11. Acknowledge the HP DButil file replacement prompt. This will complete the restoration and leave the database in a loading state. In the loading state, the database is ready to accept a restore of the transaction logs.

 **NOTE:**

The `DButil -restoresns` command uses the metadata file to put the database in a loading state. If you restore a SQL Server database to a new SQL server with a different SQL server name, you need to edit the name of the metadata file to match the new name of the SQL server.

12. Use Enterprise Manager to verify that the database is in a loading state and apply sequential transaction logs. You can use Enterprise Manager, Query Analyzer, or a third-party backup application. When using Enterprise Manager or Query Analyzer, be sure to use the No Recovery option for each log backup but the last, and choose the With Recovery option for the last log backup to make the database operational.
13. Delete the HP Business Copy EVA/MA/EMA or HP Replication Solutions Manager job that created the snapclone.

 **CAUTION:**

Delete the job that created the snapclone so that it cannot be inadvertently undone. Undoing the job would delete the snapclone.

Recovery scenario 2

Situation: A source volume has failed and needs to be replaced. SQL Server is on a standalone system. Data and logs need to be restored. Multiple transaction log backups need to be applied to roll forward the database to the most recent point in time. Tape backups are available.

Solution: You decide to restore the database, using volume replacement and an intermediate device and to apply transaction logs.

1. Create virtual disks of the same size, configuration, and VRAID level as the original source volumes that are to be replaced.
2. Present the newly created volumes to the backup server.
3. Scan for new devices using the Disk Management utility on the backup server, and partition and format the new volumes using the same configuration as the source volumes.
4. Initiate a restore from the tape backup to the new volumes.
5. Unpresent the new volumes from the backup server. (For example, using HP Command View EVA, select the appropriate virtual disk and click the **Unpresent** tab)
6. Present the new volumes to the SQL Server host. (For example, using HP Command View EVA, select the virtual disk and click the **Present** tab. Present the virtual disk to the SQL Server host.)
7. Scan for new devices on the SQL Server host using the Disk Management utility.
8. Take the bad SQL database(s) offline using Enterprise Manager or Query Analyzer.
9. Run the `DButil -restoresns` command to restore the database. You will receive a prompt to replace the files.
10. Unpresent the failed database disks from the SQL Server host. (For example, using HP Command View EVA, select the appropriate virtual disk and click the **Unpresent** tab.)
11. Change the drive letters or mount points on the new volumes to match those of the original volumes, using the Disk Management utility.
12. Acknowledge the HP DButil file replacement prompt. This completes the restoration and leaves the database in a loading state. In the loading state, the database is ready to accept the transaction logs.

**NOTE:**

The `DButil -restoresns` command uses the metadata file to put the database in a loading state. If you are restoring a SQL Server database to a new SQL server with a different SQL server name, you need to edit the name of the metadata file to match the new name of the SQL server.

13. Use Enterprise Manager to verify that the database is in a loading state and apply the sequential transaction logs. You can use Enterprise Manager, Query Analyzer, or a third-party backup application. When using Enterprise Manager or Query Analyzer, be sure to use the No Recovery option for each log backup but the last, and choose the With Recovery option for the last log backup to make the database operational.

Recovery scenario 3

Situation: SQL Server is on a standalone system. Database volumes are intact but a logical corruption has occurred in one of the databases. Multiple transaction log backups need to be applied to roll forward the database to the most recent point in time. Snapshot copies of the volumes that contain the database files are available.

Solution. You decide to restore the database using the selective file replacement method.

1. Present the snapshot copies to the SQL Server host. (For example, using HP Command View EVA, select the snapshot virtual disk and click the **Present** tab. Present the snapshot to the SQL Server host.)
2. Run the Disk Management utility on the SQL Server host and scan for new disks. Repeat as needed until all snapshot disks have been found. Assign drive letters to the snapshot volumes as needed.
3. Take the bad SQL database offline using Enterprise Manager or Query Analyzer.
4. Run the `DButil -restoresns` command to restore the database. You will receive a prompt to replace the files.
5. Manually copy the data and log files for the database to be restored from the snapshot volumes to the original volumes.
6. Acknowledge the HP DButil file replacement prompt. This completes the restoration and leaves the database in a loading state. In the loading state, the database is ready to accept transaction logs.

**NOTE:**

The `DButil -restoresns` command uses the metadata file to put the database in a loading state. If you are restoring a SQL Server database to a new SQL server with a different SQL server name, you need to edit the name of the meta file to match the new name of the SQL server.

7. Use Enterprise Manager to verify that the database is in a loading state and apply the sequential transaction logs. You can use Enterprise Manager, Query Analyzer, or a third party backup application. When using Enterprise Manager or Query Analyzer, be sure to use the No Recovery option for each log backup but the last, and choose the With Recovery option for the last log backup to make the database operational.

Recovery scenario 4

Situation: SQL Server is on a standalone system. Database volumes are intact but a logical corruption has occurred in one of the databases. The time of the corruption is not known precisely but is known to have occurred at some point after the last snapshot backup. Snapshot copies of the volumes that contain the database files are available.

Solution: You decide to restore the database using selective file replacement. Because the time that the corruption was introduced is not known, you decide not to apply transaction logs and to fall back to the last known good point which corresponds to the time of the last snapshot backup.

1. Present the snapshot copies to the SQL Server host. (For example, using HP Command View EVA, select the snapshot virtual disk and click the **Present** tab. Present the snapshot to the SQL Server host.)
2. Run the Disk Management utility on the SQL Server host and scan for new disks. Repeat as needed until all of the snapshot disks have been found. Assign drive letters to the snapshot volumes as needed.
3. Take the bad SQL database offline using Enterprise Manager or Query Analyzer.
4. Run the `DButil -recover` command to restore the database. You will receive a prompt to replace the files.
5. Manually copy the data and log files for the database to be restored from the snapshot volumes to the original volumes.
6. Acknowledge the HP DButil file replacement prompt. This completes the restoration and recovers the database at the point in time of the snapshot backup.

4 Troubleshooting HP DButil

Status and error messages are written to an HP DButil log file each time the HP DButil utility is run. These log files can help you diagnose and correct common problems encountered when using HP DButil.

The HP DButil log file is given the name `sqlserver_dbname.log` where `sqlserver` is the name of the SQL Server and `dbname` is the name of the database.

The log file is located in the same directory as HP DButil and gets overwritten each time the application is run.

Table 2 Troubleshooting HP DButil error messages

Error Message	Possible Cause(s)	Recommended Ready
SQL -pre step failed	The process that carries out the -pre step has gone away. This could be due to manual intervention or failure to get a virtual device.	This is always accompanied by some other message that gives the true cause of the problem. There is no need to do anything, as the process terminated due to the failure of a child process-as it should.
Error: Could not create component: x80040154 Check registration of SQLVDI.DLL and value of IID	SQLVDI.DLL was not registered when SQL Server was installed on the machine that you are trying to connect to.	Locate the SQLVDI.DLL file and register it using the regsvr32 command.
[Microsoft][ODBC SQL Server Driver][SQL Server] Exclusive access could not be obtained because the database is in use.	HP DButil requires exclusive access to the SQL Server during the restore step.	Stop any processes that are using the database before attempting to restore it. If necessary, take the database offline before initiating the restore.
VDS::OpenDevice fails.	HP DButil was unable to open the Virtual Device. This could be caused by a mismatch between the server name and dbname parameters given to HP DButil and the metadata file previously created.	Check the directory that HP DButil is installed in to make sure the *.meta file exists. Rename the meta file to match the sqlserver_dbname.meta. This can only be done if the meta file is the correct file with an incorrect name.
Connect fails	HP DButil was unable to connect to the SQL Server due to incorrect Server name/DB name or username/password combination.	Check the capitalization of all the input parameters. Use lowercase characters for all inputs (providing they have been similarly specified on the SQL Server/database).
Failed to open sqlserver_db-name.meta	HP DButil cannot find the metadata file, probably due to it having a different name than that specified by the server_dbname combination.	Check the directory that HP DButil is installed to make sure the *.meta file exists. Rename the metadata file to match the sqlserver_dbname.meta. This can only be done if the metadata file is the correct file with an incorrect name.
VDS::Create fails	HP DButil has been the wrong instance name for the SQL Server.	Check that the instance specified exists on the specified SQL Server.
Check on SQL failed	The SQL command has failed. This is followed by a more specific error message.	If the more specific message does not provide a cause then repeating this step may be sufficient to overcome the error.

Index

A

audience, 7

C

conventions
document, 8

D

DButil commands, 11
DButil -post, 12
DButil -pre, 11
DButil -restoresns, 12
document
conventions, 8
providing feedback, 9
related documentation, 7
documentation
HP website, 7

F

feedback
documentation, 9
product, 9

H

help
obtaining, 8
HP technical support, 8

I

installing DButil, 11

P

product
providing feedback, 9

R

recovery scenarios, 22, 23, 24, 24
examples, 22
related documentation, 7
running DButil, 11

S

SQL Server
recovery approaches, 21
Subscriber's Choice, HP, 9

T

technical support
HP, 8
service locator website, 9
troubleshooting DButil
error messages (table), 28

W

websites
HP, 9
HP Subscriber's Choice for Business, 9
product manuals, 7