

Veritas Storage Foundation™ for DB2 Graphical User Interface Guide

AIX, Linux, Solaris

5.0

Veritas Storage Foundation™ for DB2 Graphical User Interface Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version 5.0

PN: N18468C

Legal Notice

Copyright © 2006 Symantec Corporation.

All rights reserved.

Federal acquisitions: Commercial Software - Government Users Subject to Standard License Terms and Conditions.

Symantec, the Symantec Logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

DB2 is a registered trademark of IBM Corporation.

Windows is a registered trademark of Microsoft Corporation.

Veritas Storage Foundation™ is a licensed product. See the Veritas Storage Foundation™ Installation Guide for license installation instructions.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation 20330 Stevens Creek Blvd. Cupertino, CA 95014 USA

<http://www.symantec.com>

Technical Support

For technical assistance, visit <http://support.veritas.com> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1

Using the Veritas Storage Foundation for DB2 graphical user interface

| | |
|--|----|
| About GUI components | 9 |
| About GUI functions | 13 |
| Using Veritas Enterprise Administrator Service | 15 |
| Adding users to the Veritas Enterprise Administrator | 15 |
| Removing users from the Veritas Enterprise Administrator | 19 |
| Starting the Veritas Enterprise Administrator Service | 21 |
| Shutting down the Veritas Enterprise Administrator Service | 22 |
| Starting the DBED agent | 23 |
| Stopping the DBED agent | 23 |
| Opening and closing the Veritas Storage Foundation for DB2 GUI | 24 |
| Opening the Veritas Storage Foundation for DB2 Java GUI from a UNIX client | 24 |
| Opening the Veritas Storage Foundation for DB2 Java GUI from a Windows client | 25 |
| Opening the Veritas Storage Foundation for DB2 Web GUI | 25 |
| Closing the Veritas Storage Foundation for DB2 Java GUI | 25 |
| Closing the Veritas Storage Foundation for DB2 Web GUI | 25 |
| Performing tasks in the Veritas Storage Foundation for DB2 Java GUI | 26 |
| Accessing database management tasks in the Java GUI | 26 |
| Accessing clone database tasks in the Java GUI | 27 |
| Accessing Storage Checkpoint tasks in the Java GUI | 27 |
| Accessing FlashSnap tasks in the Java GUI | 28 |
| Accessing tiered storage tasks in the Java GUI | 29 |
| Accessing scheduled tasks in the Java GUI | 31 |
| Accessing system maintenance tasks in the Java GUI | 32 |
| Performing tasks in the Storage Foundation for DB2 Web GUI | 34 |
| Accessing database management tasks in the Web GUI | 34 |
| Accessing database tasks in the Web GUI | 35 |
| Accessing system maintenance tasks in the Web GUI | 36 |

Chapter 2 Managing your database

| | |
|--|----|
| Database management tasks in the Veritas Storage Foundation for DB2 GUI | 39 |
| About setting up a disk group | 41 |
| Disk group configuration guidelines | 41 |
| About selecting a volume layout | 42 |
| How to choose appropriate stripe unit sizes | 43 |
| How to choose between mirroring and RAID-5 | 43 |
| Volume configuration guidelines | 43 |
| File system creation guidelines | 44 |
| Controlling fragmentation | 45 |
| Fragmentation utilities in VxFS | 46 |
| Types of fragmentation | 46 |
| Defragmenting a file system | 46 |
| Starting a DB2 instance | 48 |
| Creating a DB2 snapshot database | 48 |
| Unmounting a DB2 snapshot database | 49 |
| Shutting down a DB2 instance | 50 |
| Restarting a DB2 instance | 51 |

Chapter 3 Using a clone database

| | |
|---|----|
| Creating a clone database | 53 |
| Creating a clone database using Checkpoints | 53 |
| Creating a clone database using Database FlashSnap | 55 |
| Creating a new snapplan for a clone database | 55 |
| Validating or modifying a snapplan for a clone database | 56 |
| Creating snapshot for a clone database | 58 |
| Start up snapshot for a clone database | 59 |
| Create snapshot for clone database | 60 |
| Removing a clone database | 61 |
| Removing a Storage Checkpoint clone database | 62 |
| Removing a database snapshot clone database | 62 |

Chapter 4 Managing Storage Checkpoints

| | |
|--|----|
| About Storage Checkpoints | 65 |
| Creating a Storage Checkpoint | 66 |
| Viewing Storage Checkpoint details | 67 |
| Mounting a Storage Checkpoint | 67 |
| Unmounting a Storage Checkpoint | 68 |
| Removing a Storage Checkpoint | 69 |
| Rolling back to a Storage Checkpoint | 69 |

| | | |
|------------------|---|-----|
| | Rolling back the database to a Storage Checkpoint | 70 |
| | Rolling back datafiles to a Storage Checkpoint | 71 |
| Chapter 5 | Managing snapshots using Database FlashSnap | |
| | Database FlashSnap in the GUI | 73 |
| | Creating a snapplan | 74 |
| | Validating or modifying a snapplan | 76 |
| | Removing a snapplan | 78 |
| | Creating a snapshot | 79 |
| | Creating a clone database with a snapshot | 80 |
| | Resynchronizing a snapshot to a database | 81 |
| | Resynchronizing a database to a snapshot | 82 |
| | Aborting the reverse resynchronization operation | 83 |
| | Viewing the snapplan log | 84 |
| Chapter 6 | Tiering storage for databases | |
| | About Database Dynamic Storage Tiering in the GUI | 87 |
| | Setting up Database Dynamic Storage Tiering parameters | 88 |
| | Setting up Database Dynamic Storage Tiering parameters | 88 |
| | Identifying tiered storage directories | 90 |
| | Classifying volumes | 91 |
| | Setting automatic policies for Database Dynamic Storage Tiering | 91 |
| | Listing policies | 97 |
| | Presetting policy for a directory | 92 |
| | Presetting policy for SMS tablespace | 93 |
| | Presetting policy for automatic storage path | 93 |
| | Setting archive1/archive 2 policy | 94 |
| | Setting external files policy | 95 |
| | Moving volumes and files | 95 |
| | Moving tablespaces | 95 |
| | Moving containers | 96 |
| | Getting reports for Database Dynamic Storage Tiering | 96 |
| | Showing activities | 97 |
| | Listing policies | 97 |
| | Showing volume usage | 97 |
| | Starting and stopping datafile statistics collection | 98 |
| | Showing statistics chart | 99 |
| Chapter 7 | Automating database tasks | |
| | About automating tasks using the Java GUI scheduler | 101 |
| | Adding a scheduled task | 102 |

| | |
|---|-----|
| Viewing scheduled jobs | 102 |
| Modifying a scheduled job task | 103 |
| Automating creation of a database checkpoint | 104 |
| Automating creation of a database snapshot | 104 |
| Automating cloning of a database using a checkpoint | 105 |
| Automating cloning of a database using FlashSnap | 105 |
| Set the schedule parameters for an automated task | 106 |

Chapter 8 Maintaining your system configuration

| | |
|--|-----|
| System configuration maintenance | 109 |
| Resynchronizing the SFDB repository | 110 |
| Rescanning the instance | 110 |
| Displaying and rescanning tablespace information | 110 |
| Viewing Storage Mapping | 111 |
| Viewing I/O statistics | 111 |
| Converting regular containers to Quick I/O Files | 112 |
| Converting Quick I/O files to regular containers | 113 |
| Updating rescan intervals | 113 |
| Checking your system configuration | 114 |
| Saving the system configuration | 116 |

Appendix A Veritas Database FlashSnap status information

| | |
|--|-----|
| Obtaining database FlashSnap snapshot status and database status | |
| from the GUI | 117 |
| Database FlashSnap snapshot status details | 117 |
| Snapshot database status details | 122 |

Appendix B Accessibility and Veritas Storage Foundation for DB2

| | |
|---|-----|
| About Veritas Storage Foundation for DB2 accessibility | 123 |
| Keyboard conventions | 124 |
| Keyboard navigation in the GUI | 125 |
| Keyboard navigation and shortcuts in Veritas Storage foundation for | |
| DB2 | 127 |
| General keyboard navigation within the GUI | 129 |
| Keyboard shortcuts | 131 |
| Keyboard navigation | 132 |
| Menu hot keys | 132 |
| Support for accessibility settings and assistive technologies | 136 |

Index

Using the Veritas Storage Foundation for DB2 graphical user interface

This chapter includes the following topics:

- [About GUI components](#)
- [About GUI functions](#)
- [Using Veritas Enterprise Administrator Service](#)
- [Starting the DBED agent](#)
- [Stopping the DBED agent](#)
- [Opening and closing the Veritas Storage Foundation for DB2 GUI](#)
- [Performing tasks in the Veritas Storage Foundation for DB2 Java GUI](#)
- [Performing tasks in the Storage Foundation for DB2 Web GUI](#)

About GUI components

You can access Veritas Storage Foundation for DB2, Veritas Volume Manager, and Veritas File System functions through the Veritas Storage Foundation for DB2 graphical user interface (GUI). Veritas Storage Foundation for DB2, supports a Java GUI and Web version of the Veritas Enterprise Administrator.

The graphical user interfaces (GUI) enable you to perform storage management duties for DB2, such as monitoring the database.

The GUI runs in a client-server environment. The server is located on a host that runs Veritas Storage Foundation for DB2. The client can run on any UNIX platform, Windows NT, Windows 98, Windows 2000, Windows Me, and Windows XP machine that supports the Java Runtime Environment.

Within the Java GUI, you can perform tasks from the main menu bar or you can right-click an object in the navigational pane on the left side of the screen.

Within the Web GUI, you can perform tasks from the main menu bar or you can select an option in the task lists on the left side of the screen.

Figure 1-1 Veritas Storage Foundation for Databases Java GUI


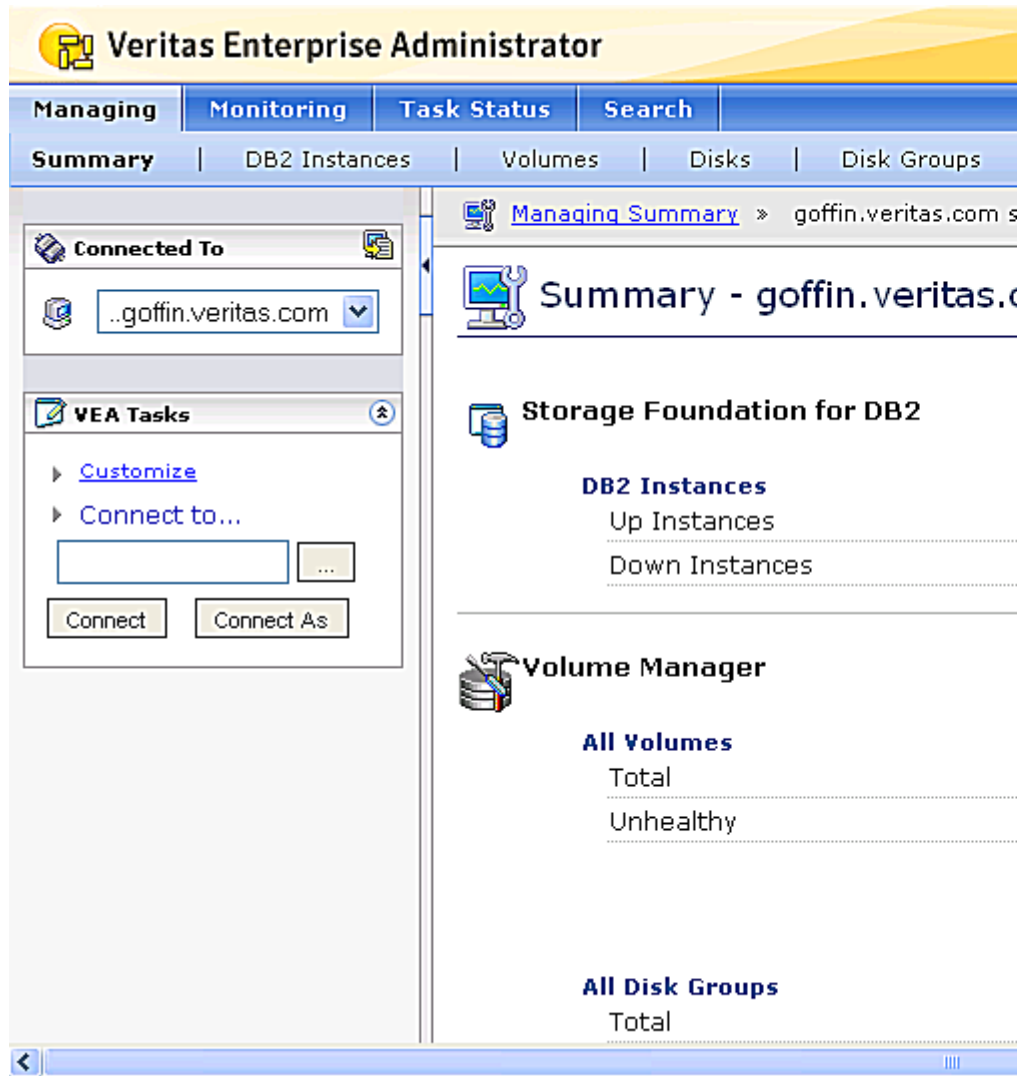


Figure 1-2 Veritas Storage Foundation for Databases Web GUI



Sections of the main window for the Java GUI:

- **Menu bar:** enables you to perform various SFDB operations. The options in the Menu Bar vary according to the object in the object tree that you have selected. To access online help from the Menu Bar, click **Help > Contents**.
- **Tool bar:** provides shortcuts to various operations available in the Menu Bar. The Tool Bar is icon-based and dynamically changes when you select something

from the Object Tree. When you use your mouse to point at an icon, a description of the icon appears.

- **Information tabs:** enable you to view different information about the same object in the Object Tree. For example, if you are viewing details about a database, you can click a different tab to view different information about that database.
- **Object tree:** is a dynamic hierarchical display of Veritas Storage Foundation for DB2 DBED agent, Veritas Volume Manager, and Veritas File System objects, Storage agent, and other objects on the system.
- **Status fields:** indicate the status of the object you are viewing. To change your view, click one of the Information Tabs at the top of the window.
- **GUI status area:** displays GUI status, which is provided through Veritas Enterprise Administrator. See the Veritas Enterprise Administrator documentation for more information.
- **Host Information bar:** enables you to toggle between different server connections.

Within the GUI, options that are not available are display the same as other options. When you click a task that is not available to you, an "Access denied" error displays. Pointing to an object on the screen displays a description of the object in a pop-up field.

Sections of the Summary page for the Web GUI:

- **Managing menu bar:** provides shortcuts to various operations available.
- **Summary menu bar:** provides shortcuts to various summary pages available.
- **Task lists:** provide drop down lists of task options available in each page. Generally these are located in the left pane of summary pages.
- **Connected to:** enables you to toggle between different server connections.

About GUI functions

The Java GUI enables you to perform many storage management tasks for DB2. The tasks that you can perform dynamically from the main menu bar and pop-up menu depend on what is highlighted on the object tree. For example, if you have the instance or database highlighted in the tree, you can open it from the DB2 menu.

The Web GUI provides summary pages enabling you to perform tasks from the drop down lists.

To access the GUI functions:

- Your user or group name must be added to the Veritas Enterprise Administrator (VEA) Authorization Database (AZDB). To make the necessary changes: See [“Using Veritas Enterprise Administrator Service”](#) on page 15. In an AIX environment, `VRTSdb2ed` is requested to be set as the DB2 Administrator during installation.
- Some operations require that you must be logged in as `root`. Please read all “Prerequisites” and “Usage notes” before starting a procedure.

GUI functionality includes:

- Instance operations
- Database operations
You can start or stop a DB2 instance through the GUI.
- Storage Checkpoint Management
You can create and roll back to Storage Checkpoints. You can also mount, unmount, and remove Storage Checkpoints. Storage Checkpoints can be used to clone the database.
- Database FlashSnap
Database FlashSnap functionality is available through the GUI. With Database FlashSnap, you can create online or offline snapshots of your database, which can be used as backups. You can also use these snapshots to recover your database if it becomes corrupt.
- Database Dynamic Tiered Storage
You can set up tiered storage parameters and policy using the GUI. With tiered storage set up, you can move files and volumes, set policy for automated movement, and get reports on storage.
- Automatic scheduling
The automated scheduler functions can be set up using the GUI. You can schedule automated checkpoints and snapshots.
- System configuration and maintenance
You can view storage topology and storage statistic information. You can also check and save your system configuration to a file. In a Solaris or AIX environment, you can convert datafiles to or from Quick I/O files. The Quick I/O feature is not available on Linux.
- Rescan System Information
You can rescan, or refresh, your system information, such as instances, databases, and tablespaces. You can also determine the rescan intervals so that automatic rescans happen as regularly as you determine.

Using Veritas Enterprise Administrator Service

To use the Veritas Storage Foundation for DB2 GUI, the DBED agent must be running on the server. The Veritas Enterprise Administrator (VEA) Service is started when you install the software.

- Prerequisites
- Use the `/opt/VRTS/bin/vxdb2edusr` utility to create login names for any user, other than `root`, who needs to run the GUI. To run `/opt/VRTS/bin/vxdb2edusr`, you must have superuser (`root`) privileges. For more information:
See [“Adding users to the Veritas Enterprise Administrator”](#) on page 15.
 - You must have superuser (`root`) privileges to execute the `vxsvctrl` command.
 - Before you can manage the database(s) through the GUI, you must update the database at least once on each database.
To update a DB2 EE database, use the `db2ed_update` command and to update a DB2 EEE database, use the `db2ed_update_all` command.
 - You must have superuser (`root`) privileges to run the `startup.sh` command.

Adding users to the Veritas Enterprise Administrator

You may want to add users to the VEA authorization Database (AZDB) to allow access to the interface to users other than `root`. You also have the option to give database administrators `root` privileges.

Having `root` privileges means that you can access the operations for volume, disk, and file system objects in the system. Use the commands in the table to add users.

Web GUI users must have a user account on the VEA authorization Database (AZDB) as well as a user profile on the Web server to access the Web GUI version of the Veritas Enterprise Services.

| | |
|----------------------------------|---------------------------------------|
| <code>-a</code> | Adds a user or group to the registry. |
| <code>user -n user_name</code> | Specifies the user to be added. |
| <code>group -n group_name</code> | Specifies the group to be added. |

| | |
|---|---|
| <code>-A</code> | Grants the user <code>root</code> access, making all SF DB2, Veritas File System, and Veritas Volume Manager objects accessible. If the <code>-A</code> is not specified, the user or group will not be able to perform any Veritas Volume Manager or Veritas File System functions, and will not be able to access Storage Foundation for DB2. |
| | Allows the user access to the GUI if the user's primary group is different than the group owner. |
| <code>-h fully_qualified_host_name</code> | Specifies a fully qualified host name on which you want to add a user. For a stand alone installation, the command is run on the local host and the host name is optional. For a centrally managed server installation, you must be logged in as <code>root</code> and specify the host name. |
| <code>-d domain_name</code> | Specifies the domain to which the user belongs. By default the domain is the host name specified with <code>-h</code> . |
| <code>-t domain_type</code> | Specifies the type of domain to which the user belongs. Valid values are <code>nis</code> , <code>nisplus</code> , <code>Idap</code> , <code>unixpwd</code> , and <code>gssapi</code> . By default the domain type is assumed to be a UNIX login. |

To add users other than root to the Veritas Enterprise Administrator AZDB

- 1 If you installed packages manually, make sure that the GUI package was installed.

| OS | Command |
|---------|---------------------------|
| AIX | lslpp -L VRTSd2gui |
| Linux | rpm -qi VRTSd2gui -common |
| Solaris | pkginfo -l VRTSd2gui |

- 2 Check if VxSVC is online. Bring it online if it is offline.

```
# /opt/VRTSob/bin/vxsvcctl status
# /opt/VRTSob/bin/vxsvcctl start
```

- 3 To give root privileges within the GUI to the database administrator, use the `vxdb2edusr` command:

```
# /opt/VRTS/bin/vxdb2edusr -a user [-A] [-f] -n user_name
[-h fully_qualified_host_name -d domain_name -t domain_type]
```

See the table for command usage details. For example:

To add a database administrator with the name “db2inst1” as a user with root privileges, enter:

```
# /opt/VRTS/bin/vxdb2edusr -a user -A -f -n db2inst1
```

To add a UNIX login with the name “db2inst1” into a stand alone installation:

```
# /opt/VRTS/bin/vxdb2edusr -a user -n db2inst1 -h test.example.com
```

To add an NIS user with the name “db2inst1” in example.com domain with a centrally managed server installation of VEA:

```
# /opt/VRTS/bin/vxdb2edusr -a user -n db2inst1 -h test.example.com
test.example.com
```

- 4 To add a user without `root` privileges, use the `vxdb2edusr` command:

```
# /opt/VRTS/bin/vxdb2edusr -a user -n user_name
```

See the table for command usage details. For example:

To add "db2inst1" as a user, enter the following:

```
# /opt/VRTS/bin/vxdb2edusr -a user -n db2inst1
```

- 5 To add a group to the AZDB, use the `vxdb2edusr` command:

```
# /opt/VRTS/bin/vxdb2edusr -a group [-A] [-f] -n group_name  
[-h fully_qualified_host_name -d domain_name -t domain_type]
```

See the table for command usage details. For example:

To add group "dba" to the AZDB on the local host:

```
# /opt/VRTS/bin/vxdb2edusr -a group -A -f -n dba
```

To add UNIX group "dba" to the AZDB on the local host:

```
# /opt/VRTS/bin/vxdb2edusr -a group -A -f -n dba -h test.example.com
```

To add NIS group "dba" to the centrally managed server installation of VEA:

```
# /opt/VRTS/bin/vxdb2edusr -a group -A -f -n dba -h test.example.com -d exa
```

To add a user profile to the Web server

- 1 Create the user name and password on the Veritas Enterprise Administrator AZDB.
See ["To add users other than root to the Veritas Enterprise Administrator AZDB"](#) on page 17.
- 2 Use your browser to connect to the Web server at `https://host:8443/VEAWeb/Login`.
See ["Opening and closing the Veritas Storage Foundation for DB2 GUI"](#) on page 24.
- 3 Select the **Create Profile** option.
- 4 Enter the user name and password. Confirm the password.
- 5 Click **Create Profile**.

Removing users from the Veritas Enterprise Administrator

You may need to restrict access to the VEA authorization database (AZDB). You can remove users or user groups from the AZDB if they have been previously added. Use the commands in the table to remove users.

| | |
|---|---|
| <code>-r</code> | Removes a user or group to the registry. |
| <code>user -n user_name</code> | Specifies the user to be added. |
| <code>group -n group_name</code> | Specifies the group to be added. |
| <code>-A</code> | Grants the user <code>root</code> access, making all SF DB2, Veritas File System, and Veritas Volume Manager objects accessible. If the <code>-A</code> is not specified, the user or group will not be able to perform any Veritas Volume Manager or Veritas File System functions, and will not be able to access Storage Foundation for DB2. |
| | Allows the user access to the GUI if the user's primary group is different than the group owner. |
| <code>-h fully_qualified_host_name</code> | Specifies a fully qualified host name on which you want to add a user. For a stand alone installation, the command is run on the local host and the host name is optional. For a centrally managed server installation, you must be logged in as <code>root</code> and specify the host name. |
| <code>-d domain_name</code> | Specifies the domain to which the user belongs. By default the domain is the host name specified with <code>-h</code> . |
| <code>-t domain_type</code> | Specifies the type of domain to which the user belongs. Valid values are <code>nis</code> , <code>nisplus</code> , <code>Idap</code> , <code>unixpwd</code> , and <code>gssapi</code> . By default the domain type is assumed to be a UNIX login. |

You cannot remove `root` from the VEA Service console registry.

To remove users other than root from the Veritas Enterprise Administrator AZDB

- 1** Make sure that the optional GUI package was installed.

| OS | Command |
|-----------|---------------------------|
| AIX | lspp -L VRTSd2gui |
| Linux | rpm -qi VRTSd2gui -common |
| Solaris | pkginfo -l VRTSd2orgui |

- 2** Check if VxSVC is online. Bring it online if it is offline.

3 To remove a user, use the `vxdb2edusr` command as follows.

```
# /opt/VRTS/bin/vxdb2edusr -r user -n user_name  
[-h fully_qualified_host_namee -d domain_name -t domain_type]
```

See the table for command usage details. For example:

To remove the user “db2inst1” from the local host, enter the following:

```
# /opt/VRTS/bin/vxdb2edusr -r user -n db2inst1
```

To remove a UNIX login “db2inst1” from the local host, enter:

```
# /opt/VRTS/bin/vxdb2edusr -r user -n db2inst1 -h test.example.com
```

To remove a NIS login “db2inst1” in domain example.com from the centrally managed server install of VEA enter:

```
# /opt/VRTS/bin/vxdb2edusr -r user -n db2inst1 -h test.example.com -d e
```

4 To remove a group use the command.

```
# /opt/VRTS/bin/vxdb2edusr -r group -n group_name
```

See the table for command usage details. For example:

To remove the group “dba” from the local host, enter:

```
# /opt/VRTS/bin/vxdb2edusr -r group -n dba
```

To remove UNIX group "dba" from the local host,:

```
# /opt/VRTS/bin/vxdb2edusr -r group -A -f -n dba -h test.example.com
```

To remove NIS group "dba" from the centrally managed server installation of VEA:

```
# /opt/VRTS/bin/vxdb2edusr -r group -A -f -n dba -h test.example.com -d
```

Starting the Veritas Enterprise Administrator Service

The DBED agent and VEA Service must be running before you access the Veritas Storage Foundation for DB2 GUI. You may need to start VEA manually.

To manually start Veritas Enterprise Administrator Service

1 Verify the status of the VEA Service:

```
# /opt/VRTSob/bin/vxsvcctl status  
  
Current state of server : NOT RUNNING
```

2 Start the VEA Service:

```
# /opt/VRTSob/bin/vxsvcctl start  
  
Initializing Storage Foundation Provider 5.0 for DB2
```

If you are using AIX, press to continue after the output is displayed.

3 Again verify the status of the VEA Service:

```
# /opt/VRTSob/bin/vxsvcctl status  
  
Current state of server : RUNNING
```

4 If the DB2 instances or databases do not show up in the GUI after starting or restarting the VEA service, you may need to run the `db2ed_update`(for DB2 EE) or `db2ed_update_all`(for DB2 EEE) command. The databases should be displayed after a few seconds.

5 To use the Web GUI, start the Web Console: for the

```
/opt/VRTSob/webgui/veaw start
```

Shutting down the Veritas Enterprise Administrator Service

Although the VEA Service should remain running, you may need to shut it down manually.

To manually shut down Veritas Enterprise administrator Service

1 Verify the status of the VEA Service:

```
# /opt/VRTSob/bin/vxsvcctrl status  
  
Current state of server : RUNNING
```

2 Stop the VEA Service:

```
# /opt/VRTSob/bin/vxsvcctrl stop  
  
Successfully unloaded the Storage Foundation Provider 5.0 for DB2
```

3 Again verify the status of the VEA Service:

```
# /opt/VRTSob/bin/VxSVCctrl status  
  
Current state of server : NOT RUNNING
```

4 To stop the Web Console: for the

```
/opt/VRTSob/webgui/veaw stop
```

Starting the DBED agent

The DBED agent should be running before you access the Veritas Storage Foundation for DB2 GUI. You may need to start the DBED agent manually.

To start the DBED agent

1 Verify if the DBED Agent is running:

```
/opt/VRTSobc/pal33/bin/vxpal -a DBEDAgent -c status  
  
Agent State: RUNNING
```

2 Start the DBED Agent

```
/etc/rc2.d/S75vxpal.DBEDAgent /opt/VRTS/orgui/config/startup.sh start
```

Stopping the DBED agent

Although the DBED agent should remain running, you may need to shut it down manually.

To stop the DBED agent

- ◆ Stop the DBED Agent:

```
/etc/rc2.d/S75vxpal.DBEDAgent /opt/VRTS/orgui/config/startup.sh stop
```

Opening and closing the Veritas Storage Foundation for DB2 GUI

You can run the GUI from a Windows or UNIX client machine.

For the Java GUI, you must have the client software installed before you can use the Java GUI.

For the Web GUI, supported browsers are:

| | |
|---------|--|
| Solaris | Mozilla 1.6 |
| HP-UX | Mozilla 1.6 |
| AIX | Mozilla 1.6 |
| Linux | Netscape 7.2 or greater, Firefox 1.07 or greater, Mozilla 1.6 |
| Windows | Internet Explorer 5.5, Netscape 7.2 or greater, Firefox 1.07 or greater, Mozilla 1.6 |

Opening the Veritas Storage Foundation for DB2 Java GUI from a UNIX client

The VEA Service and DBED agent must be running on the server. If you need to start the VEA Service or DBED agent:

See [“Starting the Veritas Enterprise Administrator Service”](#) on page 21.

See [“Starting the DBED agent”](#) on page 23.

To start the Veritas Storage Foundation for DB2 Java GUI from a UNIX client

- 1 From an open terminal window, type `/opt/VRTSob/bin/vea` and press **Enter**.
- 2 In the VEA window, select **Connect to a Host or Domain**.
- 3 In the Connection pop-up window, enter the host name for the server to which you are connecting and click **Connect to Host or Domain**.
- 4 Enter your login name and password. Then click **OK**.

Opening the Veritas Storage Foundation for DB2 Java GUI from a Windows client

The VEA Service and DBED agent must be running on the server. If you need to start the VEA Service or DBED agent:

See “[Starting the Veritas Enterprise Administrator Service](#)” on page 21.

See “[Starting the DBED agent](#)” on page 23.

To start the Veritas Storage Foundation for DB2 Java GUI from a Windows client

- 1 Click **Start**, then select **Programs > Symantec > Veritas Enterprise Administrator**.
- 2 In the VEA window, select **Connect to a Host or Domain**.
- 3 In the Connection pop-up window, enter the host name for the server to which you are connecting and click **Connect**.
- 4 Enter your login name and password. Then click **OK**.

Opening the Veritas Storage Foundation for DB2 Web GUI

Starting the Web GUI is the same for a Windows or UNIX client.

To start the Veritas Storage Foundation for DB2 Web GUI

- ◆ In a supported browser, navigate to the host where the Web Console is online:

```
http://<host_name>:8181/VEAWeb/Login
```

Closing the Veritas Storage Foundation for DB2 Java GUI

Stopping the Java GUI is the same for a Windows or UNIX client.

To close the Veritas Storage Foundation for DB2 Java GUI

- 1 From the menu bar, select **File > Exit**.
- 2 The GUI displays a message indicating that you will be disconnected from the host if you continue. Click **Yes** to continue or, click **No** to keep the GUI running.

Closing the Veritas Storage Foundation for DB2 Web GUI

Stopping the Java GUI is the same for a Windows or UNIX client.

To start the Veritas Storage Foundation for DB2 Web GUI

- ◆ In the Web GUI upper right corner, click **Logout**.

Performing tasks in the Veritas Storage Foundation for DB2 Java GUI

You can perform the following tasks in the Veritas Storage Foundation for DB2 Java GUI.

- Managing your database
- Using a clone database
- Managing Storage Checkpoints
- Managing snapshots
- Using tiered storage
- Automating tasks
- Maintaining your configuration

Accessing database management tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI for setting up a database.

- Starting a DB2 instance
- Creating a DB2 snapshot database instance
- Unmounting a DB2 snapshot database
- Shutting down a DB2 instance
- Restarting a DB2 instance

To access DB2 database wizards in the Java GUI

- 1 Click the **DB2 Instances** icon in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Select one of the following methods to start the instance.
 - From the menu bar, select **DB2**.
 - Right click the **DB2 Instances** icon to display a pop-up menu.

Select the appropriate wizard from the list presented.

See [“Starting a DB2 instance”](#) on page 48.

See [“Creating a DB2 snapshot database”](#) on page 48.

See [“Unmounting a DB2 snapshot database”](#) on page 49.

See [“Shutting down a DB2 instance”](#) on page 50.

See [“Restarting a DB2 instance”](#) on page 51.

Accessing clone database tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI for:

- Creating a clone database
- Removing a clone database

To access the Create Clone Database wizard in the Java GUI

- 1 Click the database in the object tree.
- 2 Select one of the following methods to access the Create Clone Database wizard:
 - In the menu bar, click **DB2 > Create Clone Database**.
 - Right click the database object to display a pop-up menu. Then, click **Create Clone Database**.
In the Create Clone Database wizard, click **Next**.
- 3 Information about the current database instance displays for verification. Select Storage Checkpoint or Database FlashSnap for creating a clone database and click

See [“Creating a clone database”](#) on page 53.

To access the Remove Clone Database wizard in the Java GUI

- 1 In the object tree, click a clone database that was created using a Storage Checkpoint. You may need to expand the object tree to find the clone database.
- 2 Select one of the following methods to remove the clone database:
 - From the menu bar, click **DB2 > Remove Clone Database**.
 - Right click the database object to display a pop-up menu. Then, click **Remove Clone Database**.

See [“Removing a clone database”](#) on page 61.

Accessing Storage Checkpoint tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI for:

- Creating a Storage Checkpoint
- Viewing a Storage Checkpoint
- Mounting a Storage Checkpoint
- Unmounting a Storage Checkpoint

- Removing a Storage Checkpoint
- Rolling back a Storage Checkpoint

To access the Storage Checkpoint wizards in the Java GUI

- 1 Click the **Storage Checkpoints** icon in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Select one of the following methods to access the Create Storage Checkpoint wizard.
 - From the menu bar, select **Storage Checkpoints**.
or
 - Right-click the **Storage Checkpoints** icon to display a pop-up menu.
- 3 Select the appropriate Storage Checkpoint wizard.
 - See [“Creating a Storage Checkpoint”](#) on page 66.
 - See [“Viewing Storage Checkpoint details”](#) on page 67.
 - See [“Mounting a Storage Checkpoint”](#) on page 67.
 - See [“Unmounting a Storage Checkpoint”](#) on page 68.
 - See [“Removing a Storage Checkpoint”](#) on page 69.
 - See [“Rolling back to a Storage Checkpoint”](#) on page 69.

Accessing FlashSnap tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI for:

- Creating a snapplan
- Modifying or validating a snapplan
- Removing a snapplan
- Creating a snapshot
- Creating a clone using a snapshot
- Resynchronizing a snapshot to a database
- Resynchronizing a database to a snapshot
- Aborting a reverse resynchronization
- viewing a snapplan log

To access the snapplan wizards in the Java GUI

- 1 Click the **Snapplans** icon in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Select one of the following methods to access the Create Snapshot wizard.
 - From the menu bar, select **Snapplans**.
 - Right click the Snapshot Plans icon to display a pop-up menu.
- 3 Select a wizard option.
 - See [“Creating a snapplan”](#) on page 74.
 - See [“Validating or modifying a snapplan”](#) on page 76.
 - See [“Removing a snapplan”](#) on page 78.
 - See [“Resynchronizing a snapshot to a database”](#) on page 81.

To access the snapshot wizards in the Java GUI

- 1 Click a snapplan in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Select one of the following methods to access the snapshot wizards:
 - From the menu bar, select **DB2**.
 - Right click the **DB2 instances** icon to display a pop-up menu.
- 3 Select the appropriate wizard option.

Accessing tiered storage tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI for:

- Setting up parameters
- Setting policy for automated movement
- Managing file and volume movement
- Getting reports

To access tiered storage parameters wizards in the Java GUI

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.
- 2 Select one of the following methods to access the Database Dynamic Storage Tiering parameters wizards:
 - In the menu bar, click **DB2 > Database Dynamic Storage Tiering Management**.

- Right click the database object to display a pop-up menu. Then, click **Database Dynamic Storage Tiering Management**.

Select the appropriate wizard from the list presented. The wizard displays.

See “[To set up Database Dynamic Storage Tiering parameters](#)” on page 88.

See “[To modify Database Dynamic Storage Tiering parameters](#)” on page 89.

See “[To classify volumes](#)” on page 91.

To access the tiered storage policy wizards in the Java GUI

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.

- 2 Select one of the following methods to access the Database Dynamic Storage Tiering policy wizard:

- In the menu bar, click **DB2 > Database Dynamic Storage Tiering Management**.

- Right click the database object to display a pop-up menu. Then, click **Database Dynamic Storage Tiering Management**.

Select the appropriate wizard from the list presented. The wizard displays.

See “[Listing policies](#)” on page 97.

See “[Presetting policy for a directory](#)” on page 92.

See “[To preset policy for SMS tablespace](#)” on page 93.

See “[To preset policy for automatic storage path](#)” on page 93.

See “[To set policy for moving archive logs](#)” on page 94.

See “[To set policy for moving external files](#)” on page 95.

To access the tiered storage movement wizards in the Java GUI

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.

- 2 Select one of the following methods to access the Database Dynamic Storage Tiering move wizards:

- In the menu bar, click **DB2 > Database Dynamic Storage Tiering Management**.

- Right click the database object to display a pop-up menu. Then, click **Database Dynamic Storage Tiering Management**.

Select the appropriate wizard from the list presented. The wizard displays.

See “[To move tablespaces](#)” on page 96.

See [“To move containers”](#) on page 96.

To access the tiered storage report wizards in the Java GUI

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.
- 2 Select one of the following methods to access the Database Dynamic Storage Tiering report wizards:
 - In the menu bar, click **DB2 > Database Dynamic Storage Tiering Management**.
 - Right click the database object to display a pop-up menu. Then, click **Database Dynamic Storage Tiering Management**.

Select the appropriate wizard from the list presented. The wizard displays.

See [“To show activities”](#) on page 97.

See [“To show volume usage”](#) on page 98.

See [“To create a datafiles Statistics chart”](#) on page 99.

Accessing scheduled tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI to automate tasks for:

- Adding a task to the scheduler
- Modifying a scheduled task
- Deleting a scheduled task
- Viewing scheduled tasks
- Automating creation of clone database using checkpoint
- Automating creation of a clone database using FlashSnap
- Setting up scheduling parameters

To access the scheduler wizard in the Java GUI

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.
- 2 To add an automated task using the scheduler, open the database object folder, right-click **Scheduled Tasks**, and select **Add Task**.

The scheduler wizard displays.

See [“To add a task to the scheduler”](#) on page 102.

- 3 Existing tasks are listed in the scheduler directory. Select the task in the scheduler directory to display it in the right pane or right-click it to display additional options.
- 4 To access modification options for existing tasks, select the task in the scheduler directory and right-click it. Select the appropriate option.
See [“To modify a task to the scheduler”](#) on page 103.
See [“Viewing scheduled jobs”](#) on page 102.
See [“Automating creation of a database checkpoint”](#) on page 104.
See [“Automating creation of a database snapshot”](#) on page 104.
See [“Automating cloning of a database using a checkpoint”](#) on page 105.

Accessing system maintenance tasks in the Java GUI

You can use the Veritas Storage Foundation for DB2 Java GUI for:

- Resynchronizing the SFDB repository
- Rescanning the instance
- Displaying and rescanning tablespace information
- Viewing Storage Mapping
- Viewing I/O statistics
- Converting regular containers to Quick I/O Files
- Converting Quick I/O files to regular containers
- Updating rescan intervals
- Checking your system configuration
- Saving the system configuration

Note: Quick I/O is not supported on Linux.

To access the maintenance options in the Java GUI

- 1 Click the DB2 database icon in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Select one of the following methods.
 - From the menu bar, select **DB2**.

- Right click the DB2 database icon to display a pop-up menu.
- 3 Select the appropriate maintenance option.
 - See “Resynchronizing the SFDB repository” on page 110.
 - See “Rescanning the instance” on page 110.
 - See “Displaying and rescanning tablespace information” on page 110.
 - See “Viewing Storage Mapping” on page 111.
 - See “Viewing I/O statistics” on page 111.
 - See ??? .
 - See “Converting regular containers to Quick I/O Files” on page 112.
 - See “Converting Quick I/O files to regular containers” on page 113.
 - See “Updating rescan intervals” on page 113.
 - See “Checking your system configuration” on page 114.
 - See “Saving the system configuration” on page 116.

To access the Storage Mapping topology in the Java GUI

- 1 Start the GUI and connect to the desired host.
- 2 Expand the DB2 Instances icon in the object tree and then expand the desired database.
- 3 Select one of the following methods to generate datafile statistics:
 - From the menu bar, select **Container > Topology**.
 - Right click the container to display a pop-up menu. Then click **Topology**.

To access the I/O statistics wizard in the Java GUI

- 1 Start the GUI and connect to the desired host.
- 2 Expand the DB2 Instances icon in the object tree and then expand the desired database.
- 3 Select one of the following methods to generate datafile statistics:
 - From the menu bar, select **Container > Statistics**.
 - Right click the container to display a pop-up menu. Then click **Statistics**. The **Statistics** page displays.

Performing tasks in the Storage Foundation for DB2 Web GUI

You can perform the following tasks in the Veritas Storage Foundation for DB2 Web GUI.

- Managing your database
- Using Storage Mapping
- Using a clone database
- Managing Storage Checkpoints
- Managing snapshots
- Maintaining your configuration

Tasks are accessible through drop down lists in the right pane, or under the database objects on the summary pages.

Accessing database management tasks in the Web GUI

You can use the Veritas Storage Foundation for DB2 Web GUI for setting up a database.

- Starting a DB2 instance
- Creating a DB2 snapshot database instance
- Unmounting a DB2 snapshot database
- Shutting down a DB2 instance
- Restarting a DB2 instance

To access DB2 database wizards in the Web GUI

- 1 Start the Web GUI and connect to the desired host.
- 2 In the Summary View, select Storage Foundation for DB2.

- 3 Click the DB2 instance.
- 4 Select the database task in the drop down list in the left pane to display the corresponding wizard.

See [“To start a DB2 instance”](#) on page 48.

See [“To create a DB2 snapshot database”](#) on page 48.

See [“To unmount a DB2 snapshot database”](#) on page 49.

See [“To shut down a DB2 instance”](#) on page 50.

See [“To restart a DB2 instance”](#) on page 51.

Accessing database tasks in the Web GUI

You can use the Veritas Storage Foundation for DB2 Web GUI for:

Clone tasks:

- Creating a clone database
- Removing a clone database

Storage Checkpoint tasks:

- Creating a Storage Checkpoint
- Viewing a Storage Checkpoint
- Mounting a Storage Checkpoint
- Unmounting a Storage Checkpoint
- removing a Storage Checkpoint
- Rolling back a Storage Checkpoint

Flashsnap tasks:

- Creating a snapplan
- Modifying or validating a snapplan
- Removing a snapplan
- Creating a snapshot
- Creating a clone using a snapshot
- Resynchronizing a snapshot to a database
- Resynchronizing a database to a snapshot
- Aborting a reverse resynchronization
- viewing a snapplan log

To access the wizards in the Web GUI

- 1 Start the Web GUI and connect to the desired host.
- 2 In the Summary View, select Storage Foundation for DB2.
- 3 Click the DB2 instance.
- 4 Select a database task in the clone drop down list to display the corresponding wizard.
 - See [“To create a clone database using a Storage Checkpoint”](#) on page 54.
 - See [“Creating a clone database using Database FlashSnap”](#) on page 55.
 - See [“To remove a Storage Checkpoint clone database”](#) on page 62.
 - See [“To remove a Storage Checkpoint clone database”](#) on page 62.
- 5 Select a database task in the Storage Checkpoint drop down list to display the corresponding wizard.
 - See [“Creating a Storage Checkpoint”](#) on page 66.
 - See [“Viewing Storage Checkpoint details”](#) on page 67.
 - See [“Mounting a Storage Checkpoint”](#) on page 67.
 - See [“Unmounting a Storage Checkpoint”](#) on page 68.
 - See [“Removing a Storage Checkpoint”](#) on page 69.
 - See [“Rolling back to a Storage Checkpoint”](#) on page 69.
- 6 Select a database task in the FlashSnap drop down list to display the corresponding wizard.
 - See [“Creating a snapplan”](#) on page 74.
 - See [“Validating or modifying a snapplan”](#) on page 76.
 - See [“Removing a snapplan”](#) on page 78.
 - See [“Creating a snapshot”](#) on page 79.
 - See [“Creating a clone database with a snapshot”](#) on page 80.
 - See [“Resynchronizing a snapshot to a database”](#) on page 81.
 - See [“Resynchronizing a database to a snapshot”](#) on page 82.
 - See [“Aborting the reverse resynchronization operation”](#) on page 83.
 - See [“Viewing the snapplan log”](#) on page 84.

Accessing system maintenance tasks in the Web GUI

You can use the Veritas Storage Foundation for DB2 Web GUI for:

- Resynchronizing the SFDB repository
- Rescanning the instance
- Displaying and rescanning tablespace information
- Viewing Storage Mapping
- Viewing I/O statistics
- Viewing DB2 container topology or statistics
- Converting regular containers to Quick I/O Files
- Converting Quick I/O files to regular containers
- Updating rescan intervals
- Checking your system configuration
- Saving the system configuration

Note: Quick I/O is not supported on Linux.

To access the system maintenance tasks in the Web GUI

- 1 Start the Web GUI and connect to the desired host.
- 2 In the Summary View, select Storage Foundation for DB2.
- 3 Click the DB2 instance.
- 4 Select a database task in the drop down list to display the corresponding wizard.
 - See [“Resynchronizing the SFDB repository”](#) on page 110.
 - See [“Rescanning the instance”](#) on page 110.
 - See [“Displaying and rescanning tablespace information”](#) on page 110.
 - See [“Viewing Storage Mapping”](#) on page 111.
 - See [“Viewing I/O statistics”](#) on page 111.
 - See ??? .
 - See [“Converting regular containers to Quick I/O Files”](#) on page 112.
 - See [“Converting Quick I/O files to regular containers”](#) on page 113.
 - See [“Updating rescan intervals”](#) on page 113.
 - See [“Checking your system configuration”](#) on page 114.
 - See [“Saving the system configuration”](#) on page 116.

Managing your database

This chapter includes the following topics:

- [Database management tasks in the Veritas Storage Foundation for DB2 GUI](#)
- [About setting up a disk group](#)
- [About selecting a volume layout](#)
- [File system creation guidelines](#)
- [Controlling fragmentation](#)
- [Starting a DB2 instance](#)
- [Creating a DB2 snapshot database](#)
- [Unmounting a DB2 snapshot database](#)
- [Shutting down a DB2 instance](#)
- [Restarting a DB2 instance](#)

Database management tasks in the Veritas Storage Foundation for DB2 GUI

If you are using Veritas Storage Foundation for DB2 to set up a new database, complete these tasks in the order listed below:

Determine the number and sizes of file systems you need for the database you want to create. See the *Veritas File System Administrator's Guide*.

- Create volumes to meet your file system needs. You can use disk mirroring as a safeguard against disk failures and striping for better performance. See [“About selecting a volume layout”](#) on page 42.
- If you plan to create volume snapshots for the database and use them on either the same host or a secondary one, ensure that your volume layout is consistent with Database FlashSnap requirements. See [“Database FlashSnap in the GUI”](#) on page 73.
- Create the VxFS file systems you need on the volumes. See [“File system creation guidelines”](#) on page 44.
- Install and configure your database. For best OLTP performance, use Quick I/O. You must create Quick I/O files before creating the tablespaces. See the *Veritas Storage Foundation for DB2 Administrator's Guide*.
- Note:** The Quick I/O feature is not available on the Linux operating system.
- If you would like the ability to view detailed storage stack topology information to ensure your storage stack configuration is optimized for the database, configure and use Storage Mapping. See [“Viewing Storage Mapping”](#) on page 111.
- If you are using Quick I/O, convert all database files to Quick I/O files. See [“Converting regular containers to Quick I/O Files”](#) on page 112.
- If using ODM, link the ODM library. See the *Veritas Storage Foundation for DB2 Administrator's Guide*.
- If you are not currently running on VxVM and VxFS, make sure Veritas Storage Foundation for DB2 is installed and convert your existing database configuration. See the *Veritas Storage Foundation for Oracle DB2Installation Guide*.
See the *Veritas Storage Foundation for DB2 Administrator's Guide*.

For backup and recovery on the same host, you can use the Storage Checkpoint facility to create file system snapshots of the database. A Storage Checkpoint creates an exact image of a database instantly and provides a consistent image of the database from the point in time the Storage Checkpoint was created. See [“Creating a Storage Checkpoint”](#) on page 66.

For off-host processing or backup, you can use the Database FlashSnap feature to create a volume snapshot for the database. Database FlashSnap lets you capture an online image of an actively changing database at a given instant, known as a snapshot. You can perform backups and off-host processing tasks on snapshots while providing continuous availability of your critical data. See [“Database FlashSnap in the GUI”](#) on page 73.

About setting up a disk group

Before creating file systems for a database, set up a disk group for each database.

A disk group lets you group disks, volumes, file systems, and files that are relevant to a single database into a logical collection for easy administration. Because you can move a disk group and its components as a unit from one machine to another, you can move an entire database when all the configuration objects of the database are in one disk group. This capability is useful in a failover situation.

Disk group configuration guidelines

Follow these guidelines when setting up disk groups:

- Only disks that are online and do not already belong to a disk group can be used to create a new disk group.
- Create one disk group for each database.
- The disk group name must be unique. Name each disk group using the DB2 database name specified by the environment variable `$DB2DATABASE` and a `dg` suffix. The `dg` suffix helps identify the object as a disk group. Also, each disk name must be unique within the disk group.

Note: Users should not share a disk group between different DB2 instances. Although it is not recommended, sharing a disk group among all databases in the same instance may make sense if the instance contains several small databases. In this case, name the disk group using the DB2 instance name specified by the environment variable `$DB2INSTANCE` and a `dg` suffix.

- Never create container files using file systems or volumes that are not in the same disk group.

In earlier releases of Veritas Volume Manager, a system installed with VxVM was configured with a default disk group, `rootdg`, that had to contain at least one disk. VxVM can now function without any disk group having been configured. Only when the first disk is placed under VxVM control must a disk group be configured.

Note: Most VxVM commands require superuser or equivalent privileges.

See the “Tuning for Performance” chapter of the *Veritas Storage Foundation for DB2 Administrator's Guide*.

For more about disk groups and disk group procedures, see the *Veritas Enterprise Administrator Guide*.

For more about disk groups and disk group procedures, see the *Volume Manager Web GUI Online Help*.

About selecting a volume layout

Veritas Volume Manager offers a variety of layouts that allow you to configure your database to meet performance and availability requirements. The proper selection of volume layouts provides optimal performance for the database workload.

An important factor in database performance is the tablespace placement on the disks.

Disk I/O is one of the most important determining factors of your database's performance. Having a balanced I/O load usually means optimal performance. Designing a disk layout for the database objects to achieve balanced I/O is a crucial step in configuring a database.

When deciding where to place tablespaces, it is often difficult to anticipate future usage patterns. VxVM provides flexibility in configuring storage for the initial database set up and for continual database performance improvement as needs change. VxVM can split volumes across multiple drives to provide a finer level of

granularity in data placement. By using striped volumes, I/O can be balanced across multiple disk drives. For most databases, ensuring that different containers or tablespaces, depending on your database, are distributed across the available disks may be sufficient.

Striping also helps sequential table scan performance. When a table is striped across multiple physical devices, a high transfer bandwidth can be achieved by closely matching several DB2 parameters to ensure that database extents match the stripe size for the device. Another very important consideration when using the DB2 database, which by default performs striping at the tablespace container level, is setting the `DB2_STRIPED_CONTAINERS` variable.

See “Tuning for Performance” in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

How to choose appropriate stripe unit sizes

When creating a striped volume, you need to decide the number of columns to form a striped volume and the stripe unit size. You also need to decide how to stripe the volume. You may stripe a volume across multiple disk drives on the same controller or across multiple disks on multiple controllers. By striping across multiple controllers, disk I/O can be balanced across multiple I/O channels. The decision is based on the disk and controller bandwidth and the database workload. In general, for most OLTP databases, use the default stripe unit size of 64 K or smaller for striped volumes and 16 K for RAID-5 volumes.

How to choose between mirroring and RAID-5

VxVM provides two volume configuration strategies for data redundancy: mirroring and RAID-5. Both strategies allow continuous access to data in the event of disk failure. For most database configurations, we recommend using mirrored, striped volumes. If hardware cost is a significant concern, but having higher data availability is still important, use RAID-5 volumes.

RAID-5 configurations have certain performance implications you must consider. Writes to RAID-5 volumes require parity-bit recalculation, which adds significant I/O and CPU overhead. This overhead can cause considerable performance penalties in online transaction processing (OLTP) workloads. If the database has a high read ratio, however, RAID-5 performance is similar to that of a striped volume.

Volume configuration guidelines

Follow these guidelines when selecting volume layouts:

- Put the database log files on a file system created on a striped and mirrored (RAID-0+1) volume separate from the index or data tablespaces. Stripe multiple devices to create larger volumes if needed. Use mirroring to improve reliability. Do not use VxVM RAID-5 for redo logs.
- When normal system availability is acceptable, put the tablespaces on file systems created on striped volumes for most OLTP workloads.
- Create striped volumes across at least four disks. Try to stripe across disk controllers. For sequential scans, ensure that the `NUM_IOSEEVERS` and `DB2_PARALLEL_IO` settings are tuned to match the number of disk devices used in the stripe.
- For most workloads, use the default 64 K stripe-unit size for striped volumes and 16 K for RAID-5 volumes.
- When system availability is critical, use mirroring for most write-intensive OLTP workloads. Turn on Dirty Region Logging (DRL) to allow fast volume resynchronization in the event of a system crash.
- When system availability is critical, use RAID-5 for read-intensive OLTP workloads to improve database performance and availability. Use RAID-5 logs to allow fast volume resynchronization in the event of a system crash.
- For most decision support system (DSS) workloads, where sequential scans are common, experiment with different striping strategies and stripe-unit sizes. Put the most frequently accessed tables or tables that are accessed together on separate striped volumes to improve the bandwidth of data transfer.

See “Tuning for Performance” in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

For more about volumes and volume procedures, see the *Veritas Enterprise Administrator Guide*.

For more about volumes and volume procedures, see the *Volume Manager Web GUI Online Help*.

File system creation guidelines

Follow these guidelines when creating VxFS file systems:

- Specify the maximum log size when creating file systems for databases. In a Solaris environment, also specify the maximum block size.
- Except for specifying the maximum log size and support for large files as required, use the VxFS defaults when creating file systems for databases.
- Never disable the intent logging feature of the file system.

- For database logs, create a single file system using a simple (and mirrored, if necessary) volume. Put the other tablespaces and database files on separate file systems created on striped, striped and mirrored, or RAID-5 volumes.
- When using the command line, use the mount points to name the underlying volumes. For example, if a file system named `/db01` is to be created on a mirrored volume, name the volume `db01` and the mirrors `db01-01` and `db01-02` to relate to the configuration objects. If you are using the `vxassist` command or the GUI, this is transparent.

Choose a file system block size that matches or is a multiple of the `PAGESIZE` parameter in the `create database` or `create tablespace` statement for your DB2 database or tablespace. The `PAGESIZE` parameter is defined in the `create database` or `create tablespace` statement.

It is possible to have a file system block size that is smaller than the database page size because the database page-size limit can be bigger than the file system block size. It is fine if the file system block size is smaller than the database page size because VxFS will not perform multiple I/O operations for each database I/O operation. VxFS is capable of performing I/Os with multiple blocks. For example, if your database page size is 8k and your file system block size is 4K, VxFS can put two 4k blocks together to perform one 8k database I/O operation. The DB2 instance will also need data in `EXTENTSIZE`, which is a multiple of `PAGESIZE`. These page size rules also apply for extent size.

When creating the file system, set the number of file system blocks in the intent log so that the log size is 16MB. For example, if the file system block size is 8K (that is, 8192), it will take 2000 blocks to make a 16MB log ($2000 \times 8192 = \sim 16\text{MB}$). If the file system block size is 4K (that is, 4096), then twice as many blocks as in the 8K case would need to be allocated (4000 in this example).

For more about file systems and file system procedures, see the *Veritas Enterprise Administrator Guide*.

For more about file systems and file system procedures, see the *Volume Manager Web GUI Online Help*.

Controlling fragmentation

When free resources are initially allocated to files in a Veritas file system, they are aligned in the most efficient order possible to provide optimal performance. On an active file system, the original order is lost over time as files are created, removed, or resized. As space is allocated and deallocated from files, the available free space becomes broken into fragments. This means that space must be assigned to files in smaller and smaller extents. This process is known as fragmentation.

Fragmentation leads to degraded performance and availability. The degree of fragmentation depends on file system usage and activity.

Fragmentation utilities in VxFS

Allocation units in VxFS are designed to help minimize and control fragmentation. Over time, however, file systems eventually become fragmented.

VxFS provides online reporting and optimization utilities to enable you to monitor and defragment a mounted file system. These utilities are accessible through the file system administration command, `fsadm`. Using the `fsadm` command, you can track and eliminate fragmentation without interrupting user access to the file system.

Types of fragmentation

VxFS addresses two types of fragmentation:

- **Directory fragmentation**
As files are created and removed, gaps are left in directory inodes. This is known as directory fragmentation. Directory fragmentation causes directory lookups to become slower.
- **Extent fragmentation**
As files are created and removed, the free extent map for an allocation unit changes from having one large free area to having many smaller free areas. Extent fragmentation occurs when files cannot be allocated in contiguous chunks and more extents must be referenced to access a file. In a case of extreme fragmentation, a file system may have free space that cannot be allocated.

Defragmenting a file system

You can use the online administration utility `fsadm` to defragment or reorganize file system directories and extents.

The `fsadm` utility defragments a file system mounted for read/write access by:

- Removing unused space from directories.
- Making all small files contiguous.
- Consolidating free blocks for file system.

Use the following options with the `fsadm` utility:

- | | |
|--------------|---|
| Requirements | <ul style="list-style-type: none"> ■ To reorganize a file system using the <code>fsadm</code> command, you must have superuser (root) privileges . |
| Usage notes | <ul style="list-style-type: none"> ■ If you specify <code>-d</code> and <code>-e</code>, directory reorganization is always completed first. ■ If you use both <code>-D</code> and <code>-E</code> with the <code>-d</code> and <code>-e</code> options, the fragmentation reports are produced both before and after reorganization. ■ The <code>-t</code> and <code>-p</code> options control the amount of work performed by <code>fsadm</code>, either in a specified time or by a number of passes. By default, <code>fsadm</code> runs five passes. If both <code>-t</code> and <code>-p</code> are specified, <code>fsadm</code> exits if either of the terminating conditions are reached. |
| Options | <ul style="list-style-type: none"> ■ <code>-d</code>: Reorganizes directories. Directory entries are reordered to place subdirectory entries first, then all other entries in decreasing order of time of last access. The directory is also compacted to remove free space. ■ <code>-a</code>: Use in conjunction with the <code>-d</code> option to consider files not accessed within the specified number of days as “aged” files. Aged files are moved to the end of the directory. The default is 14 days. ■ <code>-e</code>: Reorganizes extents. Files are reorganized to have the minimum number of extents. ■ <code>-D -E</code>: Produces reports on directory and extent fragmentation, respectively. ■ <code>-v</code>: Specifies verbose mode and reports reorganization activity. ■ <code>-l</code>: Specifies the size of a file that is considered large. The default is 64 blocks. ■ <code>-t</code>: Specifies a maximum length of time to run, in seconds. ■ <code>-p</code>: Specifies a maximum number of passes to run. The default is five. ■ <code>-s</code>: Prints a summary of activity at the end of each pass. ■ <code>-r</code>: Specifies the pathname of the raw device to read to determine file layout and fragmentation. This option is used when <code>fsadm</code> cannot determine the raw device. |

To defragment a file system on a volume

- 1 Select the volume containing the file system to be defragmented.
- 2 Choose **Actions > File System > Defrag File System**.
 See “[Accessing database management tasks in the Java GUI](#)” on page 26.
 See “[Accessing database management tasks in the Web GUI](#)” on page 34.
- 3 Select **Yes** in the displayed dialog box.

Starting a DB2 instance

You can start a DB2 instance from the GUI. You must know the password of the instance owner. The steps to restart a DB2 database instance are the same.

To start a DB2 instance

- 1 Access the DB2 Start Instance wizard.
See [“To access DB2 database wizards in the Java GUI”](#) on page 26.
See [“Accessing database management tasks in the Web GUI”](#) on page 34.
- 2 Enter the instance name and home directory. Verify your UNIX user name and enter your password. The UNIX user name is the same as the instance name.
- 3 Click .
- 4 At the confirmation prompt, click **Yes** to confirm that you want to start the database.

If the DB2 instance starts successfully, you will receive a confirmation message. Click **OK** to continue.

Creating a DB2 snapshot database

If you created a snapshot using the Database FlashSnap feature, you can create a DB2 snapshot database from the secondary host. This is the equivalent of cloning the database from the primary host. You must know the password of the instance owner and there must be an existing snapshot.

See [“Database FlashSnap in the GUI”](#) on page 73.

- Prerequisites
- Make sure you have enough space to create a clone database on your system.
 - You must have an existing snapshot.

To create a DB2 snapshot database

- 1 Access the DB2 Create Snapshot wizard.
See [“To access DB2 database wizards in the Java GUI”](#) on page 26.
See [“Accessing database management tasks in the Web GUI”](#) on page 34.
- 2 In the Authentication section, enter the password for the UNIX user name. The user name is a read-only field.

- 3 For the “Primary database information” section, enter the DB2 Instance and DB2 Database information.
- 4 Create or restart the database by clicking the appropriate radio button:
 - Select **Startup database** to start a new snapshot database based on a snapshot.
 - Select **Restart database** to restart an existing snapshot database.
- 5 Enter the snapshot database information in the appropriate fields:
 - The New DB2 Database.
 - Snapplan file.
 - The volume name field is required to start an off-host snapshot database.
 - Snapshot disk group name
 - The Relocate path is required to start an on-host snapshot database.
 - You can obtain this information by viewing the details after you create a snapshot. If you did not retain the information from the snapshot creation, you can use the `db2ed_vmchecksnap` command or use View Log via the GUI to retrieve the information.
See [“Creating a snapplan”](#) on page 74.
- 6 Click **Start Database**.
- 7 At the confirmation prompt, click **Yes** to confirm that you want to start the database.

If the snapshot database was successfully started, you will receive a confirmation message. Click **OK** to continue.

Unmounting a DB2 snapshot database

You can unmount a DB2 snapshot database using the GUI. For example, you are required to unmount the snapshot database before resynchronizing the snapshot with the primary database.

To unmount a DB2 snapshot database

- 1 Access the DB2 Unmount Snapshot Database wizard.
 - See [“To access DB2 database wizards in the Java GUI”](#) on page 26.
 - See [“Accessing database management tasks in the Web GUI”](#) on page 34.
- 2 In the Authentication section, enter the password for the UNIX user name. The user name field is read-only.

- 3 Enter the DB2 Database name in the Primary Database Information section.
- 4 In the Snapshot database section, enter the name of the new DB2 database, Snapplan file, and enter the relocate path.
- 5 Click **Unmount Database** at the bottom of the screen.
- 6 At the confirmation prompt, click **Yes** to confirm that you want to unmount the snapshot database.

If the snapshot database was successfully unmounted, you receive a confirmation message. Click **OK** to continue.

Shutting down a DB2 instance

The GUI enables you to shut down a DB2 instance. For example, you must shut down the database to perform a Storage Rollback of an entire database.

To shut down a DB2 instance

- 1 Access the DB2 Shutdown Instance wizard.
 - See [“To access DB2 database wizards in the Java GUI”](#) on page 26.
 - See [“Accessing database management tasks in the Web GUI”](#) on page 34.
- 2 Verify the DB2 instance information, such as the instance name, home directory, and instance owner. Then, click **Next** to continue.
- 3 Verify your UNIX user name, then enter your password in the **Password** field.
- 4 In the dialog box, select the type of shutdown you want to perform:
 - **Normal**

Use this option to shut down the DB2 instance in normal situations. If you shut down using this option, no new instance connections are allowed. DB2 waits for all currently connected users to disconnect from the instance, and then closes and dismounts the database before shutting down the instance. The next instance start up does not require an instance recovery.
 - **Force**

Use this option to forcibly shut down the DB2 database when there are existing connections to the database. When this operation is selected, no new database connections are allowed. DB2 immediately terminates all sessions and shuts down the database. The next database startup may require an instance recovery.

- 5 Click **Shutdown instance** at the bottom of the screen.
- 6 At the confirmation prompt, click **Yes** to confirm that you want to shut down the instance.

If the instance was successfully shut down, you will receive a confirmation message. Click **OK** to continue.

Restarting a DB2 instance

You can restart a DB2 instance from the GUI. You must have a UNIX user name and password.

To restart a DB2 instance

- 1 Access the DB2 Start Instance wizard.
See [“To access DB2 database wizards in the Java GUI”](#) on page 26.
See [“Accessing database management tasks in the Web GUI”](#) on page 34.
- 2 Verify the DB2 instance information, such as the instance name, home directory, and instance owner. Then, click **Next** to continue.
- 3 Enter your password then click **Start Instance**.
- 4 At the confirmation prompt, click **Yes** to confirm that you want to start the database.

If the DB2 instance was successfully restarted, you will receive a confirmation message. Click **OK** to continue.

Using a clone database

This chapter includes the following topics:

- [Creating a clone database](#)
- [Removing a clone database](#)

Creating a clone database

Cloning a database enables you to perform operations without affecting your production database. The Storage Foundation for DB2 GUI provides the Create Clone Database wizard to enable you to create a clone of your database. Choose the method to create the clone database:

- Storage Checkpoint
- Database FlashSnap

If storage is more important than performance, use Storage Checkpoint to create a clone database. A clone database created with Storage Checkpoint requires half as much storage as a clone database created with Database FlashSnap. However, over a longer period of time, the storage will slowly increase for a clone created with Storage Checkpoint.

If performance is important, use Database FlashSnap to create a clone database. A clone database created with Storage Checkpoint will slow the performance of the database.

Creating a clone database using Checkpoints

Use the Create Clone Database wizard to create a clone of your database using Storage Checkpoints for optimal storage.

To create a clone database using a Storage Checkpoint

- 1 Access the Create Clone Database wizard and select Storage Checkpoint as the method for creating a clone.

See [“Accessing clone database tasks in the Java GUI”](#) on page 27.

See [“Accessing database tasks in the Web GUI”](#) on page 35.

- 2 In the object tree, click a specific Storage Checkpoint in the Storage Checkpoints folder. You may need to expand the tree view to find the Storage Checkpoint.

- 3 Select one of the following methods to create a clone database:

- Click a Storage Checkpoint and click **Create Clone Database** in the menu bar.
- Right click the Storage Checkpoint object to bring up a pop-up menu. Then, click **Create Clone Database** .
The Create Clone Database wizard is displayed. Click **Next**.

- 4 The **Create Clone Database** screen displays with read-only information about the current database. Review and click **Next**.

- 5 The **New Database Information** dialog is displayed. Enter the target information. This information is the basis for your duplicated database. You need to enter or select the following:

- Instance name
- Password
- Database name
- Mount point
- Redo log directory (Optional. If you do not specify any options, minimum recovery will be performed.)

Check Restart Database if you want to restart a Storage Checkpoint Clone Database that is already created. Otherwise, you are creating a new Storage Checkpoint Clone Database.

The mount point in the target information is not required for restarting a Storage Checkpoint Clone Database.

- 6 Enter the source information. This is the information from your original database, which is the one that the Storage Checkpoint represents. You need to enter or verify the following:

- Instance name (This field is read-only.)
- Database name (This field is read-only.)

- Database log directory (By default, this field is populated. Modify it only if the information is incorrect.)

7 Click **Create Clone Database**.

If the database was successfully cloned, you will receive a confirmation message. Click **OK** to continue.

Creating a clone database using Database FlashSnap

Use the Create Clone Database wizard to create a clone of your database using Database FlashSnap for optimal performance.

The Create Clone Database wizard for using Database FlashSnap goes through the following dialogs:

- Create Snapplan
- Validate/Modify Snapplan
- Create Snapshot
- Startup Clone Database

Choose to create a clone database from a new or existing snapplan.

If you create a clone database using an existing snapplan, you use a portion of the Create Clone Database wizard depending on the status of the snapplan you select:

- `SNAPSHOT_END` status: picks up at Create Snapshot point in the process. This status uses the default snapplan settings.
- `INIT_FULL` status: picks up at Validate/Modify Snapplan. This status enables selection of forced snapshot creation, authentication settings, and primary database settings.
- `INIT_DB` status: picks up at Create Snapplan. This status enables the review and selection of all settings.

Creating a new snapplan for a clone database

If you are creating a clone database from a new snapplan, you must create the new snapplan.

To create a clone database with a new snapplan

- 1 Access the Create Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 Choose **Database FlashSnap**.

- 3 In the **Database FlashSnap Clone Database** dialog, select **New Snapplan**. Click **Next**.
The **Create Snapplan** dialog is displayed.
- 4 Confirm the Database Name. This is a read-only field.
- 5 Enter the following values:
 - Full path of the snapplan file
 - Name of the secondary host
 - Snapshot plex tag
- 6 Click **Next** to continue. The **Validate/Modify Snapplan** screen displays with default values set.
See [“To validate or modify a snapplan using the Create Clone Database wizard”](#) on page 57.
- 7 After the **Validate/Modify Snapplan** dialog, use the following dialog procedures to complete the clone database:
 - See [“To create a snapshot using the Create Clone Database wizard”](#) on page 58.
 - See [“To start up a snapshot using the Create Clone Database wizard”](#) on page 59.

Validating or modifying a snapplan for a clone database

After creating a new snapplan or selecting an existing snapplan, validate or modify it before proceeding to create a snapshot.

To validate or modify a snapplan using the Create Clone Database wizard

- 1 Access the Validate/Modify dialog of the Create Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 In the Validate/Modify screen, modify any incorrect settings. Click **Next** to validate the snapplan.

For more information regarding the snapplan parameters:

See *Creating and working with Snapplans using db2ed_vmchecksnap in the Veritas Storage Foundation for DB2 Administrator's Guide*.

If the snapplan was successfully validated, a confirmation message displays.

- 3 To see the snapplan details, click the **Show details** checkbox. The snapplan details are displayed in the pop-up window.

When you have finished reviewing them, click **OK**.

The **Create Snapshot** screen appears.

See [“To create a snapshot using the Create Clone Database wizard”](#) on page 58.

A sample of the snapplan details:

```
SNAPSHOT_MODE is online_snapshot
PRIMARY_HOST is host1
SECONDARY_HOST is host2
The version of PRIMARY_DG-DB2dg is 120.
SNAPSHOT_DG is SNAP_DB2dg
SNAPSHOT_PLAN_FOR is database
Examining DB2 volume and disk layout for snapshot.
Volume testvol01 on DB2dg is ready for snapshot.
Original plex and DCO log for testvol01 is on DB2dg01 DB2dg02.
Snapshot plex and DCO log for testvol01 is on DB2dg03.
Volume testvol02 on DB2dg is ready for snapshot.
Original plex and DCO log for testvol02 is on DB2dg01 DB2dg02.
Snapshot plex and DCO log for testvol02 is on DB2dg03.
Volume testvol03 on DB2dg is ready for snapshot.
Original plex and DCO log for testvol03 is on DB2dg01 DB2dg02.
Snapshot plex and DCO log for testvol03 is on DB2dg03.
Volume testvol04 on DB2dg is ready for snapshot.
Original plex and DCO log for testvol04 is on DB2dg01 DB2dg02.
Snapshot plex and DCO log for testvol04 is on DB2dg03.
Volume testvol05 on DB2dg is ready for snapshot.
Original plex and DCO log for testvol05 is on DB2dg01 DB2dg02.
```

```
Snapshot plex and DCO log for testvol105 is on DB2dg03.  
Volume udb_home on DB2dg is ready for snapshot.  
Original plex and DCO log for udb_home is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for udb_home is on DB2dg03.  
SNAP_DB2dg for snapshot will include: DB2dg03  
ALLOW_REVERSE_RESYNC is yes  
The snapplan sp1 has been created.
```

Creating snapshot for a clone database

After your snapplan is validated, you can create a snapshot for your clone database.

To create a snapshot using the Create Clone Database wizard

- 1 Access the Create Snapshot dialog of the Create Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 In the **Create Snapshot** dialog, verify the snapplan name.
- 3 Click **Next** to continue.
- 4 If the snapshot creation was successful, a confirmation message appears.

- 5 Click the **Show details** checkbox to view the actions taken.

A sample of the snapplan details:

```
db2ed_vmsnap started at 2004-06-08 11:41:39
DB20000I The SET WRITE command completed successfully.
DB20000I The SET WRITE command completed successfully.
A snapshot of DB2DATABASE UDB is in DG SNAP_DB2dg.
Snapplan spl is used for the snapshot.
```

If `-r <relocate_path>` is used in `db2ed_vmclonedb`, make sure `<relocate_path>` is created and owned by DB2 Instance Owner. Otherwise, the following mount points need to be created and owned by DB2 Instance Owner:

```
/db2/testvol01.
/db2/testvol02.
/db2/testvol03.
/db2/testvol04.
/db2/testvol05.
/db2/udb_home.
```

```
db2ed_vmsnap ended at 2004-06-08 11:42:14
```

- 6 Click **OK** to continue. You can now use the snapshot as a database backup. Save this information for when you start your snapshot database.
- 7 The Startup Snapshot dialog displays. proceed to the procedure for starting up a snapshot.
See [“To start up a snapshot using the Create Clone Database wizard”](#) on page 59.

Start up snapshot for a clone database

You can start up the clone after creating a clone snapshot.

To start up a snapshot using the Create Clone Database wizard

- 1 Access the Startup Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 For the primary database information section, verify the DB2 Instance and DB2 Database information.
- 3 Select a startup option by clicking the appropriate radio button:

- **Create database** if you are creating a new snapshot database
 - **Restart database** to restart an existing snapshot database.
- 4 For the Snapshot database information section, enter the New DB2 Database. Verify the Snapplan file, Snapshot disk group name, and Relocate path in the appropriate fields.

See creating and working with snapplans using `db2ed_vmchecksnap` in the *Veritas Storage Foundation for DB2 Administrator's Guide*.
 - 5 Click **Create Clone Database**.

If the snapshot database was successfully started, you will receive a confirmation message. Click **OK** to continue.

Create snapshot for clone database

If you are using an existing snapplan to create the clone database, some parameters are required.

To create a clone database with an existing snapplan using default settings

- 1 Access the Create Clone Database wizard.

See [“Accessing clone database tasks in the Java GUI”](#) on page 27.

See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 In the **Create Clone Database** dialog, select create a clone database using Database FlashSnap. Click **Next**.
- 3 In the **Database FlashSnap Clone Database** dialog, select **Existing Snapplan**. The snapplan displays its status. For example, `sp00 [init_db]` `/ora_home/snapplan/sp00`, where `INIT_DB` is the status of the snapplan.
- 4 Select a snapplan from the drop-down list `SNAPSHOT_END` status. The **Startup Snapshot Database** dialog displays.
- 5 Use the dialog procedure for Startup Clone Database:

See [“To start up a snapshot using the Create Clone Database wizard”](#) on page 59.

To create a clone database with an existing snapplan using `INIT_FULL`

- 1 Access the Create Clone Database wizard.

See [“Accessing clone database tasks in the Java GUI”](#) on page 27.

See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 In the **Create Clone Database** dialog, select **Existing Snapplan**. The snapplan displays its status. For example, `sp00 [init_full]` `/ora_home/snapplan/sp00`, where `INIT_FULL` is the status of the snapplan.

- 3 Select a snapplan from the drop-down list with `INIT_FULL` status. The **Create Snapshot** dialog displays.
- 4 Use the following dialog procedures in order:
 - See [“To create a snapshot using the Create Clone Database wizard”](#) on page 58.
 - See [“To start up a snapshot using the Create Clone Database wizard”](#) on page 59.

To create a clone database with an existing snapplan using `INIT_DB`

- 1 Access the Create Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 In the **Create Clone Database** dialog, select **Existing Snapplan**. The snapplan displays its status. For example, `sp00 [init_db] /ora_home/snapplan/sp00`, where `INIT_DB` is the status of the snapplan.
- 3 Select a snapplan from the drop-down list with `INIT_DB` status. The **Validate/Modify Snapplan** dialog displays with default values set.
- 4 Use the following dialog procedures in order:
 - See [“To validate or modify a snapplan using the Create Clone Database wizard”](#) on page 57.
 - See [“To create a snapshot using the Create Clone Database wizard”](#) on page 58.
 - See [“To start up a snapshot using the Create Clone Database wizard”](#) on page 59.

Removing a clone database

The Veritas Storage Foundation for DB2 provides the Remove Clone Database wizard, enabling you to:

- shut down the clone database
- unmount the clone database
- remove the Storage Checkpoint
- resynchronize the snapshot

Removing a Storage Checkpoint clone database

Removing a Storage Checkpoint clone database enables you to unmount the clone database with the option of removing the checkpoint.

To remove a Storage Checkpoint clone database

- 1 Access the Remove Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 The **Remove Clone Database** screen displays. Verify the read-only information. Click **Next** to continue.
- 3 In the **Storage Checkpoint Clone Database** section, verify the Instance Name and enter the Password. Also, verify the Database Name and Mount Point.
In the **Primary Database Information** section, verify the Instance Name and Database Name.
- 4 Check the **Remove Storage Checkpoint** option if you want to remove the Storage Checkpoint along with removing the cloned database.
- 5 Click the **Remove Clone Database** button.
- 6 In the **Confirmation** dialog, click **Yes** to proceed.

Removing a database snapshot clone database

Removing a Database FlashSnap clone database unmounts the FlashSnap database with the option of resynchronizing the snapshot.

To remove a Database FlashSnap clone database

- 1 Access the Remove Clone Database wizard.
See [“Accessing clone database tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 The **Remove Clone Database** dialog displays. In the Authentication section, verify the UNIX user name and enter the password.
In the **Snapshot Database Information** section, verify the Snapplan File name and the Relocate path.
You may check the **Resync Snapshot** option to resynchronize the snapshot. If you check the **Resync Snapshot** option, you must enter a password.
In the **Primary Database Information** section, verify the following information:

- Primary Host Name
 - DB2 Database Name
 - Instance Owner Name
- 3 Click the **Remove Clone Database** button.
 - 4 In the **Confirmation** dialog, click **Yes** to proceed with removing the clone database.

Managing Storage Checkpoints

This chapter includes the following topics:

- [About Storage Checkpoints](#)
- [Creating a Storage Checkpoint](#)
- [Viewing Storage Checkpoint details](#)
- [Mounting a Storage Checkpoint](#)
- [Unmounting a Storage Checkpoint](#)
- [Removing a Storage Checkpoint](#)
- [Rolling back to a Storage Checkpoint](#)

About Storage Checkpoints

A Storage Checkpoint is like an online backup of a database, including partitioned databases, that contains a point-in-time database image. Storage Checkpoints can later be used to restore the image of the entire database to any earlier state recorded by the Storage Checkpoints. You cannot recover a container or a tablespace. Storage Checkpoints are used in a DB2 UDB EE environment. DB2 UDB EEE (SMP only) supports data partitioning across clusters of massively parallel computers.

Veritas Storage Foundation for DB2 uses the SFDB repository to determine the list of tablespaces, containers, and file systems for Storage Checkpoint creation and removal.

For more information about Storage Checkpoints:

See the chapter on using Storage Checkpoints and Storage Rollback in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

Creating a Storage Checkpoint

You can create a Storage Checkpoint that is:

- offline
- online

To create an offline Storage Checkpoint, the instance can be either running or stopped and the database must be inactive. To create an online Storage Checkpoint, the instance must be running and the database can be either active or inactive, but it must be in archive log enabled mode.

- Prerequisites
- Enable ARCHIVELOG mode before taking online Storage Checkpoints.
See the chapter on using Storage Checkpoints and Storage Rollback in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

To create a Storage Checkpoint

- 1 Access the Create a Storage Checkpoint wizard.
See [“Accessing Storage Checkpoint tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 Verify the database name. This is a read-only field.
- 3 Select one of the following options:
 - To create an online Storage Checkpoint, click the **Online** button.
 - To create an offline Storage Checkpoint, click the **Offline** button.
- 4 If you want to remove the Storage Checkpoint when the file system becomes full, click the **Remove This Storage Checkpoint** button.
If you want to retain the Storage Checkpoint when the file system becomes full, click the **Retain This Storage Checkpoint** button.

- 5 Click **Create** to continue.
- 6 At the prompt, click **Yes** to proceed with creating the Storage Checkpoint.
If the Storage Checkpoint is successfully created, a confirmation message displays. To see details, click the **Show Details** checkbox. The details are displayed in a pop-up window.
Click **OK** to continue.

Viewing Storage Checkpoint details

You can view the details of a Storage Checkpoint.

To view the details of a Storage Checkpoint

- 1 Click a specific **Storage Checkpoint** in the object tree. (You may need to expand the tree view to find the Storage Checkpoint.) The Storage Checkpoint information is displayed on the right side of the window.
- 2 To view file system quota information for the Storage Checkpoint, click the **File System Quota** tab at the top of the window, just above the Storage Checkpoint details.

The file system quota information is only available for disk layout version 6.

Mounting a Storage Checkpoint

You can mount and write to Storage Checkpoints just as you can do with any file system.

See “Using Storage Checkpoints and Storage Rollback” in the *Veritas Storage Foundation for DB2 Administrator’s Guide*.

- Prerequisites
- The directory containing the mount point must be writable by the database administrator group. You should have created this group during installation. If not, create the group before mounting the Storage Checkpoint.

To mount a Storage Checkpoint

- 1 Access the Mount a Storage Checkpoint wizard.
See [“Accessing Storage Checkpoint tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 Verify that you are mounting the correct Storage Checkpoint and click **Next** to continue. If you selected the wrong Storage Checkpoint, click **Cancel**. The information on this screen is read-only.
- 3 On the second screen, enter the mount point (absolute path) where the Storage Checkpoint should be mounted.
- 4 To mount the Storage Checkpoint as read-only, click **Read Only**.
or
To mount the Storage Checkpoint as read-write, click **Read/Write**. This will allow you to make changes to the Storage Checkpoint.
When you select the **Read/Write** option, the GUI creates an identical Storage Checkpoint with the same name plus a `wr<001>` suffix, where `<001>` is a sequential number. The GUI mounts the new Storage Checkpoint and leaves the original Storage Checkpoint unmounted. This allows you to roll back to the original Storage Checkpoint.
- 5 Click **Mount** to mount the Storage Checkpoint.
- 6 At the prompt, click **Yes** to proceed with mounting the Storage Checkpoint.

Unmounting a Storage Checkpoint

You can unmount a Storage Checkpoint just as you can do with any file system.

To unmount a Storage Checkpoint

- 1 Click a specific Storage Checkpoint in the object tree. (You may need to expand the tree view to find the Storage Checkpoint.)
If you want to unmount a Storage Checkpoint that was originally mounted with the **Read/Write** option, you should unmount the new Storage Checkpoint that was automatically created by the GUI, which is the Storage Checkpoint that contains the `wr<001>` suffix, where `<001>` is a sequential number, at the end of the name.
- 2 Access the Unmount a Storage Checkpoint wizard.
See [“Accessing Storage Checkpoint tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.

- 3 Verify that you are unmounting the correct Storage Checkpoint and click **Unmount** to continue. If you selected the wrong Storage Checkpoint, click **Cancel**. The information on this screen is read-only.
- 4 At the prompt, click **Yes** to proceed with unmounting the Storage Checkpoint.
- 5 A confirmation dialog is displayed. Click **OK** to continue.

Removing a Storage Checkpoint

Occasionally, you may need to manually remove Storage Checkpoints that are no longer needed. For example, you can remove a Storage Checkpoint on a file system to free up needed space.

- Prerequisites
- Before you can remove a mounted Storage Checkpoint, you must first unmount it.

To remove a Storage Checkpoint

- 1 Click a specific Storage Checkpoint in the object tree. (You may need to expand the tree view to find the Storage Checkpoint.)
- 2 Access the Unmount a Storage Checkpoint wizard.
See [“Accessing Storage Checkpoint tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 At the prompt, click **Yes** to remove the Storage Checkpoint.
If the Storage Checkpoint was successfully removed, a confirmation message displays. Click **OK** to continue.

Rolling back to a Storage Checkpoint

You can roll back the entire database to a Storage Checkpoint.

The GUI does not automatically roll back any logs associated with a Storage Checkpoint.

See “Guidelines for DB2 Recovery” in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

You must be the Database Administrator to perform Storage Rollback operations. For an online Storage Rollback, the instance should be running, but the database must be in archive mode (either active or inactive). For an offline Storage Rollback, the instance can be running or inactive, but the database must be inactive.

Storage Checkpoints can only be used to roll back files that are damaged due to a software error or a human error (for example, accidental deletion of a table). Because Storage Checkpoints reside on the same physical disks as the primary file system, when a file is corrupted due to a media failure, the file on the Storage Checkpoints will not be available either. In this case, you need to restore files from a tape backup.

After the files are rolled back, you may need to follow the recovery procedure described in your DB2 manuals to recover the database before the database can be used.

Rolling back the database to a Storage Checkpoint

Rolling back the entire database rolls back all the containers used by the database, to a Storage Checkpoint.

While the Storage Rollback process is running, it creates a temporary file, `/filesystem/.VRTSstrb.lock`, in each file system. Do not remove these temporary lock files.

To rollback the database to a Storage Checkpoint

- 1 Make the DB2 database inactive by either shutting down the instance or disabling all user connections.
- 2 Click the a specific Storage Checkpoint in the object tree. (You may need to expand the tree view to find the Storage Checkpoint.)
- 3 Access the Roll Back a Storage Checkpoint wizard.
See [“Accessing Storage Checkpoint tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 4 Verify that you are rolling back to the correct Storage Checkpoint and click **Next** to continue. If you selected the wrong Storage Checkpoint, click **Cancel**. The information on this screen is read-only.
- 5 On the second screen, use the drop-down menu to select the appropriate buffer size in the Rollback Buffer Size field. The default buffer size is 128K. The buffer size configured for reads and writes when performing a Storage Rollback can affect performance. Vary the size to determine the optimal setting for your system.
- 6 Use the drop-down menu to select the appropriate number of threads in the Number of Threads field. The default number of threads is four.

Depending on the number of CPUs available on your system and the type of volume on which the file system is located, this default setting may specify too few or too many threads.

- 7 Click the **Roll back a database** button to indicate that you are rolling back the entire database to the Storage Checkpoint.
- 8 Click **Next** to continue. Click **Roll Back** to continue.
If the Storage Rollback was successful, a confirmation message is displayed.
- 9 To see the Storage Rollback details, click the **Show details** checkbox. The details are displayed in the pop-up window.
When you have finished viewing the details, click **OK**.
- 10 Click **Yes** to roll back the tablespace or tablespaces. Perform any necessary DB2 recovery. (You cannot recover your database through the GUI.)
- 11 Activate the database again.

Rolling back datafiles to a Storage Checkpoint

Rolling back datafiles rolls back database files to a Storage Checkpoint. You can also use this operation to roll back more than one datafile.

This option is only available for offline and online Storage Checkpoints. You cannot roll back a datafile to an instant Storage Checkpoint.

To roll back datafiles to a Storage Checkpoint

- 1 Verify that the datafile to which you want to roll back is offline.
- 2 Access the Rollback a Storage Checkpoint wizard.
See [“Accessing Storage Checkpoint tasks in the Java GUI”](#) on page 27.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 If the database is online, you will receive a prompt asking you if you want to continue. Click **Yes** to continue the Storage Rollback.
- 4 In the Rollback a Storage Checkpoint wizard, verify that you are rolling back to the correct Storage Checkpoint and click **Next** to continue. If you selected the wrong Storage Checkpoint, click **Cancel**. The information on this screen is read-only.
- 5 On the second screen, use the drop-down menu to select the appropriate buffer size in the Rollback Buffer Size field. The default buffer size is 128K. The buffer size configured for reads and writes when performing a Storage Rollback can affect performance. Vary the size to determine the optimal setting for your system.

- 6 Use the drop-down menu to select the appropriate number of threads in the Number of Threads field. The default number of threads is four.

Depending on the number of CPUs available on your system and the type of volume on which the file system is located, this default setting may specify too few or too many threads.
- 7 Click the **Roll back a Datafile** button to indicate that you are rolling back the datafile to the Storage Checkpoint.
- 8 Click **Next** to continue.
- 9 Select the datafile, or datafiles, you would like to roll back by clicking on the appropriate checkbox in the Datafile list.
- 10 Click **Roll Back** to continue.
- 11 If the Storage Rollback was successful, a confirmation message displays. Click **OK** to continue. You are then returned to the rollback window.
- 12 Perform any necessary Oracle media recovery.
- 13 Put the recovered datafile back online.

Managing snapshots using Database FlashSnap

This chapter includes the following topics:

- [Database FlashSnap in the GUI](#)
- [Creating a snapplan](#)
- [Validating or modifying a snapplan](#)
- [Removing a snapplan](#)
- [Creating a snapshot](#)
- [Creating a clone database with a snapshot](#)
- [Resynchronizing a snapshot to a database](#)
- [Resynchronizing a database to a snapshot](#)
- [Aborting the reverse resynchronization operation](#)
- [Viewing the snapplan log](#)

Database FlashSnap in the GUI

With Veritas Database FlashSnap, you can create a point-in-time copy of a database for backup and off-host processing. For more information:

See *Using Database FlashSnap for Backup and Off-Host Processing in the Veritas Storage Foundation for DB2 Administrator's Guide*.

From the GUI, you can:

- Create snapshots of your database using snapplans.

- Resynchronize snapshots to your database.
- Resynchronize your database to a snapshot.

Database FlashSnap functionality is not supported on duplicated (cloned) databases.

Database FlashSnap is available for DB2 EE instances only.

Creating a snapplan

A snapplan specifies snapshot scenarios for the DB2 database (such as online, offline, and instant). The snapplan is used as a basis for creating a snapshot. You must either create or validate a snapplan before you can create the snapshot image.

| | |
|---------------|---|
| Prerequisites | <ul style="list-style-type: none">■ You must be logged in as the DB2 database administrator.■ The disk group must be version 110 or higher.■ Persistent FastResync must be enabled.■ To set up your snapplan for online snapshots, the database must be in archive log mode. |
| Usage notes | <ul style="list-style-type: none">■ Database FlashSnap functionality is not supported on cloned (duplicated) databases.■ The snapplan name is user-defined.■ Each entry in the snapplan is a line in <code>parameter=argument</code> format. See the <code>db2ed_vmsnapplan(1M)</code> and <code>db2ed_vmchecksnap(1M)</code> manual page for more information. |

To create a snapplan

- 1 Access the Create Snapplan wizard.
 - See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
 - See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 Confirm the Database Name. This is a read-only field.
- 3 Enter the full path of the snapplan file.
- 4 Enter the name of the secondary host.
- 5 Enter the snapshot plex tag.
- 6 Click **Next** to continue.

The Validate/Modify Snapplan screen is displayed with default values set.

- 7 If needed, modify any incorrect settings. Then, click **Validate** to ensure that the settings have been configured correctly.

For more information regarding the snapplan parameters:

See “Creating and Working with Snapplans Using dbed_vmchecksnap” in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

- 8 At the confirmation prompt, click **Yes** to confirm that you want to validate the snapplan.

If the snapplan was successfully validated, you will receive a confirmation message.

- 9 To see the snapplan details, click the **Show details** checkbox. The snapplan details are displayed in the pop-up window.

When you have finished reviewing them, click **OK**.

Snapplan details:

SNAPSHOT_MODE is online_snapshot

PRIMARY_HOST is host1

SECONDARY_HOST is host2

The version of PRIMARY_DG-DB2dg is 110.

SNAPSHOT_DG is SNAP_DB2dg

SNAPSHOT_PLAN_FOR is database

Examining DB2 volume and disk layout for snapshot.

Volume testvol01 on DB2dg is ready for snapshot.

Original plex and DCO log for testvol01 is on DB2dg01 DB2dg02.

Snapshot plex and DCO log for testvol01 is on DB2dg03.

Volume testvol02 on DB2dg is ready for snapshot.

Original plex and DCO log for testvol02 is on DB2dg01 DB2dg02.

Snapshot plex and DCO log for testvol02 is on DB2dg03.

Volume testvol03 on DB2dg is ready for snapshot.

Original plex and DCO log for testvol03 is on DB2dg01 DB2dg02.

Snapshot plex and DCO log for testvol03 is on DB2dg03.

```
Volume testvol104 on DB2dg is ready for snapshot.  
Original plex and DCO log for testvol104 is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for testvol104 is on DB2dg03.  
  
Volume testvol105 on DB2dg is ready for snapshot.  
Original plex and DCO log for testvol105 is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for testvol105 is on DB2dg03.  
  
Volume udb_home on DB2dg is ready for snapshot.  
Original plex and DCO log for udb_home is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for udb_home is on DB2dg03.  
  
SNAP_DB2dg for snapshot will include: DB2dg03  
  
ALLOW_REVERSE_RESYNC is yes  
  
The snapplan sp1 has been created.
```

Validating or modifying a snapplan

Before creating a snapshot, it is necessary to validate your snapplan to make sure the information is correct and that the snapshot will succeed. Occasionally, you may also need to modify a snapplan. For example, you would need to modify your snapplan if you wanted to change the primary disk group that is included in the snapshot.

- | | |
|---------------|---|
| Prerequisites | <ul style="list-style-type: none">■ You must be logged in as the DB2 database administrator.■ A DCO log must be attached to the snapshot plex. |
| Usage notes | <ul style="list-style-type: none">■ Each entry in the snapplan is a line in <code>parameter=argument</code> format. |

To validate or modify a snapplan

- 1 Click a specific snapplan in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Access the Validate/Modify Snapplan wizard.
See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 In the Modify/Validate Snapplan wizard, correct the path and file name of the snapplan in the **Snapplan file** field. Then, click **Next** to continue.

- 4 If you are modifying your snapplan, enter any new parameters that should be updated. You are not required to do anything if no changes are required.
 For more information regarding the snapplan parameters:
 See “Creating and Working with Snapplans Using dbed_vmchecksnap” in the *Veritas Storage Foundation for DB2 Administrator’s Guide*.
- 5 Click **Validate** to check that the settings have been configured correctly.
- 6 At the confirmation prompt, click **Yes** to confirm that you want to validate the snapplan.
 If the snapplan was successfully validated, you will receive a confirmation message.
- 7 To see the snapplan details, click the **Show details** checkbox. The snapplan details are displayed in the pop-up window.
 When you have finished reviewing them, click **OK**.

Snapplan details:

SNAPSHOT_MODE is online_snapshot

PRIMARY_HOST is host1

SECONDARY_HOST is host2

The version of PRIMARY_DG-DB2dg is 110.

SNAPSHOT_DG is SNAP_DB2dg

SNAPSHOT_PLAN_FOR is database

Examining DB2 volume and disk layout for snapshot.

Volume testvol01 on DB2dg is ready for snapshot.

Original plex and DCO log for testvol01 is on DB2dg01 DB2dg02.

Snapshot plex and DCO log for testvol01 is on DB2dg03.

Volume testvol02 on DB2dg is ready for snapshot.

Original plex and DCO log for testvol02 is on DB2dg01 DB2dg02.

Snapshot plex and DCO log for testvol02 is on DB2dg03.

Volume testvol03 on DB2dg is ready for snapshot.

Original plex and DCO log for testvol03 is on DB2dg01 DB2dg02.

Snapshot plex and DCO log for testvol03 is on DB2dg03.

```
Volume testvol104 on DB2dg is ready for snapshot.  
Original plex and DCO log for testvol104 is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for testvol104 is on DB2dg03.
```

```
Volume testvol105 on DB2dg is ready for snapshot.  
Original plex and DCO log for testvol105 is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for testvol105 is on DB2dg03.
```

```
Volume udb_home on DB2dg is ready for snapshot.  
Original plex and DCO log for udb_home is on DB2dg01 DB2dg02.  
Snapshot plex and DCO log for udb_home is on DB2dg03.
```

```
SNAP_DB2dg for snapshot will include: DB2dg03
```

```
ALLOW_REVERSE_RESYNC is yes
```

```
The snapplan sp1 has been created.
```

Removing a snapplan

You may need to remove a snapplan that is no longer necessary. However, you cannot create a snapshot if you do not have a snapplan.

Prerequisites ■ You must be logged in as the DB2 database administrator.

To remove a snapplan

- 1 Click the snapplan you want to remove in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Access the Remove Snapplan dialog:
See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 At the prompt, click **Yes** to continue with removing the snapplan.
If the snapplan was successfully removed, you will receive a confirmation message similar to the following:
- 4 Click the **Show details** checkbox to view the actions taken. You should see information similar to the following:

```
The snapplan <filename> has been removed.
```

Click **OK** to continue.

Creating a snapshot

After having created and validated a snapplan, you can create a snapshot of your database. You can use the snapshot as a database backup or as a test database to perform operations without affecting your production database.

- | | |
|---------------|---|
| Prerequisites | <ul style="list-style-type: none">■ Before you can create a snapshot, a snapshot mirror of a volume must exist.■ You must have a validated snapplan before creating a snapshot image.■ You must be logged in as the DB2 database administrator. |
| Usage notes | <ul style="list-style-type: none">■ See the <code>db2ed_vmsnap(1M)</code> manual page for more information.■ If you are creating an offhost snapshot, perform these steps on the secondary host. |

To create a snapshot

- 1 Access the Create Snapshot wizard.
See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 2 Verify the snapplan location. (This is a read-only field.)
- 3 If you need to force the snapshot creation, click the **Force snapshot creation** checkbox.

Force snapshot creation can be used after a snapshot operation has failed and the problem was fixed without using Veritas Storage Foundation for DB2 commands. (That is, the volumes were synchronized without using Veritas Storage Foundation for DB2 commands.)

In this situation, the status of the snapplan will appear as unavailable for creating a snapshot. The **Force snapshot creation** option ignores the unavailable status, checks for the availability of volumes, and creates the snapshot after the volumes pass the availability check.

- 4 Click **Finish** to continue.

- 5 If the snapshot creation was successful, you will receive a confirmation message.

Click the **Show details** checkbox to view the actions taken. You should see information similar to the following:

- 6 Click **OK** to continue.

You can now use the snapshot as a database backup.

You will need this information when you start your snapshot database.

Creating a clone database with a snapshot

After having created a snapshot, you can use the snapshot to create a clone database. The cloned database can be used for decision-making and testing that cannot be done on your production database.

- Prerequisites
- Make sure you have enough disk space to create a clone database on your system.
 - You must have an existing snapshot.

To create a clone database with a snapshot

- 1 Click the database object DB2 icon in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Access the Create Snapshot Database wizard.
See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 In the “Authentication” section, enter the Unix user name and password.
- 4 For the primary database, enter the DB2 Instance and DB2 Database information.
- 5 Select whether you are recovering the database or restarting the database by clicking the appropriate radio button.

Select **Startup database** if you are restarting a new snapshot database based on a snapshot.

Select **Restart database** if you are restarting a snapshot database that has already been created.

- 6 For the snapshot database, enter the new DB2 Database, Snapplan file, volume name, Snapshot disk group name, and Relocate path in the appropriate fields.

To start an on-host snapshot database, the Relocate path is required. To start an off-host snapshot database, the volume name field is required. You can obtain this information by viewing the details after you create a snapshot. If you did not retain the information from the snapshot creation, you can use the `db2ed_vmchecksnap` command or use Viewlog via the GUI to retrieve the information.

See “Creating and Working with Snapplans Using `db2ed_vmchecksnap`” in the *Veritas Storage Foundation for DB2 Administrator's Guide*.

- 7 Click **Start Database**.
- 8 At the confirmation prompt, click **Yes** to confirm that you want to start the database.

If the snapshot database was successfully started, confirmation message displays. Click **OK** to continue.

Resynchronizing a snapshot to a database

Resynchronizing a snapshot to a database will refresh the snapshot so that it contains the most recent changes made to your production database.

- | | |
|---------------|---|
| Prerequisites | <ul style="list-style-type: none"> ■ You must be logged in as the DB2 database administrator. ■ You must shut down the clone database and unmount the file systems. |
| Usage notes | <ul style="list-style-type: none"> ■ See the <code>db2ed_vmsnap(1M)</code> manual page for more information. |

To resynchronize a snapshot to a database

- 1 Click a snapplan, located under the **Snapplans** icon, in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Access the Resync Snapshot wizard.
 - See “[Accessing FlashSnap tasks in the Java GUI](#)” on page 28.
 - See “[Accessing database tasks in the Web GUI](#)” on page 35.
- 3 At the confirmation prompt, click **OK** to continue resynchronizing the snapshot to the database.

- 4 If the resynchronization is successful, you a confirmation message displays. Click **OK** to continue.
- 5 Status information similar to the following example is displayed. Click **OK** when you have finished viewing the information.

```
db2ed_vmsnap started at 2005-05-13 17:20:05
```

```
The option resync has been completed.
```

To see the details, click the **Show details** checkbox. The details are displayed in a pop-up window.

Resynchronizing a database to a snapshot

Resynchronizing your database to a snapshot, also known as reverse resynchronization, reverts your database to a snapshot. Use this option if your database becomes corrupted and you need to restore your database to a previous point-in-time.

- | | |
|---------------|--|
| Prerequisites | <ul style="list-style-type: none">■ You must be logged in as the DB2 database administrator.■ You must shut down the primary database and the clone database; also, unmount the file systems. |
| Usage notes | <ul style="list-style-type: none">■ See the <code>db2ed_vmsnap(1M)</code> manual page for more information. |

To resynchronize a database to a snapshot

- 1 Click a snapplan, located under the **Snapplans** icon, in the object tree. (You may need to expand the tree view to find the icon.)
- 2 Access the Reverse Resync wizard.
See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 Verify the Database name and snapplan information. These are read-only fields.
- 4 Click the Begin button, then click **Finish** to start the reverse resynchronization process. To view details, click the **Show details** checkbox. The details are displayed in a pop-up window.
- 5 At the confirmation prompt, click **Yes** to continue.

The Begin option performs the following actions:

- Imports the disk group that was deported from the secondary host and joins it back to the original disk group.
- Mounts the snapshot volumes.
- Mounts the file systems that are configured for the primary database.
- Brings up the database snapshot image as the primary database.

The primary database must be offline to perform this action.

- 6 If the begin action was successful, a confirmation message displays. Click **Show details** to see the actions completed. When you are through, click **OK** to continue.
- 7 Again, click the snapplan on which you want to perform the reverse resynchronization.
- 8 Select one of the following methods to access the Reverse Resync wizard:
 - From the menu bar, select **Snapplan > Reverse Resync Snapshot**.
 - Right-click the snapplan to bring up a pop-up menu. Then click **Reverse Resync Snapshot**.

The Reverse Resync wizard displays.
- 9 Click the **Commit** button, then click **Finish** to commit the reverse resynchronization process.
- 10 At the confirmation prompt, click **Yes** to continue.

The Commit option performs the following actions:

 - Commits the reverse resynchronization changes.
 - Resynchronizes the original volume from the data in the snapshot and then discards the content of the original database.

Warning: This action cannot be undone.

- 11 If the commit action was successful, a confirmation message displays. Click **Show details** to see the actions completed. When you are through, click **OK** to continue.

Aborting the reverse resynchronization operation

Occasionally, you may need to stop the reverse resynchronization process after you have begun. You can only abort the reverse resynchronization process after

you have completed the Begin operation and before performing a Commit operation.

To abort the reverse resynchronization operation

- 1 Click the snapplan for which the reverse resynchronization was begun.
- 2 Select one of the following methods to access the Reverse Resync wizard.
Access the Reverse Resync Snapshot wizard.
See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.
See [“Accessing database tasks in the Web GUI”](#) on page 35.
- 3 Verify the Database name and snapplan information. These are read-only fields.
- 4 Click **Abort**, then click **Finish** to abort the reverse resynchronization process.
- 5 At the confirmation prompt, click **Yes** to continue.

The Abort option performs the following actions:

- Unmounts the snapshot volumes.
- Mounts the original volumes back with the file systems that are configured to use the volume.

This action can only be performed after a “begin” action has been completed and cannot be used after a reverse resynchronization has been committed.

- 6 If the abort action was successful, you will receive a confirmation message. Click **Show details** to see the actions completed. When you are through, click **OK** to continue.

Viewing the snapplan log

The snapplan log displays information about the snapplan, disk group, snapshot Plex tag. It also displays the volume name and the snapshot status.

To view a snapplan log

- 1 In the object tree, click the snapplan whose log you wish to see. (You may need to expand the tree view to find the icon.)
- 2 Access the View Log wizard.

See [“Accessing FlashSnap tasks in the Java GUI”](#) on page 28.

See [“Accessing database tasks in the Web GUI”](#) on page 35.

A pop-up window appears displaying the path of the Snapplan and the Log. The information displayed is similar to the following output:

```
SNAPSHOT_VERSION=4.0
PRIMARY_HOST=exmoor
SECONDARY_HOST=exmoor
PRIMARY_DG=DB2dg
SNAPSHOT_DG=SNAP_DB2dg
DB2DATABASE=UDB
DB2HOME=/db2/udb_home
REDOLOG_DEST=/db2/udb_home/inst01/NODE0000/SQL00001/SQLLOGDIR/
SNAPSHOT_MODE=online_snapshot
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=db2ed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=yes
STORAGE_INFO
DB2dg03
SNAP_PLEX=testvol01-02 testvol02-02 testvol03-02 testvol04-02 testvol05
STATUS_INFO
SNAP_STATUS=init_full
DB_STATUS=init
LOCAL_SNAPPLAN=/home/inst01/snapplans/sp1
```

- 3 Click **OK** to continue.

Tiering storage for databases

This chapter includes the following topics:

- [About Database Dynamic Storage Tiering in the GUI](#)
- [Setting up Database Dynamic Storage Tiering parameters](#)
- [Setting automatic policies for Database Dynamic Storage Tiering](#)
- [Moving volumes and files](#)
- [Getting reports for Database Dynamic Storage Tiering](#)

About Database Dynamic Storage Tiering in the GUI

The Veritas Storage Foundation for DB2 GUI enables you to set up automated policies for Database Dynamic Storage Tiering as well as perform manual operations. The Veritas Storage Foundation for DB2 GUI enables you to:

- Set up Database Dynamic Storage Tiering parameters
- Manage volume and file movement by setting automatic policy
- Operate Database Dynamic Storage Tiering by moving volumes and files manually
- Get reports on Database Dynamic Storage Tiering
- Get reports on file system statistics across MVS file system

Note: Tiering storage is supported in the Java GUI only for Storage Foundation for DB2 5.0.

Setting up Database Dynamic Storage Tiering parameters

Setting up Database Dynamic Storage Tiering requires some planning in advance. You will need to determine and specify:

- Directories where storage tiering policies will apply
- Frequency of sweeping and purging
- Storage class options

Setting up Database Dynamic Storage Tiering parameters

You must set up the Database Dynamic Storage Tiering parameters before you can apply automatic management policies.

To set up Database Dynamic Storage Tiering parameters

- 1 Access the Database Dynamic Storage Tiering Parameters wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 Set up the parameters for the Database Dynamic Storage Tiering object or the GUI fills in default values. Set:
 - Minimum number of storage classes for file systems. The range is 2 to 99.
 - Maximum number of storage classes for file systems. The range is 2 to 99.
 - Statistic gathering interval (minutes).
 - Sweep interval (days). The range is 1 to 10.
 - Sweep time (hours, minutes, am/pm). You can enter one sweep time by specifying hour, minute, and AM or PM. Specify time according to the server time, not local time.
 - Purge interval (days). The range is 1 to 120.
 - Purge time (hours, minutes, am/pm). You can enter one purge time by specifying hour, minute, and AM or PM. Specify time according to the server time, not local time.

After setting the parameters click **Next**.

- 3 Set the storage classes or the GUI fills in default values. Click **Add Class** to specify each storage class in a separate row of the table.
 To sort table rows by Storage Class or Description, click the column headers. Empty classes are sorted to the bottom of the table.
 Pre-defined Storage Classes can not be modified or removed.
- 4 Click **Finish** when all storage classes are entered.
- 5 If the parameter updates are successfully implemented, a confirmation message displays. Click **OK** to continue.

To modify Database Dynamic Storage Tiering parameters

- 1 Access the Database Dynamic Storage Tiering wizard.
 See “[To access tiered storage parameters wizards in the Java GUI](#)” on page 29.
- 2 Modify the parameters for the Database Dynamic Storage Tiering object. Set:
 - Minimum number of storage classes for file systems. The range is 2 to 99.
 - Maximum number of storage classes for file systems. The range is 2 to 99.
 - Statistic gathering interval (minutes).
 - Sweep interval (days). The range is 1 to 10.
 - Sweep time (hours, minutes, am/pm). You can enter one sweep time by specifying hour, minute, and AM or PM. Specify time according to the server time, not local time.
 - Purge interval (days). The range is 1 to 120.
 - Purge time (hours, minutes, am/pm). You can enter one sweep time by specifying hour, minute, and AM or PM. Specify time according to the server time, not local time.

After setting the parameters click **Next**.

- 3 Set the storage classes or the GUI fills in default values. Click **Add** to specify each storage class in a separate row of the table.
 To sort table rows by Storage Class or Description, click the column headers. Empty classes are sorted to the bottom of the table.
 To modify the description of a class, click the class name and modify appropriately.
 To remove a class, select the class and click **Remove**.
 Pre-defined Storage Classes can not be modified or removed.

- 4 Click **Finish** when all storage classes are entered.
- 5 If the parameter updates are successfully implemented, a confirmation message displays. Click **OK** to continue.

Identifying tiered storage directories

Set the Database Dynamic Storage Tiering parameters before identifying the directories to which they apply, or the action will fail. Once you have set the tiered storage parameters, you can access the Database Dynamic Storage Tiering Convert/Add/Remove wizard.

To identify directories covered by tiered storage policies

- 1 Access the Database Dynamic Storage Tiering Convert/Add/Remove wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 At the prompt, enter the directory name and click **Next**.
- 3 You are prompted to convert the mount device or add a new mount point.
To convert the mount device:
 - Right-click the volume name.
 - Click **Convert Mount Device**.
 - Enter the volume name.
 - Click **OK**.To add a new mount point:
 - In the Mount point field, enter the mount point.
 - Click **Get Information**.
- 4 To add volumes:
 - Right-click the volume set's name.
 - Select **Add Volumes**.
 - Enter the volume name.
 - Set the storage class.
 - Click **OK**.
- 5 To delete volumes:
 - Right-click the volume name.
 - Select **Remove Volumes**.

- Click **Yes** in confirmation dialog.
- 6 Click **Close**.
 - 7 If the directory is successfully converted, a confirmation message displays. Click **OK** to continue.

Classifying volumes

Classifying volumes enables you to set specific policies for them. You must define Database Dynamic Storage Tiering parameters before you can classify a volume.

To classify volumes

- 1 Access the Database Dynamic Storage Tiering Classify Volumes wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 At the prompt, enter the directory name and click **Next**.
- 3 Select a storage class and click **Set Class**.
To leave volumes unclassified, select **None** as the storage class.
- 4 Click **Finish**.
- 5 If the directory is successfully classified, a confirmation message displays. Click **OK** to continue.

Setting automatic policies for Database Dynamic Storage Tiering

You can manage volume and file movement by setting automatic policies. Policies for automatically moving data from one storage tier to another can be set based on:

- Directory
- SMS tablespace
- Automatic storage path
- Archive1/Archive2
- External files

Listing policies

You can display a list of all the policies that have been set for Database Dynamic Storage Tiering.

To list policies

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.
- 2 Select one of the following methods to access the Database Dynamic Storage Tiering List Policy wizard:
 - In the menu bar, click **DB2 > Database Dynamic Storage Tiering Management > List Policy**.
 - Right click the database object to bring up a pop-up menu. Then, click **Database Dynamic Storage Tiering Management > List Policy**.
- 3 Click **Yes** in the confirmation dialog.
All policies are displayed.
- 4 When finished, click **OK**.

Presetting policy for a directory

You can preset movement policy based on directories.

To preset policy for a directory

- 1 Access the Database Dynamic Storage Tiering Preset Policy wizard.
See [“To access the tiered storage policy wizards in the Java GUI”](#) on page 30.
- 2 In the **Preset Policy** page, enter the directory name.
- 3 To create or modify a policy, select **Create or modify a policy**.
- 4 Click **Next**. The file name patterns and storage classes display for the selected directory.
- 5 To add a file pattern to a storage class, double-click the cell in the **File Pattern** column, and enter a file pattern such as *.log, or *.dbf. Use commas to separate each file pattern.
- 6 To modify a file pattern for a storage class, double-click the cell corresponding to the file pattern and edit the file pattern.
- 7 To remove a file pattern for a storage class, double-click the cell corresponding to the file pattern and click **Remove File Pattern**.
- 8 To remove the policy, select **Remove Policy**.
- 9 To set the policy, click **Finish**.
- 10 Click **Yes** to confirm.

Presetting policy for SMS tablespace

You can preset policy based on directories.

To preset policy for SMS tablespace

- 1 Access the Database Dynamic Storage Tiering Preset Policy for SMS Tablespace wizard.
See [“To access the tiered storage policy wizards in the Java GUI”](#) on page 30.
- 2 In the **Preset Policy for Directory for SMS Tablespace** page, select one SMS Tablespace from the drop-down list.
- 3 Select **Create or modify** policy.
- 4 Click **Next**. The storage classes and file name patterns display for the selected SMS tablespace.
- 5 To add a file pattern to a storage class, double-click the cell in the **File Pattern** column, and enter a file pattern such as *.log, or *.dbf. Use commas to separate each file pattern.
- 6 To modify a file pattern for a storage class, double-click the cell corresponding to the file pattern and edit the file pattern.
- 7 To remove a file pattern for a storage class, double-click the cell corresponding to the file pattern and click **Remove File Pattern**.
- 8 To remove the policy, select **Remove Policy**.
- 9 To set the policy, click **Finish**.
- 10 Click **Yes** to confirm.

Presetting policy for automatic storage path

You can preset policy for automatic storage path.

To preset policy for automatic storage path

- 1 Access the Database Dynamic Storage Tiering Preset Policy for Automatic Storage Path wizard.
See [“To access the tiered storage policy wizards in the Java GUI”](#) on page 30.
- 2 Select **Create or modify** policy.
- 3 Click the **Next**. The storage classes and file name patterns display for the selected automatic storage path.
- 4 To add a file pattern to a storage class, double-click the cell in the **File Pattern** column, and enter a file pattern such as *.log, or *.dbf. Use commas to separate each file pattern.

- 5 To modify a file pattern for a storage class, double-click the cell corresponding to the file pattern and edit the file pattern.
- 6 To remove a file pattern for a storage class, double-click the cell corresponding to the file pattern and click **Remove File Pattern**.
- 7 To remove the policy, select **Remove Policy**.
- 8 To set the policy, click **Finish**.
- 9 Click **Yes** to confirm.

Setting archive1/archive 2 policy

You can set policy for moving archive logs. You must define Database Dynamic Storage Tiering parameters before you can set movement policies for archive logs.

To set policy for moving archive logs

- 1 Access the Database Dynamic Storage Tiering Archive1/Archive2 Policy wizard.

See [“To access the tiered storage policy wizards in the Java GUI”](#) on page 30.
- 2 In the **Policy of Archives**, enter the archive log destination number.
- 3 Select **Create or modify** policy.
- 4 Click **Next**. The storage classes and file name patterns display for the selected archive log.
- 5 To add a file pattern to a storage class, double-click the cell in the **File Pattern** column, and enter a file pattern such as *.log, or *.dbf. Use commas to separate each file pattern.
- 6 To modify a file pattern for a storage class, double-click the cell corresponding to the file pattern and edit the file pattern.
- 7 To remove a file pattern for a storage class, double-click the cell corresponding to the file pattern and click **Remove File Pattern**.
- 8 To remove the policy, select **Remove Policy**.
- 9 To set the policy, click **Finish**.
- 10 Click **Yes** to confirm.

Setting external files policy

You can set policy for moving external files. You must define Database Dynamic Storage Tiering parameters before you can set movement policies for external files.

To set policy for moving external files

- 1 Access the Database Dynamic Storage Tiering External Files Policy wizard.
See [“To access the tiered storage policy wizards in the Java GUI”](#) on page 30.
- 2 Select **Create or modify** policy.
- 3 Enter the mount point.
- 4 Click **Next**.
- 5 Enter the directory/file names, separated by a new line.
- 6 For each storage class, select the number of days the file is not accessed before it will be moved to that storage class. You can leave the Days field blank if you do not want to set policy for a storage class.
- 7 To set the policy, click **Finish**.
- 8 To remove the external files policy:
 - Select **Remove Policy**.
 - Click **Next**.
 - Click **Finish**.
 - Click **Yes** to confirm.

Moving volumes and files

You can operate Database Dynamic Storage Tiering by moving volumes and files manually as needed from one storage class to another. You can move:

- Tablespaces
- Containers

Moving tablespaces

You can move tablespaces from one storage class to another. You must define Database Dynamic Storage Tiering parameters before you can move tablespaces.

To move tablespaces

- 1 Access the Database Dynamic Storage Tiering Move Tablespaces wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
The tablespaces and their assigned storage classes display in the Move Tablespaces wizard. To sort, click on column headings.
- 2 Select the tablespaces to move.
- 3 Select the storage class to which the tablespaces will be moved.
- 4 To assign the selected tablespaces to the selected storage class, click **Set class**.
- 5 Repeat as needed for each class of storage until all tablespaces are assigned correctly. A blank storage class indicates the tablespace is not moved.
- 6 To move the tablespaces, click **Finish**.

Moving containers

You can move containers from one storage class to another. You must define Database Dynamic Storage Tiering parameters before you can set movement policies for external files.

To move containers

- 1 Access the Database Dynamic Storage Tiering Move Containers wizard.
See [“To access the tiered storage movement wizards in the Java GUI”](#) on page 30.
The containers and their assigned storage classes display in the Move Containers wizard. To sort, click on column headings.
- 2 Select the containers to move.
- 3 Select the storage class to which the containers will be moved.
- 4 To assign the selected containers to the selected storage class, click **Set class**.
- 5 Repeat as needed for each class of storage until all datafiles are assigned correctly. A blank storage class indicates the tablespace is not moved.
- 6 To move the datafiles, click **Finish**.

Getting reports for Database Dynamic Storage Tiering

Reports available in the GUI for Database Dynamic Storage Tiering include:

- Show activities

- Show volume usage
- Show datafile statistics chart

Showing activities

You can display activities for selected date parameters.

To show activities

- 1 Access the Database Dynamic Storage Show Activities wizard.
See [“To access the tiered storage report wizards in the Java GUI”](#) on page 31.
- 2 Select the time period.
- 3 Click **Next**. All Database Dynamic Storage Tiering activities are displayed for the period of time selected.
- 4 To close the page when finished, click **Close**.

Listing policies

You can display a list of all the policies that have been set for Database Dynamic Storage Tiering.

To list policies

- 1 In the object tree, select the database object. You may need to expand the tree view to find the database object.
- 2 Select one of the following methods to access the Database Dynamic Storage Tiering List Policy wizard:
 - In the menu bar, click **DB2 > Database Dynamic Storage Tiering Management > List Policy**.
 - Right click the database object to bring up a pop-up menu. Then, click **Database Dynamic Storage Tiering Management > List Policy**.
- 3 Click **Yes** in the confirmation dialog.
All policies are displayed.
- 4 When finished, click **OK**.

Showing volume usage

You can show volume usage for the selected mount point.

To show volume usage

- 1 Access the Database Dynamic Storage Tiering Show Volume Usage wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 Enter the mount point.
- 3 Click **Get Volume Usage**. Volume usage displays for the mount point entered.
- 4 Enter additional mount points as required.
- 5 To close the page when finished, click **Close**.

Starting and stopping datafile statistics collection

Starting statistics collection for a datafile means the datafile is included when generating statistics charts. Stopping statistics collection for a datafile means the datafile will be excluded when generating statistics charts.

Collecting statistics on all datafiles could cause the repository to grow very fast for large numbers of datafiles, since collection is daily. To monitor important datafiles, use the Start/Stop Datafile Statistics Collection wizard to select the appropriate datafiles.

See [“Showing statistics chart”](#) on page 99.

Starting datafiles statistics collection

- 1 Access the Database Dynamic Storage Tiering Start/Stop Datafile Statistics Collection wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 Select a datafile for which statistics will be collected.
- 3 Click **Start**.
- 4 Click **Finish** when all datafiles are set appropriately.

Stopping datafiles statistics collection

- 1 Access the Database Dynamic Storage Tiering Start/Stop Datafile Statistics Collection wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 Select the datafiles for which statistics will no longer be collected.
- 3 Click **Stop**.
- 4 Click **Finish** when all datafiles are set appropriately.

Showing statistics chart

Before creating datafile statistics charts, you must enable the collection of datafiles statistics by using the Start/Stop Datafile Statistics Collection wizard.

See [“Starting and stopping datafile statistics collection”](#) on page 98.

You can create datafile statistics charts based on:

- Read/write statistics
- Time period
- Sorting parameters
- Datafiles

The statistics display in a bar chart.

To create a datafiles Statistics chart

- 1 Access the Database Dynamic Storage Tiering Show Datafiles Statistics Chart wizard.
See [“To access tiered storage parameters wizards in the Java GUI”](#) on page 29.
- 2 Select the report statistic:
 - Total read/write operations
 - Total read/write blocks
 - Total read/write bytes
 - Average read/write response time
- 3 Select the time period.
- 4 Select the sorting parameters.
- 5 Select the datafiles.
- 6 Click **Next**. The statistics are displayed in a bar chart.
- 7 To display the statistical trend, click **Show Statistic Trend**. The Show Statistic Trend line chart displays.
- 8 To move the containers to different storage classes, click **Move Container Classes**.
- 9 To close the chart pages when finished, click **Close**.

Automating database tasks

This chapter includes the following topics:

- [About automating tasks using the Java GUI scheduler](#)
- [Adding a scheduled task](#)
- [Viewing scheduled jobs](#)
- [Modifying a scheduled job task](#)
- [Automating creation of a database checkpoint](#)
- [Automating creation of a database snapshot](#)
- [Automating cloning of a database using a checkpoint](#)
- [Automating cloning of a database using FlashSnap](#)
- [Set the schedule parameters for an automated task](#)

About automating tasks using the Java GUI scheduler

Tasks that can be automated by using the Java GUI scheduler:

- Creating a database using checkpoint
- Creating a database using snapshot
- Cloning a database using checkpoint
- Cloning a database using snapshot

Note: Automating database tasks is supported in the Java GUI only for Storage Foundation for DB2 5.0.

Adding a scheduled task

To automate a database task, you must add it to the scheduler.

To add a task to the scheduler

- 1 Access the scheduler wizard in GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Enter the required authorization information for the host on which you want to run the schedule:
 - User name
 - Password

By default, authorization is set to the database owner. If you are logged in as the database owner, you do not need to specify a password.
- 3 Enter the task information:
 - Task name
 - Task status: Normal, Hold, or Invalid
 - Host name
- 4 Select a task:
 - Create a new database checkpoint
 - Create a database snapshot
 - Clone a database using checkpoint
 - Clone a database using snapshot
- 5 Click **Next** to add the appropriate task.
See [“To create a database checkpoint”](#) on page 104.
See [“To create a database snapshot”](#) on page 104.
See [“To clone a database using checkpoint”](#) on page 105.
See [“To clone a database using snapshot”](#) on page 105.

Viewing scheduled jobs

You can view all existing automated tasks.

To view a sortable summary of scheduled jobs

- 1 Access a scheduled job in the GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Right-click the scheduled job and select **View Jobs**. The View Jobs page displays.
- 3 Select jobs to view by status, name, and date. Click **Refresh** to get the jobs that meet the conditions.
- 4 To view additional details, click **Details**. By default, the Job Details tab is collapsed. To display job attributes and values, expand the **Detail** tab.
- 5 To view command line interface output for a job, click the **CLI Output** tab in the Detail view.
- 6 Click **Close** when finished.

Modifying a scheduled job task

You can modify an automated database task using the scheduler wizard.

To modify a task to the scheduler

- 1 Access the scheduler wizard in GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Enter the required authorization information:
 - User name
 - Password
- 3 Modify the task information:
 - Task name
 - Task status: Normal, Hold, or Invalid
 - Host name
- 4 Select a task:
 - Create a database using checkpoint
 - Create a database using snapshot
 - Clone a database using checkpoint

- Clone a database using snapshot
- 5 Click **Next** to modify the appropriate task.
See [“To create a database checkpoint”](#) on page 104.
See [“To create a database snapshot”](#) on page 104.
See [“To clone a database using checkpoint”](#) on page 105.
See [“To clone a database using snapshot”](#) on page 105.

Automating creation of a database checkpoint

You can automate the creation of a database checkpoint using the scheduler.

To create a database checkpoint

- 1 Access the scheduler wizard in GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Select the Create database checkpoint option.
- 3 Click **Next**.
- 4 Choose whether to take the Checkpoint when the database is:
 - Online
 - Offline
- 5 Choose the appropriate action if the file system is full:
 - Remove this Storage Checkpoint
 - Retain this Storage Checkpoint
- 6 Click **Next** to set the schedule parameters.
See [“Set the schedule parameters for an automated task”](#) on page 106.

Automating creation of a database snapshot

You can automate the creation of a database snapshot using the scheduler.

To create a database snapshot

- 1 Access the scheduler wizard in GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Select the Create database snapshot option.

- 3 Select the name of the snapplan that will be used for creating the snapshot.
- 4 Click **Next** to set the schedule parameters.
See [“Set the schedule parameters for an automated task”](#) on page 106.

Automating cloning of a database using a checkpoint

You can automate the cloning of a database from a checkpoint using the scheduler.

To clone a database using checkpoint

- 1 Access the scheduler wizard in GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Select the Clone database checkpoint option.
- 3 Click **Next**.
- 4 To restart the database, check the Restart database checkbox.
- 5 Enter the target database information:
 - Instance name
 - Database name
 - Mount Point
 - Redo log directory
- 6 Enter the source database information:
 - Instance name
 - Database name
 - Database log directory
- 7 Click **Next** to set the schedule parameters.
See [“Set the schedule parameters for an automated task”](#) on page 106.

Automating cloning of a database using FlashSnap

You can automate the creation of a database from a snapshot using the scheduler.

To clone a database using snapshot

- 1 Access the scheduler wizard in GUI.
See [“Accessing scheduled tasks in the Java GUI”](#) on page 31.
- 2 Select the Clone database snapshot option.

- 3 Click **Next**.
- 4 Select an existing snapplan to use for creating the clone database.
- 5 Click **Next**.
- 6 Specify whether the snapshot database will be created or restarted.
- 7 In the Create Snapshot Database page, set the secondary host name.
- 8 Enter the target database information:
 - New DB2 database
 - Snapplan name
 - Snapshot diskgroup name
 - Relocate path
- 9 Click **Next** to set the schedule parameters.

See [“Set the schedule parameters for an automated task”](#) on page 106.

Set the schedule parameters for an automated task

- Specific routine
- Week days
- Days of the month
- Specific dates

The scheduler enables you to set or modify schedule parameters for automated tasks based on:

To set schedule parameters

- 1 Select and configure a task.
 - See [“To create a database checkpoint”](#) on page 104.
 - See [“To create a database snapshot”](#) on page 104.
 - See [“To clone a database using checkpoint”](#) on page 105.
 - See [“To clone a database using snapshot”](#) on page 105.After you configure the task, you must set the schedule parameters.
- 2 To set a specific routine, click **Specific Routine** and set the parameters as required:
 - Start time
 - End time

- Date for scheduled effects
 - Restart task option for run day
 - Days before repeating task
 - Start date
- 3 To set weeks days, click **Week Days** and check the days you want the automated task to run.
 - 4 To set days of the month, click **Days of Month** and check the days you want the automated task to run.
 - 5 To include or exclude specific days, click **Specific Dates** and select the dates you want to include or exclude running the automated task. Use the right and left arrows as needed to adjust the Include dates and Exclude dates.
 - 6 When the scheduling information is complete, click **Finish**. The automated task is created.

Maintaining your system configuration

This chapter includes the following topics:

- [System configuration maintenance](#)
- [Resynchronizing the SFDB repository](#)
- [Rescanning the instance](#)
- [Displaying and rescanning tablespace information](#)
- [Viewing Storage Mapping](#)
- [Viewing I/O statistics](#)
- [Converting regular containers to Quick I/O Files](#)
- [Converting Quick I/O files to regular containers](#)
- [Updating rescan intervals](#)
- [Checking your system configuration](#)
- [Saving the system configuration](#)

System configuration maintenance

VERITAS Storage Foundation for DB2 maintains a repository that stores the pertinent information needed to display configuration information. This repository is located at `/etc/vx/DB2.$DB2INSTANCE.$DB2DATABASE`. When the database configuration changes, the information stored in the repository may not be up-to-date. You can resynchronize the repository, if needed. You can also rescan,

or refresh, instances and tablespaces to make sure system information is up-to-date.

The GUI also allows you to check and save the configuration of each DB2 instance on your system. Information on all volumes, file systems and their types, and disk groups can be displayed.

Resynchronizing the SFDB repository

You can resynchronize the SFDB repository to make sure system information is up-to-date.

To resynchronize the SFDB repository

- 1 Access the wizard for this task.
See [“Accessing system maintenance tasks in the Java GUI”](#) on page 32.
See [“Accessing system maintenance tasks in the Web GUI”](#) on page 36.
- 2 Select option to resynchronize the repository.
- 3 In the confirmation dialog, click **OK** to continue.

Rescanning the instance

You can rescan database instances to make sure system information is up-to-date.

To rescan the instance

- 1 Access the wizard for this task.
See [“Accessing system maintenance tasks in the Java GUI”](#) on page 32.
See [“Accessing system maintenance tasks in the Web GUI”](#) on page 36.
- 2 Select the rescan option.
- 3 In the confirmation dialog, click **OK** to continue.

Displaying and rescanning tablespace information

You can rescan, or refresh, tablespaces to make sure system information is up-to-date.

To display tablespace information

- 1 Access the wizard for this task.
See [“Accessing system maintenance tasks in the Java GUI”](#) on page 32.
See [“Accessing system maintenance tasks in the Web GUI”](#) on page 36.
- 2 Expand your view to display the tablespace information.
- 3 Select the rescan option.
- 4 Click **OK** to continue.

Viewing Storage Mapping

The Storage Foundation for DB2 GUI enables you to view Storage Mapping topology. For a more detailed description of the Topology page:

See using Storage Mapping in the Veritas Storage Foundation for DB2 Administrator's Guide.

To view DB2 container Storage Mapping topology

- 1 Access the DB2 container Storage Mapping topology wizard.
See [“Accessing system maintenance tasks in the Java GUI”](#) on page 32.
See [“Accessing system maintenance tasks in the Web GUI”](#) on page 36.
- 2 When you are through, click **Close** to quit.

Viewing I/O statistics

The Storage Foundation for DB2 GUI enables you to view Storage Mapping I/O statistics.

Note: Quick I/O is not supported on Linux.

To view DB2 container I/O statistics

- 1 Access the DB2 container I/O statistics wizard.
See [“Accessing system maintenance tasks in the Java GUI”](#) on page 32.
See [“Accessing system maintenance tasks in the Web GUI”](#) on page 36.
- 2 To enable the **Get Statistics** button, enter the sampling information:
 - The number of times you wish to sample the data. You can enter a value from 0 to 9999.

- The interval, in minutes, between sampling points. You can enter a value from 0 to 9999.

The **Total Sampling Period (in hours)** displays the time used to perform the sampling of the data.

The **Display accumulated data** checkbox is checked by default.

The **Time Elapsed (in minutes)** displays the amount of time it took to sample the data. The progress status bar displays the percentage of sampling that is being completed.

- 3 After successful completion, a Confirmation dialog is displayed, indicating that the data was successfully collected for the specified time. Click **OK**, to continue.
- 4 When you are finished, click **Close**.

Converting regular containers to Quick I/O Files

VERITAS Storage Foundation for DB2 provides an option to convert your regular datafiles to Quick I/O files to improve performance.

Prerequisites

Files you want to convert must be regular containers on VxFS file systems or links that point to regular VxFS files.

Usage notes

Converting existing containers to be Quick I/O files may not be the optimal thing to do if these files are fragmented.

You can only convert DMS containers. This operation is not supported on SMS containers.

Note: Quick I/O is not supported on Linux.

To convert DB2 containers to Quick I/O files

- 1 Make the database inactive.
- 2 Select the container in the GUI.
- 3 Select the conversion option in the GUI to generate container statistics.
- 4 At the prompt, click **Yes** to convert to a Quick I/O file.
- 5 In the confirmation dialog, click **OK** to view the information.

Converting Quick I/O files to regular containers

VERITAS Storage Foundation for DB2 provides an option to convert your Quick I/O files to regular DB2 containers.

Prerequisites:

Files you want to convert must be Quick I/O files on VxFS file systems or links that point to Quick I/O files.

Note: Quick I/O is not supported on Linux.

To convert Quick I/O files to regular containers

- 1 Make the database inactive.
- 2 Select the container in the GUI.
- 3 Select the conversion option in the GUI to generate container statistics.
- 4 At the prompt, click **Yes** to convert to a regular container.
- 5 In the confirmation dialog, click **OK** to view the information.

Updating rescan intervals

VERITAS Storage Foundation for Oracle periodically scans the system for updated information. You can modify the rescan intervals to be faster or slower. A partial scan is a scan of existing known objects and a full scan is a scan of known and unknown objects, meaning that a search for new information is performed.

To update a rescan interval

- 1 Select the database icon in the GUI.
- 2 Select the **Update Rescan Intervals** option in the GUI to update the rescan interval times.
- 3 To update a partial rescan interval, enter a new time (in seconds) in the appropriate field. To update a full rescan interval, enter a new time (in seconds) in the appropriate field.
- 4 Click **Save** to save your changes. A confirmation dialog displays.
- 5 Click **OK** to continue.

Checking your system configuration

You can rescan, or refresh, database instances, datafiles, and tablespaces to make sure system information is up-to-date.

You can check the System Configuration only when using VERITAS Storage Foundation for DB2 Enterprise Edition.

To scan the system configuration of a database

- 1 Select the database icon in the GUI.
- 2 Select the **Check System Configuration** option in the GUI to scan the database.
- 3 In the confirmation dialog, click **Yes** or **No** to continue.
- 4 Click **OK** to return to the main window.

Sample rescan output:

```
Examining File System and DB2 Container attributes.
```

```
Total of 0 containers over 0 file systems.
```

```
WARNING: No file systems to examine.
```

```
Examining Quick I/O settings.
```

```
Examining Cached Quick I/O settings.
```

```
NOTICE: No file systems have Cached Quick I/O enabled.
```

```
The database has:
```

```
3 SMS Containers
```

```
1 DMS File Containers
```

```
0 DMS Device Containers
```

```
Examining DB2 container fragmentation.
```

```
NOTICE: Could not examine DB2 container fragmentation.
```

```
Examining File System tunable settings.
```

```
NOTICE: Parameters for all VxFS file systems used by PROD.
```

```
Examining DB2 Volume layout and attributes.
```

```
WARNING: Data for database PROD is spread over multiple volume groups.
```

```
Examining DB2 internal information.
```

DB2 Version is 8.1.

Examining DB2 logging mode.

The database has transaction logs in directory
 /udb_home/prod/inst01/NODE0000/SQL00001/SQLOGDIR/.

WARNING: Transaction log directory is not mirrored using VxVM.

The database is in circular log mode (not archivelog).
 WARNING: Database recovery is affected by this mode.

The database is archiving logs in the default location.
 All archived logs remain under
 /udb_home/prod/inst01/NODE0000/SQL00001/SQLOGDIR/.

Examining DB2 Database Free Space.

DB20000I The SQL command completed successfully.

| | |
|-------------------|------------------------|
| Name | = SYSCATSPACE |
| Type | = System managed space |
| Total pages | = 5570 |
| Used pages | = 5570 |
| Free pages | = Not applicable |
| Page size (bytes) | = 4096 |

| | |
|-------------------|------------------------|
| Name | = TEMPSPACE1 |
| Type | = System managed space |
| Total pages | = 1 |
| Used pages | = 1 |
| Free pages | = Not applicable |
| Page size (bytes) | = 4096 |

| | |
|-------------------|------------------------|
| Name | = USERSPACE1 |
| Type | = System managed space |
| Total pages | = 1 |
| Used pages | = 1 |
| Free pages | = Not applicable |
| Page size (bytes) | = 4096 |

| | |
|------|--------|
| Name | = TBS1 |
|------|--------|

| | |
|-------------------|--------------------------|
| Type | = Database managed space |
| Total pages | = 5000 |
| Used pages | = 160 |
| Free pages | = 4800 |
| Page size (bytes) | = 4096 |

Saving the system configuration

You can rescan, or refresh, database instances, datafiles, and tablespaces to make sure system information is up-to-date.

To save the system configuration of a database

- 1 Select the database icon in the GUI.
- 2 Select the **Save System Configuration** option in the GUI to view the database configuration information.
- 3 In the System Configuration wizard, enter a path name, or directory, in the Path field to indicate where you would like to save the system configuration information.
- 4 Click **Save** to save the configuration information.
- 5 In the confirmation dialog, click **OK** to continue.

Veritas Database FlashSnap status information

This appendix includes the following topics:

- [Obtaining database FlashSnap snapshot status and database status from the GUI](#)

Obtaining database FlashSnap snapshot status and database status from the GUI

You can obtain both the snapshot status and the snapshot database status from the GUI. The tables in this section provide detailed information regarding the various status values.

Database FlashSnap snapshot status details

To view snapshot status information from the GUI, click on a specific snapplan in the object tree. The snapshot status can be seen on the right side of the window in the **Snapplan State** field. For information regarding the various status values, see the following table:

Table A-1

| Snapshot status (as seen in the Snapplan State field) | Completed Operations | Allowed Operations |
|---|---|---|
| init_full | <ul style="list-style-type: none"> ■ Modify/Validate Snapplan (successful) ■ Resync Snapshot (successful) ■ Reverse Resync Snapshot with the commit option (successful) | <ul style="list-style-type: none"> ■ Create Snapshot |
| snapshot_start | <ul style="list-style-type: none"> ■ Create Snapshot (failed) | <ul style="list-style-type: none"> ■ If the Create Snapshot operation failed, contact your system administrator for help. You can use the VxVM utilities to create a snapshot and resynchronize the snapshot volumes, then use the Create Snapshot operation with the Force snapshot creation option for the subsequent snapshot. |
| snapshot_end | <ul style="list-style-type: none"> ■ Create Snapshot (successful) | <ul style="list-style-type: none"> ■ Resync Snapshot ■ Reverse Resync Snapshot with the begin option ■ Create Snapshot Database with the Create database option |

Table A-1 (continued)

| Snapshot status (as seen in the Snapplan State field) | Completed Operations | Allowed Operations |
|---|---|---|
| resync_start | <ul style="list-style-type: none"> ■ Resync Snapshot (failed) | <ul style="list-style-type: none"> ■ If the Resync Snapshot operation failed, contact your system administrator for help. You can use the VxVM utilities to resynchronize the snapshot volumes, then use the Create Snapshot operation with the Force snapshot creation option for the subsequent snapshot. |
| reverse_resync_begin_start | <ul style="list-style-type: none"> ■ Reverse Resync Snapshot with the begin option (failed) | <ul style="list-style-type: none"> ■ Contact VERITAS support. |
| reverse_resync_begin_end | <ul style="list-style-type: none"> ■ Reverse Resync Snapshot with the begin option (successful) | <ul style="list-style-type: none"> ■ Reverse Resync with the commit option ■ Reverse Resync with the abort option |
| reverse_resync_abort_start | <ul style="list-style-type: none"> ■ Reverse Resync Snapshot with the abort option (failed) | <ul style="list-style-type: none"> ■ Contact VERITAS support. |
| reverse_resync_abort_end | <ul style="list-style-type: none"> ■ Reverse Resync Snapshot with the abort option (successful) | <ul style="list-style-type: none"> ■ Reverse Resync Snapshot with the begin option ■ Resync Snapshot ■ Create Snapshot Database with the Restart database option |
| reverse_resync_commit_start | <ul style="list-style-type: none"> ■ Reverse Resync Snapshot with the commit option (failed) | <ul style="list-style-type: none"> ■ Contact VERITAS support. |
| restartdb_start | <ul style="list-style-type: none"> ■ Start Up Snapshot Database with the Restart database option (failed) | <ul style="list-style-type: none"> ■ Try to start the snapshot database manually. |

Table A-1 (continued)

| Snapshot status (as seen in the Snapplan State field) | Completed Operations | Allowed Operations |
|---|---|---|
| restartdb_end | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Restart database option (successful) | <ul style="list-style-type: none"> ■ Shut Down Database with the unmount option |
| restartdb_end | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Restart database option (successful) | <ul style="list-style-type: none"> ■ Unmount Snapshot Database with the unmount option |
| mountdb_start | <ul style="list-style-type: none"> ■ <code>dbed_vmclonedb -o mountdb</code> command failed from the CLI Note: This option is not supported in the GUI. | <ul style="list-style-type: none"> ■ Recover the snapshot database manually, then run the <code>dbed_vmclonedb -o update_status</code> command from the CLI Note: This option is not supported in the GUI. |
| mountdb_end | <ul style="list-style-type: none"> ■ <code>dbed_vmclonedb -o mountdb</code> command from the CLI was successful Note: This option is not supported in the GUI. | <ul style="list-style-type: none"> ■ Unmount Database FlashSnap ■ <code>dbed_vmclonedb -o update_status</code> command from the CLI Note: This option is not supported in the GUI. |
| recoverdb_start | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Restart database option (failed) | <ul style="list-style-type: none"> ■ Recover the snapshot database manually, then run the <code>dbed_vmclonedb -o update_status</code> command from the CLI Note: This option is not supported in the GUI. |

Table A-1 (continued)

| Snapshot status (as seen in the Snapplan State field) | Completed Operations | Allowed Operations |
|---|--|--|
| recoverdb_start | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Restart database option (failed) | <ul style="list-style-type: none"> ■ Recover the snapshot database manually, then run the <code>db2ed_vmclonedb -o update_status</code> command from the CLI Note: This option is not supported in the GUI. |
| recoverdb_end | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Restart database option (successful) | <ul style="list-style-type: none"> ■ Shut Down Database with the umount option |
| recoverdb_end | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Restart database option (successful) | <ul style="list-style-type: none"> ■ Unmount Snapshot Database with the umount option |
| umount_start | <ul style="list-style-type: none"> ■ <code>dbed_vmclonedb -o umount</code> command failed from the CLI | <ul style="list-style-type: none"> ■ Verify that your file system(s) are not busy and retry the command. |
| umount_start | <ul style="list-style-type: none"> ■ <code>db2ed_vmclonedb -o umount</code> command failed from the CLI Note: This option is not supported in the GUI. | <ul style="list-style-type: none"> ■ Verify that your file system(s) are not busy and retry the command. |
| umount_end | <ul style="list-style-type: none"> ■ <code>dbed_vmclonedb -o umount</code> command from the CLI was successful Note: This option is not supported in the GUI. | <ul style="list-style-type: none"> ■ Start Up Snapshot Database with the restart database option ■ Resync Snapshot ■ Reverse Resync Snapshot with the begin option |

Table A-1 (continued)

| Snapshot status (as seen in the Snapplan State field) | Completed Operations | Allowed Operations |
|---|--|--|
| umount_end | <ul style="list-style-type: none"> ■ db2ed_vmclonedb -o umount command from the CLI was successful <p>Note: This option is not supported in the GUI.</p> | <ul style="list-style-type: none"> ■ Restart Snapshot Database with the restart database option ■ Resync Snapshot ■ Reverse Resync Snapshot with the begin option |

Snapshot database status details

To view snapshot database status information from the GUI, click on a specific snapplan in the object tree. The database status can be seen on the right side of the window in the **Database Status** field. For information regarding the various status values, see the following table:

Table A-2

| Database status (as seen in the Database State field) | Completed operations |
|---|---|
| init | <ul style="list-style-type: none"> ■ Modify/Validate Snapplan (successful) ■ Create Snapshot (successful) ■ Reverse Resync Snapshot (successful) |
| init Db | <ul style="list-style-type: none"> ■ Modify/Validate Snapplan (successful) ■ Create Snapshot (successful) ■ Reverse Resync Snapshot (successful) |
| database_recovered | <ul style="list-style-type: none"> ■ Start Up Snapshot Database with the startup database option (successful) |
| database_recovered | <ul style="list-style-type: none"> ■ Create Snapshot Database with the Create database option (successful) |

Accessibility and Veritas Storage Foundation for DB2

This appendix includes the following topics:

- [About Veritas Storage Foundation for DB2 accessibility](#)
- [Keyboard conventions](#)
- [Keyboard navigation and shortcuts in Veritas Storage foundation for DB2](#)
- [Keyboard shortcuts](#)
- [Menu hot keys](#)
- [Support for accessibility settings and assistive technologies](#)

About Veritas Storage Foundation for DB2 accessibility

Veritas products meet federal accessibility requirements for software as defined in Section 508 of the Rehabilitation Act:

<http://www.access-board.gov/508.htm>

Veritas Storage Foundation for DB2 Graphical User Interface (GUI) inherits the font size and color settings from the operating system it is running on.

Keyboard shortcuts are available for all major GUI operations and menu items. Veritas products are compatible with operating system accessibility settings as well as a variety of assistive technologies. All manuals also are provided as accessible PDF files, and the online help is provided as HTML displayed in a compliant viewer.

The following topics detail accessibility features and compliance in Veritas Storage Foundation for DB2:

- [Keyboard conventions](#)
- [Keyboard shortcuts](#)
- [Menu hot keys](#)
- [Support for accessibility settings and assistive technologies](#)

Keyboard conventions

All program functions and menu items are accessible using the keyboard exclusively. Veritas Storage Foundation for DB2 uses standard operating system navigation keys and keyboard shortcuts.

To see a table of the standard Microsoft navigation keys and keyboard shortcuts, select your version of Microsoft Windows from the drop-down listbox at:

<http://www.microsoft.com/enable/products/keyboardassist.aspx>

Veritas Enterprise Administrator (VEA) has the following exceptions and additions to the Microsoft Keyboard standard:

Table B-1 VEA exceptions to Microsoft keyboard conventions

| Convention | Mnemonic | Accelerator |
|------------|----------|-------------|
| Cancel | C | ALT+C |
| OK | O | ALT+O |
| Refresh | e | None |
| Yes | Y | ALT+Y |
| No | N | ALT+N |
| New | N | ALT+N |

Table B-2 VEA additions to Microsoft keyboard conventions

| Convention | Mnemonic | Accelerator |
|------------|----------|-------------|
| Connect | C | None |
| Contents | C | None |
| Delete Now | D | ALT+ D |

Table B-2 VEA additions to Microsoft keyboard conventions (*continued*)

| Convention | Mnemonic | Accelerator |
|--------------------------------|----------|-------------|
| Delete Temporary Files on Exit | e | ALT+ E |
| Disconnect | D | None |
| Help | H | ALT+ H |
| Number of Items in History | N | ALT+ N |
| Rescan | R | None |
| Set to defaults | S | ALT+ S |
| Remember Password | R | ALT+ R |
| Password | P | ALT+P |
| Username | U | ALT+ U |

Table B-3 VEA Help additions to Microsoft keyboard conventions

| Convention | Mnemonic | Accelerator |
|----------------------|----------|-------------|
| Find in Topic.. | F | Ctrl + F |
| Find Next | n | F3 |
| Find Previous | P | None |
| Hide Navigation Tabs | H | None |
| Index | I | None |
| Print | P | None |
| Search | S | None |
| Synchronize | n | None |
| Close | o | ALT+o |

Keyboard navigation in the GUI

You can navigate and use Veritas Storage Foundation for DB2 with only the keyboard. In the GUI, the current active tree or table has a dark blue highlight, and the current active tab, radio button, or checkbox is enclosed within a rectangle formed by dotted lines. These areas are said to have focus and will respond to commands.

All Veritas GUIs use the following keyboard navigation standards:

- Tab moves the focus to the next active area, field, or control, following a preset sequence. Shift+Tab moves the focus in the reverse direction through the sequence.
- Ctrl+Tab exits any Console area that you internally navigate with Tab.
- Up and Down arrow keys move focus up and down the items of a list.
- ALT in combination with the underlined mnemonic letter for a field or command button shifts the focus to that field or button.
- Enter activates your selection. For example, after pressing Tab to select the **N**ext button in a wizard panel, press Enter to display the next screen.

Navigation in dialog boxes

Dialog boxes contain groups of controls necessary to set options or settings for programs.

Here are some general rules about dialog box navigation:

- Tab moves focus between controls within the dialog box along a preset sequence.
- Controls displaying a mnemonic (an underlined letter) can be selected regardless of focus by typing ALT and the underlined letter.
- A dark border indicates the default command button. Press Enter at any time to choose the button with a dark border.
- ALT C chooses the **C**ancel button if one exists.
- Spacebar chooses a control you select with Tab.
- Spacebar changes the state of a checkbox or radio button that has focus. Typing a mnemonic (if one is available) will move the focus to the checkbox or radio button and change its state.
- Arrow keys move focus within listboxes, sliders, groups of option controls, or groups of page tabs.
- Items that cannot be changed are not visited by the Tab key sequence. Options that are unavailable are grayed-out and can neither be selected nor given focus.

While the controls described here are typically found in dialog boxes, they also can occur in other contexts. The same navigation standards will apply.

Tabbed dialog boxes

Some dialog boxes use tabbed pages to subcategorize groups of many options. Each tabbed page contains different groups of controls. Use Tab to move the focus between tabs within a dialog box. Typing the mnemonic for a tab also moves the focus to the tab and displays its page of controls.

The following table lists keyboard navigation rules within tabbed dialog boxes:

Table B-4 Keyboard navigation in tabbed dialog boxes

| Keyboard input | Result |
|---------------------------|--|
| Ctrl+Page Down | Switches to the next tab and displays the page |
| Ctrl+Page Up | Switches to the previous tab and displays the page |
| Right arrow or Left arrow | When the focus is on a tab selector, chooses the next or previous tab in the current row and displays the page |

Listboxes

Listboxes display a column of available choices.

There are different kinds of listboxes with additional navigation conventions:

- Drop-down listboxes by default show only the selected item. A small button to the right of the control shows a downward-pointing arrow. Select the arrow to display more items from the listbox. If there are more choices than can fit in the preset listbox area, a slider appears along the side of the listbox. Show or hide the list using F4. Enter selects or deselects an item.
- Extended selection listboxes support selecting single items, blocks of items, or combinations of the two. After selecting an item, hold down Shift+ or Ctrl+navigation keys to select or deselect additional items or blocks of items.

Keyboard navigation and shortcuts in Veritas Storage foundation for DB2

All program functions and menu items are accessible using the keyboard exclusively. Veritas Storage Foundation for DB2 Adaptive Server Enterprise (ASE) uses standard operating system navigation keys and keyboard shortcuts.

To see a table of the standard Microsoft navigation keys and keyboard shortcuts, select your version of Microsoft Windows from the drop-down listbox at:

<http://www.microsoft.com/enable/products/keyboard.aspx>

Veritas Enterprise Administrator (VEA) has the following exceptions and additions to the Microsoft Keyboard standard:

Table B-5 VEA additions to Microsoft keyboard conventions

| Convention | Mnemonic | Accelerator |
|--------------------------------|----------|-------------|
| Connect | C | None |
| Contents | C | None |
| Delete Now | D | ALT+ D |
| Delete Temporary Files on Exit | e | ALT+ E |
| Disconnect | D | None |
| Help | H | ALT+ H |
| Number of Items in History | N | ALT+ N |
| Rescan | s | None |
| Set to defaults | S | ALT+ S |
| Remember Password | R | ALT+ R |
| Username | U | ALT+ U |

Table B-6 VEA exceptions to Microsoft keyboard conventions

| Convention | Mnemonic | Accelerator |
|------------|----------|-------------|
| Cancel | C | ALT+C |
| OK | O | ALT+O |
| Refresh | e | None |
| Yes | Y | ALT+Y |
| No | N | ALT+N |
| Shift+F10 | None | None |

Table B-7 VEA Help additions to Microsoft keyboard conventions

| Convention | Mnemonic | Accelerator |
|-----------------|----------|-------------|
| Find in Topic.. | F | Ctrl + F |

Table B-7 VEA Help additions to Microsoft keyboard conventions (*continued*)

| Convention | Mnemonic | Accelerator |
|----------------------|----------|-------------|
| Find Next | n | F3 |
| Find Previous | P | None |
| Hide Navigation Tabs | H | None |
| Index | I | None |
| Print | P | None |
| Search | S | None |
| Synchronize | n | None |

General keyboard navigation within the GUI

You can navigate and use Veritas Storage Foundation for DB2 with only the keyboard. In the GUI, the active tree or table node has a highlight, and the current active tab, radio button, or checkbox is enclosed within a rectangle formed by dotted lines. These areas are said to have focus and will respond to commands.

All Veritas GUIs use the following keyboard navigation standards:

- Tab moves the focus to the next active area, field, or control, following a preset sequence. Shift+Tab moves the focus in the reverse direction through the sequence.
- Ctrl+Tab exits any Console area that you internally navigate with Tab.
- Up and Down arrow keys move focus up and down the items of a list.
- ALT in combination with the underlined mnemonic letter for a field or command button shifts the focus to that field or button.
- Enter activates the default push button. For any other other button, tab to it and press Enter. For example, after pressing Tab to select the **N**ext button in a wizard panel, press Enter to display the next screen. Press Enter will activate the default push button. For the other button, tab to it and press Enter.

Keyboard navigation within dialog boxes

Dialog boxes contain groups of controls necessary to set options or settings for programs.

Here are some general rules about dialog box navigation:

- Tab moves focus between controls within the dialog box along a preset sequence.
- Controls displaying a mnemonic (an underlined letter) can be selected regardless of focus by typing ALT and the underlined letter.
- A dark border indicates the default command button. Press Enter at any time to choose the button with a dark border.
- ALT C chooses the **Cancel** button if one exists. The ESC key also cancels an action. You can press ESC any time even if there is no **Cancel** button.
- Spacebar chooses a control you select with Tab.
- Spacebar changes the state of a checkbox or radio button that has focus. Typing a mnemonic (if one is available) will move the focus to the checkbox or radio button and change its state.
- Arrow keys move focus within listboxes, sliders, groups of option controls, or groups of page tabs.
- Items that cannot be changed are not visited by the Tab key sequence. Options that are unavailable are grayed-out and can neither be selected nor given focus.

While the controls described here are typically found in dialog boxes, they also can occur in other contexts. The same navigation standards will apply.

Tabbed dialog boxes

Some dialog boxes use tabbed pages to subcategorize groups of many options. Each tabbed page contains different groups of controls. Use Tab to move the focus between tabs within a dialog box. Typing the mnemonic for a tab also moves the focus to the tab and displays its page of controls.

Table B-8 Keyboard navigation within tabbed dialog boxes

| Keyboard input | Result |
|---------------------------|--|
| Ctrl+Page Down | Switches to the next tab and displays the page |
| Ctrl+Page Up | Switches to the previous tab and displays the page |
| Right arrow or Left arrow | When the focus is on a tab selector, chooses the next or previous tab in the current row and displays the page |

Listboxes

Listboxes display a column of available choices.

There are different kinds of listboxes with additional navigation conventions:

- Drop-down listboxes by default show only the selected item. A small button to the right of the control shows a downward-pointing arrow. Select the arrow to display more items from the listbox. If there are more choices than can fit in the preset listbox area, a slider appears along the side of the listbox. Show or hide the list using F4. Enter selects or deselects an item.
- Extended selection listboxes support selecting single items, blocks of items, or combinations of the two. After selecting an item, hold down Shift+ or Ctrl+navigation keys to select or deselect additional items or blocks of items.

Keyboard shortcuts

All menu items can be selected by using accelerator or mnemonic keyboard shortcuts. An accelerator is a key combination that provides shortcut access to a GUI function. A mnemonic (sometimes referred to as a “hot key”) is a single-key equivalent (used in combination with the ALT key) for selecting GUI components such as menu items. The mnemonic “hot key” letter is underlined in the GUI.

For example:

- ALT to go into menu pull-down mode
- F key to access the **F**ile menu
- O key to activate the open command

Mnemonics are case-insensitive. Keys can be pressed sequentially instead of simultaneously.

Table B-9 Keyboard shortcuts

| Keyboard input | Action |
|---|--|
| Tab, Shift-Tab (for reversing the action) | Navigates between main components of the user-interface |
| Shift-F10 | Display Context-sensitive menu |
| Ctrl-A | Selects all items in list |
| F3 | Find Next |
| Enter, Return | Activates default button (does not require keyboard focus) |

Routine functions such as opening, saving, and printing files can be performed using the standard Microsoft keyboard shortcuts.

Keyboard shortcuts are not case-sensitive. Mnemonic keystrokes may be pressed either sequentially or simultaneously. All menu items have mnemonics, but not all menu items have accelerators.

Keyboard navigation

The following table lists some of the keys frequently used to navigate with the keyboard:

Table B-10 Keyboard navigation

| Keyboard input | Result |
|--------------------|---|
| TAB | Move forward between panes in the active Console window. |
| SHIFT+TAB | Move backwards between panes in the active Console window. |
| SHIFT+ UP ARROW | Move up one item in the tree view. |
| SHIFT+DOWN ARROW | Move down one item in the tree view. |
| SHIFT+PAGE UP | Move to the top item visible in the tree view. |
| HOME | Move to the first item in the tree view. |
| END | Move to the last item in the tree view. |
| RIGHT ARROW | Expands the highlighted item. If the highlighted item does not contain hidden items, behaves like DOWN ARROW. |
| LEFT ARROW | Collapses the highlighted item. If the highlighted item does not contain expanded items, behaves like UP ARROW. |

Menu hot keys

The following table lists the hot keys associated with the different menus. In some cases, the options available depend upon the object selected.

Table B-11 Menu hot keys

| Action | Keyboard Input |
|-------------------|---|
| File | ALT+F to open the menu, then: <ul style="list-style-type: none"> ■ C- Display the connection dialog window ■ D-Display the disconnection dialog window ■ n-Create new window ■ r-Display the properties of the connected host ■ u-Page setup for print ■ w-Print preview ■ P-Print the page ■ x -Exit |
| ToolsTools | ALT+T to open the menu, then: <ul style="list-style-type: none"> ■ P-Set the display preferences ■ M-Manage user profiles ■ E-Display the error console ■ S-Search for storage devices |
| Actions | ALT+ A to open the menu, then: <ul style="list-style-type: none"> ■ e-Refresh ■ s-rescan |
| View | ALT+ V to open the menu, then: <ul style="list-style-type: none"> ■ c-Select host connection view ■ g-View alert logs ■ o-View control panel ■ s-Return to summary page view |
| Window | ALT+ W to open the menu, then: <ul style="list-style-type: none"> ■ T-Tear off a component from being shown ■ 1-To tear off details ■ 2-To tear off the system details |
| Help | ALT+H to open the menu, then: <ul style="list-style-type: none"> ■ C-Display the contents ■ b-Display information about the Help viewer |

Table B-11 Menu hot keys (*continued*)

| Action | Keyboard Input |
|-------------------------|--|
| DB2 Instances | ALT+D to open the menu, then: <ul style="list-style-type: none"> ■ S-Start the DB2 instance ■ U-Update the rescan intervals to be faster or slower ■ E-Rescan system information ■ R-View properties |
| DB2 Instance | ALT+D to open the menu, then: <ul style="list-style-type: none"> ■ C-Create Snapshot Database ■ D-Shut down the DB2 Instance ■ E-Rescan system information ■ S-Start the DB2 instance ■ R-View properties |
| DB2 Database | ALT+D to open the menu, then: <ul style="list-style-type: none"> ■ o-Resync the database repository ■ h-Check system configuration ■ a-Save system Configuration ■ E-Rescan system information ■ C-Create Clone Database ■ R-View properties |
| Tablespaces | <ul style="list-style-type: none"> ■ R-Rescan all tablespaces ■ R-View properties |
| Container | Alt+C to open the menu, then: <ul style="list-style-type: none"> ■ C-Conversion, to convert oracle datafiles to Quick I/O files. ■ T-Topology, display mapping of the topology tree ■ S-Statistics, generates sampling data statistics ■ R-View properties |
| Oracle Databases | ALT+O to open the menu, then: <ul style="list-style-type: none"> ■ S- Start the Oracle Database ■ D-Start the Snapshot Database ■ v-Reverse resync a snapshot ■ U- Update the rescan intervals ■ R-Rescan the databases |

Table B-11 Menu hot keys (*continued*)

| Action | Keyboard Input |
|----------------------------|---|
| Oracle Instance | ALT+ O to open the menu, then: <ul style="list-style-type: none"> ■ S- Startup the database instance ■ D- Shut down the database ■ C-Create Clone Database ■ o-Resync Repository ■ h- Check System configuration ■ a- Save system configuration ■ R-Rescan Oracle instance |
| Monitoring Agent | ALT+M to open the menu, then: <ul style="list-style-type: none"> ■ S-Start the monitoring agent ■ o-Stop the monitoring agent ■ E-Enable the agent at boot time ■ D-Disable the agent at boot time |
| Snapplans | ALT+ S to open menu, then: <ul style="list-style-type: none"> ■ C-Create a snapplan ■ M-Modify/Validate a snapplan ■ R-Rescan Snapplans |
| Snapplan | ALT+ S to open menu, then: <ul style="list-style-type: none"> ■ M-Modify/Validate a snapplan ■ R-Remove a Snapplan ■ C-Create a snapshot using snapplan ■ y-Resync a snapshot ■ v-Reverse Resync a snapshot ■ L-View log for the Snapplan |
| Storage Checkpoints | ALT+S to open menu, then: <ul style="list-style-type: none"> ■ C-Create a Storage Checkpoint ■ P-Create a Storage Checkpoint policy ■ R- Rescan system information |

Table B-11 Menu hot keys (*continued*)

| Action | Keyboard Input |
|---------------------------|---|
| Storage Checkpoint | ALT+S to open menu, then: <ul style="list-style-type: none"> ■ B-Rollback a storage checkpoint ■ M-Mount a storage checkpoint ■ U-Unmount a storage checkpoint ■ R-Remove a storage checkpoint ■ C-Create a Clone Database |
| Tablespaces | ALT+ S to open the menu, then: <ul style="list-style-type: none"> ■ R-Rescan tablespaces |
| Datafile | ALT+D to open the menu, then: <ul style="list-style-type: none"> ■ T- Topology/Statistics to generate datafile statistics ■ C-Conversion (convert Oracle datafiles to Quick I/O files) |

Support for accessibility settings and assistive technologies

Veritas software responds to operating system accessibility settings.

Veritas products are compatible with Microsoft's accessibility utilities. In Windows 2000, accessibility options involving keyboard responsiveness, display contrast, alert sounds, and mouse operation can be set through the Control Panel (**Start > Settings > Control Panel > Accessibility Options**) and through the Accessibility Wizard (**Start > Programs > Accessories > Accessibility > Accessibility Wizard**).

Note: Though all graphics in Veritas documentation can be read by screen readers, setting your screen reader to ignore graphics may improve performance.

Index

A

- accessibility
 - assistive technology support 136
 - dialog boxes 126, 129
 - keyboard navigation 125, 129
 - keyboard shortcuts 131
- ARCHIVELOG mode 66
- assistive technology support 136
- automating tasks
 - adding a scheduled task 102
 - creating checkpoint database 104
 - creating clone database using checkpoint 105
 - creating clone database using FlashSnap 105
 - creating FlashSnap database 104
 - modifying a scheduled task 103
 - set schedule parameters 106
 - viewing scheduled tasks 102

B

- browsers, supported for Web GUI 24

C

- clone database
 - creating 53
 - creating a clone with a snapshot 80
 - creating a new snapplan for clone 55
 - creating a snapshot for clone 58, 60
 - creating with Database FlashSnap 55
 - creating with Storage Checkpoints 53
 - removing a clone 61
 - removing a clone database 62
 - removing a snapshot database clone 62
 - startup clone snapshot 59
 - validating a snapplan 56
- configuration maintenance
 - checking your system configuration 114
 - converting Quick I/O files to regular containers 113
 - converting regular containers to Quick I/O 112
 - displaying tablespace information 110

- configuration maintenance (*continued*)
 - maintaining system configurations 109
 - rescanning tablespace information 110
 - rescanning the database instance 110
 - resynchronizing the SFDB repository 110
 - saving the system configuration 116
 - updating rescan intervals 113
 - viewing I/O statistics 111
 - viewing storage mapping 111

D

- Database FlashSnap
 - aborting reverse resynchronization 83
 - creating a clone 55
 - creating a clone snapshot clone 58
 - creating a clone snapshot for clone 60
 - creating a clone with a snapshot 80
 - creating a new snapplan for clone 55
 - creating a snapplan 74
 - creating a snapshot 79
 - modifying a snapplan 76
 - prerequisites 73
 - removing a snapplan 78
 - removing a snapshot clone database 62
 - resynchronizing a database to a snapshot 82
 - resynchronizing a snapshot to a database 81
 - startup clone snapshot 59
 - validating a clone snapplan 56
 - validating a snapplan 76
 - viewing the snapplan 84
- database information
 - displaying tablespace information 110
 - rescanning 110
 - rescanning tablespace information 110
- databases
 - restoring using Storage Checkpoints 70
 - setting up 39
- DB2
 - creating a snapshot database 48
 - restarting an instance 51
 - shutting down an instance 50

DB2 (*continued*)
 starting an instance 48
 unmounting a snapshot database 49

defragmenting a file system 46

dirty region logging 44

disk group
 naming a disk group 41

disk groups
 configuration guidelines 41

DRL 44

DSS workloads
 guidelines 44

Dynamic Database Storage Tiering
 getting reports 96
 moving files and volumes 95
 setting automatic policy 91
 setting up parameters 88

F

fast recovery 44

file locks
 during Storage Rollback 70

file systems
 configuration guidelines 44

fragmentation 45
 controlling 46
 defragmenting a file system 46
 types 46

free space 69

G

guidelines
 creating file systems 44
 disk groups 41
 for DSS workloads 44
 for OLTP workloads 44
 striped volumes 44
 volumes 43

L

lock files
 during Storage Rollback 70

M

maintaining your configuration
 about configuration maintenance 109
 checking your system configuration 114

maintaining your configuration (*continued*)
 converting Quick I/O files to regular
 containers 113
 converting regular containers to Quick I/O 112
 displaying tablespace information 110
 rescanning tablespace information 110
 rescanning the database instance 110
 resynchronizing the SFDB repository 110
 saving the system configuration 116
 updating rescan intervals 113
 viewing I/O statistics 111
 viewing storage mapping 111

menus
 Storage Rollback Administration 69

mirroring
 choosing 43

O

OLTP workloads
 guidelines 44

Q

Quick I/O Files
 converting Quick I/O files to regular
 containers 113
 converting regular containers to Quick I/O 112

R

RAID-5 44
 choosing 43
 performance 43

RAID-5 log 44

redo logs
 configuration guidelines 45
 creating a file system 45

restoring
 databases 69
 files 69
 tablespaces 69

resynchronization
 using DRL logs 44
 using RAID-5 logs 44

rolling back
 databases to a Storage Checkpoint 70

S

scheduling tasks
 adding a scheduled task 102

- scheduling tasks (*continued*)
 - creating checkpoint database 104
 - creating clone database using checkpoint 105
 - creating clone database using FlashSnap 105
 - creating FlashSnap database 104
 - modifying a scheduled task 103
 - set schedule parameters 106
 - viewing scheduled tasks 102
 - selecting volume layouts 42
 - SFDB repository
 - resynchronizing 110
 - snapplan
 - abort reverse resync 83
 - creating 74
 - modifying 76
 - removing 78
 - validating 76
 - view the snapplan log 84
 - snapshot
 - creating 79
 - creating a clone with a snapshot 80
 - resynchronizing a snapshot to a database 81
 - reverse resync 82
 - Storage Checkpoint
 - removing a clone Storage Checkpoint clone 62
 - Storage Checkpoints
 - concepts 65
 - creating 66
 - creating a clone 53
 - databases 69
 - mounting 67
 - removing 69
 - restoring files 69
 - rolling back 69
 - tablespaces 69
 - unmounting 68
 - viewing details 67
 - Storage Mapping
 - viewing I/O statistics 111
 - viewing topology 111
 - Storage Rollback
 - databases 69
 - restoring files 69
 - tablespaces 69
 - Storage Rollback Administration menu 69
 - Storage_agent 39
 - stripe unit sizes
 - choosing 43
 - striped volumes 44
 - configuration guidelines 44
- T**
- tiered storage
 - getting reports 96
 - moving files and volumes 95
 - setting automatic policy 91
 - setting up parameters 88
- V**
- Veritas Storage Foundation for DB2
 - accessibility
 - assistive technology support 136
 - dialog boxes 126, 129
 - keyboard navigation 125, 129
 - keyboard shortcuts 131
 - adding VEA users 15
 - closing the Java GUI 25
 - GUI components 9
 - GUI functions 13
 - Java GUI prerequisites 24
 - Java GUI supported tasks 26
 - opening the Java GUI (UNIX) 24
 - opening the Java GUI (Windows) 25
 - removing VEA users 19
 - shutting down the DBED agent 23
 - shutting down VEA 22
 - starting the DBED agent 23
 - starting VEA 21
 - VEA prerequisites 15
 - Web GUI supported tasks 34
 - volume layouts
 - selecting 42
 - volumes
 - configuration guidelines 43
- W**
- workloads
 - write-intensive 44