



**SIOS Protection Suite for Linux
SAP MaxDB Recovery Kit
v8.3.2**

Administration Guide

Nov 2014

This document and the information herein is the property of SIOS Technology Corp. (previously known as SteelEye® Technology, Inc.) and all unauthorized use and reproduction is prohibited. SIOS Technology Corp. makes no warranties with respect to the contents of this document and reserves the right to revise this publication and make changes to the products described herein without prior notification. It is the policy of SIOS Technology Corp. to improve products as new technology, components and software become available. SIOS Technology Corp., therefore, reserves the right to change specifications without prior notice.

LifeKeeper, SteelEye and SteelEye DataKeeper are registered trademarks of SIOS Technology Corp.

Other brand and product names used herein are for identification purposes only and may be trademarks of their respective companies.

To maintain the quality of our publications, we welcome your comments on the accuracy, clarity, organization, and value of this document.

Address correspondence to:
ip@us.sios.com

Copyright © 2014
By SIOS Technology Corp.
San Mateo, CA U.S.A.
All rights reserved

Table of Contents

Chapter 1: Introduction	1
SAP DB / MaxDB Recovery Kit Administration Guide	1
Document Contents	1
SIOS Protection Suite Documentation	1
SAP DB Documentation	2
Chapter 2: SAP DB / MaxDB Recovery Kit Requirements	3
Hardware Requirements	3
Software Requirements	3
Chapter 3: Overview	4
SAP DB / MaxDB Resource Hierarchy	4
Chapter 4: Configuration Considerations	6
Using Raw I/O	6
Using Mirrored File Systems with DataKeeper	6
Active/Standby Considerations	6
Active/Standby Scenarios	6
Scenario 1	6
Scenario 2	7
Active/Standby Configuration Example	8
Active/Active Considerations	8
Active/Active Scenario	8
Active/Active Configuration Example	9
Chapter 5: Configuring SAP DB / MaxDB with SPS	11
Install the SAP DB / MaxDB Software	11
Create the SAP DB / MaxDB Database	11
Create the User_Key	12

Install the SPS Software	14
Chapter 6: Resource Configuration Tasks	15
Creating an SAP DB Resource Hierarchy	16
Extending an SAP DB Resource Hierarchy	17
Unextending an SAP DB Resource Hierarchy	18
Deleting an SAP DB Resource Hierarchy	19
Testing Your Resource Hierarchy	19
Performing a Manual Switchover from the LifeKeeper GUI	19
Chapter 7: Resource Hierarchy Administration	20
Modifying User_Keys	20
Modifying OS User	20
Updating Parameters	20
Chapter 8: Troubleshooting	21
General Tips	21
SAP DB / Max DB Recovery Kit Error Messages	22
Action: Verify that the underlying device is a shared device.	22
Appendix: Creating Device Spaces Using Raw I/O	24
Requirements	24
Naming Conventions	24
Raw I/O Setup Steps	24
Adding a Device Space after Creating Hierarchy	25

Chapter 1: Introduction

SAP DB / MaxDB Recovery Kit Administration Guide

SAP DB / MaxDB (SAP DB) is a SQL-based, industrial-strength database system that can be deployed for a wide array of purposes. It is highly scalable, platform-independent and provides full transaction support. The database system was originally owned by SAP but has since been released to the Open Source community.

The SAP DB / MaxDB Recovery Kit provides fault resilient protection for SAP DB databases in an SPS for Linux environment.

Document Contents

This guide includes the following topics to help you successfully define and manage your SAP DB hierarchy:

- [SAP DB / Max DB Recovery Kit Requirements](#). Lists the hardware and software necessary to properly set up, install and operate the SAP DB / Max DB Recovery Kit.
- [Overview](#). Describes the SAP DB / Max DB Recovery Kit's features and functionality.
- [Configuration Considerations](#). Contains information to consider before you install and configure the SAP DB / Max DB Recovery Kit.
- [Configuring SAP DB with SAP](#). Provides instructions for installing and configuring the SAP DB software and SAP software.
- [Resource Configuration Tasks](#). Describes the various functions you may perform on your hierarchies using the LifeKeeper GUI: create, extend, delete and unextend.
- [Hierarchy Administration](#). Provides important recommendations for ongoing administration of the SAP DB hierarchy.
- [Troubleshooting](#). Lists and describes the error messages associated with the SAP DB / Max DB Recovery Kit.
- [Appendix](#). Provides requirements and instructions for setting up raw devices for use with the SAP DB / Max DB Recovery Kit.

SIOS Protection Suite Documentation

The following SPS product documentation is available from the SIOS Technology Corp. website:

- SPS for Linux Release Notes
- SPS for Linux Technical Documentation

- SPS for Linux Installation Guide
- Optional Recovery Kit Documentation

This documentation, along with documentation associated with optional LifeKeeper Recovery Kits, is provided on the SIOS Technology Corp. website at:

<http://docs.us.sios.com/>

and from the Help menu in the LifeKeeper GUI.

SAP DB Documentation

You can find the SAP DB/MaxDB documentation, including the Installation Guide, User Manual and Reference Manual, at the following locations on the web:

SAP DB 7.5

<http://mysql.netvisao.pt/downloads/maxdb/7.5.00.html>

SAP DB 7.4

http://www.sapdb.org/7.4/sap_db_documentation.htm

SAP DB 7.3

http://www.sapdb.org/sap_db_documentation.htm

Chapter 2: SAP DB / MaxDB Recovery Kit Requirements

Your SPS configuration must meet the following requirements prior to the installation of SPS for Linux SAP DB / MaxDB Recovery Kit. Please see the SPS for Linux Installation Guide for specific instructions regarding the installation and configuration of your SPS hardware and software.

Hardware Requirements

- **Servers** - Servers should be configured in accordance with the requirements described in the SPS for Linux Installation Guide and the SPS for Linux Release Notes.
- **IP Network Interface Cards** - Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that an SPS cluster requires at least two communications paths. Two separate LAN-based communication paths using dual independent sub-nets are recommended for heartbeats, and at least one of these should be configured as a private network. Using a combination of TCP and TTY heartbeats is also supported.

Software Requirements


- **TCP/IP Software** – Each server in your SPS configuration requires TCP/IP Software.
- **SAP DB / MaxDB Software** – SPS supports versions 7.3.00 (7.3.0.21 and higher) and 7.4.3 (7.4.3.10) of the SAP DB software. These versions can be downloaded directly from www.sapdb.org/sap_db_downloads.htm. SPS also supports versions 7.5.x, 7.6.x, 7.7.x and 7.8.x of the MaxDB software. The MaxDB version can be downloaded from <http://www.sdn.sap.com/irj/scn/maxdb-downloads>.
Note: The same version of the SAP DB software must be installed on all servers in the cluster.
- **SPS Software** – It is imperative that you install the same version of the SPS software and apply the same versions of the SPS software patches to each server in your cluster.
- **SPS for Linux SAP DB / Max DB Recovery Kit** – The SAP DB / Max DB Recovery Kit is provided on the SPS for Linux Installation Image File (Steeleye-lkSAPDB). It is packaged, installed and removed via Red Hat Package Manager, rpm.

Chapter 3: Overview

The SPS for Linux SAP DB / Max DB Recovery Kit provides a mechanism for protecting SAP DB instances within SPS. The SAP DB software, SPS Core and SAP DB / Max DB Recovery Kit are installed on two or more servers in a cluster. Once the SAP DB database instance is under SPS protection, clients connect to the database using an SPS protected IP address. The SPS protected IP address must be created separately and a dependency made manually between the parent SAP DB resource instance and the child IP address resource. In the event that the SAP DB server fails, SPS will first attempt to recover it on the local server. If the local recovery fails, then SPS will fail over to a backup server.

SAP DB / MaxDB Resource Hierarchy

The following example shows a typical SAP DB / MaxDB resource hierarchy:



The dependencies in the above example correspond to the following protected resources:

Resource	SAP DB Software Component
LKIP.EXAMPLE.COM	protects the switchable IP address used for client connections
sapdb_EXAMPLE /data	protects the database data device space for the EXAMPLE database
sapdb_EXAMPLE /log	protects the database log device space for the EXAMPLE database
sapdb_EXAMPLE /rundirectory	protects the database RUNDIRECTORY for the EXAMPLE database

Resource	SAP DB Software Component
sapdb_EXAMPLE /sys	protects the database system device space for the EXAMPLE database
sapdb_ik/data	protects the independent data path
sapdb_ik/dep	protects the dependent path
sapdb_ik/prog	protects the independent program path

In the event of failover, SPS will bring the file system, IP address and database resources (including all the resource dependencies) in service on a backup server. Clients will be disconnected and will need to re-connect to the server. Any SQL statement that has not been committed will need to be re-entered.

Chapter 4: Configuration Considerations

This section contains information that you should consider before you start to configure and administer the SAP DB / Max DB Recovery Kit.


Using Raw I/O

If you plan to use SAP DB / MaxDB with raw devices, you must install the SPS Raw I/O Recovery Kit from the SPS Installation Image file. You must also properly set up the raw I/O devices prior to use. See the [Appendix](#) for instructions.

Using Mirrored File Systems with DataKeeper

The SAP DB / Max DB Recovery Kit supports the use of DataKeeper as a device space. In addition, the SAP DB / MaxDB software can be installed on mirrored file systems.

For example, a dependent file system for an SAP DB / MaxDB resource would look similar to the following, which shows a file system for the system device space and its dependency, the DataKeeper resource mirror.



Active/Standby Considerations

In an Active/Standby configuration, the backup server is not actively running the SAP DB / MaxDB but stands by in case the primary server experiences a failure. The following scenarios provide specific requirements that must be adhered to when protecting an SAP DB resource instance in Active/Standby configurations.

Active/Standby Scenarios

The typical Active/Standby configurations are explained below in Scenarios 1 and 2.

Scenario 1

The SAP DB *IndepDataPath*, *IndepProgPath* and *DependPath* are installed to **one or more shared file systems on the primary server**.

- The paths *IndepDataPath*, *IndepProgPath*, and *DependPath* must be shared between all servers that will protect the resource instance.

Scenario 2

- The directory structure under `/usr/spool/sql` must be replicated manually to each server in the cluster. This directory structure should not be located on shared storage since it must be accessible from the target server during resource extend operations. (*Please disregard for MaxDB 7.8 as this directory no longer exists.*)
- The registry file `/etc/opt/sdb` must exist on each server in the cluster for MaxDB 7.5.x versions. This file should not be located on shared storage since it must be accessible from the target server during resource extend operations.
- The database instance data device spaces (`data devspaces`), log device spaces (`log devspaces`) and system device spaces (`sys devspaces`) must reside on a shared disk (either shared file system or shared raw device).

Scenario 2


The SAP DB `IndepDataPath` and `IndepProgPath` are installed locally on both servers. The SAP DB `DependPath` can be installed locally or on a shared file system on the primary server.

- The directory structure under `/usr/spool/sql` must exist on all servers with the same permissions as well as the same owner and group. (*Please disregard for MaxDB 7.8 as this directory no longer exists.*)
- The registry file `/etc/opt/sdb` must exist on each server in the cluster for MaxDB 7.5.x versions. This file should not be located on shared storage since it must be accessible from the target server during resource extend operations.
- The database instance data device spaces (`data devspaces`), log device spaces (`log devspaces`) and system device spaces (`sys devspaces`) must reside on a shared disk (either shared file system or shared raw device).
- The database instance run directory (`RUNDIRECTORY`) must be located on shared storage. The value of `RUNDIRECTORY` can be modified via the DBMCLI command `param_directput`. If the value of `RUNDIRECTORY` is modified after the database is created, the database instance must be stopped and restarted to complete the parameter update.
- The database instance config (`<IndepDataPath>/config`) directory structure must exist in the same location on all servers in the cluster where the database instance will be protected. In addition, the parameter files for the database instance must be copied from the template (or primary) server to all backup servers in the cluster. The parameter files must be redistributed to all servers in the cluster after any parameter has been updated. The required files are:

```
config/<db instance>
config/<db instance>.<01>...<N>      (Note: There may be
multiple .<number> files.)
config/<db instance>.cfg
config/<db instance>.pab
config/<db instance>.upc
```

Active/Standby Configuration Example

IndepData = /shr1/data
IndepPrograms = /shr1/programs
DependPath = /shr1/depend



Configuration Notes:

- Both servers use the *IndepProgPath*, *DependPath* and *IndepDataPath* on the shared storage.
- The database instance *DB1* is located on the shared storage. This includes all log device spaces, data device spaces and system device spaces.
- The directory structure */usr/spool/sql* has been replicated to *Server 2*. All entries for the *SAP_DBTech.ini* file have been added. (Please disregard for *MaxDB 7.8* as this directory and file no longer exist.)
- *Server 2* cannot access files and directories on the shared disk while *Server 1* is active.

Active/Active Considerations

In an Active/Active configuration, each server is actively running one SAP DB instance while acting as a backup for the other server in case of failure. The following scenario provides specific requirements that must be adhered to in sequential order when protecting an SAP DB resource instance in an Active/Active configuration.

Active/Active Scenario

The SAP DB *IndepDataPath*, *IndepProgPath* and *DependPath* are installed locally on both servers.

- The directory structure under */usr/spool/sql* must exist on all servers with the same permissions as well as the same owner and group. (Please disregard for *MaxDB 7.8* as this directory no longer exists.)

Active/Active Configuration Example

- The registry file `/etc/opt/sdb` must exist on each server in the cluster for MaxDB 7.5.x versions. This file should not be located on shared storage since it must be accessible from the target server during resource extend operations.
- The database instance data device spaces (`data devspaces`), log device spaces (`log devspaces`) and system device spaces (`sys devspaces`) must reside on a shared disk (either shared file system or shared raw device).
- The database instance run directory (`RUNDIRECTORY`) must be located on shared storage. The value of `RUNDIRECTORY` can be modified via the DBMCLI command `param_directput`. If the value of `RUNDIRECTORY` is modified after the database is created, the database instance must be stopped and restarted to complete the parameter update.
- The database instance config (`<IndepDataPath>/config`) directory structure must exist in the same location on all servers in the cluster where the database instance will be protected. In addition, the parameter files for the database instance must be copied from the template (or primary) server to all backup servers in the cluster. The parameter files must be redistributed to all servers in the cluster after any parameter has been updated. The required files are:

```
config/<db instance>
config/<db instance>.<01>...<N>      (Note: There may be
multiple .<number> files.)
config/<db instance>.cfg
config/<db instance>.pab
config/<db instance>.upc
```


Active/Active Configuration Example

```
IndepData = /usr/sapdb/data
IndepPrograms = /usr/sapdb/prog
DependPath = /usr/sapdb/dep
```

```
IndepData = /usr/sapdb/data
IndepPrograms = /usr/sapdb/prog
DependPath = /usr/sapdb/dep
```

Configuration Notes:

- The *IndepDataPath*, *IndepProgPath* and *DependPath* are locally installed on both servers.
- Each database is configured on separate shared disks. The database instance includes all log device spaces, system device spaces and data device spaces.
- The `RUNDIRECTORY` for each database instance is also on a shared disk.
- The database configuration files for `DB1` have been copied to *Server 2* and the database configuration files for `DB2` have been copied to *Server 1*. The configuration files are located at `<IndepDataPath>/-config/<db instance>`.
- Initially, *Server 1* runs `DB1` and *Server 2* runs `DB2`. In a switchover situation, one server can run both databases.



Chapter 5: Configuring SAP DB / MaxDB with SPS

The following sequence is recommended for installing and configuring the SAP DB / MaxDB database and SPS software. Each of these steps links to detailed tasks that follow.

[Install the SAP DB / MaxDB software](#)

[Create the SAP DB / MaxDB database](#)

[Create the User_Key](#)

[Install the SPS Core and SAP DB / Max DB Recovery Kit](#)

After you have performed these tasks, you will be ready to create the SPS resource hierarchy to protect your SAP DB / MaxDB database.

Install the SAP DB / MaxDB Software

Install the SAP DB /MaxDB software on all servers in the cluster using identical parameters/settings. Refer to the *SAP DB Installation Guide* for details. The following are additional recommendations to ensure that SPS will work with SAP DB / MaxDB:

- A non-root system user (OS User) must exist on all servers as follows:
 - This OS User should be designated as the owner of the SAP DB / MaxDB software installation and subdirectories or have adequate permissions on the software installation path and subdirectories as required to manage a database instance.
 - This OS User must have authority to use the DBMCLI and x_server utilities. The OS User must be able to start and stop the vserver using the x_server commands.
 - The OS User name should contain alpha-numeric characters only.
 - The User ID and Group ID of this OS User must be identical on all servers.
- The SAP DB / MaxDB client software packages must be installed. These packages must include the SAP DB / MaxDB DBMCLI client utility, and the SAP DB / MaxDB x_server utility.
- Each SPS server containing an SAP DB resource hierarchy must have identical service entries in the */etc/services* file for the SAP DB instance.

Create the SAP DB / MaxDB Database

Follow the instructions in your *SAP DB / MaxDB User Manual* to create your database. In addition, please note the following recommendations:

Create the User_Key

- There must be a DBM operator with authority for starting, stopping, obtaining status and obtaining database parameters via client utilities.
- The database instance data device spaces (`data devspaces`), log device spaces (`log devspaces`) and system device spaces (`sys devspaces`) must reside on a shared disk (either shared file system, or shared raw device).
- The SAP DB / MaxDB database name should contain alphanumeric characters only.
- A User_Key is required for use by the SAP DB / Max DB Recovery Kit during operation with the DBMCLI utility. See [Create the User_Key](#) for required parameters.
- After creating your database, you should disable automatic startup of the SAP DB / MaxDB database instance. Once under SPS protection, SPS will handle the start and stop of the database.
- The `SAP_DBTech.ini` file must exist on all servers in `/usr/spool/sql/ini`. If this file does not exist, several SAP DB / MaxDB utilities may return erroneous results. This will also affect the behavior of SPS during resource create and extension. In an Active/Standby configuration, you must manually copy this file to the backup server. (*Please disregard for MaxDB 7.8 as this directory and file no longer exist.*)
- For version 7.5.x, verify the `databases.ini` file exists on all servers in the `IndepDataPath/-config` directory.

Create the User_Key

The SAP DB / MaxDB instance requires several options for a user to successfully access a database instance. These required pieces of information must be passed in to the SAP DB / MaxDB tool being used to access the database instance. The SAP DB / MaxDB software includes the **xuser** tool for simplifying the specification of many required options to SAP DB / MaxDB tools. The **xuser** tool allows you to predefine and save user data. Once this data has been saved, it can be used when you call the **DBMCLI** or other tools requiring user options. This predefined user data is stored in a user key (User_Key). An individual user can manage and maintain several user keys for the same or multiple databases. Each key includes a combination of options including username/password, database name as well as database server name.

The SAP DB / Max DB Recovery Kit requires a valid User_Key for each database instance under protection. This User_Key must be created and accessible by the OS User that owns the database instance. The user information specified for each User_Key must be for a DBM operator with the following permissions:

- DBStart
- DBStop
- DBInfoRead
- ParamRead

The User_Key can be generated using the command:

```
xuser -b <file name>
```

where `<file name>` is the name of a file containing the valid XUSER entries as follows:

Create the User_Key

Parameter	Parameter Definition
USERKEY	Unique name for the User_Key
USERID	User name of the dbm operator
PASSWORD	Password of the user
SERVERDB	Name of the database instance that this key will refer to
SERVERNODE	The name of the server where the database is running (this should be the DNS or host file entry for the SPS protected IP)
SQLMODE	This determines what SQL dialects are compatible
CACHELIMIT	This determines the cache limits for a given session
TIMEOUT	Time in seconds before terminating an inactive session (-1 is the default)
ISOLATION	Determines the isolation level used for locks that affect the user (-1 is the default)
DB_LOCAL	Specifies the database locale

Refer to the *SAP DB / MaxDB User Manual* for more information on parameters. Once proper entries have been specified, use the **xuser** tool to generate the `.XUSER.62` file in the OS User home directory. A sample XUSER file is included below containing two keys (an entry must exist for the DEFAULT User_Key).

```

DEFAULT
NULLDB
NULLDB
NULLDB
NULLDB
LKIP.example.com
INTERNAL
-1
-1
-1
my_locale
LK_USERKEY
LKDBMOPER
LKDBMPASSWD
DB1
LKIP.example.com
INTERNAL
-1
-1
-1
en_US

```

This example XUSER file specifies that two user keys be created, `DEFAULT` and `LK_USERKEY`. Once the **xuser** tool has been run to generate the User_Key(s), the file specified for use by the **xuser** tool should be deleted.

Install the SPS Software

Once you have installed the SAP DB / MaxDB software, created your database and created the User_Key, you are ready to install the SPS Core software and any required patches followed by the SAP DB / Max DB Recovery Kit. Also, if you plan to use SAP DB / MaxDB with raw devices, you must install the SPS Raw I/O Recovery Kit from the SPS Installation Image file. See the [Appendix](#) for requirements and instructions on setting up raw devices.

Refer to the SPS for Linux Installation Guide for details on installing the SPS packages.

Chapter 6: Resource Configuration Tasks

Once you have completed the setup tasks described in the previous section, you are ready to create and extend your SAP DB resource hierarchies.

The following tasks are available for configuring the SPS for Linux SAP DB / Max DB Recovery Kit:

- **Create Resource Hierarchy** - Creates an SAP DB resource hierarchy
- **Delete Resource Hierarchy** - Deletes an SAP DB resource hierarchy
- **Extend Resource Hierarchy** - Extends an SAP DB resource hierarchy from the primary server to the backup server
- **Unextend Resource Hierarchy** - Unextends (removes) an SAP DB resource hierarchy from a single server in the SPS cluster

Please refer to your SPS for Linux Technical Documentation located on the SIOS Technology website for instructions on configuring your LifeKeeper Core resource hierarchies.

The following tasks are described in the Administration section within the SPS for Linux Technical Documentation because they are common tasks with steps that are identical across all Recovery Kits.

- **Create a Resource Dependency**. Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- **Delete a Resource Dependency**. Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- **In Service**. Brings a resource hierarchy into service on a specific server.
- **Out of Service**. Takes a resource hierarchy out of service on a specific server.
- **View/Edit Properties**. View or edit the properties of a resource hierarchy on a specific server.

Note: Throughout the rest of this section, configuration tasks are performed using the **Edit** menu. You may also perform most of the tasks:

- from the toolbar
- by right-clicking on a global resource in the left pane of the status display
- by right-clicking on a resource in the right pane of the status display

Using the right-click method allows you to avoid entering information that is required using the **Edit** menu.

Creating an SAP DB Resource Hierarchy

Perform the following steps on the primary server:

1. On the **Edit** menu, select **Server**, then **Create Resource Hierarchy**.
The **Create Resource Wizard** dialog will appear.
2. Select **SAP DB Database** from the drop-down list and click **Enter**.
3. You will be prompted for the following information. When the **Back** button is active in any of the dialog boxes, you can go back to the previous dialog box. This is helpful should you encounter any error requiring you to correct the previously entered information. You may click **Cancel** at any time to cancel the entire creation process.

Field	Tips
Switchback Type	Choose either intelligent or automatic. This determines how the SAP DB resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. Intelligent switchback requires administrative intervention to switch the resource back to the primary server, while automatic switchback occurs as soon as the primary server is back on line and reestablishes SPS communication paths. Note: The switchback strategy must match that of the dependent resources to be used by the SAP DB resource.
SAP DB Programs Directory	This field contains by default the SAP DB Program Path found in the <i>SAP_DBTech.ini</i> file on the corresponding server. You may type in another directory path. The valid characters allowed for the pathname are letters, digits and the following special characters: - _ . / (Please disregard for MaxDB 7.8 as <i>SAP_DBTech.ini</i> no longer exists.)
SAP DB Instance	This field contains by default the name of the first SAP DB instance found on the system for which no SPS hierarchy exists. The drop-down list shows other database instances that may be available on your system.
SAP DB System User	This is the System User that owns or has permission to execute SAP DB commands. This user must exist on the corresponding server. Enter a valid user name in the selection window.
User_Key	This field contains a default value for the XUSER User_Key. The User_Key is used to store database user data for use with SAP DB Tools. Enter a valid User_Key for the corresponding server, OS User and database instance combination.
SAP DB Database Tag	This is a unique tag name for the new SAP DB database resource on the primary server. The default tag name consists of the SAP DB instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits and the following special characters: - _ . /

4. Click **Create**. The **Create Resource Wizard** will then create your SAP DB resource hierarchy. SPS will validate the data entered. If SPS detects a problem, an error message will appear in the information box.
5. You should see a message indicating that you have successfully created an SAP DB resource hierarchy and you must extend that hierarchy to another server in your cluster to achieve failover protection. Click **Next**.
6. Click **Continue**. SPS will then launch the **Pre-Extend Wizard**. Refer to Step 2 under [Extending an SAP DB Resource Hierarchy](#) for details on how to extend your resource hierarchy to another server.

Extending an SAP DB Resource Hierarchy

This operation can be started from the **Edit** menu or initiated automatically upon completing the **Create Resource Hierarchy** option, in which case you should refer to Step 2 below.

1. On the **Edit** menu, select **Resource**, then **Extend Resource Hierarchy**. The **Pre-Extend Wizard** appears. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the SPS Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The **Pre-Extend Wizard** will prompt you to enter the following information.
Note: The first two fields appear only if you initiated the Extend from the **Edit** menu.

Field	Tips
Template Server	Select the server where your SAP DB resource is currently in service.
Tag to Extend	Select the SAP DB resource you wish to extend.
Target Server	Enter or select the server you are extending to.
Switchback Type	This determines how the SAP DB resource will be switched back to the primary server after it comes in-service (active) on the backup server following a failover. You can choose either intelligent or automatic. The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box. Note: Remember that the switchback strategy must match that of the dependent resources to be used by the SAP DB resource.
Template Priority	Select or enter a Template Priority. This is the priority for the SAP DB hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. Note: This selection will appear only for the initial extend of the hierarchy.

Field	Tips
Target Priority	This is the priority for the new extended SAP DB hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. Note that SPS assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

3. After receiving the message that the pre-extend checks were successful, click **Next**.
4. Depending upon the hierarchy being extended, SPS will display a series of information boxes showing the Resource Tags to be extended, some of which cannot be edited.
5. The **Extend Wizard** will prompt you to enter the following information.

SAP DB Programs Directory	This field contains by default the SAP DB Program Path found in the <code>SAP_DBTech.ini</code> file on the corresponding server. The valid characters allowed for the pathname are letters, digits and the following special characters: <code>-_./</code> (Please disregard for MaxDB 7.8 as <code>SAP_DBTech.inino</code> longer exists.)
User_Key	This field contains a default value for the XUSER User_Key. The User_Key is used to store database user data for use with SAP DB Tools. Enter a valid User_Key for the corresponding server, OS User and database instance combination.
SAP DB Database Tag	This is a unique tag name for the new SAP DB database resource on the target server. The default tag name consists of the SAP DB instance. You may type in another unique tag name. The valid characters allowed for the tag are letters, digits and the following special characters: <code>-_./</code>

6. After receiving the message "Hierarchy extend operations completed", click **Next Server** to extend the hierarchy to another server, or click **Finish** if there are no other extend operations to perform.
7. After receiving the message "Hierarchy Verification Finished", click **Done**.

Unextending an SAP DB Resource Hierarchy

To remove a resource hierarchy from a single server in the SPS cluster, do the following:

1. On the **Edit** menu, select **Resource**, then **Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the SAP DB resource. It cannot be the server where the resource is currently in service. (This dialog box will not appear if you selected the **Unextend** task by right-clicking on a resource instance in the right pane.) Click **Next**.
3. Select the SAP DB hierarchy to unextend and click **Next**. (This dialog will not appear if you selected the **Unextend** task by right-clicking on a resource instance in either pane).
4. An information box appears confirming the target server and the SAP DB resource hierarchy you have chosen to unextend. Click **Unextend**.

5. Another information box appears confirming that the SAP DB resource was unextended successfully. Click **Done** to exit the **Unextend Resource Hierarchy** menu selection.

Deleting an SAP DB Resource Hierarchy

To delete an SAP DB resource from all servers in your SPS configuration, complete the following steps:

1. On the **Edit** menu, select **Resource**, then **Delete Resource Hierarchy**.
2. Select the name of the **Target Server** where you will be deleting your SAP DB resource hierarchy. *(This dialog will not appear if you selected the **Delete Resource** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance.)*
3. Select the **Hierarchy to Delete**. *(This dialog will not appear if you selected the **Delete Resource** task by right-clicking on a resource instance in the left or right pane.)* Click **Next**.
4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Next**.
5. Another information box appears confirming that the SAP DB resource was deleted successfully.
6. Click **Done** to exit.

Testing Your Resource Hierarchy

You can test your SAP DB resource hierarchy by initiating a manual switchover that will simulate a failover of the resource instance from the primary server to a backup server.

Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource and InService**. For example, an in-service request executed on a backup server causes the SAP DB resource hierarchy to be placed in service on the backup server and taken out of service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

If you execute the Out of Service request, the resource hierarchy is taken out of service without bringing it in service on the other server.



IMPORTANT: After bringing your resource hierarchy in service on the backup server, you should attempt to connect to the databases, especially when using raw devices as device spaces. This is necessary to ensure that all disk partitions are visible on the backup servers and the raw bindings are being established correctly.

If the raw bindings have not been established on the backup servers, it is most likely caused by the fact that new partitions were created on the primary server and added to the configuration, but the partition tables have not yet been updated on the backup servers.

The solution is to reboot the backup servers so that the partition tables are updated correctly.

Chapter 7: Resource Hierarchy Administration

The following tasks may be required after your resource hierarchies have been created.

Modifying User_Keys

If the User_Key for an existing hierarchy needs to be changed, the hierarchy must be deleted and recreated.

Modifying OS User

If the OS User that owns the database instance needs to be changed, the hierarchy must be deleted and recreated.

Updating Parameters

When database parameters are updated for an SAP DB instance, it is necessary to ensure that the updated parameter files are redistributed to all servers protecting the instance. If the `IndepDataPath` is on a shared disk, then all servers protecting the instance will be updated automatically.

Chapter 8: Troubleshooting

General Tips

The following error messages are not generated by the SAP DB / Max DB Recovery Kit but may be encountered while using the recovery kit.

Error Message	Solution
Unable to create pipe <code>/usr/spool/sql/fifo/<db instance></code>	The directory <code>/usr/spool/sql</code> must have proper permissions to allow access for system user that owns the database instance.
Open device space <code><dev></code> permission denied	The device spaces on the backup and primary must have the same owner as well as the same user and group permissions.
Unable to set uid on startup	The setuid bit on <code><DependPath>/pgm/dbmsrv</code> must be set and the owner of the file must be the SAP DB system user.
runtime environment error	There are several possible causes with different solutions: <ul style="list-style-type: none"> • The database instance parameter and configuration files do not exist. Create the database parameter files or copy the files from the template server. • The database has encountered a library problem. The server and software installation combination may require the use of the library <code>libpthread-0.8.so</code>. Consult the SAP DB documentation for instructions. • The database instance environment has been corrupted. The processes must be manually killed. Then attempt to restore the resource to the in-service state.
open Registry: Permission denied	The directory <code>/usr/spool/sql/ini</code> should be owned by the system user and group that owns the SAP DB software. In addition, the user and group must also have read/write permissions on the directory.
ERR_USRREAD: could not read user data	The config files from <code><IndepDataPath>/config/<db instance></code> do not exist on the server or do not have the correct permissions. Verify that the files exist with the correct permissions for the system user that owns the database instance.

SAP DB / Max DB Recovery Kit Error Messages

Error Number	Message
111000	Usage: %s independent_program_path <validate:value_1:....:value_n:>
111001	Usage: %s %s %s
111002	No value specified to script %s for input argument %s.
111003	User %s with User_Key %s cannot access instance %s. Action: Specify a User_Key for the given user with database access rights.
111004	The user %s does not exist on the server %s.
111005	The SAP DB instance %s is not running on server %s.
111006	Database Manager Utilities were not found in the specified path %s.
111007	An SPS internal error occurred in utility %s. Action: Retry operation.
111008	Unable to obtain %s device space information for SAP DB instance %s for user %s and User_Key %s. Action: Verify that the user and User_Key are valid for the corresponding database instance.
111009	Unable to create raw resource hierarchy for %s. Action: Verify that the underlying device is a shared device.
111010	Unable to create filesystem resource hierarchy for %s. Action: Verify that the underlying device is a shared device.
111011	Unable to determine the type of the dev space or install path %s. Action: Valid dev space types include file system and/or raw devices.
111012	The path %s is not on a shared filesystem .
111013	The SAP DB instance %s is already under SPS protection on server %s.
111014	The SAP DB instance %s has been successfully started on server %s.
111015	The SAP DB instance %s has been successfully stopped on server %s.
111016	Unable to start SAP DB instance %s on server %s.
111017	Unable to stop SAP DB instance %s on server %s.
111018	Attempting db_warm for database instance %s after db_start failure.
111019	The SAP DB x_server has been successfully started on server %s.
111020	The SAP DB x_server has been successfully stopped on server %s.

Action: Verify that the underlying device is a shared device.

Error Number	Message
111021	The SAP DB x_server is not running on server %s
111022	Unable to start SAP DB x_server on server %s. Action: A problem has occurred using the x_server utility, check the SAP DB logs and correct the problem.
111023	Unable to stop SAP DB x_server on server %s. Action: A problem has occurred using the x_server utility; check the SAP DB logs and correct the problem.
111024	The SAB DB file SAP_DBTech.ini was not found on server %s. Action: Verify that SAP DB is installed correctly on the specified server.
111025	The user id for user %s is not the same on server %s and %s.
111026	The group id for user %s is not the same on server %s and %s.
111027	The service file entries for are not the same on server %s and %s.
111028	One or more of the SAP DB service file entries do not exist on server %s.
111029	No dependents were found for resource %s on server %s.

Appendix: Creating Device Spaces Using Raw I/O

Requirements

In order to use the SAP DB / Max DB Recovery Kit with raw I/O, the following requirements must be met:

- The Linux OS must support raw I/O devices. For most distributions, this support was included in the 2.4 kernel, but there are some distributions that support raw I/O on a 2.2 kernel.
- All raw I/O devices must be bound to a shared disk partition. The number of device spaces (*devspaces*) that will be located on raw I/O devices determines the exact number of raw devices and shared disk partitions required. Refer to the *SAP DB Manual* for guidelines for creating *devspaces* on raw devices.
- The version of the SAP DB / MaxDB software must support the use of raw I/O devices.

Naming Conventions

The naming of raw devices and controller varies by Linux distribution.

- On Red Hat, the device name is `/dev/raw/raw<number>` and the controller is `/dev/rawctl`
- On SuSE SLES 11 versions, the device name is `/dev/raw/raw<number>` and the controller is `/dev/raw/rawctl`

Raw I/O Setup Steps

1. Select a shared disk partition of appropriate size for the SAP DB device space.
2. Bind an unused raw device node to this partition. Since this needs to be done every time the machine is rebooted and requires root access, you may want to add the raw bindings to a system initialization file (i.e. `rc.local` or `boot.local`). These bindings must be removed from the file once the hierarchy is under SPS protection. SPS will re-establish the raw bindings for raw I/O devices that are under SPS protection. Use the command `raw -qa` to see which raw device nodes are already in use. For example:

```
# raw -qa  
  
# raw /dev/raw/raw1 /dev/sda1
```

3. Set global read permissions on both the raw device controller (`/dev/rawctl` or `/dev/raw`) and the disk partition on all servers that will protect the database instance.

Adding a Device Space after Creating Hierarchy

```
# chmod a+r /dev/rawctl (or chmod a+r /dev/raw )
```

4. Set group and user read/write permissions on the raw device on all servers that will protect the database instance.

```
# chmod 664 /dev/raw/raw1
```

5. Change the owner of the raw device to the `SAP DB OS User` for the given database instance on all servers that will protect the database instance.

```
# chown -R sapdb:sapdb /dev/raw/raw1
```

6. Add the device space to the database using `param_adddevspace` or `db_adddevspace`. Refer to the *SAP DB User Manual* and/or the *Database Manager CLI Manual*.

Adding a Device Space after Creating Hierarchy

If a tablespace is added on a raw I/O device or shared file system after the SAP DB hierarchy has been created in SPS, you must manually create a resource hierarchy for the raw device or file system via the LifeKeeper GUI. The newly created resource hierarchy must then be made a dependent (child) of the SAP DB resource hierarchy. The updated parameter files must be redistributed if necessary to all servers that protect the database instance (*this is not required if the `IndepDataPath` is located on a shared disk*).