

Serial Storage Architecture



Spare Tool

Serial Storage Architecture



Spare Tool

Note

Before using this information and the product it supports, be sure to read the general information under "Product Warranties and Notices" included with your system unit.

First Edition (January 2000)

The following paragraph does not apply to any country where such provisions are inconsistent with local law: THIS PUBLICATION IS PRINTED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication could contain technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. The manufacturer may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time, without notice.

It is possible that this publication may contain reference to, or information about, products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that such products, programming, or services will be offered in your country. Any reference to a licensed program in this publication is not intended to state or imply that you can use only the licensed program indicated. You can use any functionally equivalent program instead.

Requests for technical information about products should be made to your authorized reseller or marketing representative. ©Copyright International Business Machines Corporation 2000. All rights reserved. Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication, or disclosure is subject to restrictions set forth in the GSA ADP Schedule Contract with IBM Corp.

Contents

About this book	v
If you need more information	v
ISO 9000	v
Trademarks	v
Chapter 1. Introduction	1
Chapter 2. Installation	3
Installation requirements	3
Obtaining the SSA Spare Tool	3
Installing the SSA Spare Tool	5
Operating requirements	5
Removing the SSA Spare Tool	5
Chapter 3. Configuration	7
Setting up the hot spare client	7
Setting up a hot spare disk drive	7
Setting up the hot spare server	8
Chapter 4. Checking communications between the hot spare server and its clients.	11
Chapter 5. Starting and stopping the SSA Spare Tool	13
Chapter 6. Automated SSA Spare Tool scripts	15
Chapter 7. Manually running a hot spare replacement script	19
Chapter 8. Controlling the SSA Spare Tool with smit	21
Chapter 9. Handling problems	23
Reporting problems	24
Chapter 10. Testing scenario	25
Chapter 11. Messages	27
Index	29

About this book

This book provides information about installing and using the SSA Spare tool.

If you need more information

When you install the Advanced SerialRAID Adapter or the Advanced SerialRAID Plus Adapter, refer to the configuration documentation for your hardware and software.

If you are installing the Fast-Write Cache feature, refer to the *Advanced SerialRAID Adapter: Fast-Write Cache Option Card: Installation Guide*.

For further details on configuring your system with the Advanced SerialRAID Adapter or the Advanced SerialRAID Plus Adapter, check the SSA World Wide Web page at:

<http://www.hursley.ibm.com/ssa/>

Refer to the *Advanced SerialRAID Adapters: User's Guide and Maintenance Information* manual, SA33-3285, for information about how to update your system configuration after you have installed the Advanced SerialRAID Adapter.

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacture of this product.

Trademarks

RS/6000 is a registered trademark of International Business Machines Corporation.

Chapter 1. Introduction

This document explains how to install and enable the Serial Storage Architecture (SSA) Spare Tool. It is intended for storage network administrators who are familiar with SSA, the Logical Volume Manager (LVM), and general storage network principles.

The SSA Spare Tool is an aid for managing IBM SSA storage networks. It works with the LVM to identify stale partitions or missing physical volumes automatically in LVM *mirrored* volume groups. If the SSA Spare Tool finds any stale partitions or missing physical volumes, it:

- Automatically resynchronizes the stale partition if no hardware replacement is needed.
- Issues an e-mail message that indicates what has been done, or whether a disk drive must be replaced.
- Logically exchanges, for a preassigned spare disk drive, the failed disk drive that is in the volume group.
- Resynchronizes the new disk drive with the disk drives that remain in the volume group.
- Notifies the user when it has logically replaced a failed disk drive and resynchronized the replacement disk drives with the other disk drives in the volume group.

When the failed disk drive is physically exchanged, and the data is resynchronized, other scripts can:

- Move the data from the temporary spare disk drive to the new replacement disk drive.
- Prepare the spare disk drive so that it can be used as a spare for any disk drive that fails in the future.

The SSA Spare Tool maintains a protected, software-mirrored environment, and automatically moves data from a failed disk drive to a spare disk drive when required.

The SSA Spare Tool supports all SSA disk drive units.

To enable the SSA Spare Tool, you must install, onto the hot spare server and its related clients, the filesets that are associated with the tool.

These filesets enable you to set up hot sparing for mirrored SSA disk drives that are in a clustered environment. A hot spare disk drive is a disk drive that is made available, dynamically, when a mirrored disk drive fails. When the mirrored disk drive fails, the scripts that are contained in the filesets exchange the failed disk drive with the hot spare disk drive. The mirrored data is then rebuilt onto the hot spare disk drive. One machine, which is designated as server, manages hot sparing and disk drive exchange for all its client machines. This server can also be a client.

Chapter 2. Installation

This chapter describes:

- Installation requirements
- How to obtain the SSA Spare Tool
- How to install the SSA Spare Tool
- How to remove the SSA Spare Tool

Installation requirements

The SSA Spare Tool contains two different **installp** packages: one for servers, and one for clients. To use the SSA Spare Tool, you must install the server fileset onto the hot spare server that has **rsh** access to all the clients. You must also install the client fileset onto all the clients.

Access the **AIX Fix Distribution** (<http://service.boulder.ibm.com/aix.www/aixfixes>) web site and ensure that you have the latest versions of the SSA device drivers. If you do not have the latest versions of these device drivers, download any of the following device drivers that need to be updated. You can search for these device drivers by name, or search for the authorized program analysis report (APAR) number that is associated with each version of AIX. The AIX versions that are compatible with the SSA Spare Tool are:

- Version 4.3.2
- Version 4.3.3

Obtaining the SSA Spare Tool

You can download the SSA Spare Tool by accessing the following uniform resource locator (URL) address:

<http://www.hursley.ibm.com/ssa/rs6k/>

When you have selected your adapter and level of AIX, you can download the SSA Spare Tool.

The SSA Spare Tool is in a **tar** format. When you have downloaded the **tar** file:

1. Put the **tar** file into an empty directory.
2. Change directory (**cd**) to the directory that now contains the **tar** file.
3. Run **tar -xvf sparetool.tar** to obtain the **devices.ssa.sparetool** fileset.

The **devices.ssa.sparetool** fileset has four sub filesets:

- **devices.ssa.sparetool.common** - required both for client and for server installations
- **devices.ssa.sparetool.server** - installs the server code
- **devices.ssa.sparetool.client** - installs the client code

- **devices.ssa.sparetool.smit** - installs the SMIT support (server only)

Installing the SSA Spare Tool

To install the SSA Spare Tool:

1. Remove any previous versions of the SSA Spare Tool from your system. Refer to “Removing the SSA Spare Tool” for more information.
2. Install the **installp** packages by using **smit**.

Operating requirements

For correct operation of the SSA Spare Tool, ensure that:

- All the monitored SSA disk drives that are in the storage networks have the logical volumes 100% mirrored and are on separate physical volumes.
- The volume group is not varied on in concurrent mode.
- The hot spare server must have **rsh** access to all clients. It uses **shortname (hostname -s)** to address those clients.
- This tool does not support **boot** or **dump** logical volume types.
- These scripts handle single failures in a volume group. The scripts do not start any action if two or more disk drives fail at the same time in any one volume group. If the sparing process is running when another disk drive fails, the SSA Spare Tool completes this sparing process before it addresses the second failure.
- Suitable spare SSA disk drives are available when needed.

Notes:

1. The SSA Spare Tool does not support failures of Redundant Array of Independent Disks (RAID) arrays or of disk drives that are in RAID arrays.
2. Disk drive failures that occur concurrently in separate volume groups are handled serially.
3. Any volume groups that are imported on other client machines are updated automatically if disk drive sparing and exchange are successful.

Removing the SSA Spare Tool

To remove this product, use the **installp** command. For example **installp -ug devices.ssa.sparetool.client** removes the client code.

To determine whether you have removed the tool, checking whether the Spare Tool directory is still on your system.

Chapter 3. Configuration

Setting up the hot spare client

To install the SSA Spare Tool onto the client:

1. Install the client code. To do this, use `smit`, or run:

```
installp -qacd devices.ssa.sparetool devices.ssa.sparetool.  
common devices.ssa.sparetool.client
```

2. Configure the client/server relationship. To do this, change the entry in the `/usr/ssa/sparetool/bin/notify_server` script to tell the client the server name. That is:

```
SERVER=<SERVERNAME>
```

For example:

```
SERVER=bugsbunny
```

Setting up a hot spare disk drive

To make spare disk drives available as hot spare disk drives, add the spare disk drives to volume groups and use the naming rule of `spareXX`, where `XX` is a unique, cluster-wide ID.

Note: You should set **autovaryon** for **boot**, although you cannot set **autovaryon** for all the systems.

Because all systems do not have access to all the disk drives, you might need to have multiple volume groups with spares; for example, `spare01`, `spare02`, and `spare03`, with at least one disk drive in each.

The spare disk drives that you want the scripts to use must belong to a separate volume group, and must **not** contain logical volumes. Create enough spare volume groups. The spare volume groups need not be on the nodes that contain the active volume groups. The spare disk drives, however, must be accessible to the nodes that have active volume groups where a failure might occur.

If a failure occurs on a loop that contains host systems but no hot spare disk drives, the SSA Spare Tool searches for a spare disk drive on another loop. This action occurs if all the nodes on the failed loop can access the spare disk drive.

Vary on all the spare volume groups that you have created. If you plan to use the Spare Tool in a high-availability environment, import all the spare volume groups into other host systems that can access them. If a sparing occurs, the Spare Tool scripts ensure that the other host systems are updated with any Physical volume changes.

Setting up the hot spare server

To install the SSA Spare Tool onto a server:

1. Install the server code. To do this, use `smit`, or run:

```
installp -qacd devices.ssa.sparetool devices.ssa.sparetool.  
common devices.ssa.sparetool.server devices.ssa.sparetool.smit
```

Note: The `smit` portion is optional but recommended.

2. If you have not installed the `smit` portion, use the `odmget/odmchange` command to change the `rdsk` application settings for mail. Alternatively, use the `smit chg_rdsk` command, which is equivalent to the selections:

```
smit  
  Devices  
    SSA Spare Tool  
      Change/Show Application Settings
```

The following information is displayed:

```
Change/Show Application Settings  
  
Type or select values in entry fields.  
Press Enter AFTER making all desired changes.  
  
                                [Entry Fields]  
* Rdsk Chkpt                      reset  
* Rdsk Alert ID                    [root]  
* Rdsk Alert Domain                [host.subdomain.subd>  
* Automation                       yes      +  
* StalePP Replacement              yes      +  
  
F1=Help      F2=Refresh      F3=Cancel      F4=List  
F5=Reset     F6=Command     F7=Edit       F8=Image
```

You can reset mail settings by using the `Rdsk Alert ID` option. All the `smit` options have online help available.

3. To create the `SSAdisks` and `SSAloops` for the following object database manager classes, type:

```
/usr/ssa/sparetool /bin/load_diskdb -h client1 client2
```

4. Create the cron task for the `check_spare` script. This script periodically tests the disk drives that are in the spare volume group.

- a. Run `crontab -e` as the root user on the server machine.
- b. Enter the new cron at the end of the file:

```
minute hour day_of_month month weekday /usr/ssa/sparetool /bin/check_spare 2>  
/dev/null 1>/dev/null
```

For example:

```
30 10 * * * /usr/ssa/sparetool /bin/check_spare 1>/dev/null 2>/dev/null
```

This cron causes the `check_spare` script to run at 10:30 daily.

Notes:

1. The client1 client2 ... values must be the same as those that you use when you run the **hostname -s** command on the host systems that you are supplying to the **load_diskdb** script.
2. Run the **load_diskdb** script *only after you have created spare volume groups*. Run also the **load_diskdb** script if:
 - You have added any new host systems to the configuration.
 - Sparing is complete and you have used a method other than the **replacing_failed_disk** script to replace the failed disk drive.
 - You want the spare disk drive that is currently "in use" to be the permanent replacement for the failed disk drive.
3. Run the **reload_db** script if:
 - You have added any new SSA disk drives to any systems that you are maintaining with the SSA Spare Tool.
 - You have added, removed, or changed the names of, any spare volume groups.
 - You have added disk drives to, or removed disk drives from, any volume groups; or you have changed the names of disk drives in any volume groups.
 - You have changed the topology and want to maintain the systems with the SSA Spare Tool.
4. Do *not* run the **load_diskdb** or **reload_db** scripts if:
 - You have any broken loops.
 - Any host systems that are supported by the SSA Spare Tool are down.
 - Sparing is now running.
 - Sparing is complete and you have not used the **replacing_failed_disk** script to replace the failed disk drive.

If the host list that you are supplying for the **load_diskdb** script is larger than the command line limit, you can run the **load_diskdb** script by entering:

```
load_diskdb -h 'cat filename'
```

(filename is the file that has client host names).

Chapter 4. Checking communications between the hot spare server and its clients

To check if communications exist between the client and the server, add the following to the **/etc/syslog.conf** file on the server and the clients:

```
*.debug /tmp/syslog.out
*.daemon /tmp/syslog.out
*.user /tmp/syslog.out
```

Check also whether the **/tmp/syslog.out** file exists and whether the **syslogd** daemon is running on the server and the clients.

If the **/tmp/syslog.out** file does not exist, enter:

```
touch /tmp/syslog.out
```

You can verify whether the **syslogd** daemon is running by entering:

```
lssrc -s syslogd
```

If the **syslogd** is not running, you can start the daemon by entering:

```
startsrc -s syslogd
```

If the **syslogd** is not in an active state, you can refresh the syslogd daemon by entering:

```
refresh -s syslogd
```

When you have completed these steps, you can check the communications by entering the following command on the client from the **/usr/ssa/sparetool/bin** directory:

```
dd_exe -h <hot spare server name> -m "client_name test
message"
```

You can then check the following to verify that the communications work correctly:

- **/tmp/syslog.out** file on the client from which the message was sent
- **/tmp/syslog.out** file on the server
- **/tmp/rdssa/invoke_replace.log** file on the server

Chapter 5. Starting and stopping the SSA Spare Tool

When you install the SSA Spare Tool, the **dd_serv** daemon is started automatically. If you want to stop the daemon, edit the **/etc/inittab** file and remove the entry for the daemon and end the existing **dd_serv** process.

If you want to restart the SSA Spare Tool, enter:

```
mkitab ssaspares:2:respawn:/usr/ssa/sparetool /bin/dd_serv
```

If the **devices.ssa.Spare Tool .smit** portion of the Spare Tool was installed, you can use the **smit** fastpath **chg_rdisk**. This menu allows you to modify the application settings. If you set the Automation flag to **no**, automatic replacements do not occur.

Chapter 6. Automated SSA Spare Tool scripts

Note: You can use these scripts with Quorum on or off. If, however, you have Quorum on while running these scripts, and you lose Quorum, the volume group is varied off. When this happens, these scripts no longer work for those volume groups.

check_spares

This script checks to ensure that the spare disk drives are in a usable condition. This script creates a logical volume on a spare disk drive, then removes the logical volume. This test verifies the usable condition of all spare disk drives. You can set up this script to run periodically from a cron task.

If the **devices.ssa.Spare Tool.smit** portion of the Spare Tool was installed, you can use the smit fastpath **health_check**.

choose_spare

This script chooses a spare disk drive. It attempts to find a disk drive that is the same size as the failed disk drive and in the same SSA loop. If the script cannot find a disk drive of the same size, it searches for a larger disk drive in the same SSA loop. If the script cannot find a suitable spare disk drive on the same loop, it searches for a spare disk drive on other SSA loops that include the host systems of the original loop.

cust_notify

This script produces different types of customer notification, such as e-mail and pager notification. Although basic function and mail text are provided, you can customize the **/usr/ssa/sparetool/bin/cust_notify** file as required.

invoke_replace_disk

When the client notifies the server of an LVM error, the daemon runs this script on the server. The script checks whether a physical volume is missing. If a physical volume is missing, this script attempts to recover the physical volume by running **varyonvg** on the active volume group. If the script cannot start a recovery, it chooses a suitable spare disk drive and calls the **replace_disk** script.

load_diskdb

Run this script on the hot spare server to add the SSAloops and SSAdisks object classes to the object database manager. If you add new hosts to the system, run this script to update the object database manager. For example, run **load_diskdb -h client1 client2 ..** to create the SSA disk drives and SSA loops for the object database manager classes.

Ensure that the spare volume group is set up before you run **load_diskdb**.

If the **devices.ssa.sparetool.smit** portion of the Spare Tool was installed, you can use the smit fastpath **load_diskdb**.

reload_db

This script reinitializes the SSAdisks and SSAloops databases. Run this script if you have changed the system cabling, or added or removed disk drives from the configuration. For example, run **reload_db** to refresh the SSA disk drives and SSA loops for the object database manager classes.

You cannot run this script if a spare disk drive is in use or other Spare Tool activity is running.

If the server machine is the control workstation of an SP2 configuration, the **reload_db** script uses the SDR to get information about the attached nodes. The **reload_db** script:

- Gets the control workstation name from the SDR.
- Gets a list of nodes where **host_responds=1**.
- Gets the **reliable_hostname** for each node where **host_responds=1**.
- Checks network access via **rsh** and checks whether the SSA Spare Tool software install has occurred. The **reload_db** script does these actions for each **reliable_hostname**.
- Uses **rsh** to get the shortname (hostname -s) for each node that has the code installed and **rsh** access.
- Checks network access for each shortname via **rsh** to the node that is using the shortname.

The output from these actions is used to call the **load_diskdb** script to refresh the databases. It works correctly only if the number of nodes generated is less than, or equal to, 8 nodes. To ensure that you can use reload_db as required:

- Install the SSA Spare Tool software only onto nodes that need to be monitored by it.
- Ensure that the SDR on the control workstation is correctly set up and that it gives the correct responses.
- Ensure that the rsh command works correctly when it uses either the shortname, or the reliable_hostnames, for all nodes.

If you cannot meet these requirements, you must use the load_diskdb script every time you need to refresh the Spare Tool databases.

If the **devices.ssa.sparetool.smit** portion of the Spare Tool was installed, you can use the smit fastpath reload_menu or reload_db .

notify_server

The AIX **errornotify** function calls this script when **LVM_SA_PVMISS**, **LVM_SA_STALEPP**, **LVM_MWCWFAIL** or **LVM_SA_WRERR** events occur on a configured client.

replace_disk

This script removes the failed disk drive from the active volume group and adds the spare disk drive to the volume group. The script then recreates and synchronizes the data that was on the failed disk drive. This action returns the volume group to a fully protected status.

replacing_failed_disk

This script provides prompts for the physical exchange procedure for failed disk drives. It moves the data from the hot spare disk drive to the location of the original disk drive in the SSA loop. The script then prepares the hot spare disk drive for use during the next disk drive failure. After you have exchanged the failed disk drive, the script requests the serial number of the new disk drive. The script checks whether the number is a valid serial number or external label. A valid serial number is one that does not belong to another volume group. When entering the serial number of the disk drive, do not enter the complete number. For example, for a disk drive with serial number 000629C848C700D, enter 29C848C7.

You can exchange the failed disk drive for a disk drive that is:

- Physically swapped with the failed disk drive
- Installed in another slot in an enclosure, but is not yet configured as a pdisk or hdisk
- Installed in another slot in an enclosure and already configured as a pdisk or hdisk
- Installed in another slot in an enclosure, is already configured as a pdisk or hdisk, and is configured into the SSAdisks and SSALoops databases (that is, it was present when reload_db or load_diskdb was last run)
- The disk drive that failed but has now been repaired

Notes:

1. For the exchange to be successful, the replacement disk drive must meet the size requirements and must be accessible to those same host systems that had access to the failed disk drive. It is not necessary for the replacement disk drive to be in the loop that contained the failed disk drive, but it is recommended.
2. If the spare disk drive that replaced the original failed disk drive has also failed and been replaced by another spare, you must exchange those failed disk drives in sequence, starting with the most recent failure. This procedure keeps the sizes and spare volume groups consistent.

Attention: If you receive the following message, do not proceed until you have notified, and received instructions from, your system administrator.

This disk <diskname> belongs to another volume group <vgname>. PROCEEDING CAN CAUSE POTENTIAL DATA LOSS AND ALSO CAN CAUSE SYSTEM CODE PROBLEMS.

DO NOT PROCEED WITHOUT THE ONSITE SYSTEM ADMINISTRATOR. DO NOT PRESS Y/N WITHOUT THE SYSTEM ADMINISTRATOR'S CONSENT.

Chapter 7. Manually running a hot spare replacement script

If you do not want the SSA Spare Tool to run automatically, you can run the **manual_replace_disk** script. This script identifies a failed disk drive that is online, then prompts you to exchange the failed disk drive with another disk drive. The script exchanges a failed SSA disk drive that is used as a system disk and is connected to 1 to 8 systems. ***It does not exchange a RAID array.*** The volume group must be varied on for this script to run. You can start the script from an SP control workstation or from the host system that has the failed disk drive.

To start this script, you need **rsh** access to the host systems that are connected to the failed disk drive. If the failure is on the host system on which you are running the script, you do not need **rsh** access to that host system. Note that this script does not work for disk drives that have boot or dump logical volumes. The replacement disk drive that you choose must have a capacity that is equal to, or greater than, the capacity of the failed disk drive. The replacement disk drive must not be large enough to violate the 1016 PPs/PV limit when it is added to the volume group.

Do not use this script for disk drives that are part of a concurrent volume group. Before starting this script, you must know:

- The name of the hdisk that failed on the system that has the volume group varied on
- The hostname where the failure occurred
- The names of all systems that are connected to the failed disk drive

This script must be run on a system that is permitted to run **rsh**, **rcp**, and other similar commands on the systems that are connected to the disk drive for **root**. If you run this script from the host system, and the volume group is varied on, the script does not require **rsh** access for this host system only. A host system does not need **rsh** access to itself.

To run this script, enter:

```
manual_replace_disk -d -i -c -F -f <failed hdisk> -h <host1> [host2... host8]
```

where:

- -d = Turn on debugging.
- -i = Interactive mode.
- -c = Certify the disk drive; this takes some time.
- -F = Force the replacement of an active PV.
- -f <failed disk> = The disk drive that has failed on host1.
- -h <host1> = Host with the varied on VG containing the failed disk drive.
- [host2 ... host8] = Up to seven additional host systems are connected to this disk drive.

If the Interactive flag is **not** set, either on the command line or from smit, you are **not** prompted for input. In this mode, the script assumes that the new disk drive has already been exchanged with the failed disk drive.

If the **devices.ssa.sparetool.smit** portion of the Spare Tool was installed, you can use the smit fastpath **run_recovery**.

Notes:

1. Ensure that the SSA Spare Tool does not automatically start while you are manually running the hot spare replacement script. To do this, stop the **dd_serv** daemon on the server. Also, edit the **/etc/inittab** file so that the daemon does not start automatically.
2. If the **devices.ssa.sparetool.smit** portion of the Spare Tool was installed, you can use the smit fastpath **chg_rdisk**. This menu allows you to modify the application settings by setting the Automation flag to **no**. Automatic exchanges do not then occur.

Chapter 8. Controlling the SSA Spare Tool with smit

The SSA spare tool can be controlled via smit. The SSA Spare Tool menu is in the **devices** section of smit. Alternatively, you can access it directly. To do this, type:

```
smit ssapare
```

The menu structure is:

1. Change/Show Spare Tool Databases
 - 1.1 View the Spare Tool Databases
 - 1.1.1 View the SSA disks Database
 - 1.1.2 View the SSA loops Database
 - 1.2 Create the Spare Tool Databases
 - 1.3 Refresh the Spare Tool Databases
 - 1.3.1 Reload the Spare Tool Databases
 - 1.3.2 View Log
 - 1.4 Show the SSA Physical Disks Configuration
 - 1.5 List all Failed SSA Disks
 - 1.6 List all Spare SSA Disks in use
2. Change/Show Application Settings
 - 2.1 Rdisk Chkpt
 - 2.2 Rdisk Alert ID
 - 2.3 Rdisk Alert Domain
 - 2.4 Automation
 - 2.5 Stale PP Replacement
3. Manual Disk Recovery
 - 3.1 Replace a Failed SSA disk
 - 3.2 View Log
4. Automated Disk Recovery
 - 4.1 View Automated Sparing Logs
 - 4.1.1 View the Invoke Replace Disk Log
 - 4.1.2 View the Replace Disk Log
 - 4.1.3 View the CE Recovery Log
 - 4.2 Recover a Failed SSA Disk (Hot Spared)
5. SSA Health Check
 - 5.1 Check SSA Hot Spares
 - 5.2 View the Spare Check Log

These menus allow you to inspect and change the configuration and settings, perform disk recovery, view log files, and check the health of the hot spare disk drives.

Because all the menus that are in the spare tool panels for smit have full online help available, they provide what is probably the easiest and most natural way to use the tool.

Chapter 9. Handling problems

- Ensure that your system meets the following conditions:
 - Your machine meets the requirements that are stated in Installation requirements and Operating requirements.
 - TCP/IP is installed and active, and you can communicate with the host systems by using TCP/IP.
 - You have applied the appropriate SSA PTFs that are needed to support the system you are using.
- Inspect the SSA Spare Tool log files: **/tmp/syslog.out** and **/tmp/rdssa/invoke_replace_disk.log.\$\$** to determine whether:
 - A disk drive is missing.
 - A spare disk drive has been chosen.
 - Sparing is complete, or where it stopped and why.
 - Sparing is not disabled via the rdsd Object settings.
- Verify **rsh** access from the server to the clients.
- Inspect the status of the active volume groups on the clients with **lsvg <vgname>** and **lsvg --p <vgname>** to determine whether:
 - A physical volume is missing.
 - Any stale physical volumes or partitions exist.
- Inspect the AIX error logs by using the **errpt -a** command. The error logs tell whether the system has generated any LVM error labels.
- Ensure that the error log daemon, **errdemon**, is running. Ensure also that the error log is not full.

The following log files exist on the server. They record information about events that occur and the action that is taken. You might want to examine the contents of these log files:

- **/tmp/rdssa/invoke_replace.log.\$\$**: This file contains a running log of every time **invoke_replace_disk** was run.
- **/tmp/rdssa/replace_disk.log.\$\$**: This file contains the output from the **replace_disk** script.
- **/tmp/rdssa/choose_spare.log**: This file contains the output from the **choose_spare** script.
- **/tmp/rdssa/remove_spare.log**: This file contains the output from the **remove_spare** script.
- **/tmp/rdssa/manual_replace_disk.log.\$\$**: This file contains the output from the **manual_replace_disk** script.

Reporting problems

Before you report a problem, collect the following information:

- The version of AIX that you are using
- The SSAdisks and SSAloops information in the object database manager on the server
- The error logs in the **/tmp/rdssa** file
- The **~/smit.log** file

The AIX Support Line provides direct customer support. This fee service provides voice and electronic access to the IBM support organization. The AIX Support Line addresses your questions about how to use the product and your reports of suspected software defects in this product.

If you do not have an AIX support contract, you can report suspected defects by sending an e-mail to:

aixsupt@aix.boulder.ibm.com

Chapter 10. Testing scenario

Note: Run this procedure only in a test environment. Use a temporary logical volume that is only for test purposes.

Create a data volume group, with logical volumes and a file system, and mount the file system. Create a spare volume group that includes a spare disk drive of suitable size. Perform an I/O operation to the logical volume, and remove a disk drive from your system while the I/O operation is running. This action identifies the disk drive as missing, and starts the sparing process. Monitor the log files in **/tmp/rdssa** in the server machine.

Chapter 11. Messages

While running, the SSA Spare Tool sends messages to the user that is defined in the rdsd Object. These messages are related to the status of the sparing process

The user can be set up via the smit fastpath **chg_rdsd**, or can be hard coded in the **/usr/ssa/sparetool /bin/cust_notify** file.

Note: Text in parenthesis indicates dynamic information that is included in the messages.

MSG1(disk serial number) failed at (date; time) on (host name)

Explanation: The scripts have found a failed disk drive.

MSG2Failed disk (disk serial number) has been recovered at (date; time) on (host name).

Explanation: The disk drive has been recovered and no sparing is required.

MSG3No spares available for failed disk (disk serial number) at (date; time) on (host name).

Explanation: The scripts cannot find a spare for the failed disk drive.

MSG4Failed disk replacement on (host name) for (disk serial number) as hdisk1 is not a member on host1 on (date)

Explanation: The failed disk drive is not a part of a volume group and no sparing is required.

MSG5Failed disk replacement on (host name) for (disk serial number) as Concurrent vg not supported on (date)

Explanation: The scripts do not support concurrent volume groups.

MSG6Failed disk replacement on (host name) for (disk serial number) as failed disk shows active on (date)

Explanation: The failed disk drive is active and sparing is not required.

MSG7Failed disk replacement on (host name for (disk serial number) as host1 does not have vg varied on.

Explanation: On the failed host system, the volume group of the failed disk drive is not varied on.

MSG8Failed disk replacement on (host name) for (disk serial number) as only one copy of lv on (date)

Explanation: The physical volume is not 100 % mirrored.

MSG9Failed disk replacement on (host name) for (disk serial number) as more than one pv is non active (date)

Explanation: More than one physical volume are in a missing state in the volume group.

MSG10Failed disk replacement on (host name) for (disk serial number) as more than one stale pv for vg on (date)

Explanation: More than one stale physical volume are in that volume group.

MSG11aFailed disk replacement on (host name) for disk serial number) as hot spare ACXXXXX is a member of svg on (date)

Explanation: The spare chosen belongs to a volume group.

MSG11bFailed disk replacement on (host name) for (disk serial number) as cannot perform I/O with spare on (date)

Explanation: The spare disk drive is not good.

MSG12Disk (disk serial number) successfully replaced on (host name) with spare (disk serial number) (date & time)

Explanation: The sparing process is complete for the failed disk drive.

MSG13 replacing_failed_disksuccessful (date & time) on (host name) for failed disk drive (disk drive serial number) new sn is AXZXZZZZ

Explanation: The service representative or customer has run the **replacing_failed_disk** script, and the process is successful.

Index

D

deleting the SSA Spare Tool 5

H

handling problems 23

I

installation requirements 3

installing the SSA Spare Tool 5

M

messages 27

 general 27

O

operating requirements 5

P

problem handling 23

problem reporting 24

Q

Quorum 15

R

removing the SSA Spare Tool 5

reporting problems 24

requirements

 installation 3

 operating 5

S

scripts

 check_spare 15

 choose_spare 15

 controlling the SSA Spare Tool with smit 21

 cust_notify 15

 invoke_replace_disk 15

 load_diskdb 15

 load_diskdb -h client1 client2 15

 notify_server 16

 reload_db 16

 replace_disk 16

 replacing_failed_disk 17

 smit 21

setting up a hot spare disk drive 7

setting up the hot spare client 7

setting up the hot spare server 8

T

testing scenario 25

U

URL (AIX Fix) 3

URL (SSA Spare Tool) 3



Printed in the U.S.A.

SA33-3307-00

