

# SavWareHA

High Availability Software for SCO UNIX, OpenServer and UnixWare

**Reference Manual Revision 2.04**

**For**

**SavWareHA for OpenServer5 Release 2.04d**

**SavWareHA for SCO Unix Release 2.04d**

**SavWareHA for UnixWare 7 Release 2.05d**

**Copyright 2000 Avnet Computer Marketing, Hall-Mark Division  
All Rights Reserved**

# TABLE OF CONTENTS

<b>1. INTRODUCTION</b> .....	<b>7</b>
OPERATIONAL OVERVIEW .....	7
<i>No SCO Driver Modifications</i> .....	7
<i>Network Connection</i> .....	8
<i>SCSI Bus Implementation</i> .....	8
<i>SavWareHA Monitor (smon)</i> .....	8
GLOSSARY .....	9
SYSTEM REQUIREMENTS .....	10
<i>Operating Systems</i> .....	10
<i>Disk Space</i> .....	10
<i>Memory</i> .....	10
<i>Network Cards</i> .....	10
<i>Serial Device Connection</i> .....	10
CUSTOMER SUPPORT .....	11
<b>2. INSTALLING SAVWAREHA</b> .....	<b>13</b>
INSTALLATION PROCEDURE .....	13
<i>Licensing Types</i> .....	13
<i>Initial Installation – SCO OpenServer5 and SCO Unix</i> .....	13
<i>Initial Installation – SCO UnixWare 7</i> .....	14
CONFIGURING SAVWAREHA.....	15
<i>SCO TCP/IP</i> .....	15
<i>Host Equivalence</i> .....	15
<i>Primary / Standby Designation</i> .....	15
<i>Configuration Process</i> .....	15
CONFIGURING THE SAVWAREHA MONITOR.....	18
<i>smon</i> .....	18
<i>/etc/sentinel.d/links</i> .....	18
REGISTRATION - ‘REGGIE’ .....	19
<i>Registration Procedure</i> .....	20
<i>Central Registration System Information</i> .....	20
SAVWAREHA FILE AND PROGRAM DESCRIPTIONS.....	20
<i>/etc/sentinel.d/administrators</i> .....	20
<i>/etc/sentinel.d/dkinfo</i> .....	21
<i>/etc/sentinel.d/dkuse</i> .....	21
<i>/etc/sentinel.d/filesys</i> .....	21
<i>/usr/adm/dmx.log</i> .....	22
<i>/etc/sentinel.d/mode</i> .....	22
<i>/etc/sentinel.d/roles</i> .....	22
<i>/etc/sentinel.d/sentinel</i> .....	22
<i>/etc/sentinel.d/sentinel.sh</i> .....	23
<i>/etc/sentinel.d/sentinel.strt</i> .....	23
<i>/etc/sentinel.d/terminals</i> .....	23
<i>/etc/default/netdisk</i> .....	23
<i>/etc/default/mirror</i> .....	24
<b>3. OPERATING SAVWAREHA</b> .....	<b>27</b>
BOOTING SAVWAREHA .....	27
<i>Standby System Bootup</i> .....	27
<i>Primary System Bootup</i> .....	27

---

<i>Access of Filesystems</i> .....	27
SHUTDOWN & MAINTENANCE .....	27
<i>Normal Shutdown</i> .....	27
<i>Taking the Primary System Off-Line</i> .....	28
<i>Placing the Primary System Back On-Line</i> .....	28
<i>Taking the Standby System Off-Line</i> .....	28
<i>Placing the Standby System Back On-Line</i> .....	28
ERROR CONDITIONS, DETECTION, & RECOVERY .....	29
<i>Resolving SavWareHA Mode Discrepancies</i> .....	29
<i>Switch to Fallback Mode</i> .....	29
<i>Recover to Normal Mode</i> .....	30
<i>Regeneration</i> .....	30
SAVWAREHA CHECKOUT .....	31
<i>Error Handling &amp; Notification</i> .....	31
<i>Final Testing</i> .....	31
SAVWAREHA ADMINISTRATION MENU.....	32
<i>SavWareHA Menu Keystrokes</i> .....	32
<i>SavWareHA Administration Screen</i> .....	33
<i>SavWareHA Screen Definitions</i> .....	33
<i>Mirror Status Codes</i> .....	34
SAVWAREHA SUB-MENUS .....	35
<i>Mirror Sub-Menu</i> .....	35
<i>Display Mirror Status</i> .....	35
<i>Fix Broken Mirror</i> .....	35
<i>Regenerate Mirror</i> .....	36
<i>Turn Off Standby Element</i> .....	36
<i>Turn Off Primary Element</i> .....	37
<i>Check Mirror Integrity</i> .....	38
<i>Change Read Flags</i> .....	38
<i>Make Unix File System</i> .....	38
<i>Remove Mirror</i> .....	39
<i>Display Mirror Log</i> .....	40
SAVWAREHA SUB-MENUS - UTILITY SUB-MENU .....	40
<i>Utility Sub-Menu</i> .....	40
<i>Display Monitor Status</i> .....	40
<i>Build A Mirror</i> .....	41
<i>Display Mirror Log</i> .....	41
<i>Software Registration</i> .....	42
<i>Force Change To Fallback Mode</i> .....	42
<i>Recover From Fallback to Normal</i> .....	43
<b>4. ERROR CONDITIONS &amp; RECOVERY .....</b>	<b>45</b>
ERROR CLASSES .....	45
GENERAL ERROR ACTIONS .....	45
SPECIFIC ERROR RECOVERY PROCEDURES.....	46
<i>Meaning of 'Good' versus 'Active'</i> .....	46
<i>Mirror showing 'Good' / 'Good'</i> .....	46
<i>Mirror showing 'Active' / 'Active'</i> .....	46
<i>Mirror showing 'Not Closed' / 'Not Closed'</i> .....	46
<i>Mirror showing 'Good' / 'Not Setup'</i> .....	47
<i>Mirror showing 'Active' / 'Not Setup'</i> .....	47
<i>'fsck' Reports 'Cant stat filesystem'</i> .....	47
<i>Mirror Showing 'Good' / 'Off'</i> .....	48

---

---

Mirror Showing 'Active' / 'Off' .....	48
Mirror Showing 'Off' / 'Good' .....	48
Mirror Showing 'Off' / 'Active' .....	48
Mirror Showing 'Good' / 'ERROR' .....	50
Mirror Showing 'Active' / 'ERROR' .....	50
Mirror Showing 'Disk ERROR' / 'Good' .....	50
Mirror Showing 'Disk ERROR' / 'Active' .....	50
Mirror showing 'Good' / 'MISMATCH' .....	51
Mirror showing 'Active' / 'MISMATCH' .....	51
Mirror Showing 'Good' / 'Need Regen' .....	51
Mirror Showing 'Active' / 'Need Regen' .....	52
Mirror Showing 'Good' / 'REGEN' .....	52
Mirror Showing 'Active' / 'REGEN' .....	52
UNIQUE OR DIFFICULT SITUATIONS .....	53
Installation of SavWareHA, no backup .....	53
DATA RECOVERY / ARCHIVING .....	54
Moving one half of the mirror aside .....	54
Backward Regeneration .....	55
Recovery of files when mirror shows 'Active' / 'Off' .....	56
PANICS AND FILESYSTEM ERRORS .....	57
System Panic, backup uncertain .....	57
System Panic, backup uncertain, filesystem corrupted .....	58
MIRROR STATUS MEANINGS & RECOVERY .....	60
Primary Says 'Normal' / Standby Says 'Normal' .....	60
Primary Says 'Fallback' / Standby Says 'Fallback' .....	60
Primary Says 'Normal' / Standby Says 'Fallback' .....	61
Primary Says 'Fallback' / Standby Says 'Normal' .....	62
MONITOR STATUS ERRORS .....	63
Host access errors .....	63
NetDisk errors .....	63
Communications errors .....	63
Configuration errors .....	64
Fallback notifications .....	64
<b>5. MANUAL PAGES .....</b>	<b>65</b>
SENTINEL .....	65
SENTINEL.SH .....	65
SMON .....	66
NETDSK .....	66
SNTLMAIL .....	67
<b>6. EXTENDED MONITOR .....</b>	<b>69</b>
Application Process Monitoring .....	69
Network Monitoring .....	69
Administration .....	69
INSTALLATION .....	69
Primary System Only .....	69
Installation Instructions .....	69
SCO Unix 3.2v4 and SCO ODT3 .....	70
SCO OpenServer5 .....	70
Installation Continued .....	70
CONFIGURATION PARAMETERS .....	70
Application Monitoring .....	70

---

---

<i>Network Monitoring</i> .....	70
<i>Network To Monitor</i> .....	71
<i>Automatic Fallback</i> .....	71
<i>Interval (seconds)</i> .....	71
<i>Processes To Monitor</i> .....	71
ADMINISTRATION MENU .....	71
<i>X Windows Navigation</i> .....	71
CHARACTER MODE NAVIGATION.....	72
ACTIONS .....	73
<i>Application Monitoring</i> .....	73
<i>Network Monitoring</i> .....	73
<i>Fallback</i> .....	73
LOGS.....	73
<b>7. APPLICATION NOTES .....</b>	<b>75</b>
SAVWAREHA CONFIGURATION WORKSHEET .....	75
SCO LLI DRIVERS .....	76
SAVWAREHA DATA OVER A SECOND NETWORK.....	76
ON LINE BACKUPS WITH 'XA'.....	77
MRACDB ROOT DATABASE CONSISTENCY .....	78
SWITCHING NETWORK DEVICES (IP MATCHING TECHNIQUE).....	79
<i>Requirements and Limitations</i> .....	79
<i>Theory of Operation</i> .....	79
<i>Installation</i> .....	80
SWITCHING NETWORK DEVICES (IP ALIASING) .....	81
<i>Requirements:</i> .....	82
S75INITTAB .....	83
SAVWAREHA RELEASE NOTES 2.03A.....	84
SAVWAREHA RELEASE NOTES 2.04B.....	85
SAVWAREHA RELEASE NOTES 2.04C.....	86
SAVWAREHA RELEASE 2.05X6 FOR UNIXWARE7 .....	87

# 1. Introduction

**SavWareHA** was developed by Avnet Computer Marketing in response to the need for a cost-effective fault-tolerant system residing on two physically separate machines.

**SavWareHA** provides a redundant data and system security level previously relegated to extremely expensive, proprietary computers and operating systems.

Although **SavWareHA** has been installed in sites ranging from Vancouver, BC to Johannesburg, SA, it should **never** be viewed as a substitute for frequent, checksum-verified data backups to removable media, such as those provided by 'XA' from Avnet Computer Marketing. See the discussion of unattended, verified system backups in the "Application Notes" section of this manual; also, please secure a copy of 'XA' if you do not already have one.

## Operational Overview

The software package consists of a combination of tightly coupled Unix drivers and utilities collectively called **SavWareHA**. The major components of **SavWareHA** are: **SavWareHA**, **DMX Net Disk Driver**, and the **SavWareHA Monitor**.

```
Sentinel
PRIMARY SYSTEM stmain      Normal Operation
STANDBY SYSTEM stback     On stmain

> Mirror 0
/dev/mirror00
1461Mb

/dev/srp00 ACTIVE
R/W [          ]

stback:/dev/rsrp00 ACTIVE
W [          ]

Utilities
```

**SavWareHA** provides the administration utilities to configure and monitor all processes and events associated with the **SavWareHA** system.

### No SCO Driver Modifications

Unlike other products on the market today, **SavWareHA** has never modified the standard SCO device drivers that operate either the master (Primary System) or the slave (Standby System) sides of the mirror. This means that any SCO supported device can be mirrored to any other SCO supported device, using any SCO supported connectivity scheme.

**SavWareHA** imposes **no restrictions** on device choice other than those espoused by SCO itself.

### Network Connection

The **DMX Net Disk Driver** provides the network connection for the remote half of the mirrored filesystem(s). This link is referred to as the **data link**, and is used to pass disk requests between the **Primary System** and **Standby System**. An additional network or serial connection may be used to provide user connectivity; this link is referred to as the **IO link**.

### SCSI Bus Implementation

**SavWareHA** will support a **SCSI link**. The reason that Avnet Computer Marketing does not actively recommend this connectivity feature is the fact that any bus error on the exposed SCSI bus will (most probably) immediately cause both machines to crash. Also, if one machine must be removed from the **SavWareHA** setup, there is the real possibility that termination of the SCSI bus would be violated, also causing system problems.

### SavWareHA Monitor (*smon*)

The **SavWareHA Monitor**, which runs as the daemon '**smon**', is a background process that continuously monitors both the Primary and Standby Systems. All events that necessitate some change in the status of the mirrored systems will be sensed and initiated by this daemon.

**SavWareHA** (or more particularly, the **SavWareHA Monitor** '**smon**') may be considered to operate on the Standby System. The **SavWareHA Monitor** will periodically pass a request (or '**token**') between the two machines to verify that the **SavWareHA** system is fully viable. This interval may be tailored by the user. Failure of the Primary System will prevent the **SavWareHA Monitor** (running on the Standby System) from receiving the expected response in a timely manner.

Once a failure has been detected, the **SavWareHA Monitor** will initiate a transition to **fallback mode**. This transition will involve notifying specified system administrators, performing a '**fsck**' on the mirrored unit(s), and enabling login sessions on the Standby System. The automatic transition to fallback mode may be set to manual mode at any time.

When the Primary System is repaired and brought back on line, recovery to a fully fault tolerant status, or **normal mode**, is performed through a menu option. The recovery process:

- Stops all users of the mirror (mounted on the Standby System during fallback mode)
- Unmounts the mirror from the Standby System
- Starts a regeneration process from the Standby to the Primary System
- Mounts the mirror on the Primary System
- Enables user logins on the Primary System.

## Glossary

<b>Primary system</b>	The initial computer system on which the application(s) and users will be running. The Primary System is typically at least as powerful as the Standby System.
<b>Standby system</b>	The system that maintains a mirrored, concurrent copy of the user's application(s), data bases, and programs. This system is ready to take over operation of the mirror Primary System if the Primary System becomes untrustworthy.
<b>Normal mode</b>	The time when the Primary System is fully operational; the users and application(s) are running on the Primary System.
<b>Fallback mode</b>	The time when the users and application(s) have been moved to the Standby System. The only trustworthy copy of the data base and application(s) are on the Standby System.
<b>Mirroring</b>	The use of software and hardware to effect an unnoticed duplication of data across multiple disk drives. A <b>hot mirror</b> means that data is dynamically replicated across the disk drives; a <b>cold mirror</b> means that the data is replicated on demand across the disk drives. <b>SavWareHA</b> constructs a hot mirror.
<b>I/O link</b>	<b>SavWareHA</b> depends on at least two connections between the primary and Standby Systems. One of these connections is the I/O link that handles most, if not all, of the connectivity needs of a typical <b>SavWareHA</b> system. This I/O link can be implemented as a 10BaseT, thin Ethernet, RS422, or similar connection depending on the choice of serial device for connectivity.
<b>Serial link</b>	Serial ports on each of the systems can be connected together to form a link between the mirrored systems. This serial link is used by the monitor program on the Standby System as an additional way to check the status of the Primary System. The serial link allows the monitor program to distinguish between a network failure and a system failure. No user data is transmitted via the serial link.
<b>Data link</b>	<b>SavWareHA's</b> other connection between the primary and secondary system is the path through which all data is transferred between the mirrored filesystems. This link is usually a point-to-point link, implemented typically with 100BaseT or 10BaseT Ethernet; for short distances, 10Base2 (coax) may also be used, with proper termination. A SCSI connection may also be used.
<b>Regeneration</b>	The block by block duplication of one side of a mirror onto the other. This creates an exact copy of the source side onto the destination. All information is duplicated: <i>'disk'</i> and <i>'divvy'</i> tables, as well as Unix filesystem information.

## System Requirements

### Operating Systems

- **SCO Unix 3.2v4.0 (or later) with TCP/IP 1.2**
- **SCO Open Desktop Version 2.0 (or later)**
- **SCO Enterprise Server**
- **SCO OpenServer 5**
- **SCO UnixWare 7**

The general requirement is that each machine have SCO Unix 3.2v4.0 (or later), and TCP/IP version 1.2 (or later) installed.

### Disk Space

**SavWareHA** will take approximately 1 megabyte of disk space on the root filesystem of each system for the programs and configuration files. Also, disk partitions of equal block size must be defined on each machine to provide a mirror of the user data.

**Note:** Because **SavWareHA** requires two fully functional SCO Unix operating systems (one on the Primary and one on the Standby), the root filesystem cannot be mirrored between the two systems. User accounts, printer setup and other configuration information can be automatically duplicated on each system using the `/etc/mracdb` functionality. Refer to the appendix for detailed instructions.

### Memory

**SavWareHA** will use approximately 60Kb of memory during normal operation. An additional 16Kb will be used during a **regeneration**.

### Network Cards

In addition to requiring SCO TCP/IP, any network card supported by the **SCO LLI** (Link Level Interface) Drivers or **SCO TCP/IP** can be used by **SavWareHA**. Refer to the SCO **“Hardware Compatibility Guide”** for a detailed list of supported network cards. One network card is required in each system for the data mirroring. Additional network card(s) connecting to separate networks are recommended if users on the system will be using networking devices and software.

### Serial Device Connection

A terminal server (either Ethernet based, or using a proprietary connectivity scheme) is required if serial devices (printers, terminals, modems, and the like) are going to be connected to the **SavWareHA** system's IO Link and are to be switched automatically to the Standby System if a failure occurs. Any multi-host capable serial connectivity scheme, such as the Equinox SST, that is supported under SCO Unix may be used.

## Customer Support

Avnet Computer Marketing provides support for **SavWareHA** through a combination of telephone and dial-up methods. It is essential that all systems on which **SavWareHA** is installed have a secure, direct telephone line into at least a 9600 baud error-correcting modem on the Standby System. When the modem is connected in this manner, technical support can always access either machine (unless the Standby System has been brought down). It is almost impossible to provide timely support on any system not configured with such a modem.

Avnet Computer Marketing's Customer Support Department is open from 7:00am to 5:00pm PST, Monday through Friday, United States national holidays excepted. Contact Information:

Avnet Computer Marketing  
140 W. Technology, Suite 400  
Irvine, CA 92618

(949) 789-4361 Support Telephone Line  
(949) 789-4365 FAX

E-mail: [cpgsupport@avnet.com](mailto:cpgsupport@avnet.com)  
WWW: <http://svtg.avnet.com>

Requests for assistance or support on **SavWareHA** must be initiated by fax or e-mail, and may incur a charge, unless the customer has signed up for **Premier Support**. Premier Support members may use their Premier Support Hot Line number. For more information, request information on "Premier Support" through the Customer Support Department.

Avnet Computer Marketing will also provide pre-sales design support to configure **SavWareHA** systems for specific applications. Avnet Computer Marketing's staff has a wide range of experience in multi-user Unix applications, as well as substantial experience in constructing and implementing heterogeneous enterprise-wide networks.



## 2. Installing SavWareHA

### Installation Procedure

The **SavWareHA** install procedure will extract the programs and files from the distribution media, serialize the installation and install the drivers into the kernel. **SavWareHA** will not start operating until it is properly setup using the configuration procedure.

#### Licensing Types

One of three versions of **SavWareHA** are activated on installation. Choice of version depends on the presence of a serial number, and registration of the software. The versions are:

- **Demonstration:** An unlicensed version of **SavWareHA**. This will allow one 200Mb mirror, and will require that the system be rebooted every three days.
- **Evaluation:** A licensed, but unregistered, version of **SavWareHA**. This will allow full functionality of the **SavWareHA** product for 30 calendar days, after which time the mirroring will cease (no data is destroyed, just the ability to mirror is removed).

If a fully registered copy has not been purchased, an evaluation copy can be fully registered any time within the 30 day period without reloading and reconfiguring **SavWareHA**. See the section "**Registration - reggie**" in this manual.

- **Fully Registered:** A licensed copy of the software that has been fully paid for, and is then registered with the '**reggie**' registration program. This will function indefinitely on the two systems.

#### Initial Installation – SCO OpenServer5 and SCO Unix

**SavWareHA** is installed using SCO's '**custom**' program in single-user mode. All steps should be performed on both the primary and Standby Systems. The steps are:

1. Place the Primary System in maintenance (single user) mode; log on to 'root'.
2. Enter the '**custom**' command. The following steps are all run from within the '**custom**' program.
3. Choose "**Install**".
4. Choose "**A New Product**".
5. Choose "**Entire Product**".
6. Insert the **SavWareHA** diskette in the drive.

## *Installing Sentinel*

---

7. Choose *"Continue"*. The **SavWareHA** installation process will make two checks before the installation is started:
  - First, the install procedure will attempt to make a new kernel before **SavWareHA** is installed. If a new kernel cannot be successfully created, the install procedure will not start. Pre-existing errors in the link kit must be resolved before **SavWareHA** can be installed.
  - Second, the install procedure will insure that the contents of *'/etc/inittab'* match the contents of the files in *'/etc/conf/init.d/\*'*. If they do not match then the discrepancies must be resolved by the system administrator before **SavWareHA** can be installed.
8. *'custom'* prompts for insertion of the **SavWareHA** diskette to continue installation; since the **SavWareHA** diskette is still in the drive, press [RETURN]. *'custom'* extracts files and installs the **SavWareHA** programs and files.
9. The **SavWareHA** install procedure will prompt for the serial number. This is found on the card that came with the **SavWareHA** package. If a serial number is not available, a demonstration version of **SavWareHA** can be installed by pressing [RETURN]. The demonstration copy has the following restrictions: The mirrored filesystem can be a maximum of 200Mb, the system must be rebooted at least every three days and only one mirrored filesystem is allowed. The full system may be activated without reinstallation later by registration with a valid serial number.
10. When the install procedure prompts *"Rebuild the kernel ?"*, press **"y"**.
11. When the install procedure prompts *"Should this kernel boot by default ?"*, press **"y"**.
12. When the install procedure prompts *"Rebuild the kernel environment ?"*, press **"y"**.
13. At the installation's end, exit *'custom'*.
14. Remove the **SavWareHA** diskette from the drive.
15. Give the *'shutdown'* command to reboot the system. The system must be rebooted before the **SavWareHA** drivers can be accessed.
16. Perform all of the above steps on the Standby System.

## *Initial Installation – SCO UnixWare 7*

1. Insert the SavWareHA distribution diskette for UnixWare7 in the diskette drive.
2. Use the following command line installation procedure or use the Software Manager graphical interface.
3. Login to the root account and run a command shell window
4. `pkgadd -d diskette1`
5. `shutdown -g0 -y -i6`

## Configuring SavWareHA

Before configuring **SavWareHA**, the following items must be installed and functional.

### SCO TCP/IP

SCO TCP/IP must be installed on both machines. IP addresses and machine names must be uniquely determined for each card on each machine (a good discussion of Internet addresses can be found in the SCO “**TCP/IP Runtime System**” manual). Avnet Computer Marketing recommends using a 100BaseT or 10BaseT Ethernet point to point connection for the data link because of it’s ease of support.

Administration Recommendation: SCO assigns a unique system name to each network card in each system. It is recommended that the network card for the data link on the primary system be given the name “neptune” and the network card for the data link on the standby system be given the name “pluto”. This is not a requirement for SavWareHA to operate correctly but can help simplify the system administration process as well as help prevent the confusion in the role reversal situation that naming the systems “primary” and “standby” would have.

### Host Equivalence

Each machine must allow the root account of the other machine to have trusted access. This is accomplished by adding entries for both machines into the files `/etc/hosts` and `/.rhosts`: (Refer to the SCO “**TCP/IP Configuration Guide**” for more detail). Please note that `/.rhosts` must be owned by root, have read-only set for user, group and other (use `chown root /.rhosts` and `chmod 400 /.rhosts` to accomplish this). Also, a password must be assigned to `root` on each machine.

### Primary / Standby Designation

One machine must be designated the Primary System (the other must then be designated the Standby System). If both systems are of equivalent configuration and power, no special requirements are placed on the selection of the Primary System. However, if the two hosts are not equivalent in power or configuration, the more powerful (or more fully configured) host should be chosen as the primary.

### Configuration Process

**SavWareHA** must be configured using the following procedure after installation is complete on both systems. It may be useful to refer to the **SavWareHA** configuration worksheet (see *Application Notes* for this form) for an overview of the information that will be required during the configuration process.

1. Boot both machines into multi-user mode.
2. Mirroring a Unix Filesystem: Unmount the filesystems that will be mirrored on each machine. Use the `umount` command or the `sysadmsh` program (choose `filesystems`, then `unmount`). **NOTE:** All the data that might have existed on the Standby System’s half of the mirror will be **overwritten** with the Primary System’s data during the regeneration phase!

3. Mirroring a Unix Filesystem: Remove filesystem configuration entries.

**SCO OpenServer5 and SCO Unix:** Use the *'mkdev fs'* command or the *'sysadmsh'* program (choose *'Filesystems'*, then *'Remove'*; this allows you to remove existing filesystem definitions). The filesystem configuration entries are held in the files *'/etc/default/filesys'* and *'/etc/checklist'*; these files may be edited by hand.

**SCO UnixWare7:** Use the Filesystem Manager in the Graphical Interface to Remove Mount Configuration for the disk that is to be mirrored. (The information is stored in */etc/vfstab*.)

Mirroring a Raw Database Partition: The application using the raw database partition must be stopped. The database must also be shut down. Refer to the shutdown procedure in the database documentation. (After the configuration process is complete, the database application will need to be re-configured to access the device */dev/rmirror00* instead of the original SCO disk device).

4. Log in as root on the Primary System. Attempt to log into the Standby System with the *'rlogin'* command; if the *'password:'* prompt is seen, then *'.rhosts'* has not been properly configured on the Standby System. Perform the same *'rlogin'* command, but this time (since you are on the Standby System), attempt to log into the Primary System. The *'password:'* prompt should not be seen here either. To sign off both systems, enter *'exit'* twice or [CONTROL+D] twice; this will return your focus to the Primary System.
5. Give the *'sentinel'* command. The **SavWareHA** program uses the arrow keys or space bar to move between selections and the [RETURN] key to make a selection. Use the [END], [ESC] or [CONTROL-E] key to exit or cancel an operation.
6. When prompted, enter the system name for the network card of the data link on the primary system.

Administration Recommendation: It is recommended that the network card for the data link on the primary system be given the name "neptune". This must be done with netconfig before SavWareHA is configured.

(Press [F3] for a listing of defined system names). This should be the same name as that used to *'rlogin'* from the Standby System to the Primary System.

**NOTE:** [F3] here passes through the *'/etc/hosts'* file looking for valid addresses that can be reached via *'ping'*. This should serve as a further verification that the hosts are properly configured.

7. When prompted, enter the system name of the network card that is used for the data link on the standby system.

Administration Recommendation: It is recommended that the network card for the data link on the standby system be given the name "pluto". This must be done with netconfig before SavWareHA is configured

(Press [F3] for a listing of defined system names). This should be the same name as that used to *'rlogin'* from the Primary System to the Standby System.

8. Choose the "Utilities" option and press [RETURN].
9. Choose the "Build a mirror" option and press [RETURN]. This will be **non-destructive** to the data on the Primary System; but will **overwrite** all data on the Standby System's filesystem.
10. Enter the disk partition device name on the Primary System of the filesystem that will become the **SavWareHA** master mirror side (Press [F3] for a listing of all disk partitions on the Primary System).

**NOTE:** [F3] here passes through `/etc/conf/cf.d/m SCSI`, and attempts to access devices that appear to be valid disk (Sdsk) entries. If the disks do not actually exist, or are not completely configured (via the `mkdev hd` command), errors will be reported to the system console. Depending on the configuration of the system, these errors may be ignored.

11. Enter the disk partition device name on the Standby System that will contain the slave mirror side. (Press [F3] for a listing of all disk partitions on the Standby System). **Any existing data in this disk partition will be overwritten with the Primary System's data !!** This disk partition must be at least as large as the disk partition on the Primary System.
12. Mirroring Unix Filesystem: Enter the mount point for the filesystem (such as `/u`). This same mount point name will apply to both systems.

Mirroring Raw Database Partition: Enter "N" at the mount point prompt.

13. When prompted with "OK to continue?", double check the entries and press "y" to complete the mirror configuration. At this time, the regeneration process will begin. Both systems should show massive disk access activity, as will the data link.
14. (**Versions of SavWareHA before 2.02c**): If prompted "mkfs: Make a Unix Filesystem (y/n)": Answer "n" (no) if there is existing data in the disk partition on the Primary System. **An answer of "y" to this question will DESTROY all data on the mirrored disk partitions on BOTH systems!!!** Answer "y" (yes) if there was no filesystem in the disk partition on the Primary System. If in any doubt at all, entering 'n' is always the best choice.
15. Press [RETURN] to continue. The mirror generation will continue and mirrored access is now available to the users. The mirror is mounted automatically at this time on the specified mount point.

Mirroring Raw Database Partition: The database application must now be configured to access `/dev/rmirror00` instead of the original SCO device (i.e. `/dev/ru`).

16. Exit the **SavWareHA** program by pressing the "q", [END], [ESC] or [CONTROL-E] key.
17. Test your application's access to the mounted filesystem(s) on the Primary System. User applications will not be able to access the filesystem(s) on the Standby System unless **SavWareHA** is switched to fallback mode.

## Installing Sentinel

---

18. Configure the **SavWareHA** monitor program links as necessary. By default, the initial link is set to *'net'*. (See the section “**Configuring the SavWareHA Monitor**” in this manual for more information).
19. Configure the terminals (*'stty'* settings, terminal types) as required for the application. Remember that this will have to be performed on both the Primary and Standby Systems. See application note for S75inittab and mracdb for details on keeping the standby root information up to date.
20. Add installation-specific *'Start'* and *'Stop'* scripts to the directory *'/etc/sentinel.d'*. The scripts should follow the SCO default of *'S###xxxx'* and *'K###xxxx'*. *'S###xxxx'* is the startup script (*'S'*), sequence number (*'##'*), and descriptive name (*'xxxx'*). *'K###xxxx'* is the stop or kill script (*'K'*), sequence number (*'##'*), and descriptive name. Usually, one script is written to handle both start and stop, named *'S###xxxx'*, and is then linked with the *'ln'* command to the *'K###xxxx'* name.
21. Test the **SavWareHA** system as described in the “**SavWareHA Checkout**” section of this manual.

## Configuring the SavWareHA Monitor

### smon

**smon** is a program that will run on the Standby System and monitor the Primary System. In the case of a Primary System failure, **smon** can be configured to automatically switch to fallback mode and will send a message to the administrator(s). In the case of a network failure, the Primary System will *'turn off'* the slave (Standby System) side of the mirror without interrupting processing.

Please note that disruption of a TCP/IP link (not the *'LLI'* link) will cause up to a **9 minute time out**. This appears not to be configurable. Avnet Computer Marketing suggests that the *'LLI'* method be used where network packet security is not an issue. SavWareHA release 2.03a or higher should be installed before using the *'LLI'* method.

The configuration file for *'smon'* is *'/etc/sentinel.d/links'*. This file is primarily accessed on the Standby System; however, it should be copied to the Primary System so that the file is identical on both systems. Each line describes a link between the Primary and Standby System. The contents of this file is described on the following page:

### /etc/sentinel.d/links

<b>net</b> { <i>machine name</i> [ <i>time-out</i> ]}	The machine name of the Primary System. If the machine name is omitted, the Primary System name is assumed. <i>'time-out'</i> is the default number of seconds that <i>smon</i> will attempt to access the Primary System. To set this value, the “machine name” field must be filled in. As a shortcut, you can specify the “machine name” as <i>'-'</i> to default to the Primary System name.
<b>serial</b> <i>primarydev= standbydev</i> [ <i>optional</i> ]	The serial link between the Standby and Primary Systems.

<b>Manual</b> <i>[optional]</i>	Always 'true'; failure of <i>smon</i> to sense the Primary System will not cause a fallback. If this is found anywhere in the file, the system will never automatically fall back.
------------------------------------	--

'*smon*' will only cause an automatic fallback if all links fail, but will never cause a fallback if the keyword '**manual**' appears anywhere in the '*links*' file on a line by itself.

Changes to /etc/sentinel.d/links will take effect automatically within 60 seconds without requiring a system reboot. Use the sentinel menu option "Display Monitor Status" to show the current configuration and status.

For example, the file might contain:

```
net pluto 8
serial tty2a=tty1a
```

Please note there is a [TAB] after the word 'serial' in this example.

This indicates that the Primary System is "pluto", time-out is set to 8 seconds, and that there is a serial link between tty1a (COM1:) on the Standby System to tty2a (COM2:) of the Primary System.

Note: The cable required to connect two DB-9 COM ports is as follows:

```
2    ---    3
3    ---    2
5    ---    5
```

Note: The cable required to connect two DB-25 COM ports is as follows:

```
2    ---    3
3    ---    2
7    ---    7
```

The links can be monitored manually using the *sentinel* command with the '*Utilities:Display Monitor Status*' option. One line will be displayed for each link defined in the links file. The link configuration and status information will be displayed.

Serial links may also be established between ports on terminal servers. The '*smon*' program will periodically test each link and report failures to the system administrator(s). This technique allows automatic unattended notification of terminal server failures.

## Registration - 'reggie'

The **SavWareHA** Software must be registered within 30 days of installation. The serialized **SavWareHA** product will be fully functional for this grace period. Use the **SavWareHA** Registration option from the '*Utilities*' menu to display the Registration Request string and to perform the product registration. The Registration Request string must be transmitted to the central registration system and a Registration Code will be returned. The Registration Code must be entered into the system to allow the **SavWareHA** software to continue to operate.

### Registration Procedure

The registration procedure is designed to be as easy as possible. Use any of the following ways to register:

- **FAX:** The program will generate the Registration Request and send it to the customer support fax number via a Unix fax system. A fax system such as VSI\*FAX or True Fax must have been previously installed and configured on the system. When the fax is received by the central registration system, it will generate the Registration Code and return it via fax, e-mail, or voice as appropriate.
- **E-Mail:** The program will generate the Registration Request and send it via electronic mail to the central registration system at '[reggie@avnet.com](mailto:reggie@avnet.com)'. An e-mail package must have been previously installed and configured with a connection to the Internet. The Registration Code will be returned via e-mail to the originating system.
- **Print:** The program will print the Registration Request to a printer. The request may be faxed or mailed to the central registration system. The registration code will be faxed back to the fax number as displayed on the fax request. Make sure that the return fax number is correctly and clearly indicated on the initial fax or letter.
- **File:** The Registration Request can be written to a file and then transmitted to the central registration system via any means possible. The return contact information must be supplied for the Registration Code.

When the Registration Code is returned from the central registration system, it can be entered electronically (Direct Modem), entered manually through the keyboard or read in from a file.

### Central Registration System Information

Avnet Computer Marketing  
Voice (949) 789-4361; FAX (949) 789-4365  
E-mail: [reggie@avnet.com](mailto:reggie@avnet.com)  
WWW: <http://svtg.avnet.com/reggie>

## **SavWareHA File and Program Descriptions**

### /etc/sentinel.d/administrators

#### **(Text File)**

List of users on both the Primary and Standby Systems that have been designated as system administrators. Each administrator will have a login id listed on separate lines. For example:

root

mary\_r  
dmann@aa.sa.org (requires proper configuration of MMDF)

These administrators will receive mail and messages regarding the status of the **SavWareHA** system. This file must be manually edited using a text editor such as 'vi'.

Make certain that 'mail' functions properly by sending a test mail message to the proposed system administrator, especially if 'mmdf' must be modified. For instance, in the example above, the administrator should send mail to 'dmann':

```
mail dmann@aa.sa.org
Hello down there ;- )
[CONTROL-D]
```

This should initiate a 'mmdf' request, and the recipient should receive it in a timely fashion. Much more information on 'mmdf' may be found in the SCO '**System Administrator's Guide**' manual.

/etc/sentinel.d/dkinfo

**(Program)**

This program displays the disk hardware name(s), size, SCSI adapter and adapter number, and SCSI target ID of each disk partition on the system found through the file '/etc/conf/cf.d/m SCSI'. This is used in the **SavWareHA** setup when [F3] is pressed to configure the mirror.

/etc/sentinel.d/dkuse

**(Program)**

This program shows if and where a disk partition is used. This is invoked when [F3] is pressed during mirror construction.

/etc/sentinel.d/filesys

**(Text File)**

Defines the mirrored filesystems. For example, the file might read:

```
# bdev(primary) bdev(standby) mountpoint flags mountflags
/dev/mirror00 /dev/mirror00 /u A
```

This indicates that on the Primary System, the filesystem referenced by '/dev/mirror00' will be mirrored over to the Standby System onto the filesystem '/dev/mirror00'. This mirror will be mounted on the Primary System on directory '/u'.

## *Installing Sentinel*

---

The flags and mountflags columns contain the flags to be passed to the *mount* command. The flag **'A'** means to automatically mount the filesystem during the transition to multi-user mode. The flag **'R'** means that this is a raw database partition that is mirrored but not automatically mounted. This file must be identical on both systems.

This file is created and modified by the **SavWareHA** configuration utility.

*/usr/adm/dmx.log*

### **(Text File)**

Comprehensive log file for **SavWareHA**. Information regarding status changes (normal to fallback and vice versa), configuration changes, and errors will be recorded in this log.

*/etc/sentinel.d/mode*

### **(Text File)**

Text representation of the mode of the **SavWareHA** system. Its value is either *'normal'* or *'fallback'*. A copy of this is accessed on both systems.

*/etc/sentinel.d/roles*

### **(Text File)**

Consists of two lines, one per machine, that determine the current **SavWareHA** configuration. For example, the file might read:

```
neptune primary
pluto standby
```

This would indicate that the machine *'neptune'* is currently operating as the Primary System, and *'pluto'* is currently operating as the Standby System. This file must be identical on both systems.

Administration Recommendation: It is recommended that the network card for the data link on the primary system be given the name "neptune" and the network card for the data link on the standby system be given the name "pluto". This must be done with netconfig before SavWareHA is configured. If this is done, the roles file will be:

```
netptune primary
pluto standby
```

This file is created and modified by the **SavWareHA** configuration utility.

*/etc/sentinel.d/sentinel*

### **(Program)**

This is the **SavWareHA** executable program. It is linked to `/etc/sentinel`. All **SavWareHA** menu-driven functions are accessed from this program.

[/etc/sentinel.d/sentinel.sh](#)

**(Shell Program)**

This is a supporting shell script for the **SavWareHA** program. This program is linked to `/etc/sentinel.sh`.

[/etc/sentinel.d/sentinel.strt](#)

**(Shell Program)**

A link between this file, `'/etc/rc2.d/S85sentinel'` and `'/etc/rc0.d/K05sentinel'` is made to determine what processes are to be started up or shutdown during transition phases of the **SavWareHA** system.

[/etc/sentinel.d/terminals](#)

**(Text File)**

The list of terminals and serial resources that are to be enabled and disabled during transitions from normal to fallback mode. The format for this file is serial device name without the preceding `'dev'`. For example, the file might contain:

```
tty1a
ttyn0101
ttybd
```

This would denote that terminals `'dev/tty1a'`, `'dev/ttyn0101'`, and `'dev/ttybd'` are to be enabled for login on the Primary System under 'normal' mode, and on the Standby System under 'fallback' mode. This file must be manually edited using a text editor. This file should be identical on both systems.

Administration Recommendation: It is recommended that `S75inittab` be used to control the terminals if more than three terminals are going to be switched. This will reduce the amount of time required for a system switchover. See application note for `S75inittab`.

[/etc/default/netdisk](#)

This is created and modified by the **SavWareHA** program during **SavWareHA** configuration. It contains the network addressing for the data link. This file is only used on the Primary System. This file can be manually edited as necessary, but the Primary System will need to be rebooted for the changes to take effect. The format of the file is:

```
machine:device - (Size) (0) (Unit#) (Primary LLI) (Remote LLI)
```

Please note that the file requires [TAB] characters for field separators.

For example, the file might look like:

**pluto:/dev/rsrp00 - 0 0 0 /dev/wdn0 /dev/wdn0**

<b>Machine</b>	The machine name used by the data link to access the Standby System. To route <b>SavWareHA</b> data over a secondary network link, modify this machine name.
<b>Device</b>	The slave device name on the Standby System accessed by the 'netdisk' driver.
<b>Size</b>	If this field is 0, the size will be requested from the Standby System. The size can be forced by putting the number of blocks in this field.
<b>0</b>	This field is reserved.
<b>Unit #</b>	This field relates the entries in this file to the device name <i>'dev/ndsk?0'</i> . This is used to cross reference the filesystems to <i>'etc/default/mirror'</i> .
<b>Primary LLI</b>	This is an optional field used to redirect the <b>SavWareHA</b> data requests directly to the SCO LLI drivers without using TCP/IP. If present, this field must be the LLI device name of the network card on the Primary System. (i.e. <i>'dev/wdn0'</i> ).
<b>Standby LLI</b>	If present, this field must be the LLI device name of the network card on the Standby System (i.e. <i>'dev/wdn0'</i> ).

*/etc/default/mirror*

This file is used to describe all of the mirrors on the system. On the Primary System, it defines the primary element and the standby element for each mirror. On the Standby System, only the standby element is defined. This file is modified by the **SavWareHA** program and other **SavWareHA** utilities. The format of the file is:

(Block Dev) (Raw Dev) (Size) (0) (Unit#)(Flags)

For example, the file might contain:

<b>#block device</b>	<b>raw device</b>	<b>size</b>	<b>factor</b>	<b>unit</b>
<b>/dev/srp00</b>	<b>/dev/rsrp00</b>	<b>0</b>	<b>0</b>	<b>OR</b>
<b>/dev/ndsk00</b>	<b>/dev/rndsk00</b>	<b>0</b>	<b>0</b>	<b>OR</b>

<b>Block Dev</b>	Block Device name of the primary element of the mirror.
<b>Raw Device</b>	The Character Device name of the primary element of the mirror.

<b>Size</b> {usually 0}	If the size is 0, the size will be determined by accessing the raw device. The size can be forced by specifying the number of blocks on the device.
<b>0</b>	Reserved
<b>Unit #</b>	This is the mirror number. It cross references the file information to the devices <i>'/dev/mirror?0'</i> .
<b>Flags</b>	Mirror Flags R - Read from Primary only T - Read from Standby only M - Standby disabled S - Primary disabled



# 3. Operating SavWareHA

## Booting SavWareHA

### Standby System Bootup

When a **SavWareHA** system is started from a power-off or complete shutdown status, the Standby System should be started first. The Primary System will automatically look for a response from the Standby System telling in what mode (*'normal'* or *'fallback'*) the Primary System should be immediately after the Primary System comes into multi-user mode. If the Standby System is not in full multi-user mode when the Primary System comes up, the **SavWareHA** system will assume (correctly) that the Standby System is unable to handle mirroring requests, and therefore mirroring over to the Standby System (the slave side) will be disabled (the slave side is turned *'OFF'*).

### Primary System Bootup

After the Standby System has come up into multi-user mode, the Primary System should then be started. The booting sequence for the Primary System will proceed as usual. The Primary System will check the Standby System's mode during the transition to multi-user mode, and compare it with its own mode. If the Standby System is in fallback mode and the Primary System is in normal mode, the Primary System will automatically switch itself to fallback mode.

### Access of Filesystems

Access to the mirror, and to the master and slave sides of the mirror, must be done through the mirror device node (OpenServer 5 *-/dev/mirror00*, UnixWare 7 *-/dev/mirror000*). Access of either side directly while the mirror is mounted will immediately compromise the mirrored filesystem(s).

If access to one side only is desired for some reason, make certain to turn that side *'OFF'* with the *'Fix Broken Mirror: Turn Off ..... Element'* selection. After the desired side's status has been set to *'OFF'*, the side may have *'fsck'* run against it, and then may be mounted. Be aware, though, that once the mirror is brought back into full operation, the side marked as *'OFF'* will be **overwritten** with the other side's data.

## Shutdown & Maintenance

### Normal Shutdown

The proper procedure to shutdown a **SavWareHA** system is to issue the *'shutdown'* command from the Primary System. During the shutdown, the Primary System will turn off the monitor program on the Standby System so it doesn't think there is a Primary System failure. The Standby System will continue to run but it will not take control of the users or filesystems once the Primary System has been shutdown. The Standby System can then be shutdown with the normal shutdown command on its console.

### Taking the Primary System Off-Line

The Primary System can be taken off-line for maintenance, updates or diagnostics by manually issuing the *'Utilities: Force Fallback'* command from the **SavWareHA** *'Utilities'* sub-menu on either the Primary or Standby System. This will cause the Standby System to take control of the users and the filesystems. Once control has passed to the Standby System, then the Primary System may be shutdown gracefully, and taken off-line. See the *'Utilities Sub-Menu'* section.

Alternatively, if the installation has 'hot swap' disk drives, the system administrator may desire to perform maintenance on the drive or drives without interrupting processing with a fallback transition. To accomplish this, use the *'Fix Broken Mirror: Turn Off Primary Element'* option on the mirror in question. This will mark the mirror as 'OFF', and allow removal / insertion of drives. Bring the mirror back into compliance with the *'Fix Broken Mirror: Regenerate Mirror'* option on the mirror in question.

### Placing the Primary System Back On-Line

When all off-line maintenance is complete on the Primary System, boot it into multi-user mode. The Primary System will access the Standby System, see that it should be in fallback mode, and set its mode accordingly. Log in to *'root'* on the Primary System and enter the command **sentinel**. Use the *'Utilities: Recover from Fallback to normal'* option. Since this is "planned downtime", the system administrator should first notify all active users that a transition back to normal mode will take place shortly. All users will be evicted from the Standby System during this procedure; once the transition is complete and a regeneration started, the users may log in to the Primary System and continue without further interruption.

### Taking the Standby System Off-Line

The Standby System can be taken off-line to perform maintenance, updates or diagnostics by turning off the Standby System mirror element with the *'Fix Broken Mirror: Turn Off Standby Element'* option, and then issuing the *'shutdown'* command on the Standby System. Data mirroring will be stopped and the users will continue uninterrupted on the Primary System. Applications running on the Primary System will not be affected in any manner.

### Placing the Standby System Back On-Line

When the off-line maintenance is complete on the Standby System, boot it in to multi-user mode. Log in to *'root'* on the Primary System and run the command **sentinel**. Select the appropriate mirror number, and use the *"Fix Broken Mirror: Regenerate Mirror"* selection from the mirror sub-menu. This will start a regeneration of the data to the Standby System. Users will continue uninterrupted on the Primary System during the regeneration process.

## Error Conditions, Detection, & Recovery

### Resolving SavWareHA Mode Discrepancies

After both systems are booted to multi-user mode, administrator action may be required if the SavWareHA modes get out of sequence between the two systems. Refer to the following table:

Primary System	Standby System	Actions
Normal	Normal	None required. The Primary and Standby Systems are both in normal operation. Mirroring continues.
Normal	Fallback	Communication has been lost on all channels between the systems. The filesystem data is out of sync. Both systems may have taken control of their respective disk partitions. Restore communication and follow the procedure(s) shown in <i>'Unique or Difficult Situations'</i> section of this manual.
Fallback	Fallback	The Primary System has or had failed. The filesystem and users are controlled by the Standby System. The Primary System can be tested and repaired while the users are running on the Standby System. When the system administrator is confident that the Primary System is operational. The <i>"Recover to normal mode"</i> option of <b>SavWareHA</b> should be used.
Fallback	Normal	An unrecoverable error has occurred on both systems. Neither system has control of the users or disk partitions. The filesystem integrity may be compromised. Test all communications links and follow the procedure(s) shown in <i>'Unique or Difficult Situations'</i> section of this manual.

### Switch to Fallback Mode

In the case of a Primary System failure, the Standby System's monitor program detects a failure by the lack of access acknowledgments across all entries specified in the *'/etc/sentinel.d/links'* file. When this occurs, **SavWareHA** will automatically switch to fallback mode -- **if** the keyword *'manual'* is **not** present in the *'links'* file. When **SavWareHA** switches to fallback mode either by the **SavWareHA** monitor program or manually from the **SavWareHA** utility menu, the following events occur:

1. The Standby System will take control of the mirrored filesystem(s). This is accomplished by the **'S04mount'** routine.
2. **SavWareHA** will cause *'fsck'* to be run on any mirrored filesystems if necessary. This is also accomplished by the **'S04mount'** routine.
3. **SavWareHA** will then allow login on those terminals specified in *'/etc/sentinel.d/terminals'*. This is accomplished by the **'S85terminals'** routine.
4. The fallback system will execute a series of application "start" procedures. These procedures will run any application-specific programs to recover transactions or restart daemons. Application-specific "start" and "stop" procedures can be added by

## Operating Sentinel

---

creating a file in the directory `/etc/sentinel.d` starting with the letter "**S**" for start and the letter "**K**" for stop (or kill), followed by a sequence number and procedure name. The permissions for the file must be set for read and execute on user, group, and other (accomplished by `chmod 555 <filename>`). Examples of these files can be found in the `/etc/sentinel.d` directory. These procedures may be modeled after the `'X75user'` startup routine.

5. The `/etc/sentinel.d/mode` file is changed on the Standby System to "fallback", thus marking the successful transition off of the Primary System. This is accomplished by the `'S90mode'` routine.
6. **SavWareHA** cleans up any temporary files created on the Standby System. This is accomplished by the `'S91cleanup'` routine.

### Recover to Normal Mode

In fallback mode, the Primary System has failed and users have been transferred over to the Standby System. After the Primary System has been repaired, brought back on-line, and tested, use the following procedure to recover to normal mode:

- Be sure all users are logged off from the Standby System. Since this is "planned" downtime, the system administrator should notify all users to log off temporarily.
- On the Primary System, log in as `'root'`.
- Give the **sentinel** command.
- Choose "*Utilities*".
- Choose "*Recover from fallback to normal*". Answer `'y'` (yes) to the question `'OK to recover to normal mode?'`; the Primary System will notify the Standby System that **SavWareHA** transitions back to normal mode. The mirrored unit will be unmounted from the Standby System, and control of the mirror will pass to the Primary System. A regeneration from the Standby System to the Primary System is then begun immediately.
- Exit from the **SavWareHA** program.

### Regeneration

The users can immediately log back into the Primary System and continue operation. **SavWareHA** is in 'normal' mode and the Primary System has control of the users and the filesystems. A regeneration process has been started in the 'background' to copy the data from the Standby System to the Primary System. This process is a low priority job that will only run when there are no outstanding disk block write requests to the mirror.

During the regeneration, all reads occur from the Standby System over the data link. Because of the fact that both reads and writes are being served over the data link, users may notice that the system is operating slightly slower than usual. When the regeneration is complete, the reads will be routed to the Primary System automatically with no operator

intervention. No further administrator action is required. An entry into the log file is made that the regeneration has completed.

## SavWareHA Checkout

After **SavWareHA** has been installed and configured, a trial fallback transition should be tested.

Use the **SavWareHA** program to verify that the system is in normal mode (upper right hand corner of screen). Next, go into the *'Utilities'* menu and display the monitor status. The monitor status should be normal. **Automatic fallback will only occur if the monitor status is normal!** The monitor will be PAUSED during a regeneration and should return to NORMAL when the regeneration is complete. The status of the mirror should be ACTIVE / REGEN (regeneration is in process).

Use the mount command on the Primary System and insure that the user filesystems are mounted.

### Error Handling & Notification

Whenever any error condition occurs, mail is sent to each user listed in the *'/etc/sentinel.d/administrators'* file. Users can be at a remote location if the e-mail system is correctly configured for the system to pass messages in this manner. In addition, pagers may be activated automatically if *'mmdf'* and *'uucp'* have been properly configured.

The administration message handling should be tested once **SavWareHA** configuration is complete. This can be done by using the *'Utilities: Force Fallback Mode'* selection to manually force the system into fallback mode. Once the switch to fallback mode is complete, the system administrator(s) should receive mail describing the transition to fallback mode.

### Final Testing

Terminals, printers, and network devices on the IO link should be tested on both the Primary System in normal mode and the Standby System in fallback mode to verify operation under normal and fallback modes.

To perform a fallback and recovery test, wait for the regeneration to complete. On the Standby System, log on to *'root'*, enter the **'sentinel'** command and observe the status of the system; it should state that the mirror *is 'Controlled by Primary System'*. Issue two *'sync'* commands on the Primary System, then press the **'reset'** button on the Primary System (This is the preferred method in place of powering off the Primary System).

Within 30 seconds, the reader should see that the status of the Standby System has changed to *'Fallback Mode'*; the mirrored unit will be undergoing an *'fsck'*; and finally the mirror should be mounted on the Standby System, and all devices and logins should be available from the Standby System. Correct any login, printer, or other device errors at this time.

To recover, reboot the Primary System. The operating system will automatically prompt the user for an *'fsck'* of the root filesystem; perform this and clear the possibly many zero

## Operating Sentinel

---

length files that *'fsck'* prompts for clearance. These zero length files are temporary 'pipe' files that TCP/IP uses during operation, and may be safely discarded.

Enter *'sentinel'* on the Primary System. When ready to recover, select *'Utilities: Recover from Fallback to Normal Mode'*. This will prompt the user with another question to authorize the recovery; enter *'y'*. The Primary System then:

- Sends a message to the Standby System
- Instructs it to stop all users on the mirrored unit
- Unmounts the mirror from the Standby System
- Remounts the mirror on the Primary System
- Starts up daemons and users on the Primary System

The **SavWareHA** system will be regenerating data from the Standby over to the Primary system; all reads will be done from the Standby System until regeneration is complete.

When regeneration is complete, the read flag is turned *'off'* on the Standby and *'on'* on the Primary System. The **SavWareHA** system is back in compliance.

## SavWareHA Administration Menu

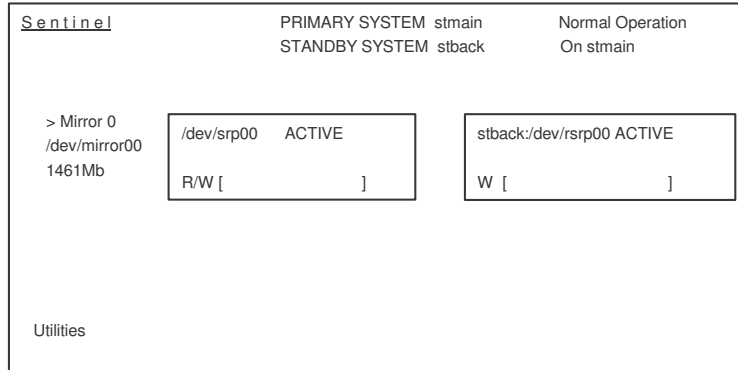
The **SavWareHA** administration menu can be started by logging on to the *'root'* account and entering the command *'sentinel'*. If no mirrors are configured, only the *'Utilities'* option will be displayed. Sub-menus can be selected for each mirror and for *'Utilities'* by using either the arrow keys or [SPACE] (space bar) to move the ">" pointer to the desired selection and pressing [RETURN].

### SavWareHA Menu Keystrokes

The following keystrokes can be used while in the **SavWareHA** administration menu:

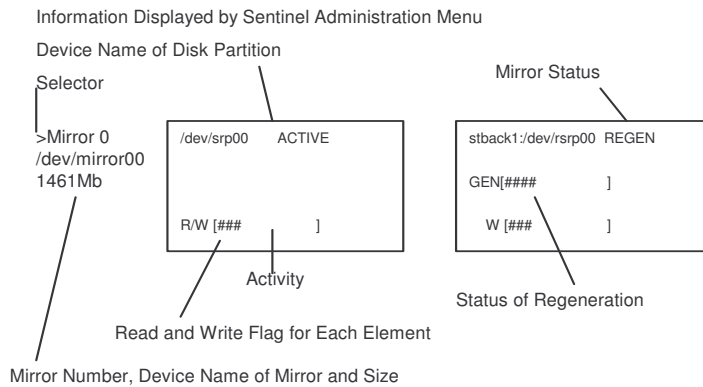
“↑”, “↓” (Up/Down Arrow)	Move Selection Pointer to Next/Previous Option
[SPACE] (spacebar)	Move Pointer to Next Option
[RETURN]	Execute Option at Pointer
[END], [ESC], [CONTROL-E]	Exit Current Menu, Abort Current Entry
[CONTROL-R]	Repaint Current Screen
“←”, “→” (Left/Right Arrow)	Move current sub-window left/right
“U”, “u”, “D”, “d” keys	Move current sub-window up/down
[BACK-TAB]	Move to previous input
[F3]	Generate and choose a response from a list of options

SavWareHA Administration Screen



Each mirror will consist of a pair of mirror element boxes, the Primary System shown on the left, and the Standby System shown on the right. Underneath the mirror description ('Mirror 0' above) is shown the device name ('/dev/mirror00'), and its size (1461Mb). The system status ('Normal Operation', 'On stmain') is shown in the upper right hand corner of the screen.

SavWareHA Screen Definitions



### Mirror Status Codes

The mirror status codes are located in the upper right hand corner of each mirror side's display box. See the *'Specific Error Recovery Procedures'* section in this manual for more information.

<b>ACTIVE</b>	Mirror is GOOD and in use. The mirror is currently mounted.
<b>CHECK</b>	An unknown error was encountered while getting the status of the mirror. Usual indications are network or device errors or misconfiguration of mirror devices. Removal and rebuilding of the mirror (not <b>SavWareHA</b> itself), or testing of the network connection is the suggested remedy.
<b>COLD</b>	The element data is valid as of the date displayed. The element can be updated manually via the Regenerate Mirror option. Usually reserved for either root mirrors or for inactive (archival) data mirroring.
<b>DISK ERROR</b>	A disk error has been reported on the mirror device. The mirror has been broken. Correct the hardware problem and regenerate the mirror with the <i>'Fix Broken Mirror: Regenerate Mirror'</i> selection.
<b>ERROR</b>	An error was encountered while trying to access the mirror device. Correct the network communication error(s) and regenerate the mirror with the <i>'Fix Broken Mirror: Regenerate Mirror'</i> selection.
<b>GOOD</b>	Mirror is GOOD but is not in use. The mirror is not currently mounted.
<b>MISMATCH</b>	The date & time codes are different on the mirror elements. This usually occurs when some of the controlling files in either <i>'etc/default'</i> or <i>'etc/sentinel.d'</i> have become corrupted; or may reflect misconfiguration of mirror devices. Removal and rebuilding of the mirror (not <b>SavWareHA</b> itself) is the suggested remedy.
<b>NEED REGEN</b>	Mirror side requires regeneration. No reads are enabled on the destination mirror element.
<b>NOT CLOSED</b>	The system was not shutdown properly. The mirror must have <i>'fsck'</i> run on the mirror device after regeneration has started. Until both regeneration and a successful <i>'fsck'</i> has occurred, the mirror will not be accessible.
<b>NOT SETUP</b>	The drive has not been setup for a mirror. This usually means that the <b>SavWareHA</b> configuration files expected to find a mirror stamp on the device, but did not.

---

## SavWareHA Sub-Menus

### Mirror Sub-Menu

This menu will be enabled when the selector is moved by pressing either “↑”, “↓”, or [SPACE] keys to position the highlight pointer at a mirror and [RETURN] is pressed. The options executed from this menu only affect the mirror selected (except for display mirror log which shows entries for all mirrors on the system).

> Display Mirror Status  
Fix Broken Mirror  
Check Mirror Integrity  
Change Read Flags  
Make Unix File System  
Remove Mirror  
Display Mirror Log

Each of these selections are described in detail on the following pages. Please note that only the system administrator should access these selections.

### Display Mirror Status

This selection allows the user to examine the status flags for each side of the mirror. Also, information as to the creation and last access dates and times are shown, and the size of each mirror side.

This screen will show the protocol for the data link (TCP/IP or LLI) when both sides of the mirror are active.

### Fix Broken Mirror

This selection displays another menu whose three selections are shown below:

> Fix Mirror by Copy  
Turn Off Standby Element  
Turn Off Primary Element

This menu controls the ‘fixing’, or regeneration, of a broken mirror; allows the user to turn off the Standby System element or the mirror selected. By pressing [END], [ESCAPE], or ‘q’, the user may exit this menu without making a selection.

Each of these options is explained below, with representative screen displays.

### Regenerate Mirror

#### [Use With Caution!]

Use this option to bring the mirror back into compliance after some error has either broken the mirror completely (either the master or slave side of the mirror has been turned **'OFF'** intentionally or by the action of the mirror driver), or one side of the mirror reported a read or write error (usually resulting in the offending side's read/write flag being set to **'W'** rite only), or when the Primary System is unable to reach the Standby System (the slave side of the mirror will then display **'ERROR'**).

The program will display the mirror elements again, and will prompt the user for authority to begin a regeneration from the **'GOOD'** or **'ACTIVE'** side of the mirror to the inactive side of the mirror. An affirmative (**'y'**) response will immediately begin the regeneration; any other response will abort the regeneration.

```
Mirror /dev/rmirror00  Fix A Broken Mirror
/dev/mirror00  /dev/srp00  ACTIVE  stback:/dev/rsrp00 OFF
1461Mb
R/W [##      ]  W [      ]
OK to copy /dev/srp00 to stback:/dev/rsrp00 ?
```

**Always check the direction of the copy to insure that valid data will not be inadvertently overwritten.**

### Turn Off Standby Element

#### [Use with Caution]

This option's effect is to turn off all reads and writes to the slave (Standby System) side of the mirror. No system interruption will be noticed by the users on the Primary System. All pending disk writes will be flushed, no more writes (or reads) will be accepted. The slave side's status will be set to **'OFF'**.

Some circumstances where the user may wish to use this option: The Standby System must be taken completely away; a verified backup of the mirrored unit in a quiescent state is desired without interrupting the Primary System's users; some impending Standby System, network error or major reconfiguration is anticipated in the near future.

Note that this will pause the **SavWareHA** monitor program on the Standby system. No further mirroring will be done, and no automatic fallback process will be initiated. Use the **'Fix Broken Mirror: Regenerate Mirror'** option to bring the mirror back into compliance.

```

Mirror /dev/mirror00 Turn Off Standby Element

/dev/mirror00 /dev/srp00 ACTIVE stback:/dev/rsrp00 ACTIVE
1461Mb R/W [ ] W [ ]

Mirror #0 is not Broken. Do you want to turn off slave element?
    
```

Turn Off Primary Element

This option's effect is to turn off all reads and writes to the Primary System's side of the mirror. No system interruption will be noticed by the users on the Primary System. All pending disk writes will be flushed, no more writes (or reads) will be accepted. The Primary System's status will be set to **'OFF'**.

Some circumstances where the user may wish to use this option: The Primary System's disk array must be taken completely away (assuming there are 'hot swap' drive carriers in use); some impending system threatening error is about to occur if the Primary System's side of the mirror is left active for either reading or writing.

```

Mirror /dev/mirror00 Turn Off Primary Element

/dev/mirror00 /dev/srp00 ACTIVE stback:/dev/rsrp00 ACTIVE
1461Mb R/W [ ] W [ ]

Mirror #0 is not Broken. Do you want to turn off master element?
    
```

Note that this will NOT pause the **SavWareHA** monitor program on the Standby system. No further mirroring will be done on the Primary System - all reads and writes will be performed on the Standby System over the data link connection. Use the **'Fix Broken Mirror: Regenerate Mirror'** option to bring the mirror back into compliance.

### Check Mirror Integrity

#### **[USE CAUTION!]**

**NOTE:** This option was removed in Revision 2.01a of **SavWareHA**, and is included here for those installations prior to Revision 2.01a.

This option will allow dynamic block by block comparison of the mirrored units. Please note that on an active filesystem there may be some errors reported by this option. This is due to the fact that the mirror reads on either side may be dynamically interrupted by a write buffer flush before both side's blocks have been read, thereby reporting an error that does not really exist. For this reason, only use this option when the mirror is quiescent (unmounted).

### Change Read Flags

#### **[USE CAUTION!]**

This option will allow the user to toggle on or off the read enabling flags for either side of the mirror. Care should be taken when using this option, as system performance may be impacted (in the case of turning on reads to both sides of the mirror)..

```
Mirror /dev/rmirror00  Change Read Flags

/dev/mirror00  /dev/srp00  ACTIVE  stback:/dev/rsrp00 ACTIVE
1461Mb        R/W [      ]      W [      ]

> Primary Reads      : On
  Standby Reads      : Off
  Ram Drive Read     : Off
Use arrows to move to desired selection.
Press [RETURN] to toggle selection.
Press [END] to make changes.
```

### Make Unix File System

#### **[USE WITH EXTREME CAUTION!]**

This option should only be used ONCE (if at all). Its function is to perform a 'mkfs' (Make Filesystem) function on the mirror device selected. This option is **NOT NECESSARY** for proper **SavWareHA** operation; since the regeneration from the Primary System to the Standby System duplicates the Primary System's already extant filesystem.

If this option is desired to be executed, the user will be prompted with the command 'mkfs' and all options as calculated by **SavWareHA** for the mirrored filesystem.

**ALL EXISTING DATA ON BOTH SIDES OF THE MIRROR WILL BE REMOVED. A CLEAN (EMPTY!) FILESYSTEM WILL RESULT FROM A POSITIVE RESPONSE.**

If this is what is desired (the user is creating a new, clean, filesystem), answer ‘y’ to the ‘mkfs .... ‘ question. A new filesystem will be built immediately.

The function of this command is to produce a clean filesystem on both sides of the mirror (actually, the ‘netdisk’ driver duplicates block for block the ‘mkfs ...’ instructions on the Standby System). If a filesystem already existed on the master (Primary System) side of the mirror, the initial regeneration would have exactly duplicated the filesystem structure, thereby making use of this option unnecessary.

### Remove Mirror

#### [USE CAUTION!]

This option instructs **SavWareHA** to remove all references to the mirrored filesystem from its control files. The mirrored filesystem must be quiescent (unmounted) in order to execute this command. Specifically, this command removes the mirror data from the following files:

```

/etc/default/filesys
/etc/default/mirror
/etc/default/netdisk
/etc/sentinel.d/filesys

```

Both the Primary and Standby Systems are updated simultaneously. Please note that **no data is destroyed**; only the entries in the **SavWareHA** control files that pertain to constructing and managing a mirrored system are touched. The user may individually mount the respective mirror components on each system after this command has terminated.

```

Mirror /dev/mirror00  Remove Mirror

/dev/mirror00  /dev/srp00  GOOD  stback:/dev/rsrp00  GOOD
1461Mb        R/W [          ]    W [          ]

OK to remove ?

```

### Display Mirror Log

This option displays the file *'usr/adm/dmx.log'*. The last 20 lines of the file are displayed; the command then waits for one of the following to be entered:

[RETURN]	End the display.
[ESCAPE] / [END]	End the display.
[PAGE UP]	Show the previous page.
[PAGE DOWN]	Show the next page.
[HOME]	Show the first 20 lines.

At the end, the program returns control to the **SavWareHA** sub-menu.

## SavWareHA Sub-Menus - Utility Sub-Menu

### Utility Sub-Menu

The *'Utility'* sub-menu lists several options that act on the system as a whole (are not mirror specific).

```
>Display Monitor Status
Build A Mirror
Display Mirror Log
Software Registration
Force Change To Fallback Mode
Recover From Fallback to Normal
```

### Display Monitor Status

This option will display the current status of the **SavWareHA** monitor program, *'smon'*. Under normal mode, the output of this option will look like:

```
Reply is:Smon Status Request_REPLY
Monitor status: Normal
Link 0 Net-UDP Status:Active Timeout:8 Error Retry:60 Errors:4 neptune
```

### Build A Mirror

#### **[Use Caution!]**

This option will allow construction of a mirror. When this selection is made, the user will be prompted to enter the device name for the Primary System (master) side of the mirror, followed by the device name for the Standby System (slave) side.

Sentinel - Build A New Mirror

Enter Primary Device Name :                    F3 for list

Enter Standby Device Name :

Enter Mountpoint                                :

When the mirror is built, all data will be copied from the Primary device to the Standby device. This will destroy all existing data on the Standby device.

**SavWareHA** then asks for the mountpoint. This will be a directory on which the mirror device (*/dev/mirror00*) will be mounted. By SCO convention, the first non-root filesystem is named *'u'*. Enter *'N'* if this is a raw database partition that is not to be mounted.

After the mountpoint has been acknowledged, **SavWareHA** will start a regeneration, and will mount the mirror. The mirror status should read *'Active' / 'Active'*, and **SavWareHA** will be in *'Normal'* operation.

[END], [ESCAPE], or [DELETE] will abort the build of a new mirror.

### Display Mirror Log

This selection will display the file *'usr/adm/dmx.log'*. The program displays the last 20 lines in the file. The output may look something like:

```
Sat Apr 16 24 13:20:58 PST 1999
  SavWareHA
  High Availability Software
  Release 2.04c

for SCO Unix 3.2v4.0 or Greater
'custom'
(05-Apr-99)     P/N 2206ST06

Mirror #0 Sat Apr 16 13:28:03 1999
Mirror copy started: /dev/srp00 to pluto1:/dev/rsrp00

Mirror #0 Sat Apr 16 15:27:02 1999
Mirror 0 Copy Complete
```

## Operating Sentinel

---

This particular log shows the initial installation of **SavWareHA**, and then the beginning and end of the regeneration of mirror #0.

### Software Registration

This is the 'reggie' program that will brand the software with the serial number and registration response key. All registration activities are conducted through this program.

```
STAR Technologies Software Registration

Product      Serial#  Status      Installed   Options
Sentinel     123456  Registered   Apr 16, 1995  0

> Log   Registration   Serialization
```

The '**Log**' selection will display the last 20 lines of the Registration Log file

The '**Registration**' selection will allow registration of a serialized copy of **SavWareHA**. This will display the Registration Request String which must be transmitted back to Avnet Computer Marketing to fully activate the product(s).

The '**Serialization**' selection will allow entry of a valid serial number for the **SavWareHA** product. This would convert a Demonstration version of the product into an Evaluation, and with proper registration into a fully licensed copy of **SavWareHA**.

### Force Change To Fallback Mode

#### **[USE CAUTION!]**

This will initiate a transition to fallback mode. This may be used to test the **SavWareHA** system, or to force a change when imminent system problems or maintenance is required on the Primary System.

A response of "y" will begin the transition immediately. The users of the mirror(s) on the Primary System are stopped ('kill' is used to evict them), the mirror is unmounted, the Standby System is then instructed to begin its fallback mode operations. The mirror is mounted on the Standby System, users are enabled, and daemons in 'etc/sentinel.d/S\*' are started.

Force Change to Fallback Mode

Ok to force change to fallback mode ?

*Recover From Fallback to Normal*

**[USE CAUTION!]**

This will initiate a transition to normal mode. This option must be selected from the Primary System only; a warning message will be displayed if attempted from the Standby System.

A response of "y" will begin the transition immediately. The users of the mirror(s) on the Standby System are stopped ('kill' is used to evict them), any user-defined daemons are stopped (via '/etc/sentinel.d/K\*' scripts), the mirror is unmounted, the Primary System is then instructed to begin its recovery to normal mode operations. The mirror is mounted on the Primary System, users are enabled, and daemons in '/etc/sentinel.d/S\*' are started.

Sentinel - Recover to Normal Mode

This procedure should be used after a failure that caused a switch to Fallback mode.

OK to recover to Normal mode ?



# 4. Error Conditions & Recovery

## Error Classes

Any error encountered in the **SavWareHA** system may be classified in as either a Class 1, Class 2, or Class 3 error.

**Class 1** errors are generally defined as errors on either system that go unnoticed by any user on the either the Primary or Standby Systems, with the exception of the designated system administrator and the console on the Primary System. These errors will require prompt attention, as they indicate a degradation of the system's integrity. Failure to respond appropriately may result in system disintegration. **SavWareHA's** usual action is to turn the offending side of the mirror 'OFF'.

**Class 2** errors are generally defined as major systemic errors on the Primary System which will be noticed by all users on the Primary System, including the system administrator. These errors will be sensed by the *'smon'* program, and will cause the Primary System to be regarded as untrustworthy. The Primary System mirror accesses will be shutdown (if possible - the system may already be down). The Standby System will then mount and check the mirrored data partition, perform any application-specific startup procedures, and users may then sign on to the Standby System and continue with their tasks.

Depending upon the particular application, the recovery from a Class 2 error may require some additional repairs of data. Additionally, the application may also require some amount of re-entry of work (usually the last uncompleted transaction). "Journalling" data base systems, as well as SCO supported journalling filesystems, will reduce or eliminate this possible situation.

**Class 3** errors consist of a major catastrophe affecting both the primary and Standby Systems, or major disruption of either the data or I/O links. With the class of error, neither machine is functional; or the I/O link has been severely disrupted, preventing user access to the system(s).

## General error actions

Class 1 errors will result first in the mirror read flag being turned 'OFF' for the offending half of the mirror, and after four additional errors will turn the offending half of the mirror '**OFF**'. The error count is cumulative since boot. These errors will not cause the application to receive an error; however, all Administrators will receive e-mailed notification of the error condition (see the file *'/etc/sentinel.d/administrators'*).

Class 2 errors will result in the Standby System taking control of the mirror and associated users. In this case, it is presumed that the Primary System had crashed or become wildly unpredictable. This could also be the result of a network disconnection, and having only one link specified to **SavWareHA** in the *'/etc/sentinel.d/links'* file on the Standby System. Alternatively, the Primary System could have become so overloaded that it failed to respond to the Standby System's status request within the number of seconds held in the same *'/etc/sentinel.d/links'* file.

Class 3 errors mean that neither system is functional -- and all that implies.

### Specific Error Recovery Procedures

While we cannot specify each and every type of error, some of the more common errors and recovery procedures are listed here. Please note that the status codes may be 'reversed'; that is in reality, the Primary System may actually reflect what the Standby System says in the following examples; if recovery procedures differ significantly, they are discussed separately.

#### Meaning of 'Good' versus 'Active'

'Good' refers to a mirror element that not has had errors recorded against it, and 'Active' refers to a mirror element that is currently 'Good' and is mounted or in use by an application or database.

#### Mirror showing 'Good' / 'Good'

This is the status of a completely correct, unaccessed mirror. No action is required to use this mirror except to access it. Always mount or access the mirror device ('/dev/mirror00' or '/dev/rmirror00'), **not** the mirror components.

/dev/srp00      GOOD R/W [            ]	stback:/dev/rsrp00 GOOD W [            ]
--	---

#### Mirror showing 'Active' / 'Active'

This is the status of a completely correct mirror that is either mounted or being accessed by an application. The mirror is currently in use. To release, use the 'umount /dev/mirror00' command or stop the database application.

/dev/srp00      ACTIVE R/W [            ]	stback:/dev/rsrp00 ACTIVE W [            ]
--	---

#### Mirror showing 'Not Closed' / 'Not Closed'

This is the result of a system crash (Class 3 error) on both machines; or, the result if the keyword 'Manual' was inserted in the Standby System's '/etc/sentinel.d/links' file, and the Primary System crashed and was rebooted.

/dev/srp00      NOT CLOSED R/W [            ]	stback:/dev/rsrp00 NOT CLOSED W [            ]
--	---

In essence, this is similar to the situation that arises when a standard Unix filesystem is left 'open' (or mounted), and the operating system halts. Use the following steps to get the mirror back up and running:

- Turn off the slave side of the mirror (!! see note below) by using the *'Fix Broken Mirror: Turn Off Standby Element'* option
- Start a regeneration by using the *'Fix Broken Mirror: Regenerate Mirror'*
- Perform a *'fsck'* on the mirrored device (**NOT** on either of the mirror components!)
- After *'fsck'* completes successfully, either bring the system back into multi-user mode, or remount the mirror device.

**NOTE:** The side that should be turned 'OFF' should be the side that has only the 'W'rite flag (not the 'R'ead flag) turned on. This would be the 'more correct' subcomponent to regenerate from, as the system was last instructed to perform reads from this device.

*Mirror showing 'Good' / 'Not Setup'*

- or -

*Mirror showing 'Active' / 'Not Setup'*

This is the result of not completing the mirror specification on both sides. **SavWareHA** expected mirror information to be available on the Standby System mirror component, but did not find any.

To correct this, remove the mirror, and rebuild. Make sure that both systems are accessible to each other.

/dev/srp00      GOOD R/W [            ]	stback:/dev/rsrp00 NOT SETUP
/dev/srp00      ACTIVE R/W [            ]	stback:/dev/rsrp00 NOT SETUP

*'fsck' Reports 'Cant stat filesystem'*

The mirror is most probably in a condition that the filesystem was not unmounted before the Primary System was shutdown. Perform the following steps:

- Turn off the slave (Standby) side of the mirror - **IF** the Primary System had the 'R'ead flag turned on, otherwise, turn off the master (Primary) side of the mirror
- Start a regeneration by using the *'Fix Broken Mirror' / 'Regenerate Mirror'*
- Perform a *'fsck'* on the mirrored device (**NOT** on either of the mirror components!) Use *'fsck /dev/mirror00'*.
- After *'fsck'* completes successfully, either bring the system back into multi-user mode, or remount the mirror device.

## Error Conditions & Recovery

---

### Mirror Showing 'Good' / 'Off'

- or -

### Mirror Showing 'Active' / 'Off'

This is a result of either a Class 1 error, or of the use of the *'Fix Broken Mirror: Turn Off Standby Element'* menu selection.

In the case of a Class 1 error, this means that the Primary or Standby System has become unreliable (disk errors on respective sides of the mirror). Proper maintenance is indicated on the offending side: Fix any drive errors, make certain that the offending side functions correctly. As soon as it is certain that the system has been repaired, perform the following steps to correct both error possibilities:

- Connect the offending side to power and to both the Data and I/O Links
- Boot the system into multi-user mode
- Verify that the network connections are functioning between the Primary and Standby Systems
- Enter **SavWareHA**; select the mirror in question; start a regeneration by the use of *'Fix Broken Mirror: Regenerate Mirror'*.

<pre>/dev/srp00      GOOD R/W [          ]</pre>	<pre>stback:/dev/rsrp00  OFF</pre>
<pre>/dev/srp00      ACTIVE R/W [          ]</pre>	<pre>stback:/dev/rsrp00  OFF</pre>

### Mirror Showing 'Off' / 'Good'

- or -

### Mirror Showing 'Off' / 'Active'

This is a result of either a Class 1 error, or of the use of the *'Turn Off Primary Element'* menu selection.

In the case of a Class 1 error, this means that the Primary System's disk filesystem has become unreliable (disk errors on the master side of the mirror). Proper maintenance is indicated on the Primary System. If the installation has 'hot swap' drives, replace the offending drive(s), rebuild any stripe that may have been in existence.

After any necessary repairs have been accomplished, use the *'Fix Broken Mirror: Regenerate Mirror'* selection to bring the mirror back into compliance.

If the installation does not have 'hot swap' drives, then it will become necessary to force a change to fallback mode, enabling the Standby System to take over while the Primary System is being repaired. This may be accomplished in the following manner:

- Either wait until the system is quiescent, or notify all users that work should be terminated shortly for a planned switchover to fallback mode
- Enter **SavWareHA** on the Primary System; select the *'Utilities'* sub-menu
- Highlight the command *'Force change to Fallback Mode'*, press **[RETURN]**, answer **'y'** to force the change

**SavWareHA** will automatically switch control over to the Standby System. Users may then log onto the Standby System and operate in fallback mode.

As soon as it is certain that the Primary System has been repaired, perform the following steps to transfer back to normal mode.

- Connect the Primary System to power and to both the Data and I/O Links
- Boot the Primary System into multi-user mode; **SavWareHA** will see that the Standby System is in fallback mode, and will adjust the Primary System accordingly
- Verify that the network connections are functioning between the Primary and Standby Systems
- Enter **SavWareHA** on the Primary System
- Notify all users that a planned recovery to normal operation is about to occur
- When the system is quiescent, select the *'Utilities'* sub-menu from the **SavWareHA** main screen on the Primary System
- 
- 
- Highlight the selection *'Recover from Fallback to Normal Mode'*; answer **'y'** to start the recovery process
- Since no *'fsck'* will have to be done, the switch over time is limited to the raw time **SavWareHA** takes to disable daemons and terminals on the Standby System, and to enable the same on the Primary System
- 
- A regeneration is started immediately from the Standby to the Primary System; users may log into the Primary System at any time and continue work.

/dev/srp00      OFF	stback:/dev/rsrp00      GOOD R/ W [                      ]
---------------------	---

/dev/srp00      OFF	stback:/dev/rsrp00      ACTIVE R/ W [                      ]
---------------------	---

## Error Conditions & Recovery

---

### Mirror Showing 'Good' / 'ERROR'

- or -

### Mirror Showing 'Active' / 'ERROR'

This is a result of a failure of the Data Link between the Primary and Standby machines. Test the link by *'rlogin'* or *'telnet'* or *'ping'*; these will probably not work correctly (if at all). Check the integrity of the cable (or other connection) between the machines. Once any of the above three commands function, SavWareHA should clear the 'ERROR' and replace it with 'NEED REGEN' (see below).

<code>/dev/srp00      GOOD R/W [            ]</code>	<code>stback:/dev/rsrp00    ERROR</code>
--	--

<code>/dev/srp00      ACTIVE R/W [            ]</code>	<code>stback:/dev/rsrp00    ERROR</code>
--	--

### Mirror Showing 'Disk ERROR' / 'Good'

- or -

### Mirror Showing 'Disk ERROR' / 'Active'

This is a result of a hard disk error on the Primary System. The mirror has shut off reads and writes to the local disk(s); all reads and writes are being serviced by the Standby System. As soon as possible, repairs need to be effected on the Primary System.

If the Primary System has 'hot swap' drives, the drive exchange could be completed, and a regeneration started without user interruption. If the Primary System does not have this technology, then a manual transition to fallback mode will have to be initiated as soon as possible.

/dev/srp00    DISK ERROR W [                    ]	stback:/dev/rsrp00    GOOD R/W [                    ]
--	--

/dev/srp00    DISK ERROR W [                    ]	stback:/dev/rsrp00    ACTIVE R/W [                    ]
--	--

Mirror showing 'Good' / 'MISMATCH'

- or -

Mirror showing 'Active' / 'MISMATCH'

This is the result of a mismatch in the mirror specification on the offending side. **SavWareHA** expected mirror information to be identical to the Primary System, but did not find that to be the case.

To correct this, remove the mirror, and rebuild. Make sure that both systems are accessible to each other.

/dev/srp00    GOOD R/W [                    ]	stback:/dev/rsrp00    MISMATCH
--	--------------------------------

/dev/srp00    ACTIVE R/W [                    ]	stback:/dev/rsrp00    MISMATCH
--	--------------------------------

Mirror Showing 'Good' / 'Need Regen'

- or -

## Error Conditions & Recovery

---

### Mirror Showing 'Active' / 'Need Regen'

This is a result of a reboot after the Standby System went down unexpectedly, an incomplete regeneration, a repair of a network error (see above), or the Standby System's mirror date stamp is incorrect. Use the following steps to get the mirror back up and running:

- Turn off the slave side of the mirror by using the *'Fix Broken Mirror: Turn Off Standby Element'* option
- Start a regeneration by using the *'Fix Broken Mirror: Regenerate Mirror'* option.

<pre>/dev/srp00  GOOD R/W [      ]</pre>	<pre>stback:/dev/rsrp00  NEED REGEN</pre>
--	---

<pre>/dev/srp00  ACTIVE R/W [      ]</pre>	<pre>stback:/dev/rsrp00  NEED REGEN</pre>
--	---

### Mirror Showing 'Good' / 'REGEN'

-- or --

### Mirror Showing 'Active' / 'REGEN'

This is the normal status of a mirror undergoing a regeneration. Data is being copied from the 'Good' or 'Active' side of the mirror over to the side marked as 'REGEN'. While this is taking place, full access of the mirror may be made; the regeneration process runs in the background only when there are no disk writes pending.

The mirror element being regenerated will have an additional status line marked 'GEN' (as shown below). This will give an approximate indication of the amount of regenerated data written to the target mirror element.

The regeneration process will run in the background even if the initiating user logs off the machine. The process may be terminated with a *'kill'* statement to the proper process ID number; if this happens, the target mirror element will be flagged as 'NEED REGEN', and will not be available for reads.

/dev/srp00    GOOD R/W [            ]	stback:/dev/rsrp00    REGEN GEN [#####        ] W [            ]
--	--

/dev/srp00    ACTIVE R/W [            ]	stback:/dev/rsrp00    REGEN GEN [#####        ] W [            ]
--	--

## Unique Or Difficult Situations

This section will document 'unique' situations that have happened in the field requiring knowledge not only of **SavWareHA**, but also of Unix and Unix filesystems. Most of the cases described below were encountered and corrected by members of the Avnet Computer Marketing Customer Support Department; some were reported and corrected by knowledgeable customers. Should the reader find himself or herself in a 'unique' **SavWareHA** situation, please contact us immediately at:

- **Voice:**        (949) 789-4361
- **Fax:**         (949) 789-4365
- **E-mail:**     [cpgsupport@avnet.com](mailto:cpgsupport@avnet.com)

Please send all particulars of the system's hardware and software setup, including driver and board types and revision levels; circumstances surrounding the 'unique' situation; steps that have been performed already; and access to a 9600 baud ECC modem on the Standby System is mandatory in most cases. Premier Support customers may call the Premier Support Hot Line number.

Provision of support methods are governed by Avnet Computer Marketing's current Customer Support policies.

### Installation of SavWareHA, no backup

#### **Situation:**

- **SavWareHA** desired to be installed on a Primary System
- No current or verifiable backup to data unit
- Minimal window available for installation

#### **Meaning:**

Sites that merit **SavWareHA** usually exhibit a high level of activity throughout the normal business day. In some cases, this may mean that on-line verified backups are not possible. In other cases, a backup may not have been performed for some time.

## Error Conditions & Recovery

---

Installation of **SavWareHA** into these clients becomes necessary to provide not only fault tolerance, but also to allow the ability to perform on-line backups and system maintenance.

### Possible solution:

This will require a small bit of downtime to install **SavWareHA** and relink the kernel. The installation should take no more than 15 minutes of downtime, so this could be done at lunch, near shift changes, late at night, etc. Follow these steps:

- Place the Standby System into single user mode.
- Install SavWareHA on the Standby System first; register and relink the kernel.
- Place the Standby System into multi-user mode; there may be 'error' messages stating that **SavWareHA** has not been properly configured. This is normal for this case.
- Place the Primary system into single user mode.
- Install **SavWareHA** as per the instructions in this manual.
- Remove the "about to be mirrored" filesystem entry from `'/etc/default/filesys'`. Either comment out the line with a leading `'#'` sign, or remove the line completely (the latter is the preferred method).
- Reboot the Primary system into multi-user mode. The system will not mount the desired unit; users who attempt to log on will receive errors to that effect.
- Log on to `'root'` on the Primary System. Enter `'sentinel'`; configure the Primary and Standby Systems; build the mirror, and specify the same mount point to **SavWareHA** as that which was removed from `'/etc/default/filesys'`. **DO NOT MAKE A UNIX FILESYSTEM!!!!**
- The regeneration will start; mount the mirror (use the `'mountall'` command); users may log on and begin working.

(Thanks to both John Magill and Margaret Smith of Computer Configurations and Debbie Mann of the Automobile Assn of South Africa)

## Data Recovery / Archiving

This section deals with situation in which data is desired to be either recovered, archived, or 'moved aside' in preparation for new untested application programs. Even though we specify several theoretical data recovery and archival processes here, the ultimate responsibility for data integrity is on the local system administrator.

### Moving one half of the mirror aside

#### Situation:

- A new procedure or application will be placed on-line.
- MIS wishes to provide immediate recovery if this procedure or application fails

#### Meaning:

Occasionally, new packages or substantial software upgrades are installed on MIS systems. By judicious use of **SavWareHA**, the system administrator may ease the data recovery process.

Since this procedure involves a “wrong way” regeneration, it behooves the user to make certain that **verified backups** exist before doing anything of this nature! Security of data is ultimately the user’s responsibility.

Please note that this procedure is **NOT SUPPORTED**, and that the reader undertakes this knowingly **AT HIS OWN PERIL**.

### Possible Solution:

What is desired is to suspend mirroring until such time as the new procedure or application has been verified to be functional. The longer mirroring is suspended, however, the more transaction-oriented work will have to be redone! Follow these steps:

- Turn off the Standby System (slave) side of the mirror.
- Add or access the procedure or application; allow it to perform whatever function it was designed to do.
- If the procedure or application functions correctly, bring the mirror back into compliance by use of the “*Fix Broken Mirror: Regenerate Mirror*” selection detailed elsewhere in this manual.

### Backward Regeneration

If the procedure or application “blows up”, and the resultant modifications to the data become unrecoverable or untraceable, follow these steps to recover the data on-line. Please note that it is assumed the mirror in question is the first mirror, */dev/mirror00*:

- Login to the ‘root’ account
- Stop all access to the mirrored unit (which should be on the Primary System)
- Unmount the mirror with the *'umount /dev/mirror00'* command on the Primary System
- Issue the command *'etc/mirror -z1'*. This will authorize the backward regeneration
- Use the “*Fix Broken Mirror: Regenerate Mirror*” selection for the mirror in question
- At the prompt “*OK to copy /dev/u to standby:/dev/ru ?*”, issue an exclamation mark (!)
- A detailed warning message then appears, followed by the prompt “*OK to copy standby:/dev/ru to /dev/u ?*”
- Answer in the affirmative (“**y**”) to begin the backward regeneration
- **NOTE:** You will see several error messages regarding the fact that the read flags were not set on the Standby System; this is normal for this backward regeneration process
- When the backward regeneration has finished, remount the mirror, and correct your problem(s).

## Error Conditions & Recovery

---

Please note that the mirror may be mounted while the backward regeneration is in process; however, since some unknown application error caused the data to become corrupted, it would be highly advisable to wait for the completion of the regeneration prior to activating the mirror once again!

### Recovery of files when mirror shows 'Active' / 'Off'

#### **Situation:**

- Slave (or master) side of mirror previously turned off
- Unintended *'rm'* of a file or files on 'Active' side
- No way to recover the file or files from backup

#### **Meaning:**

Unintended deletion of critical files is the bane of the Unix System Administrator. If the reader happens to have removed a file or files, and had previously turned off one side of the mirror, it is possible to recover the removed files.

Although **SavWareHA** has been installed in sites ranging from Vancouver, BC to Johannesburg, SA, it should **never** be viewed as a substitute for frequent, checksum-verified data backups to removable media, such as those provided by 'XA' from Avnet Computer Marketing. See the discussion of unattended, verified system backups in the "Application Notes" section of this manual; also, please secure a copy of 'XA' if you do not already have one.

#### **Possible solution:**

What is needed is to mount the slave (or 'OFF' side of the mirror) on the Standby System, and then *'rcp'* the file or files over to the master (Primary System) side of the mirror. For the purposes of this discussion, we will assume that the Standby System had its mirror element turned 'OFF'. Follow these steps:

- Log in to the Standby System; determine what the mirror component is called on the Standby System by looking at the **SavWareHA** screen, right hand box. We will assume it is *'/dev/srp00'*.
- Issue: *'fsck /dev/srp00'*; this will repair the filesystem. This is needed because when an element is turned 'OFF', the filesystem is still marked as 'Open'.
- Issue: *'mount /dev/srp00 /mnt'*. This will mount the filesystem on *'/mnt'*.
- Issue: *'cd /mnt'*. Then, use *'rcp'* to copy the file or files over to the Primary System.
- When finished, issue: *'cd /'* and then *'umount /mnt'*. Log off the Standby System. Perform a regeneration from the Primary System over to the Standby System when ready to bring **SavWareHA** back into compliance.

After this has been completed, **BACKUP** the mirror!

## Panics and Filesystem Errors

Unix kernel panics are rarely to be taken lightly. These exceptions to the normal operation of the system usually point to some major hardware error, if no change has been made to any of the software drivers on the system. A list of panic types and general meanings is available and can be used to diagnose what part or parts have failed.

The action of the mirror when a panic occurs under **SavWareHA** will usually occasion a transition to fallback mode, since the Primary System will almost certainly cease to function. In this case, normal **SavWareHA** repair and recovery procedures may be followed as detailed previously in this manual.

In rare cases, corrupt data has already been written to the mirror; and this in turn could compromise the entire **SavWareHA** system. In this latter case, **SavWareHA** has performed exactly correctly, as far as it was concerned: It received a block or blocks, and duly wrote them to both sides of the mirror.

How one recovers from these errors is more of an art form than science. One item is essential, though: checksum verified backups, such as those generated by 'XA'.

The errors reflected here deal with 'Class 3' errors, where both systems for some reason have gone down; other errors deal with filesystem integrity problems, and require delicate manipulation of the 'inode' table using a combination of advanced Unix commands.

### System Panic, backup uncertain

#### **Situation:**

- Panic or system crash of Class 3 (both systems down)
- Backup to tape (or other media) uncertain
- Rebooted system, mirror showing 'Not Closed' / 'Not Closed'  
-- or --
- Rebooted system, mirror showing 'Not Closed' / 'ERROR'
- Unable to mount mirror
- 'fsck' reports 'Cant stat filesystem'
- Integrity of mirrored filesystem not necessarily suspect

#### **Meaning:**

This is similar to a Unix filesystem status after a system crash, with minor exceptions. There is a real good possibility that the data will be 100% recoverable.

#### **Possible Solution:**

Turn off one side of the mirror. The choice will depend on the status of the mirror elements:

- If 'Not Closed' / 'Not Closed' - Pick either; slave side would be best
- If 'Not Closed' / 'ERROR' - Pick the slave side (or 'ERROR' side) to turn off

## Error Conditions & Recovery

---

If the reader is **really** uncertain (or nervous!) about the backup, or about the status of the filesystem, then issue:

```
mount -r /dev/mirror00 /mnt
```

(Note that `/dev/mirror00` may not be the correct mirror number! Make sure to enter the correct mirror number as shown on the system.) This mounts the mirror device as 'Read Only' on `/mnt`.

The reader can now copy the data (if possible) from `/mnt` to a valid tape or other backup device. Suggested method is to use `cpio`:

- Issue `'mount -r /dev/mirror00 /mnt'` to mount the mirror read-only on `/mnt`
- Issue `'cd /mnt'`
- Issue `'find . -print | cpio -oacvB > /dev/rStp0'` (assuming that `/dev/rStp0` is the correct tape device)
- Issue `'umount /dev/mirror00'` to unmount the mirror device

After the backup has completed (or if the last backup before the crash was verified to be correct), follow these steps:

- Issue `'fsck -n /dev/mirror00'`. This will show what errors may be expected on the mirror.
- If the output of the above does not appear to portend further disaster, issue `'fsck -s /dev/mirror00'`. This will repair the filesystem, and rebuild the free list.
- `'fsck'` will state at last: `'Set filesystem status to good?'` Answer `'y'`.
- Start a regeneration; then mount the mirror, and the users may continue their work.

### System Panic, backup uncertain, filesystem corrupted

#### **Situation:**

- Panic or system crash, especially 'Panic: CLFREE...'
- Backup to tape (or other media) uncertain
- Rebooted system, mirror showing 'Not Closed' / 'Not Closed'  
-- or --
- Rebooted system, mirror showing 'Not Closed' / 'ERROR'
- Unable to mount mirror
- `'fsck'` reports `'Can't stat root inode'` or `'Dups bad in root inode table'`
- Integrity of mirrored filesystem suspect

#### **Meaning:**

This is basically a bad situation to be in. The system has been seriously compromised; the `'Panic: CLFREE...'` notice from Unix is saying the filesystem(s) in question have been trashed either by unreliable / failed hardware, or by a rogue program or daemon. 100%

recovery is not to be expected. This makes a good point for 'XA' from Avnet Computer Marketing to ensure complete verified backups every day.

If the panic refers to the mirror, then the mirrored filesystem has been compromised. Since **SavWareHA** mirrors disk block write requests exactly, both sides of the mirror will have been compromised. Both sides will have the exact same problem(s).

In the case of 'fsck' reporting 'Can't stat root inode', most likely none of the following will work. It certainly could not hurt the filesystem any more to try, though....

### Possible Solutions:

Essentially the filesystem is unrecoverable. Having said that, the filesystem may be salvageable (which is NOT the same thing). If the reader can clear (manually) some of the 'inodes' that are listed as being duplicated, then the filesystem stands a good chance of being salvaged. **YOU WILL LOSE FILES.** The only question is which ones?

First and foremost, the reader must determine WHY this panic (especially 'CLFREE' or 'Free List Panic') occurred. In the past, we have seen:

- Bad / old / mixed firmware on drives (Micropolis 21xx series in particular)
- Bad SCSI termination or cabling (especially on newly installed systems)
- Overlapping 'divvy' or 'disk' partitions (most of this has gone away with later versions of Unix)
- Unauthorized use of restricted and undocumented filesystem commands (if you know the names of these commands, you should know how to use them correctly)
- Memory or cache problems in the CPU(s) or motherboard

If the backup is available and is good, remake the Unix filesystem on the mirror, restore all data, and re-enter the transactions from the time of the backup until the time of the panic. This is the preferred method.

If the backup is NOT good or trustworthy, or if there is no other choice:

- Turn off one side (pick the slave or Standby System side)
- Issue 'fsck -n /dev/mirror00'. Record files or inodes that are reported as bad. It is possible that 'fsck' will stop with the message 'Excessive dups for inode <inode>'. If so, you will have to repeat these steps several times.
- Issue 'clri /dev/mirror00 <inode> <inode> ... <inode>' where '<inode>' refers to the list of inodes found in the step immediately above. This 'zeros out' these inodes, destroying any link to the files (or directories!!!) that the inode used to contain. But, since the inode was duplicated anyway, those files or directories were already gone.
- Repeat the above two steps until 'fsck -n /dev/mirror00' appears to go through the inode phase without errors. You will start seeing references to 'missing files'; those come from the 'clri' instructions given previously.
- Now issue: 'fsck -s -y /dev/mirror00 | tee /tmp/err.log'. This will forever make the changes to the filesystem; but will also record all errors / missing filenames / cleared inodes into the log file '/tmp/err.log' (you may name this file anything you want).

## Error Conditions & Recovery

---

- Re-issue `'fsck -y /dev/mirror00'`. This will most probably report some minor adjustments (blocks recovered, etc.). DO NOT overwrite or re-direct output to the log file created above!
- Continue to re-issue `'fsck -y /dev/mirror00'` until all errors / notices stop appearing. At this point, the filesystem has been salvaged.
- Go through the log file (called `'/tmp/err.log'` above) and attempt to determine the names of the files or directories that had been lost. Recover these from some backup; or reconstruct by hand.

After all files have been restored or recreated, **BACKUP** the filesystem.

(Thanks for assistance with the above to Dave Button, MIS Director for Dollar RentACar, LAX)

## Mirror Status Meanings & Recovery

This section deals with the various mirror status readings that may be seen from time to time on a **SavWareHA** installation. Suggestions are given as to the nature of events that may have led up to the observed status readings; also, suggested recovery procedures are given for each.

### Primary Says 'Normal' / Standby Says 'Normal'

This is the correct, completely functional status of the **SavWareHA** system. Mirroring (if configured) is operational, and the **SavWareHA** monitor will be running on the Standby System.

### Primary Says 'Fallback' / Standby Says 'Fallback'

#### **Situation:**

- Primary System had gone down
- Users had to log back in to the Standby System

#### **Meaning:**

This is the correct status after a transition to fallback mode has been accomplished. The Standby System had not received an acknowledgment from the Primary System in the allotted time-out period. The Standby then cleaned the mirrored filesystem(s) with `'fsck'`, mounted the filesystem(s), started any necessary daemons, and enabled user logins.

#### **Possible solutions:**

Ascertain why the Primary System was unable to respond to the Standby System's acknowledgment requests from the **SavWareHA** monitor. Usually some major problem (*panic*, use of `'haltsys'` command, or unknown lockup) had crashed the Primary System.

Correct any system problems on the Primary System, run diagnostics, perform any needed maintenance. Follow these steps to recover:

- Connect the Primary System to power and to both the Data and I/O Links
- Boot the Primary System into multi-user mode; **SavWareHA** will see that the Standby System is in fallback mode, and will adjust the Primary System accordingly
- Verify that the network connections are functioning between the Primary and Standby Systems
- Enter **SavWareHA** on the Primary System
- Notify all users that a planned recovery to normal operation is about to occur
- When the system is quiescent, select the *'Utilities'* sub-menu from the **SavWareHA** main screen on the Primary System
- Highlight the selection ***'Recover from Fallback to Normal Mode'***; answer **'y'** to start the recovery process
- Since no *'fsck'* will have to be done, the switch over time is limited to the raw time **SavWareHA** takes to disable daemons and terminals on the Standby System, and to enable the same on the Primary System
- A regeneration is started immediately from the Standby to the Primary System; users may log into the Primary System at any time and continue work.

### Primary Says 'Normal' / Standby Says 'Fallback'

#### **Situation:**

- Some systemic error occurred to force the Standby System into Fallback mode
- The Primary System actually did not fail, the master side of the mirror is still mounted
- The mirror status probably says *'ACTIVE'* / *'ERROR'*; or perhaps *'ACTIVE'* / *'OFF'*
- The Standby System has mounted the slave side of the mirror
- Users are finding that they get multiple logins from both systems, or that they log off one system and then get a login from the other system.

#### **Meaning:**

In this situation, one would hope that not too much untraceable work has been performed on the system(s). What has happened is that some bizarre sequence of events has been either initiated by a user with *'root'* privilege (such as directly modifying the *'/etc/sentinel.d/mode'* file); or all links failed between the two systems, but no real hardware failure had occurred on the Primary System.

This allowed the Standby System to initiate a transition to fallback mode, thereby issuing a *'fsck'* and *'mount'* of the *'/dev/mirror00'* device; then the Standby System started all user daemons and processes and logins (which results in the multiple login problem on networked or interconnected serial devices). At the same time, the Primary System was still in complete operation, but had turned *'OFF'* reads and writes to the Standby System.

Whatever the case, the reader must now attempt to determine which side of the mirror is the more *'correct'* as far as the application is concerned.

## Error Conditions & Recovery

---

### Possible solutions:

Since modifications to the application's data may have taken place asynchronously on both sides of the mirror, the first step is to stop all further use of the **SavWareHA** system. If for some reason this is not possible, the suggestion is (generally) to halt processing on the Standby System, allowing all users to either continue on the Primary System, or to log into the Primary System. On the Primary System, turn the slave side 'OFF' by the *'Fix Broken Mirror: Turn Off Standby Element'*.

Once the **SavWareHA** system is in a known state (either completely quiescent, or only the Primary System accepting data and users), then the reader must determine whether or not to attempt to 'combine' data from the Standby System with data on the Primary System. Unfortunately, in this situation, there is no step-by-step process to be followed to secure this 'combination' of data.

When the data combination is completed on the Primary System, log in to the Standby System. Make certain that the mirrored unit is unmounted; then manually edit the file *'/etc/sentinel.d/mode'*; make this file say 'normal'. Shutdown and reboot the Standby System.

After the Standby System has rebooted into multi-user mode, enter the **SavWareHA** menu, select *'Fix Broken Mirror'*. Turn off the slave with *'Turn Off Standby Element'*, and then start regeneration with *'Regenerate Mirror'*.

### Primary Says 'Fallback' / Standby Says 'Normal'

#### Situation:

- The Primary System's mode file was manually edited and set to 'fallback'
- The Standby System was booted normally.

#### Meaning:

The mirror is probably not mounted on either side. This situation is most probably the result of manual intervention or editing of the *'/etc/sentinel.d/mode'* file.

No known real-life examples of this combination have ever been encountered.

#### Possible solutions:

Modify the *'/etc/sentinel.d/mode'* file on the Primary System, and reboot the Primary System. This will then start mirroring correctly; a regeneration may possibly have to be done to synchronize the two sides of the mirror.

As a last resort, remove and reinstall the **SavWareHA** software system on both machines.

## Monitor Status Errors

This section will show all possible monitor (*smon*) status error messages, an explanation of the error message, and suggestions as to the cause and repair of the error. Please note that some of the errors will report process ID numbers and other system status-dependent numbers and values; these are represented with an *italic bold font*.

### Host access errors

- Error (bad file number) finding host (**15**)
- Error (connection refused) on sendto
- Error (Interrupted system call)
- Another smon terminated abnormally (PID=**34**), clearing lock

These errors will result when the monitor cannot locate or communicate across the network. This may be a result of a network disconnection or media error.

### NetDisk errors

- Send time-out, x\_tb(**22448**), a\_tb(**22030**)
- Network eof - closing connection

These errors will result when a network disconnection or media error has occurred.

### Communications errors

- Error (connection refused) during recvfrom UDP
- Link **0 PrimarySystemName** Failed recvfrom UDP error - Interrupted system call
- All links between primary and standby systems have failed
- Error (No such file or directory) opening UDP server socket (cnt=**0**)
- Network read error (connection reset by peer) (will close connection)
- Network read error (interrupted system call) (will close connection)
- Serial Link 1 Read Failed (**-1** bytes) (Chan **5**), Error - Interrupted system call
- Serial Link **1** Data Mismatch

These errors reflect partial or no success in an expected response from one system to another. This may be the result of either a network disruption, someone *'kill'*ed the *'smon'* process, or a failure of one of the systems.

## Error Conditions & Recovery

---

### Configuration errors

- Error on open of drive info file /etc/default/netdisk
- Error opening /dev/rndsk00: No such file or directory

These errors happen when some misconfiguration of **SavWareHA** has been noticed by the monitor.

### Fallback notifications

- Automatically switching to Fallback mode!
- Server is not running

These are the initial messages signaling a transition to fallback mode.

# 5. Manual Pages

**SavWareHA** automatically inserts the following manual pages into the 'man' system. Enter '*man <command>*' to see these entries on-line.

## sentinel

/etc/sentinel

The sentinel program provides a menu or command line driven administration facility for **SavWareHA**.

### Options

-d	Download config information to mirror driver
-b	Build a Mirror
-r	Re-Build a Mirror
-m #	Specify the mirror number for other commands
-f opt	Specify filename of mirror config info
-k	Stop DMX Mirror Daemon
-p	Start DMX Mirror Daemon
-F	Fix a mirror if easily fixable
-G	Fix a mirror (Requires -m)
-S opt	Show the size of a device
-a	Add a "Master Only" Mirror
-R	Remove a "Master Only" Mirror
-s0	Turn off mirror Master Element (Requires -m)
-s1	Turn off mirror Slave Element (Requires -m)
-Q	Query status of mirror (optional -m)
-M1	Mode change to fallback
-M2	Mode change to normal
-u dev	Check if a device is currently in use
-K	Create kernel space files
-z1	Special option to allow reverse regeneration
-v	verbose

## sentinel.sh

/etc/sentinel.sh

Provide shell services for smon and **SavWareHA** programs

### Options

-r	Report role of current machine
-m	Report current <b>SavWareHA</b> mode
-s mode	Force a switch to a different mode
-Smode	Start a switch on this machine only (used by -s)
-l	Display log file

## Manual Pages

---

-M	Mount <b>SavWareHA</b> filesystems
-U	Unmount <b>SavWareHA</b> filesystems
-i	Execute bootup code
-b	Mirror configuration support
-d	Mirror unconfiguration support
-L dev	Make a lost+found directory
-I	Copy basic config files to standby
-N	Process data link failure notification
-F	Process monitor link failure notification

### **smon**

/etc/smon

Monitor the system of the Primary System and cause an automatic switch to fallback mode on a failure:

#### Options

-s	Run as a server
-D	Send a message to the Standby System
-P	Send a message to the Primary System
-d	Display monitor information and status
-k	Kill any smon processes that are running
-t	Start serial servers
-n srv	Change service name to srv (default is smon)
-x	Run in verbose mode
-a	Change global time-out value
-F	Do not fork on startup

### **netdsk**

/etc/netdsk

Provide disk driver access over a network.

#### Options

-k	Stop running daemons
-dp	Operate as a client
-F	Do not fork on startup
-m num	Set ndsk number
-s service	Change default service name
-B	Set TCP receive buffer size
-T	Set TCP transmit buffer size
-M sys:dev	Check if this is a valid device
-o file	Write output to a file
-n sys:dev	Give ndsk?? Name
-N ndsk??	Give sys:dev Name
-v	Verbose

-W size            Set lli window size  
-D /dev/ndsk??    Check mode of link

## sntlmail

/etc/sentinel.d/sntlmail [-ic]

When invoked without any options, checks the status of all defined mirrors and mails messages to a list of administrators if all mirrors are not completely ACTIVE or GOOD.

Administrators are listed in the file /etc/sentinel.d/administrators

Administrators can be at a remote location if the E-mail facility is configured for remote mail.

Messages will be mailed once when a problem occurs. If the problem is not resolved within 4 days, 2 repeated messages will be mailed per day until the problem is resolved or the message is cleared.

### Options

- i Initialize the sntlmail program by creating an entry in crontab if one does not already exist and creating the administrators file if necessary.
- c Clear any pending message



## ***6. Extended Monitor***

The SavWareHA Extended Monitor is an add-on package for SavWareHA that provides additional monitoring capabilities.

### ***Application Process Monitoring***

The extended monitor provides the capability to define a list of process that should be monitored. If any of these processes fail during the normal operation of the system, an automatic switch to fallback mode can be initiated.

### ***Network Monitoring***

The extended monitor also provides the capability to monitor the user network. A system fallback can be initiated in the case of a network failure.

### ***Administration***

An administration program is provided that will run as a graphical X-Windows session or on a character mode device. The administration program can be used to set all controllable monitoring parameters and will also show a log of the monitor activity.

## **Installation**

### ***Primary System Only***

The Extended Monitor only needs to be installed on the Primary system.

### ***Installation Instructions***

The Extended Monitor program should be installed while the system is in multi-user mode.

Login to the root account

Insert the "SavWareHA Extended Monitor" distribution diskette into the floppy diskette drive

Select one of the following installation procedures:

## Extended Monitor

---

### SCO Unix 3.2v4 and SCO ODT3

Run the SCO 'custom' utility and select the following options:

**custom->Install->A New Product->Entire Product**

If the SCO Visual Tool Control Language has not previously been installed on the system it will now be automatically installed.

Insert SCO Visual Tcl 1.0 for Open Server 3.0 Vol 1 and press **ENTER**

Insert SCO Visual Tcl 1.0 for Open Server 3.0 Vol 2 and press **ENTER**

Insert SCO Visual Tcl 1.0 for Open Server 3.0 Vol 3 and press **ENTER**

The Release Notes for SCO VTCL will be displayed on the screen. Press the **SPACE BAR** at the "—More—(xx%)" prompts until the installation continues.

### SCO OpenServer5

Run the SCO 'custom' utility and select the following options:

**custom->Software->Install New->From system->Floppy Disk Drive 0->Full**

### Installation Continued

The Extended Monitor Administration program will run after all of the files have been transferred from the diskette. Please refer to the detailed explanation of this program later in this manual for more information. Make any necessary changes to the configuration options and exit the program by pressing the **ESC** key. This program can be re-run at any time by running `emon.adm` from the root account.

## Configuration Parameters

### Application Monitoring

This parameter determines whether the Extended Monitor will monitor application processes. If the option is enabled, the Extended Monitor will monitor application processes. If the option is disabled the Extended Monitor will not monitor application processes.

The processes to monitor must be listed in the "**Processes To Monitor**" table at the bottom of the screen.

### Network Monitoring

This parameter determines whether the Extended Monitor will monitor the user network. If the option is enabled, the Extended Monitor will monitor the user network. If the option is disabled, the Extended Monitor will not monitor the user network.

The interface name for the user network must be specified using the “Network To Monitor” parameter.

### Network To Monitor

This is the interface name for the user network. If Network Monitoring is enabled, the Extended Monitor will monitor this network interface.

### Automatic Fallback

This parameter determines whether the Extended Monitor will initiate an automatic fallback. If this option is enabled, the Extended Monitor will be allowed to initiate an automatic fallback if it recognizes a failure condition. If this option is disabled, the Extended Monitor will not initiate an automatic fallback but will record all events in the system log.

### Interval (seconds)

The parameter is the number of seconds between each check by the Extended Monitor. A high number will lower system overhead but will take longer to recognize a failure. A low number will cause more system overhead but will recognize failures more quickly. The number cannot be set below 2.

### Processes To Monitor

This parameter is a list of the processes to monitor. The process name of each process that is to be monitored should be included in this list. The name of the process in this list should match the value that is displayed when a `ps -e` command is issued. No spaces are permitted in a process name.

## Administration Menu

### X Windows Navigation

The X Windows SavWareHA Extended Monitor Administration menu can be started by clicking on the following icons **root->etc->emon.adm** or selecting **File->Open** from the main desktop and then entering the name `emon.adm` and then click on “OK”.

The “**Application Monitoring**”, “**Networking Monitoring**” and “**Automatic Fallback**” options can be selected or deselected by clicking on the toggle buttons. If the button is solid, the option is enabled. If the button is hollow (light colored), the option is disabled.

The “**Network To Monitor**” can be selected by clicking the button. A list will then be displayed of all network interfaces configured into the system. Use the mouse to click on the interface of the network that should be monitored. The selected interface will then display in the “**Network To Monitor**” button on the main administration screen.

The **Interval** can be changed by using the mouse to position the text cursor into the field. Use the **INS**, **DEL** and number keys to change the value.

## Extended Monitor

---

The **Processes To Monitor** list shows all of the processes that are currently being monitored. A new process can be added to the list by using the mouse to click on the **“Add”** button. A new window will appear that will allow the process name to be entered. The process name should be entered as it would appear in a `ps -e` command. Use the mouse to press the **“OK To Add”** button after the process name has been typed in. A process can be deleted by using the mouse to click on the process to be deleted and then click on the **“Delete”** button.

The **“Log”** button can be used to show the events that have been recorded by the Extended Monitor since the last time it was started.

The **“Close”** button can be used to exit the administration program.

## Character Mode Navigation

The character mode administration menu can be started by running the command **emon.adm** from the root account.

The **“Application Monitoring”**, **“Networking Monitoring”** and **“Automatic Fallback”** options can be selected or deselected by using the **TAB** key to highlight the desired item and pressing the **SPACE BAR** to toggle the button. If an **“\*”** is displayed between the brackets **“[ ]”**, the option is enabled. If no **“\*”** is displayed between the brackets **“[ ]”**, the option is disabled.

The **“Network To Monitor”** can be selected by using the **TAB** key to highlight the option and then pressing the **SPACE BAR**. A list will then be displayed of all network interfaces configured into the system. Use the **ARROW KEYS** to highlight the desired interface of the network that should be monitored and then press the **SPACE BAR** to select the interface. The selected interface will then display in the **“Network To Monitor”** option on the main administration screen.

The Interval can be changed by using the **TAB** key to highlight the option. Use the **INS**, **DEL** and number keys to change the value.

The **Processes To Monitor** list shows all of the processes that are currently being monitored. A new process can be added to the list by using the **TAB** key to move the highlighted selection to one of the bottom buttons and the **ARROW** keys to highlight the **“Add”** button and then pressing the **SPACE BAR**. A new window will appear that will allow the process name to be entered. The process name should be entered as it would appear in a `ps -e` command. After the process name is entered, use the **TAB** key to move to the **“OK To Add”** button and press the **SPACE BAR**. A process can be deleted by using the **TAB** and **ARROW** keys to highlight the process to be deleted and press the **SPACE BAR**. Then use the **TAB** and **ARROW** keys to highlight the **“Delete”** button and press the **SPACE BAR**.

The **“Log”** button can be used to show the events that have been recorded by the Extended Monitor since the last time it was started.

The **“Close”** button can be used to exit the administration program.

## Actions

The Extended Monitor program will automatically start on the primary system during the boot to multi-user mode and will continue to run while the system is in multi-user mode. Only one copy of the Extended Monitor program will start at a time. The Extended Monitor will delay for the Interval (seconds) and then perform the following monitoring.

### Application Monitoring

If the “**Application Monitoring**” is enabled, the Extended Monitor will monitor the process list and insure that each of the “**Processes To Monitor**” is running. No action will occur until a processes is noticed by the Extended Monitor program and then subsequently disappears from the process list. When the Extended Monitor detects a “Process Failure” it will perform the fallback checks and initiate a fallback if appropriate.

### Network Monitoring

If the “**Network Monitoring**” is enabled, the Extended Monitor will monitor the selected network interface. When the Extended Monitor detects a “Network Failure” it will perform the fallback checks and initiate a fallback if appropriate.

### Fallback

When a fallback is requested by the Extended Monitor program, the following actions will occur. A notification will be sent to the system console. If “**Automatic Fallback**” is enabled, the Extended Monitor will then check to make sure that the mirror is properly synchronized on the standby system. If it is not, a message will be generated and normal operation will continue on the primary system. If the mirror is properly synchronized, the Extended Monitor will initiate a switch to fallback mode.

### LOGS

A log of all monitoring activity is recorded in /tmp/emon.status and can be displayed by the administration program. In addition, any information regarding failures or an attempted switch-over is also recorded in the /usr/adm/dmx.log file, in the /usr/adm/messages file and is displayed on the system console.



# 7. Application Notes

This section includes notes about certain 'applications' that have been encountered in the **SavWareHA** installations to date. 'Applications' refers to both hardware and software devices and programs. We invite the reader to e-mail your own unique 'Application Notes' to *cpgsupport@avnet.com* if they might assist others in implementing **SavWareHA**.

## SavWareHA Configuration Worksheet

In planning a SavWareHA installation, we have found that the following form has helped to identify major components, configuration needs, and serves as a hard copy document of the installation. It is placed on a sheet by itself for your use.

### SavWareHA Configuration Worksheet

Customer:

System Serial #:

<u>Primary System</u>	<u>Standby System</u>
Node Name:	Node Name:
IP Addr:	IP Addr:
/.rhosts ok?	/.rhosts ok?
/etc/hosts ok?	/etc/hosts ok?
'ping' ok to S/S?	'ping' ok to P/S?
'rlogin' ok to S/S?	'rlogin' ok to P/S?
'rcmd' ok to S/S?	'rcmd' ok to P/S?

<u>Mirror #0:</u>		<u>Mirror #1:</u>	
Primary Device:	Standby Device:	Primary Device:	Standby Device:
Mounts as:		Mounts as:	

### SCO LLI Drivers

Recommended for SavWareHA Release 2.02e and later.

**SavWareHA** can be configured to use the SCO LLI Drivers to improve the performance of the data link, eliminate the long TCP/IP timeout in the case of a data link failure and reduce the CPU overhead required for TCP/IP transfers. Use the following procedure:

Issue the command *'llistat'* from the *'root'* account on both the Primary and Standby Systems and note the device name of the LLI network device. Using a text editor, modify the file *'etc/default/netdisk'*. Add the LLI device names found previously of the data link network card on the Primary and Standby System to the device line.

For example, change:

```
standby:/dev/u - 0 0 0
```

To:

```
standby:/dev/u - 0 0 0 /dev/wdn0 /dev/wdn0
```

NOTE: Use [TAB] between each field.

Shutdown and reboot the Primary System for the change to take effect. This change has no meaning or effect on the Standby System.

If you notice performance degradation and possibly many *'<defunct>'* processes, then the passing of UDP's over the LLI link has become unreliable. Revert to the TCP/IP method, or investigate you cabling / network cards closely.

### SavWareHA Data over a Second Network

When a second network is defined to SCO, a separate Internet address and machine name are used to identify the network.

It is recommended that the network card for the data link on the primary system be named "neptune" and the network card for the data link on the standby system be named "pluto". This must be done with netconfig before SavWareHA is configured.

Install **SavWareHA** as usual: Specify "neptune" and "pluto" as the system names.

This will route data over the data link network card when **SavWareHA** starts up.

## On Line Backups with 'XA'

Occasionally, a client requires 100% uptime, but also wants checksum verified backups each night. This is a common, but often overlooked, need for **SavWareHA** customers. It also points out another benefit of having **SavWareHA** installed on a mission-critical system.

The **SavWareHA** package is NOT a substitute for frequent, checksum-verified data backups to removable media, such as those provided by 'XA' from Avnet Computer Marketing. Please secure a copy of 'XA' if you do not already have one for all of your sites, **SavWareHA** or not. This discussion will center around the scripts needed to implement 'XA'.

The idea here is to form a script command file that will be called from 'cron' on the Primary System. The script file will issue commands to both the Primary and Standby Systems to:

1. Turn off the slave side of the mirror;
2. Clean and mount (read-only) the slave side of the mirror;
3. Fire off 'XA' from the command line;
4. Unmount the slave side and start regeneration.

The following script segment makes several assumptions, as listed below. Please keep in mind that any of the assumptions may be subject to change on particular systems:

- Mirrored Unit is mirror #0, the first mirror
- The Standby System is named 'standby'
- Slave sub-component is named '/dev/u'
- Slave sub-component can be mounted on directory '/u'
- A backup of the entire Standby System is desired
- Output redirection would be specified by the calling routine

Given the above assumptions, the script might look similar to the following:

```
# /usr/local/bin/backup.slv - Backup Slave Mirror
# Turn off slave side, mirror #0:
/etc/mirror -m0 -s1
# Cause slave sub-component to be cleaned:
rcmd standby /etc/fsck -y -s /dev/u
# Cause slave sub-component to be mounted, read only:
rcmd standby /etc/mount -r /dev/u /u
# Cause 'XA' to fire off and backup entire Standby System:
rcmd standby /usr/xa/bin/start.backupx
# Now that 'XA' has finished (any result), unmount slave:
rcmd standby /etc/umount /dev/u
# And bring mirror back into compliance (note CAPITAL 'G'):
/etc/mirror -m0 -G
```

## MRACDB Root DataBase Consistency

Available on SavWareHA Release 2.02f and later.

The program `/etc/mracbd` is provided in the directory `/etc/sentinel.d` to allow important configuration information on the root filesystem of the Primary system to be periodically copied to the Standby system. This allows the system administrator to make changes to the Primary system, such as adding a new user, and have those changes made automatically on the Standby system.

The `mracdb` facility is not configured during a normal installation. The installer must copy the file `mracdb` from `/etc/sentinel.d` to `/etc` and then run the program `/etc/mracdb`. This first time this program is executed it will add the proper entry to `/etc/crontab` for the periodic scheduling. It will also create the file `/etc/sentinel.d/cdb` which contains a list of files that will be copied from the Primary system to the Standby system. The default file will contain the basic entries required to maintain users, passwords and permissions. Additional entries can be made to the file to provide additional ease of maintenance.

The `mracdb` functionality can be disabled at any time by using the command `/etc/mracdb -r` from the root account.

## Switching Network Devices (IP Matching Technique)

Applies to SCO Unix Release 3.2v4.x, ODT Release 3 and OpenServer5.

Note: Please review the complete capabilities of this technique and of IP Aliasing to determine which technique will work better in your application. Typically IP Aliasing is a better technique.

Available on SavWareHA Release 2.02f and later.

The normal installation of SavWareHA allows for the switching of serial devices in the event of a fallback. This application note describes the configuration changes that can be made to allow for network devices, such as PCs or X-Terminals, to be automatically switched between systems in the event of a fallback.

### Requirements and Limitations

This application note only applies to SCO Unix Release 3.2v4.x and ODT Release 3.

There **must** be at least two network links between the Primary and Standby system. One dedicated Data Link for mirroring disk data and one dedicated IO Link for connecting user PCs and other network devices.

In addition, there **must** be a serial link between the Primary and Standby system. This will be used by the monitor program to pass control information between the systems. Note that the IO Link can not be used for the monitor program because it will be disabled on the system that is not "in control".

The IO Link will be disabled on the system that is not "in control". This means that in normal mode the IO Link on the Standby system is disabled. This will prevent any network functionality (i.e. rlogin, rcp, etc..) on that network interface. Note that network functionality will be available to the Standby system on the Data Link.

### Theory of Operation

The MAC (Hardware Address) of the IO Link network interface card (NIC) for both systems will be set to the same pre-defined value and the interfaces will be disabled during the boot to multi-user mode (/etc/tcp). Later during the boot sequence, the program /etc/sentinel.d/S01tcpa will determine which system is "in control" and enable that IO Link interface at a common IP address and with a common system name. The system that is not "in control" will remain connected to the IO Link network with its network disabled.

In the case of a fallback, the Primary system will disable the IO Link interface (if possible) and the Standby system will enable its IO Link interface at the same IP address and system name that the Primary had previously used. Any connection oriented sessions that were already established will be lost. New connections can be made from a PC or X-Terminal to the Standby with no configuration changes because the Standby now has the same address that the Primary used to have.

## Application Notes

---

### Installation

1. Create the file `/etc/sentinel.d/iolink` on both systems. It should be created as a text file that only contains the network interface name of the IO Link (i.e. `wdn0`).
2. Run the following program on both systems. This will modify `/etc/tcp`.

```
/etc/sentinel.d/tcpopts -install
```

The first hook into `/etc/tcp` sets the MAC address to a predefined value before any of the TCP daemons start so they will see the correct MAC address.

The second hook immediately disables the `iolink` before any service daemons are started.

3. Important Note: The interface for the Data Link must be listed before the interface for the IO Link in the file `/etc/strcf`. The installer may have to manually modify the file `/etc/strcf` to insure this sequencing.
4. The IP address and system name configured in `/etc/hosts` and in `/etc/tcp` for the IO Link must be set to the same value on both the Primary and Standby systems.
5. The systems names listed in `/etc/default/netdisk`, `/etc/sentinel.d/roles` and `./rhosts` on both systems must be modified to refer to the Data Link instead of the IO Link.
6. Reboot both the Primary and Standby systems.
7. Use the common IP address and system name when referring to the SavWareHA system from a network device.

## Switching Network Devices (IP Aliasing)

Applies to SavWareHA for OpenServer5 Release 2.04c and greater and SavWareHA for UnixWare 7 release 2.05x6 and greater.

IP Aliasing is a simpler form of network device switching. It uses a new (third) IP number that always floats to the system in control.

Configuration is done by creating the file `/etc/sentinel.d/ipalias` on both systems with a line that has the floater IP number and rebooting both systems. This number must be on one of the subnets configured on the system (See Network field in the `netstat -in` command).

SavWareHA will always leave each machine at its own IP address, and just add / delete the floater IP address to the system in control. This means that both machines are always accessible to the network via their fixed IP address. In addition, a user can access the IP Alias number and automatically connect to the system in control.

Example:

Result of `netstat -in` command:

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
Net0	1500	192.168.20	192.168.20.7	171	0	171	0	0

Use a text editor to add the following entry into `/etc/sentinel.d/ipalias` and reboot

```
192.168.20.207
```

You can then access the system as 192.168.20.7. If it is the system in control, you can also access it as 192.168.20.207.

## Application Notes

---

### Using an Equinox SST with SavWareHA

The Equinox SST can be used for serial ports that must be switched between two computers under the control of SavWareHA.

#### Requirements:

- SST host cards, one for each computer.
- Equinox SST Driver with redundancy (Rel 2 or higher).
- Equinox SST redundancy cable.
- External power supply for SST panel.
- SavWareHA revision 2.02f or greater.

There must be a power supply for the SST pods, as power is normally derived from the primary system through the original cable. Power can not come from the standby system through the redundancy cable, so it must be supplied from the external power supply.

After installing the SST driver on both systems, you must edit the file `/etc/conf/pack.d/eqnx/space.c`. The line of `int EQNft` must be changed from `'=0'` to `'=1'`. After the changes have been made, the kernel must be rebuilt and the system rebooted.

The file `/etc/sentinel.d/X65eqnx` must be copied to `/etc/sentinel.d/S65eqnx`.

The file `/etc/sentinel.d/S65eqnx` must be linked to `/etc/sentinel.d/K65eqnx`.

If more than one "SST Chain" is established between the systems then the file `/etc/sentinel.d/S65eqnx` must be modified to include additional lines that call `/etc/eqnx/sstft` with the `-s` and `-q` options to start and stop the interfaces. The parameter specified for `-s` and `-q` must be the TTY device name of the first TTY device on the SST chain.

The file `/etc/sentinel.d/terminals` contains a list of TTY device names that will be enabled during a switchover.

## S75inittab

S75inittab speeds the process of enabling and disabling terminals during a switchover. It should be used in place of `/etc/sentinel.d/terminals` when more than three terminals need to be switched.

**`/etc/sentinel.d/S75inittab` will cause the file `/etc/inittab` to be recreated from the kernel environment each time the system is booted.**

**`/etc/sentinel.d/S75inittab` is enabled by default on SavWareHA release 2.03a.**

The device driver names for all serial multiplexor devices that are to be switched must be manually entered into the "SHARED\_DEV" line in `/etc/sentinel.d/S75inittab` and a line must be added to `/etc/sentinel.d/cdb` for the switchover to work properly. (mracdb must have been previously configured).

The device driver name is the name of the file in the `/etc/conf/init.d` directory that contains the inittab entries for the devices.

A new inittab file is created for each system during the system boot or switchover. The new inittab that is created will have all of the "shared devices" turned off on the system that is not in control.

This functionality can be disabled by renaming `/etc/sentinel.d/S75inittab` to `/etc/sentinel.d/X75inittab` and removing `/etc/sentinel.d/K32inittab` on both systems.

### SavWareHA Release Notes 2.03a

A problem has been corrected that would cause a SavWareHA system that was configured to operate in LLI mode to revert back to TCP/IP mode in certain circumstances. This could have caused the system to hang during the long TCP/IP time-out in the event of a standby or data link failure.

A positive indication is now available to allow the system administrator to determine if the data link is operating in TCP/IP mode or LLI mode. This is displayed in the SavWareHA mirror menu option "Display Mirror Status". The indication is only displayed when the data link is in use (Both sides of the mirror are active).

The file **/etc/inittab is now regenerated during each system boot** from the kernel environment files in /etc/conf/init.d and /etc/conf/cf.d/init.base. Any changes to /etc/inittab that are not also reflected in the corresponding kernel environment files will be lost during the system boot. Previously, the file /etc/inittab was only regenerated when a new kernel was linked. This functionality is accomplished by /etc/sentinel.d/S75inittab and can be disabled by renaming the file to /etc/sentinel.d/X75inittab or removing the file.

Monitor Link Failure Notifications are now displayed on the console and a message is mailed to the system administrator(s) in addition to the previous action which was to record the event the log file

Data Link Failure Notifications are now displayed on the console and a message is mailed to the system administrator(s) in addition to the previous action which was to record the event in the log file.

A problem has been corrected that would have allowed a regeneration to a drive that had reads enabled. Regeneration is now prohibited to a drive with reads enabled.

A problem has been corrected that could have caused the system to force reads to a drive with an incomplete filesystem under certain failure conditions. Error and Warning messages were generated to the console in the previous versions. In this new release, a PANIC is generated and the primary system is halted under these circumstances to prevent possible data corruption.

The periodic monitor program has now been modified to check the system status every 2 hours. In addition, inconsistency notices are displayed on the system console as well as mail messages to the system administrator(s).

The system time and date on the standby system are now synchronized to the primary system.

A utility program has been added to the system that will accumulate all of the support logs and statistics from both systems into a single file that can be transmitted to customer support for review. The program is /etc/sentinel.d/makelog.

## SavWareHA Release Notes 2.04b

Because of a change in the default service number of netdisk and smon, Upgrades to 2.04 from any/all lesser version must be done to both the primary and standby systems at the same time.

1. SavWareHA Release 2.04b dramatically improves the netdisk throughput running in TCP/IP mode on SCO OpenServer 5.0.x. Now running either TCP/IP or LLI mode we have seen from 500KB to 2.5MB / second throughput across a 100BaseT network. This is the throughput used for disk mirroring. LLI mode is preferred as it does not incur the CPU overhead that TCP/IP mode does. However, if the boxes are geographically separated and the data (mirror) link must be routed then use TCP/IP mode as LLI packets are not routeable.
2. Single install Diskette has both SCO 4.x and OpenServer 5 versions as well as Extended Monitor
3. Auto Recognize / Operation of Shared Disk Array (SDA)
4. Role Reversal of Sentinel System (Primary <=> Standby).

### NOTICE: Important Information

Because of the ability to now do Role Reversal, it is best to not name the machines 'primary' and 'standby' or else you will confuse operators when the system named 'standby' becomes the primary system. Therefore, the new default names for the two systems is 'Pluto' and 'Neptune' (two planets that because of their elliptical orbit, alternate having the distinction of being the furthest planet from the sun).

5. Host equivalency not needed for normal operations or switch-over, only for (re-)configuration, or if using 'mracdb'.
6. Evaluation Period of Hot Root Mirror now works instead of displaying "Registration Time Expired".
7. Configure of Sentinel Mirrors automatically setup for LLI mode of operation for the NetDisk Data Link.
8. Configure of Sentinel Mirrors automatically setup to permit Role Reversal.
9. Command 'sentinel.sh -setslave' to setup to bring fresh Operating System on-line as a replacement standby system. Drive Partitions need to be pre-allocated but can't have the filesystems configured.
10. Command 'sentinel.sh -setlli' to setup for LLI mode of operation.
11. Command 'sentinel.sh -setrev'. Must be run in advance on the primary system to permit Role Reversal.

## **SavWareHA Release Notes 2.04c**

1. Add IP Alias method for OpenServer 5
2. Correct problem with a false detection of a Shared Disk Array.
3. Correct problem with smon on serial ports after a role reversal.
4. Correct problem with smon not disarming if loss of network link.

## **SavWareHA Release 2.05x6 for UnixWare7**

Differences between SavWareHA for OpenServer5 and UnixWare7

1. Install procedure uses 'pkgadd' instead of 'custom'.
2. IP Aliasing is used instead of IP Matching (See separate application note)
3. Device names for the mirrors have 3 digits instead of 2 (i.e /dev/mirror000 instead of /dev/mirror00).
4. MRACDB functionality is not currently supported.
5. Sub-division of a mirror into partitions is not currently supported.
6. EMON functionality is not currently supported
7. Use of F3 key to list available disk devices during install is not currently supported
8. The Equinox SST switchover functionality is not currently supported (Equinox Issue).





/

**/etc/.rhosts**, 15  
**/etc/conf/init.d/\***, 14  
**/etc/default/netdisk**, 76  
**/etc/hosts**, 15  
**/etc/inittab**, 14  
**/etc/mracbd**, 78  
**/etc/rc2.d/K95sentinel**, 23  
**/etc/rc2.d/S95sentinel**, 23  
**/etc/sentinel**, 65  
**/etc/sentinel**, 19  
**/etc/sentinel.**, 23  
**/etc/sentinel.d/administrators** (Text File), 20  
**/etc/sentinel.d/dkinfo** (Program), 21  
**/etc/sentinel.d/dkuse** (Program), 21  
**/etc/sentinel.d/filesys** (Text File), 21  
**/etc/sentinel.d/iolink**, 80  
**/etc/sentinel.d/links**, 18  
**/etc/sentinel.d/mode** (Text File), 22  
**/etc/sentinel.d/roles** (Text File), 22  
**/etc/sentinel.d/S06pci**, 76  
**/etc/sentinel.d/sentinel** (Program), 23  
**/etc/sentinel.d/sentinel.sh** (Shell Program), 23  
**/etc/sentinel.d/sentinel.strt** (Shell Program), 23  
**/etc/sentinel.d/sntlmail**, 67  
**/etc/sentinel.d/terminals**, 23, 29  
**/etc/sentinel.sh**, 65  
**/etc/sentinel.sh.**, 23  
**/etc/strcf**, 80  
**/etc/tcp**, 79  
**/usr/adm/dmx.log** (Text File), 22

## A

ACTIVE, 31, 34, 36, 61, 67  
adding a new user, 78  
administration, 31  
administrator, 18, 19, 20, 31  
administrator, 29  
administrators, 67  
application "start" procedures, 29

## B

background, 30  
*boot*, 14, 28, 45  
*Boot kernel by default*, 14  
*Build a mirror*, 17

## C

cancel, 16  
COM1, 19  
COM2, 19

**Configuring Sentinel**, 15  
**Configuring the Sentinel Monitor**, 18  
connection refused, 63  
CONTROL-E, 17  
**custom**, 13, 14, 41  
custom, 13  
Customer Support, 11, 53

## D

data link, 8, 15, 17, 23, 24, 30, 37, 76  
**Data link**, 9  
demo, 14  
demonstration, 14  
device name, 17  
disabled, 23  
disk partitions, 10  
Disk Space!, 10  
*Display Monitor Status*, 19  
*divvy*, 9, 59  
**DMX**, 7, 8, 65  
drivers, 7

## E

e-mail, 31  
enabled, 23  
END, 16, 17, 32, 35, 40, 41  
Equinox ELS-16 Terminal Server, 76  
**ERROR**, 34, 36, 50, 57, 58, 61  
errors, 22  
ESC, 17  
evaluation, 13  
exit, 16

## F

F3, 17  
F3, 16, 17, 21, 32  
failed a, 30  
failure, 8, 9, 10, 18, 19, 27, 29, 50, 61, 66  
fallback, 29  
**Fallback**, 29  
fallback mode, 30  
**fallback mode**, 8, 17, 18, 23, 27, 28, 29, 30, 31, 42, 49, 50, 57, 60, 61, 64, 66  
Fallback mode, 9  
fault-tolerant, 15  
**File**, 20  
filesystem, 16  
filesystem, 8, 9, 10, 14, 16, 17, 21, 22, 27, 29, 31, 38, 39, 41, 46, 47, 48, 54, 56, 57, 58, 59, 60  
filesystem(s), 17  
filesystems, 21, 29  
*Fix Broken Mirror*, 28  
*fsck*, 8, 27, 29, 31, 34, 47, 49, 56, 57, 58, 59, 60, 61

**G**

GOOD, 34, 36, 67

**H**

host, 10, 15, 63  
Host Equivalence, 15  
hot swap, 28, 37, 48, 49, 50

**I**

**I/O link**, 9  
Installing, 13  
Installing Sentinel, 13  
integrity, 29  
Internet, 15  
Internet addresses, 15  
interrupted system call, 63  
**IO link**, 8, 31  
IP address, 79

**K**

kernel, 13, 14, 54, 57

**L**

link, 8  
link kit, 14  
link kit, 14  
links, 18, 19, 29, 45, 46, 61, 63  
links file., 19  
**LLI**, 10, 18, 23, 24, 76  
log, 13, 16, 22, 28, 30, 31, 35, 40, 41, 42,  
49, 54, 59, 60, 61, 62, 65  
log file, 22  
log file, 22, 31, 59, 60, 65

**M**

machine name, 15, 18, 24, 76  
mail, 21  
mail, 11, 20, 21, 31, 67, 75  
maintenance mode, 13  
*Make a Unix Filesystem*, 17  
Man Pages, 65  
manual, 7, 8, 13, 15, 18, 19, 21, 29, 34, 50,  
54, 55, 56, 57, 62, 65  
**Memory**., 10  
Menu, 19  
messages, 21, 31, 54, 63, 64, 67  
mirror, 8, 9, 10, 13, 15, 17, 18, 21, 24, 25,  
27, 28, 30, 31, 32, 33, 34, 35, 36, 37, 38,  
39, 40, 41, 42, 43, 45, 46, 47, 48, 49, 50,  
51, 52, 54, 55, 56, 57, 58, 59, 60, 61, 62,  
65  
Mirroring, 9

**MISMATCH**, 34, 51

**mkdev**, 16, 17  
mkdev fs, 16  
*mkfs*, 17, 38, 39  
*mmdf*, 21, 31  
monitor, 79  
mount, 17, 21, 22, 31, 39, 41, 45, 54, 56, 57,  
58, 61  
mount point, 17, 54  
mountflags, 21, 22  
multiuser mode, 27

**N**

*ndsk*, 24  
**net machinename**, 18  
netdisk, 23, 39, 64, 76  
network, 8, 9, 10, 18, 23, 24, 31, 34, 36, 45,  
48, 49, 52, 61, 63, 76  
network, 8  
Network Cards\., 10  
network devices, 79  
**Normal**, 29  
NORMAL, 31  
**normal mode**, 8, 27, 28, 29, 30, 31, 40, 43,  
49  
**Normal mode**, 9  
**NOT CLOSED**, 34  
**NOT SETUP**, 34  
notification, 19, 45

**O**

*OFF*, 27, 28, 36, 37, 45, 56, 61, 62  
operating system, 7, 10, 31, 46

**P**

panic, 57, 59, 60  
password, 15  
passwords, 78  
PCs, 79  
permissions, 78  
Placing the Primary System Back On-Line,  
28  
Placing the Standby System Back On-Line,  
28  
planned downtime, 28, 30  
Primary / Standby Designation, 15  
**primary system**, 8, 15  
**Primary system**, 9  
**Primary System**, 8, 9, 13, 15, 16, 17, 18,  
19, 21, 22, 23, 24, 27, 28, 29, 30, 31, 32,  
33, 36, 37, 38, 39, 41, 42, 43, 45, 46, 47,  
48, 49, 50, 51, 53, 54, 56, 57, 60, 61, 62,  
66, 76  
Primary System Bootup, 27  
primary system name, 16

## R

Read, 20  
*Rebuild the kernel*, 14  
Recover from fallback to normal", 30  
Recover to normal mode, 29  
Recover to Normal Mode, 30  
recvfrom, 63  
REGEN, 31, 34, 50, 52  
regeneration, 8, 10, 15, 17, 28, 30, 31, 32,  
34, 35, 36, 38, 39, 41, 42, 47, 48, 49, 50,  
52, 54, 55, 56, 58, 61, 62  
regeneration., 10  
**reggie**, 13, 19, 20, 42  
registration, 13, 14, 19, 20, 42  
registration code, 20  
registration system, 19, 20  
repaired, 29  
Requirements, 10  
*rlogin*, 16, 50  
root, 78  
root filesystem, 10

## S

S01tcpa, 79  
S06pci, 76  
**SCO LLI**, 10  
SCO LLI Drivers, 76  
**SCO TCP/IP**, 10  
SCSI, 8, 9, 21, 59  
**SCSI link**, 8  
second network, 76  
sendto, 63  
sentinel, 65  
Sentinel Data over a Second Network, 76  
Sentinel File and Program Descriptions, 20  
Sentinel Mode Discrepancy, 29  
**Sentinel Monitor**, 7, 8, 18  
sentinel.sh, 65  
serial, 23  
serial link, 9, 18, 19  
**Serial link**, 9  
serial number, 14  
**serial standbydev=primarydev**, 18  
**shutdown**, 14, 23, 27, 28, 34, 45, 47  
smon, 8, 18, 19, 40, 45, 63, 65, 66  
smon, 18  
sntlmail, 67  
socket, 63  
Software Registration, 19, 20, 42  
**Standby system**, 9

**Standby System**, 8, 9, 10, 11, 13, 14, 15,  
16, 17, 18, 19, 20, 21, 22, 23, 24, 27, 28,  
29, 30, 31, 32, 33, 35, 36, 37, 38, 39, 41,  
42, 43, 45, 46, 47, 48, 49, 50, 52, 53, 54,  
55, 56, 59, 60, 61, 62, 66, 76  
Standby System Bootup, 27  
standby system name, 16  
standby system., 8  
startech, 11, 20, 75  
**stmain**, 22, 33, 40  
Support, 11  
Switch to Fallback Mode, 29  
**sysadmsh**, 15  
**sysadmsh**, 15, 16  
system call, 63  
System Checkout, 31

## T

**TCP/IP**, 10, 15, 18, 24, 32, 76  
TCP/IP 1.2, 10  
terminal server, 10, 19  
Terminal Server:, 10  
terminals, 23  
**time out**, 18  
transition, 8, 22, 23, 27, 28, 30, 31, 42, 43,  
50, 57, 60, 61, 64  
trusted access, 15  
**tty1a**, 23  
**ttyn0101**, 23

## U

UDP, 40, 63  
**umount**, 15, 46, 56, 58  
umount, 15  
Unix, 7, 9, 10, 11, 17, 20, 38, 41, 46, 53, 56,  
57, 58, 59  
Unmount, 15  
users, 78  
*Utilities*, 17, 19  
*uucp*, 31

## V

*vi*, 21

## X

**XA**, 7, 56, 57, 59, 77  
X-Terminals, 79