



FailSafe™ Version 2 TMF Administrator's Guide

007-4422-002

CONTRIBUTORS

Written by Anita Manders, Steven Levine, Lori Johnson

Production by Karen Jacobson

Engineering contributions by Luca Bagnaschi, Dale Brantly, Dave Caliguire, Kevan Rehm, Judith Schmitz, Paddy Sreenivasan

COPYRIGHT

© 2002, 2004 Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, IRIS, and IRIX are registered trademarks and SGI, the Silicon Graphics logo, and IRIS FailSafe are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

INFORMIX is a trademark of Informix Software, Inc. Netscape is a trademark of Netscape Communications Corporation. NFS is a trademark of Sun Microsystems, Inc. Oracle is a trademark of Oracle Corporation. STK is a trademark of Storage Technology Corporation. Windows is a trademark of Microsoft Corporation.

Record of Revision

Version	Description
001	January 2002 Original printing.
002	April 2004 Incorporates information for the FailSafe TMF 2.2 release, which supports the FailSafe 2.1.x and TMF 1.4+ releases

Contents

About This Guide	vii
Audience	vii
Related Documentation	vii
Obtaining Publications	ix
Conventions	x
Reader Comments	x
1. FailSafe TMF Overview	1
FailSafe TMF Monitoring	1
FailSafe TMF Plug-In	1
Required Software	3
Overview of Configuring FailSafe for TMF	3
2. Configuring FailSafe TMF	5
Verifying that TMF is Enabled	5
Creating a TMF Resource Type	5
Configuring a TMF Resource	7
FailSafe TMF Configuration Parameters	7
Optional Configuration Specifications	8
The <code>/etc/tmf/failsafe_tmf.config</code> File	10
The <code>loader</code> Directive	10
The <code>remote_devices</code> Directive	11
Configuring Tapes and TMF	12
Executing FailSafe TMF	13
Creating a TMF Resource	15
007-4422-002	v

Contents

Creating a TMF Resource Group	16
Testing the TMF Resource	17
Testing the <i>start</i> Script	17
Testing the <i>stop</i> Script	18
Testing Resource Group Failovers	18
Index	19

About This Guide

This guide provides information about configuring FailSafe 2.1x systems with the FailSafe Tape Management Facility (TMF) 2.2 plug-in. This plug-in enables TMF and its resources to be failed over from one node to another if a component fails. This guide is intended as a supplement to the information included in the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Audience

This guide is written for system administrators who are responsible for configuring and administering a FailSafe system with the optional FailSafe TMF software. These system administrators must be able to customize several shell scripts and must be familiar with TMF configuration and TMF startup and shutdown procedures.

Related Documentation

For TMF information, see:

- *TMF Administrator's Guide*
- *TMF Release and Installation Guide*

Besides this guide, other documentation for the FailSafe system includes the following:

- *FailSafe Administrator's Guide for SGI InfiniteStorage*
- *FailSafe Programmer's Guide for SGI Infinite Storage*
- *IRIS FailSafe 2.0 INFORMIX Administrator's Guide*
- *IRIS FailSafe 2.0 Netscape Server Administrator's Guide*
- *IRIX FailSafe NFS Administrator's Guide*
- *IRIS FailSafe 2.0 Oracle Administrator's Guide*
- *IRIS FailSafe Version 2 Samba Administrator's Guide*

The FailSafe reference pages are as follows:

- `cdbBackup(1M)`
- `cdbRestore(1M)`
- `cluster_mgr(1M)`
- `crsd(1M)`
- `failsafe(7M)`
- `fs2d(1M)`
- `ha_cilog(1M)`
- `ha_cmds(1M)`
- `ha_exec2(1M)`
- `ha_fsadm(1M)`
- `ha_fsd(1M)`
- `ha_gcd(1M)`
- `ha_http_ping2(1M)`
- `ha_ifd(1M)`
- `ha_ifdadmin(1M)`
- `ha_ifmx2(1M)` (FailSafe INFORMIX option)
- `ha_macconfig2(1M)`
- `ha_srmd(1M)`
- `ha_statd2(1M)`
- `haStatus(1M)`

Release notes are included with each FailSafe product. The names of the release notes are as follows:

Release Note	Product
<code>cluster_admin</code>	Cluster administration services
<code>cluster_control</code>	Cluster node control services
<code>cluster_services</code>	Cluster services
<code>failsafe2</code>	FailSafe 2.1.x release
<code>failsafe2_dmf</code>	FailSafe DMF
<code>failsafe2_informix</code>	FailSafe INFORMIX
<code>failsafe2_nfs</code>	FailSafe NFS
<code>failsafe2_oracle</code>	FailSafe Oracle
<code>failsafe2_samba</code>	FailSafe Samba
<code>failsafe2_tmf</code>	FailSafe TMF
<code>failsafe2_web</code>	FailSafe/Netscape web
<code>patch_number</code>	FailSafe patch release

Obtaining Publications

You can obtain SGI documentation as follows:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- If it is installed on your SGI system, you can use InfoSearch, an online tool that provides a more limited set of online books, release notes, and man pages. With an IRIX system, enter `infosearch` at a command line or select **Help > InfoSearch** from the Toolchest.
- On IRIX systems, you can view release notes by entering either `grelnotes` or `relnotes` at a command line.

- On Linux systems, you can view release notes on your system by accessing the `README.txt` file for the product. This is usually located in the `/usr/share/doc/productname` directory, although file locations may vary.
- You can view man pages by typing `man title` at a command line.

Conventions

The following conventions are used throughout this guide:

Convention	Meaning
[]	Brackets enclose optional portions of a command or directive line.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
GUI	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
`techpubs@sgi.com`
- Use the Feedback option on the Technical Publications Library Web page:
`http://docs.sgi.com`

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

Technical Publications
SGI
1500 Crittenden Lane, M/S 535
Mountain View, California 94043-1351

SGI values your comments and will respond to them promptly.

FailSafe TMF Overview

This guide provides information about the components that are added to the base FailSafe 2.1.x product by the FailSafe Tape Management Facility (TMF) 2.2 plug-in. It assumes that you are familiar with the basic components of FailSafe described in the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

The chapter discusses the following:

- "FailSafe TMF Monitoring"
- "FailSafe TMF Plug-In"
- "Required Software" on page 3
- "Overview of Configuring FailSafe for TMF" on page 3

FailSafe TMF Monitoring

TMF is used to manage tapes and tape devices under IRIX. The TMF plug-in lets FailSafe monitor the availability of tape device groups as defined in the `tmf.config` file.

The TMF monitoring script monitors the availability of tape devices. If the monitor script detects a failure, it will report a local monitoring failure to FailSafe, and FailSafe will take the appropriate action.

FailSafe TMF Plug-In

The FailSafe TMF plug-in lets you create a TMF tape device group resource. A tape device resource is associated with a particular tape device group as defined in the TMF configuration file.

The following example is part of a TMF configuration file that defines a tape device group named EGLF:

```
DEVICE_GROUP
    name = EGLF
    AUTOCONFIG
{
    DEVICE
        NAME      = f9840f1 ,
        device_group_name = EGLF ,
        FILE      = /hw/tape/500104f000417a18/lun0/c4p1 ,
        status    = down ,
        access    = EXCLUSIVE ,
        vendor_address = (1,0,0,2),
        LOADER    = 1180
}
```

The FailSafe TMF plug-in performs the following functions for a particular TMF resource:

- Starts TMF if it is not already running.
- Configures the associated loader if it is not already set up.
- Configures the devices if they are not already set up.
- Lets you specify different tape device types to monitor — one type per instance of the resource you create.
- Forces dismounts of tapes from drives specified in the following file:

```
    /etc/tmf/failsafe_tmf.config
```
- Lets you specify a minimum number of devices of a specific type that must be available for use.
- Monitors the minimum number of tape devices that are available. If the number of devices currently available are less than the minimum level, a monitoring failure will occur.
- Releases previous reservations that are held by another host (if the drive firmware supports this option).
- Lets you assign a different drive type to each instance of a TMF resource.

Note: The TMF resource can take a relatively long time to start. For this reason, it starts late in the FailSafe sequence.

Required Software

The required software for TMF failover is as follows:

- TMF software. See the *TMF Administrator's Guide* for more information.
- Base FailSafe software. See the installation information in *FailSafe Administrator's Guide for SGI InfiniteStorage* for a complete list of required base software.
- The `expect` command, which is in the software subsystem `sgitcl_eoe.sw.expect`.
- FailSafe TMF software:
 - Base software subsystem: `failsafe2_tmf.sw.base`
 - Release notes subsystem: `failsafe2_tmf.man.relnotes`
 - Documentation subsystem (this guide): `failsafe2_tmf.books.book_AG`

Note: The only tape configuration supported by the FailSafe TMF software is an STK silo running ACSLS library control software version 6.0 or later.

Overview of Configuring FailSafe for TMF

To configure an FailSafe cluster for failover of TMF, follow these steps:

1. Install, configure, and test the base FailSafe software as described in the *FailSafe Administrator's Guide for SGI InfiniteStorage*.
2. Install additional software, if required. See "Required Software" on page 3.
3. Create the TMF resource type. See "Creating a TMF Resource Type" on page 5.
4. Configure the TMF resources. See "Configuring a TMF Resource" on page 7.

5. Make additional configuration changes, if desired. See "Optional Configuration Specifications" on page 8.
6. Make sure that the TMF and tape driver parameters are set correctly. See "Optional Configuration Specifications" on page 8 and "Configuring Tapes and TMF" on page 12.
7. Add the TMF resources to a resource group. See "Creating a TMF Resource Group" on page 16.
8. Test the TMF failover. See "Testing the TMF Resource" on page 17.

Configuring FailSafe TMF

This chapter provides information about configuring the FailSafe TMF plug-in:

- "Verifying that TMF is Enabled"
- "Creating a TMF Resource Type"
- "Configuring a TMF Resource" on page 7
- "Creating a TMF Resource" on page 15
- "Creating a TMF Resource Group" on page 16
- "Testing the TMF Resource" on page 17

The procedures described in this chapter assume that a cluster database that does not include TMF has already been created, installed, and tested as described in the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Verifying that TMF is Enabled

To run FailSafe TMF, the TMF software must be enabled. You should ensure that the output from `chkconfig` shows the following flag set to `on`:

```
# chkconfig | grep tmf
...
          tmf                on
```

If it is not, set it to `on`. For example:

```
# chkconfig tmf on
```

Creating a TMF Resource Type

To create a TMF resource type, this subsection assumes that you are already familiar with the concepts of resource types. Table 2-1 shows the resource attributes of a TMF resource type. "Configuring a TMF Resource" on page 7 describes how to use these parameters when configuring a TMF resource.

Table 2-1 TMF Configuration Parameters and Attributes

Parameters/Attributes	Comments
<i>resource-name</i>	The name of the TMF resource; for example <i>egft</i> .
<i>device-group</i>	The device group that is to be monitored. This device group specification is a device group that is defined in the TMF configuration file <i>tmf.config</i> .
<i>devices-minimum</i>	The minimum number of devices of the specified <i>device-group</i> that you must have available to you on a node.
<i>devices-loaned</i>	Currently unused; should be left at default value.
<i>email-addresses</i>	List of addresses to send email when the monitor script detects that devices in the <i>device-group</i> have become unavailable. This may be a comma- or white-space-separated list of names.

The TMF resource type is not created at cluster creation time. You must create the resource type before a TMF resource is created. The TMF resource type must be installed if you want to add a TMF resource to a cluster that was created before the FailSafe TMF software was installed.

You can use one of the following methods to create the TMF resource type:

- Run `cmgr` and manually create the resource type. For more information on `cmgr`, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.
- Run `cmgr` and install the resource type, as follows:

```
cmgr> install resource_type TMF in cluster eagan
cmgr> show resource_types installed

TMF
NFS
template
Netscape_web
statd_unlimited
Oracle_DB
MAC_address
IP_address
INFORMIX_DB
```

```
filesystem  
volume
```

- Use the template scripts supplied with FailSafe located in `/var/cluster/cmgr-template/ cmgr-create-resource_type`.
- Execute `/var/cluster/ha/resource_type/TMF/create_resource_type` and include the path of the cluster database argument and the cluster name.
- Run the FailSafe Manager GUI and use the **Load Resource Type** task to load the resource type. For more information on the FailSafe GUI, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Configuring a TMF Resource

The FailSafe TMF plug-in performs various functions for a TMF resource, as summarized in "FailSafe TMF Plug-In" on page 1. This section describes how to configure a resource to perform each of these functions.

FailSafe TMF Configuration Parameters

Table 2-1 summarizes the FailSafe TMF plug-in configuration parameters.

The FailSafe TMF plug-in lets you specify device groups to monitor. You specify a device group through the resource attribute `device-group`; it applies to the particular resource that is being defined. A device group refers to the tape devices that belong to a device group as defined in the TMF configuration file. This attribute is required for each resource that you create.

When you create a TMF resource, you must specify the minimum number of devices of a particular device group that must be configured and available for use. This value is specified as the resource attribute `devices-minimum` and is required for each resource. The default value for `devices-minimum` is 0, which means that no action is taken by the TMF plug-in even if no tapes are available (many sites would not want FailSafe to take action if the tapes are not available because failover or local restart will not help in that situation).

When you create a TMF resource, you must specify a list of email addresses to notify when the monitoring scripts detect that devices in the device group have become unavailable. Specify this list through the resource attribute `email-addresses` as a comma- or white-space-separated list of name.

A TMF resource includes the resource attribute `devices-loaned`. This attribute is currently unused by the FailSafe TMF plug-in and should be left at its default assigned value.

Optional Configuration Specifications

There are other optional configuration specifications associated with this resource. These specifications provide required information to the FailSafe TMF plug-in that let it communicate with the tape library and they also tell the plug-in which drives within the library on which it will force dismounts.

The FailSafe TMF plug-in can force a dismount of tapes from drives within the library. There may be various reasons why you might want to do this when a failover occurs. In the case of the data migration facility (DMF), you would want to ensure that any DMF tapes that were in use on a previous host are available to DMF on the new node after a failover. If these tapes were in drives assigned to the previous host, they must be ejected and returned to the library so that they are again accessible to DMF on the new host. You may want the FailSafe TMF plug-in to dismount only tape devices associated with a particular resource or you may not want the plug-in to dismount any tapes at all.

If you are using the `tpsc` tape driver, then in order for the plug-in to be able to force a dismount of tapes, the capabilities list specified for the device in the `/var/sysgen/master.d/scsi` file must not include the `MTCAN_PREV` capability.

The following example shows entries from this file for the STK 9840 and STK 9940 drives. The description for the 9840 drive does not include the `MTCAN_PREV` capability, but the description for the 9940 drive does include it.

```

/* STK 9840 drive */
{ STK9840, TPSTK9840, 3, 4, "STK", "9840", 0, 0, {0, 0, 0, 0},
MTCAN_BSF | MTCAN_BSR | MTCANT_RET | MTCAN_CHKRDY | MTCAN_SPEOD |
MTCAN_SEEK | MTCAN_APPEND | MTCAN_SILI | MTCAN_VAR | MTCAN_SETSZ |
    MTCAN_CHTYPEANY | MTCAN_COMPRESS,
    20, 8*60, 10*60, 3*60, 3*60, 16384, 256*1024,
    tpsc_default_dens_count, tpsc_default_hwg_dens_names,
    tpsc_default_alias_dens_names,
    {0}, 0, 0, 0,
    0, (u_char *)0 },

/* STK 9940 drive */
{ STK9840, TPSTK9840, 3, 4, "STK", "T9940A", 0, 0, {0, 0, 0, 0},
    MTCAN_BSF | MTCAN_BSR | MTCANT_RET | MTCAN_CHKRDY | MTCAN_PREV |
MTCAN_SPEOD |
    MTCAN_SEEK | MTCAN_APPEND | MTCAN_SILI | MTCAN_VAR | MTCAN_SETSZ |
    MTCAN_CHTYPEANY | MTCAN_COMPRESS,
    20, 8*60, 10*60, 3*60, 3*60, 16384, 256*1024,
    tpsc_default_dens_count, tpsc_default_hwg_dens_names,
    tpsc_default_alias_dens_names,
    {0}, 0, 0, 0,
    0, (u_char *)0 },

```

If the device does have this capability specified and you are using the `tpsc` tape driver, you must do the following:

- Remove the device from the list
- Perform an autoconfig
- Reboot in order for the plug-in to be able to force the dismounts of tapes from devices of that type

If you are using the `ts` tape driver, then the `/etc/config/tspd.config` file must not specify the following for the device:

```
PREVENT_REMOVAL pathname yes
```

If the device does have this capability specified, you must edit the `tspd.config` file and restart the appropriate `ts` personality daemon.

Some of the functions of the FailSafe TMF plug-in are performed through TMF; the plug-in issues commands to the TMF daemon to use these functions. However, the plug-in forces a dismount of a tape from a drive by issuing a command to the library

software controlling the loader/library. In the case of the Storage Technology Corporation (STK) hardware, the plug-in communicates its request to the Automated Cartridge System Library Software (ACSL) software that controls the loader. The plug-in uses an `expect` script that issues commands to login to the loader and issue a dismount request to a drive.

The `/etc/tmf/failsafe_tmf.config` File

The `/etc/tmf/failsafe_tmf.config` file lets you configure additional features of the FailSafe TMF plug-in. This file exists on all hosts in the cluster, and should be edited as necessary on each machine.

The contents of the `failsafe_tmf.config` file are dependent on the drives assigned to each host in the cluster. If all hosts in the failover domain are configured through TMF to use exactly the same drives, then this file would be the same on each host in the failover domain. You must maintain this file on each host; a change on one host is unknown to the other hosts.

There are two different types of directives that you can specify in the `failsafe_tmf.config` file: the `loader` directive and the `remote_devices` directive. These are defined in the following subsections.

The `loader` Directive

The `loader` directive provides information about a TMF loader, which controls one or more tape devices that are members of TMF device groups being managed as FailSafe resources. There may be more than one such directive in this file. The loader information is used by the FailSafe TMF plug-in to force a dismount of tapes from drives that cannot be made available (that is, have `tmstat` states other than `assn`, `free`, `conn`, or `idle`) so that those tapes can be used via other tape devices in the same device group. The information is also used to force a dismount of tapes from drives that are only connected to other hosts, not this host (as described in "The `remote_devices` Directive" on page 11). If the file does not contain a `loader` directive, then the TMF plug-in will make no attempt to force a dismount of tapes from any drives.

The directive has the following format:

```
loader lname ltype lhost luser lpswd
```

where:

lname Name of the loader as defined in the TMF config file.
ltype Type of the loader as defined in the TMF config file. (Currently only STKACS is supported for *ltype*.)
lhost Server name of the loader as defined in the TMF config file.
luser User name of the loader's administrator account. For STKACS, use *acssa*.
lpswd Password for the loader's administrator account.

The TMF command `/usr/sbin/tmmls` shows the name of the loader and the server associated with it:

```
# tmmls
loader type status m server old m_pnd d_pnd r_qd comp avg
operator OPERATOR UP A IRIX 0 0 0 0 0 0(sec)
wolfy STKACS DOWN A wolfcree 0 0 0 0 0 0(sec)
panther STKACS DOWN A stk9710 0 0 0 0 0 0(sec)
1180 STKACS UP A stk9710 0 0 0 0 0 0(sec)
```

For example, suppose you want to have the FailSafe TMF plug-in dismount drives that are in the 1180 loader/library listed above. That library has the *stk9710* server associated with it. The loader directive in the `failsafe_tmf.config` file would look like the following:

```
loader 1180 STKACS stk9710 acssa acssapassword
```

In this case, the FailSafe TMF plug-in would force a dismount for each drive that is specified in the `tmf.config` file to be in the 1180 loader/library and in the plug-in's drive group. If you do not want the plug-in to dismount any tape drives associated with a particular resource, you would not place a loader directive in the `failsafe_tmf.config` file.

The remote_devices Directive

The `remote_devices` directive provides information about one or more tape devices that are part of a TMF device group, but which are not visible on this host. An example would be where a library has four SCSI drives, and two drives are connected

to each of two FailSafe hosts. If host A should crash, host B must be able to force a dismount of any tapes in A's drives so that they can then be used from host B. Because the drives are not visible on host B, the `remote_devices` directive provides the information needed to force a dismount of unseen drives.

The directive has the following format:

```
remote_devices  rname lname droid ...
```

where:

<i>rname</i>	Name of the FailSafe TMF Resource that will dismount this drive.
<i>lname</i>	Name of the loader as defined in the TMF config file. There must be a loader directive for <i>lname</i> elsewhere in this file, or the <code>remote_devices</code> directive will be ignored.
<i>droid</i>	The vendor ID of the drive on which to force a dismount. This is the unique name by which the loader identifies the drive. In the case of STKACS, this will be a comma-separated four-digit string listing the ACS, LSM, drive panel, and drive (for example, 0,0,1,3).

Note: No blanks should exist within the ID.

Multiple vendor IDs can be specified in the same `remote_devices` directive as long as they all pertain to the same loader. If all the vendor IDs will not fit on a single line, just add additional `remote_devices` directives for the same loader. For example, to enable the FailSafe TMF plug-in to force a dismount of the remote drives (0,0,1,0), (0,0,1,1), (0,0,1,2), and (0,0,1,3) in the l180 loader/library for resource `tmf_eglf`, the directive would be:

```
remote_devices tmf_eglf1 180 0,0,1,0 0,0,1,1 0,0,1,2 0,0,1,3
```

If multiple FailSafe TMF resources are defined, only the resource named `tmf_eglf` will force a dismount of these drives.

Configuring Tapes and TMF

If drives that belong to a FailSafe TMF resource are configured on more than one machine in the FailSafe cluster, they should be configured consistently. The same tape driver (for example, `ts` or `tpsc`) should be used on each host where the drive is configured.

When configuring a FailSafe TMF resource, administrators should be aware of several parameters in the `/etc/tmf/tmf.config` file. The FailSafe TMF plug-in will try to start the loader associated with its device-group if it is not up. However, if the `tmf.config` file specifies `status = UP` for the loader, this step may not be necessary, and the devices may become available sooner.

A drive that is in a FailSafe TMF resource will be configured in the `tmf.config` file of one or more hosts within the cluster. It should be configured with `status=down`. All drives associated with the resource group must be unavailable for the exclusive script to indicate that the resource is not already running.

If the drives being used do not support persistent reserve, then they should be configured in the `tmf.config` file with `access=shared`. If the drives do support persistent reserve, then it is recommended that you use this feature when using the FailSafe TMF plug-in. To use persistent reserve, you should use the `ts` tape driver, and set `access=exclusive` in the `tmf.config` file. See the `ts(7)` man page for more information about using the `ts` tape driver. The `access` option should be consistent across all hosts in the failsafe cluster where the drives are configured.

The `-g` option of the `tmconfig` command reassigns a device to a different device group name. The FailSafe TMF software does not support reassigning a device into a FailSafe TMF device group. That is because, in case of failover, the FailSafe TMF plug-in on the machine we have failed over to would not have any knowledge of this reassigned drive. It would not be able to dismount tapes that are in the drive. Using `tmconfig -g` to move devices out of a FailSafe TMF device group will decrease the number of available drives that the monitor script sees. Also, in the case of failover or stop, the drive will be configured down.

Executing FailSafe TMF

The FailSafe TMF plug-in assumes that TMF is being used as the mounting service for tape devices associated with a tape library. Each time the plug-in is run, it will verify that TMF is up and running. If TMF is not running, the plug-in will start it. If TMF cannot be started by the plug-in, a failure will occur. Next, the plug-in will verify that the tape loader associated with the devices for a resource is up and accessible. If it is not up, the plug-in will configure it up using the `/usr/sbin/tmconfig` TMF command. If it cannot configure up the loader, a failure will occur.

The FailSafe TMF plug-in uses information supplied by the `/etc/tmf/tmf.config` file to identify what devices pertain to a particular resource. It uses that information in conjunction with the resource's `remote_devices` directive in the

`failsafe_tmf.config` file to determine what actions need to be taken on tape drives defined by the resource.

The plug-in retrieves the values of the `device-group` and `devices-minimum` attributes for a particular resource. It then examines the TMF configuration file for information pertaining to drives belonging to the same device group as specified for the resource and stores the information for processing.

It will then force a dismount of tapes from any drives that are specified in the `remote_devices` directive in the `failsafe_tmf.config` file associated with the resource. The FailSafe TMF plug-in verifies the minimum number of devices of the specified type are available for use. A device is considered available if its status displayed from the `tmstat` command is one of the following:

- `assn`
- `idle`
- `free`

If devices are in the down state, it will use `tmconfig` to configure them up and make sure that they are available. If a device cannot be configured up, and the associated loader directive is in the `failsafe_tmf.config` file, the plug-in will force a dismount of the tape from that device.

If the FailSafe TMF plug-in does not find the required minimum number of drives to be available, a failure will occur.

Note: The TMF Resource Type definition defines the number of local restarts to be a large number. This means that if a failure of the resource is detected on a host, the host will try to restart the resource before failing over. For information about changing this definition, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

Creating a TMF Resource

After you have defined the resource type, you must define the TMF resources based on the resource type. Each resource requires a unique resource name.

```
# cmgr
Welcome to SGI Cluster Manager Command-Line Interface

cmgr> set cluster eagan
cmgr> define resource egft of resource_type TMF
Enter commands, when finished enter either "done" or "cancel"

Type specific attributes to create with set command:

Type Specific Attributes - 1: device-group
Type Specific Attributes - 2: devices-minimum
Type Specific Attributes - 3: devices-loaned
Type Specific Attributes - 4: email-addresses

No resource type dependencies to add

resource egft ? set device-group EGLFT
resource egft ? set devices-minimum 4
resource egft ? set email-addresses yourname@yourcompany.com
resource egft ? done
Successfully defined resource egft
```

Note: The `devices-loaned` parameter is ignored and it should be left at its default value.

The `device-group` field is case sensitive. It must exactly match what is entered in the `tmf.config` file.

FailSafe, and the FailSafe TMF plug-in, allow you to specify unique values for each of the attributes on each of the hosts. For example, a FailSafe TMF resource could be defined so that host A had `devices-minimum=3`, but host B had `devices-minimum=2`.

Creating a TMF Resource Group

You can create a resource group by using either the FailSafe GUI or the `cmgr` command. For information, see the *FailSafe Administrator's Guide for SGI InfiniteStorage*.

To define an effective resource group, you must include all of the resources on which the TMF depends. The following example shows the creation of a typical resource group:

```
cmgr> show failover_policies

Failover Policies:
    t2
dmfadmin
dgroups
ordered

cmgr> define resource_group tmfrg in cluster eagan
Enter commands, when finished enter either "done" or "cancel"

resource_group tmfrg ? set failover_policy to dgroups
resource_group tmfrg ? add resource egft of resource_type TMF
resource_group tmfrg ? done
Successfully created resource group tmfrg

cmgr> show resource_group tmfrg

Resource Group: tmfrg
    Cluster: eagan
    Failover Policy: dgroups

Resources:
    egft (type: TMF)

cmgr> show failover_policy dgroups
Failover policy: dgroups
Version: 1
Script: ordered
Attributes: Auto_Failback Auto_Recovery
Initial AFD: guinness dublin
```

Testing the TMF Resource

To ensure that the TMF resource has been correctly configured, you can test individual actions by executing the scripts. Each script, located at `/var/cluster/ha/resource_types/TMF`, requires two arguments, an input file and an output file. The contents of these files are the resource names. The scripts will display 0 if they are successfully executed or display a positive number that indicates the error type. For more information on error codes, see the *FailSafe Programmer's Guide for SGI Infinite Storage*.

All TMF scripts assume they are being run under ksh.

In the following example, you can test the `start` script by starting the NFS resource with the resource name `tmfx`.

```
$ cd /var/cluster/ha/resource_types/TMF
$ echo "tmfx" > /tmp/ipfile
$ ./start /tmp/ipfile /tmp/opfile
```

This should start the `tmfx` instance, named by the TMF resource `tmfx`.

To view the individual script actions, you must edit the script and add the following to the `action` function:

```
set -x
```

Testing the `start` Script

Use the following procedure to test the `start` script:

1. Create a TMF resource for a device group specified in the `tmf.config` file. Do not run TMF on your test node.
2. Perform the following actions:

```
# echo "resource-name" > /tmp/ipfile
# /var/cluster/ha/resource_types/TMF/start /tmp/ipfile /dev/null
```

3. Check `/var/cluster/ha/log/script_nodename` logfile and verify from the log messages that the resource was started.
4. Verify that the TMF daemon process `tmdaemon` was started. Using the TMF commands `tmmls` and `tmstat`, verify that the tape loader and tape drives of type `resource_name` that you defined came up and are available.

Note: Running this test may force the dismount of tapes in derives as specified in the `failsafe_tmf.config` file.

Testing the `stop` Script

To stop the TMF resource, enter the following command:

```
# echo "resource-name" > /tmp/ipfile
# /var/cluster/ha/resource_types/TMF/stop /tmp/ipfile /dev/null
```

Check to see if the resource is offline by using the TMF `tmstat` command to verify that the tape devices of type `resource_name` were configured down.

Testing Resource Group Failovers

You can test the failover policy by using either `cmgr` or the FailSafe GUI to move the resource group to another node in the cluster. To ensure that the resource group correctly failed over, use `cmgr` or the GUI to display the resource group states.

The following example uses `cmgr` to test the failover policy:

```
cmgr> admin offline resource_group TMF in cluster eagan

cmgr> admin move resource_group TMF in cluster eagan to node cm2

cmgr> admin online resource_group TMF in cluster eagan
```

Index

A

ACSLs library control software, 3

C

chkconfig, 5
cmgr, 6
cmgr cluster manager, 16
cmgr-create-resource_type template script, 7
configuration overview, 3
configuration parameter
 device-group, 6, 7, 15
 devices loaned, 8
 devices minimum, 6
 devices-loaned, 6
 devices-minimum, 7
 email-addresses, 6, 7
 resource-name, 6
configuration parameters, 7

D

device-group configuration parameter, 6, 7, 15
devices-loaned configuration parameter, 6, 8
devices-minimum configuration parameter, 6, 7

E

email-addresses configuration parameter, 7
examples
 verify that chkconfig flags are on, 5
execution of the plug-in, 13
expect command, 3

007-4422-002

F

FailSafe cluster manager (cmgr), 6
FailSafe Manager GUI, 7
failsafe2_tmf.man.relnotes software subsystem, 3
failsafe2_tmf.sw.base software subsystem, 3
failsafe_tmf.config file, 2, 10
forcing a tape dismount, 8

G

GUI, FailSafe Manager, 7

L

loader directive, in failsafe_tmf.config file, 10
local restarts, 14

M

monitoring, 1
MTCAN_PREV capability, 9

P

PREVENT_REMOVAL, 9

R

release notes, 3
remote_devices directive, in failsafe_tmf.config file, 11

19

- resource
 - configuration, 7
 - creation, 15
- resource group
 - creation, 16
 - failover testing, 18
- resource type creation, 5
- resource-name configuration parameter, 6

S

- scsi file, 8
- sgitcl_eoe.sw.expect software subsystem, 3
- start script testing, 17
- STK drives, 8
- stop script testing, 18
- Storage Technology Corporation (STK)
 - hardware, 10

T

- tape configuration, 12

- tape dismount, force, 8
- template scripts, 7
- testing
 - resource group failover, 18
 - start script, 17
 - stop script, 18
 - TMF resource, 17
- TMF resource type, 5
- tmf.config file, 13
- tspd.config file, 9

V

- /var/cluster/ha/resource_types/TMF file, 17
- /var/sysgen/master.d/scsi file, 8