

IRIS FailSafe™ Version 2  
TMF Administrator's Guide

Document Number 007-4422-001

---

## CONTRIBUTORS

Written by Anita Manders, Steven Levine

Illustrated by Dany Galgani

Edited by Rick Thompson

Production by Karen Jacobson

Engineering contributions by Dave Caliguire, Kevan Rehm, Judith Schmitz, Paddy Sreenivasan

---

## COPYRIGHT

© 2002 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 electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy 2E, Mountain View, CA 94043-1351.

---

## 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.

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. Windows is a trademark of Microsoft Corporation.

---

## Record of Revision

<b>Version</b>	<b>Description</b>
001	January 2002 Original printing.



---

# Contents

<b>Figures</b>	vii
<b>Tables</b>	ix
<b>Examples</b>	xi
<b>About This Guide.</b>	xiii
Audience	xiii
Structure of This Document	xiii
Related Documentation	xiv
Conventions Used in This Guide	xvi
Reader Comments.	xvi
<b>1. IRIS FailSafe TMF Overview</b>	<b>1</b>
IRIS FailSafe TMF Monitoring	1
Overview of Configuring IRIS FailSafe for TMF	3
<b>2. Configuring IRIS FailSafe for TMF</b>	<b>5</b>
Required Software.	5
IRIS FailSafe 2.1 TMF Agent	6
Creating a TMF Resource Type	7
Configuring a TMF Resource.	9
The FailSafe TMF Agent Configuration Parameters	9
Additional Optional FailSafe TMF Agent Configuration Specifications	10
The /etc/config/failsafe_tmf.config File	11
The loader Directive	12
The remote_devices Directive	13
Configuring tapes and TMF.	14
Executing the IRIS FailSafe 2.1 TMF agent.	15
Creating a TMF Resource.	16

Creating a TMF Resource Group . . . . .	.17
Testing the TMF Resource . . . . .	.18
Testing the start Script . . . . .	.19
Testing the stop Script. . . . .	.19
Testing Resource Group Failovers . . . . .	.19
<b>Index.</b> . . . . .	<b>.21</b>

---

# Figures

<b>Figure 1-1</b>	Message Paths in IRIS FailSafe in a two-node cluster . . . . .	2
-------------------	--	---



---

# Tables

<b>Table i</b>	IRIS FailSafe Release Notes . . . . .	xv
<b>Table 2-1</b>	TMF Configuration Parameters . . . . .	7



---

# Examples

**Example 2-1**    Creating a TMF Resource using `cmgr` . . . . . 16



---

## About This Guide

This guide provides information about how to configure IRIS FailSafe 2.1 systems with the IRIS FailSafe 2.1 Tape Management Facility (TMF) option. This option enables TMF and its resources to be failed over from one node to another when a FailSafe failover occurs. This guide is intended as a supplement to the information about configuring IRIS FailSafe that is described in the *IRIS FailSafe Version 2 Administrator's Guide*.

This guide was prepared in conjunction with Release 2.1 of the IRIS FailSafe TMF option and the IRIS FailSafe 2.1 release. It describes the IRIS FailSafe TMF software for TMF release 1.3.3.2 and higher. This guide uses *FailSafe* as an abbreviation for *IRIS FailSafe*.

### Audience

This guide is written for system administrators who are responsible for configuring and administering an IRIS FailSafe system with the optional IRIS 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.

### Structure of This Document

This guide contains the following chapters:

- Chapter 1, “IRIS FailSafe TMF Overview,” describes the IRIS FailSafe components that are included in the IRIS FailSafe TMF option. It also provides an overview of the steps that are required to add TMF as a highly available service to an IRIS FailSafe cluster.
- Chapter 2, “Configuring IRIS FailSafe for TMF,” explains the procedure for adding TMF as a highly available service.

## Related Documentation

For TMF installation information, see the *IRIX TMF Administrator's Guide* and the *IRIX TMF Release and Installation Guide*.

For DMF installation information, see the *DMF Administrator's Guide for IRIX Systems* and the *DMF Release and Installation Guide for IRIX Systems*.

For CXFS installation information, see the *CXFS Software Installation and Administration Guide*.

Besides this guide, other documentation for the IRIS FailSafe system includes

- *IRIS FailSafe Version 2 Administrator's Guide*
- *IRIS FailSafe Version 2 Programmer's Guide*
- *IRIS FailSafe Version 2 DMF Administrator's Guide*
- *IRIS FailSafe 2.0 INFORMIX Administrator's Guide*
- *IRIS FailSafe 2.0 Netscape Server Administrator's Guide*
- *IRIS FailSafe Version 2 NFS Administrator's Guide*
- *IRIS FailSafe 2.0 Oracle Administrator's Guide*
- *IRIS FailSafe Version 2 Samba Administrator's Guide*

The IRIS FailSafe reference pages are as follows:

- `cdbBackup(1M)`
- `cdbRestore(1M)`
- `cluster_mgr(1M)`
- `crsd(1M)`
- `failsafe(7M)`
- `fs2d(1M)`
- `ha_cilog(1M)`
- `ha_cmsd(1M)`
- `ha_exec2(1M)`
- `ha_fsadm(1M)`

- ha\_fsd(1M)
- ha\_gcd(1M)
- ha\_http\_ping2(1M) (IRIS FailSafe Netscape Web option)
- ha\_ifd(1M)
- ha\_ifdadmin(1M)
- ha\_ifmx2(1M) (IRIS FailSafe INFORMIX option)
- ha\_macconfig2(1M)
- ha\_srmd(1M)
- ha\_statd2(1M)
- haStatus(1M)

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

**Table i** IRIS FailSafe Release Notes

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

## Conventions Used in This Guide

These type conventions and symbols are used in this guide:

<i>Italics</i>	New terms, manual/book titles, variable command line arguments, and variables to be supplied by the user in examples, code, and syntax statements
Fixed-width type	Command names and command line arguments, file names, directory names, resource types, and examples of command output and code
<b>Bold fixed-width type</b>	Commands and text that you are to type literally in response to shell and command prompts
#	IRIX shell prompt for the super user ( <i>root</i> )

## Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number can be found on the back cover.)

You can contact us in any of the following ways:

- Send e-mail to the following address:  
[techpubs@sgi.com](mailto:techpubs@sgi.com)
- Use the Feedback option on the Technical Publications Library World Wide Web page:  
<http://techpubs.sgi.com>
- Contact your customer service representative and ask that an incident be filled in the SGI incident tracking system.
- Send mail to the following address:

Technical Publications  
SGI  
1600 Amphitheatre Pkwy., M/S 535  
Mountain View, California 94043-1351

- Send a fax to the attention of Technical Publications at +1 650 932 0801.

We value your comments and will respond to them promptly.



## IRIS FailSafe TMF Overview

This chapter provides information about the components that are added to the base IRIS FailSafe 2.1 product by the IRIS FailSafe 2.1 Tape Management Facility (TMF) option. It assumes that you are familiar with the basic components of IRIS FailSafe described in Chapter 1 of the *IRIS FailSafe Version 2 Administrator's Guide*. This chapter also lists the steps you must take to add TMF to the highly available services that are failed over on an IRIS FailSafe cluster.

The Tape Management Facility (TMF) is used to manage tapes and tape devices under IRIX. The TMF option lets IRIS FailSafe monitor the availability of tape device groups as defined in the TMF configuration file `tmf.config`.

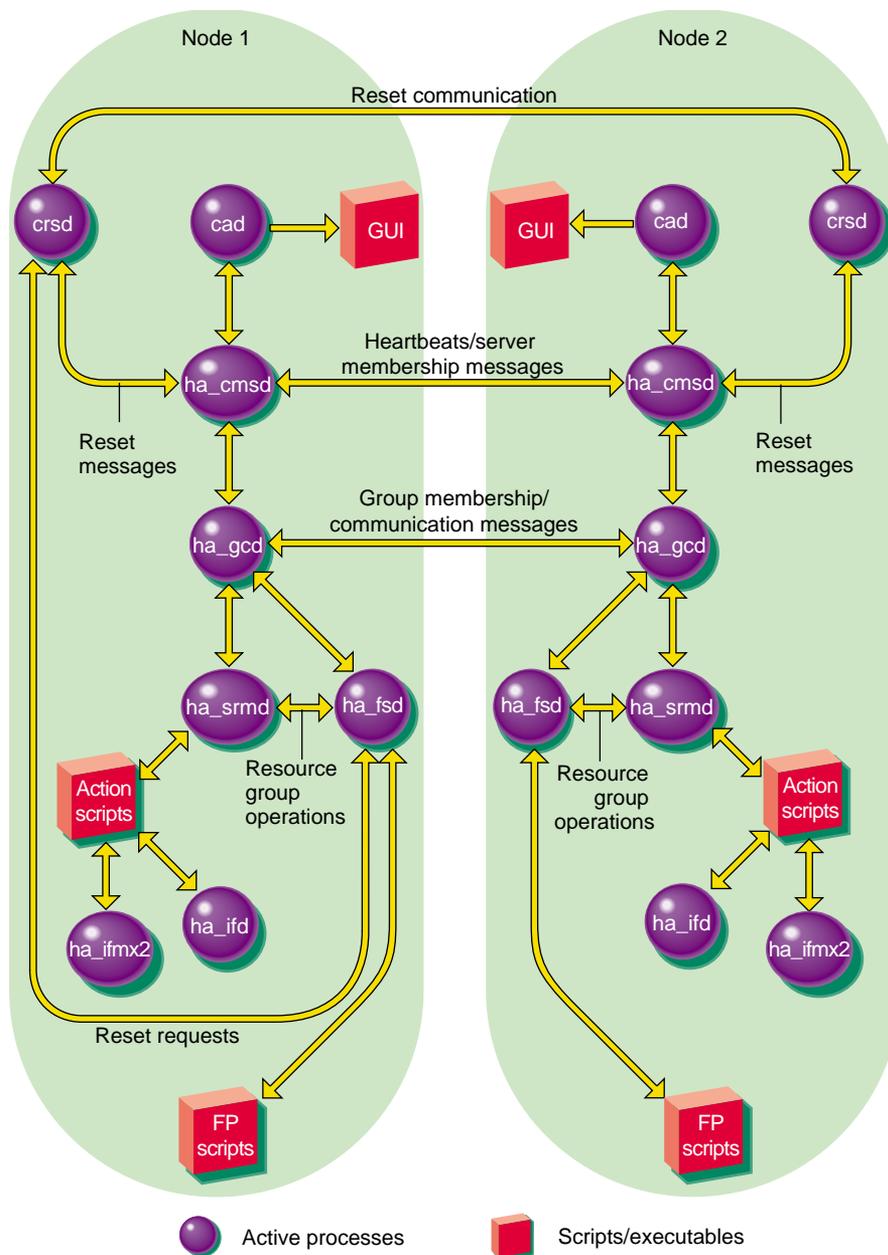
The major sections in this chapter are as follows:

- “IRIS FailSafe TMF Monitoring” on page 1
- “Overview of Configuring IRIS FailSafe for TMF” on page 3

### IRIS FailSafe TMF Monitoring

The IRIS FailSafe TMF option adds TMF device group monitoring to the IRIS FailSafe base software. 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.

IRIS FailSafe supports up to eight nodes in a cluster. Figure 1-1 shows an example of the message paths in IRIS FailSafe in a two-node cluster. Figure 1-1 also shows the interaction between the resource action scripts and the FailSafe agents (`ha_ifd`, `ha_ifmx2`). The TMF resource action scripts do not interact with any FailSafe agent directly, except for the `ha_srm` process.



**Figure 1-1** Message Paths in IRIS FailSafe in a two-node cluster

## Overview of Configuring IRIS FailSafe for TMF

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

1. Install, configure, and test the base IRIS FailSafe software as described in the *IRIS FailSafe Version 2 Administrator's Guide*.
2. Install additional software, if required. See "Required Software" on page 5.
3. Create the TMF resource type. See "Creating a TMF Resource Type" on page 7.
4. Configure the TMF resources. See "Configuring a TMF Resource" on page 9.
5. Make additional configuration changes, if desired. See "Additional Optional FailSafe TMF Agent Configuration Specifications" on page 10.
6. Make sure TMF and tape driver parameters are set correctly. See "Additional Optional FailSafe TMF Agent Configuration Specifications" on page 10 and "Configuring tapes and TMF" on page 14.
7. Add the TMF resources to a resource group. See "Creating a TMF Resource Group" on page 17.
8. Test the TMF failover. See "Testing the TMF Resource" on page 18.



## Configuring IRIS FailSafe for TMF

This chapter provides information about configuring the IRIS FailSafe TMF option for use on an IRIS FailSafe system. The major sections in this chapter are as follows:

- “Required Software” on page 5
- “IRIS FailSafe 2.1 TMF Agent” on page 6
- “Creating a TMF Resource Type” on page 7
- “Configuring a TMF Resource” on page 9
- “Creating a TMF Resource” on page 16
- “Creating a TMF Resource Group” on page 17
- “Testing the TMF Resource” on page 18

The procedures described in this chapter assume that a CDB that does not include TMF has already been created, installed, and tested as described in the *IRIS FailSafe Version 2 Administrator’s Guide*.

### Required Software

The required software for TMF failover is as follows:

- TMF software  
See the *TMF Administrator’s Guide* for more information about TMF.
- Base IRIS FailSafe software  
See the section “Installing the IRIS FailSafe Software” in Chapter 2 of the *IRIS FailSafe Version 2 Administrator’s Guide* for a complete list of required base software.
- The `expect` command  
The `expect` command is in the software subsystem `sgitcl_eoe.sw.expect`.

- IRIS FailSafe TMF software
  - The base software subsystem is `failsafe2_tmf.sw.base`.
  - The release notes subsystem is `failsafe2_tmf.man.relnotes`.
  - The *IRIS FailSafe Version 2 TMF Administrator's Guide* subsystem is `failsafe2_tmf.books.book_AG`.

---

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

---

## IRIS FailSafe 2.1 TMF Agent

The IRIS FailSafe 2.1 TMF agent 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 IRIS FailSafe 2.1 TMF agent 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/config/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.

---

## 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 9 describes how to use these parameters when configuring a TMF resource.

**Table 2-1** TMF Configuration Parameters

Resource Parameters	Comments
<i>resource-name</i>	The name of the TMF resource; for example <i>egft</i> .
<b>resource attribute</b>	
<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 <code>tmf.config</code> .

**Table 2-1 (continued)** TMF Configuration Parameters

Resource Parameters	Comments
<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. The administrator 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 cluster manager (`cmgr`) and manually create the resource type. For more information on `cmgr`, see the *IRIS FailSafe Version 2 Administrator's Guide*.
- Run cluster manager (`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

cmgr>
```

- 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 CDB 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 *IRIS FailSafe Version 2 Administrator's Guide*.

## Configuring a TMF Resource

The IRIS FailSafe 2.1 TMF agent performs various functions for a TMF resource, as summarized in “IRIS FailSafe 2.1 TMF Agent” on page 6. This section describes how to configure a resource to perform each of these functions.

### The FailSafe TMF Agent Configuration Parameters

Table 2-1 on page 7 summarizes the FailSafe TMF agent configuration parameters.

The IRIS FailSafe 2.1 TMF agent lets you specify device groups to monitor. A device group is specified through the resource attribute `device-group` and 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.

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. This list is specified 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 agent and should be left at its default assigned value.

## Additional Optional FailSafe TMF Agent Configuration Specifications

There are other optional configuration specifications associated with this resource. These specifications provide required information to the IRIS FailSafe 2.1 TMF agent that let it communicate with the tape library and they also tell the agent which drives within the library on which it will force dismounts.

The IRIS FailSafe 2.1 TMF agent 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 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 IRIS FailSafe 2.1 TMF agent to dismount only tape devices associated with a particular resource or you may not want the agent to dismount any tapes at all. If you are using the tpsc tape driver, then in order for the agent to be able to force a dismount of tapes, the capabilities list specified for the device in the file `/var/sysgen/master.d/scsi` 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_BSFF | 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 remove the device from the list, perform an autoconfig, and reboot in order for the agent to be able to force the dismounts of tapes from devices of that type.

If you are using the ts tape driver, then the file `/etc/config/tspd.config` must not specify `PREVENT_REMOVAL` pathname `yes` for the device. 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 IRIS FailSafe 2.1 TMF agent are performed through TMF and the agent issues commands to the TMF daemon to use them. However, the agent 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 agent communicates its request to the Automated Cartridge System Library Software (ACSL) software that controls the loader. The agent uses an `expect` script that issues commands to login to the loader and issue a dismount request to a drive.

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

The file `/etc/config/failsafe_tmf.config` provides a way for administrators to configure additional features of the TMF FailSafe agent. 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. It must be maintained 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 2.1 TMF agent to force a dismount of tapes from drives that cannot be made available (i.e. 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 13). If the file does not contain a `loader` directive, then the TMF agent will make no attempt to force a dismount of tapes from any drives. The directive has the following format:

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

where:

<i>lname</i>	name of the loader as defined in the TMF config file.
<i>ltype</i>	type of the loader as defined in the TMF config file. (Currently only ‘STKACS’ is supported for <i>ltype</i> .)
<i>lhost</i>	server name of the loader as defined in the TMF config file.
<i>luser</i>	user name of the loader’s administrator account. For STKACS, use <code>acssa</code> .
<i>lpswd</i>	password for the loader’s administrator account.

The TMF command `/usr/tmf/bin/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)
wolffy	STKACS	DOWN	A	wolfcree	0	0	0	0	0	0(sec)
panther	STKACS	DOWN	A	stk9710	0	0	0	0	0	0(sec)
l180	STKACS	UP	A	stk9710	0	0	0	0	0	0(sec)

For example, suppose you want to have the TMF FailSafe agent dismount drives that are in the l180 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 l180 STKACS stk9710 acssa acssapassword
```

In this case, the FailSafe TMF agent would force a dismount for each drive that is specified in the `tmf.config` file to be in the l180 loader/library and in the agent's drive group. If you do not want the agent 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 second type of directive that can be specified in the `failsafe_tmf.config` file is the `remote_devices` directive. A `remote_devices` directive provides information about one or more tape devices which 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 needs to be able to force a dismount of any tapes in A's drives so that they can then be used from host B. Since 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:

*rname*            name of the TMF FailSafe Resource that will dismount this drive.

*lname*            name of the loader as defined in the TMF config file. There must be a loader directive for *lname* elsewhere in this file, or the `remote_devices` directive will be ignored.

*droid* 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 (e.g 0,0,1,3). Note that 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 won't fit on a single line, just add additional `remote_devices` directives for the same loader. For example, to enable the TMF FailSafe agent 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_eglfl`, the directive would be:

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

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

## Configuring tapes and TMF

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

When configuring a TMF FailSafe resource, administrators should be aware of several parameters in the `/etc/config/tmf.config` file. The TMF FailSafe agent 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 TMF FailSafe 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 TMF FailSafe Agent. 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 TMF FailSafe agent 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. If `tmconfig -g` is used to move devices out of a FailSafe TMF device group, that will decrease the number of available drives that the monitor script sees. Also, in the case of failover or stop, the drive will be downed.

## Executing the IRIS FailSafe 2.1 TMF agent

The IRIS FailSafe 2.1 TMF agent assumes that TMF is being used as the mounting service for tape devices associated with a tape library. Each time the agent is run, it will verify that TMF is up and running. If TMF is not running, the agent will start it. If TMF cannot be started by the agent, a failure will occur. Next, the agent will verify that the tape loader associated with the devices for a resource is up and accessible. If it is not up, the agent will configure it up using the `/usr/tmf/bin/tmconfig` TMF command. If it cannot configure up the loader, a failure will occur.

The IRIS FailSafe 2.1 TMF agent uses information supplied by the TMF configuration file `/etc/config/tmf.config` 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 agent 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 IRIS FailSafe 2.1 TMF agent 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 either `assn`, `idle`, or `free`. If devices are in the down state it will `tmconfig` 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 agent will force a dismount of the tape from that device.

If the IRIS FailSafe 2.1 TMF agent 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. This may be changed by the administrator, if desired. See the *IRIS FailSafe Version 2 Administrator's Guide* for more information.

---

## Creating a TMF Resource

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

### Example 2-1 Creating a TMF Resource using `cmgr`

```
# 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@yourcomany.com
resource /disk5 ? done
Successfully defined resource egft
```

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

## Creating a TMF Resource Group

To create a resource group, you must first become familiar with the terms and concepts of FailSafe. A resource group can be created either by the GUI or the cluster manager (`cmgr`). For information on the GUI and `cmgr` see the *IRIS FailSafe Version 2 Administrator's Guide*.

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

```
cmgr> cshow 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

cmgr>
```

### 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 *IRIS FailSafe Version 2 Programmer's Guide*.

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 “`set -x`” to the `action` function.

## Testing the start Script

You can use the following procedure to test the `start` script.

1. You can create a TMF resource for a device group specified in your TMF configuration file `tmf.config`. Do not have TMF up and running on the node that you are going to test on.
2. Perform the following actions:  

```
echo "resource-name" > /tmp/ipfile  
  
/var/cluster/ha/resource_types/TMF/start /tmp/ipfile /dev/null
```
3. Check the failsafe logfile `/var/cluster/ha/log/script_nodename` 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

You can use the following procedure to test the `stop` script.

1. 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 TMF using the command `tmstat` 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 the `cmgr` or 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, 6

## C

cluster manager (cmgr), 8, 16, 17  
cmgr-create-resource\_type template script, 8  
configuration overview, 3  
configuration parameter  
  device-group, 7, 9, 17  
  devices loaned, 9  
  devices minimum, 8  
  devices-loaned, 8  
  devices-minimum, 9  
  email-addresses, 8, 9  
  resource-name, 7  
configuration parameters, 7, 9

## D

device-group configuration parameter, 7, 9, 17  
devices-loaned configuration parameter, 8, 9  
devices-minimum configuration parameter, 8, 9

## E

email-addresses configuration parameter, 9  
expect command, 5

## F

FailSafe agents, 1  
FailSafe cluster manager (cmgr), 8  
FailSafe Manager GUI, 9  
failsafe\_tmf.config file, 7, 11  
failsafe2\_tmf.man.relnotes software subsystem, 6  
failsafe2\_tmf.sw.base software subsystem, 6  
forcing a tape dismount, 10

## G

GUI, FailSafe Manager, 9

## I

IRIS FailSafe message paths, 1

## L

loader directive, in failsafe\_tmf.config file, 12  
local restarts, 16

## M

message paths, IRIS FailSafe, 1, 2

**R**

release notes, xv, 6  
remote\_devices directive, in failsafe\_tmf.config file,  
13  
resource  
  action scripts, 1  
  configuration, 9  
  creation, 16  
resource group  
  creation, 17  
  failover testing, 19  
resource type creation, 7  
resource-name configuration parameter, 7

**S**

scsi file, 10  
sgitcl\_eoe.sw.expect software subsystem, 5  
start script testing, 19  
stop script testing, 19  
Storage Technology Corporation (STK) hardware, 11

**T**

tape configuration, 14  
tape dismount, force, 10  
template scripts, 8  
testing  
  resource group failover, 19  
  start script, 19  
  stop script, 19  
  TMF resource, 18  
TMF  
  configuration parameters, 7  
  execution, 15  
  monitoring, 1

resource configuration, 9  
resource creation, 16  
resource group creation, 17  
resource testing, 18  
resource type, creation, 7  
tmf.config file, 14  
tspd.config file, 11

**V**

/var/cluster/ha/resource\_types/TMF file, 18  
/var/sysgen/master.d/scsi file, 10