IRIS FailSafe™ DMF Administrator's Guide

Document Number 007–3906–001

Record of Revision

Version	*Description*
1.2	December 1998 Original printing to support the initial release of IRIS FailSafe DMF.

Contents

Page

About This Guide

This guide provides information about configuring IRIS FailSafe systems with the IRIS FailSafe Data Migration Facility (DMF) option. This option enables DMF and its resources to be failed over from one node to another when a FailSafe fail over occurs. This guide is intended as a supplement to information about configuring IRIS FailSafe included in the *IRIS FailSafe Administrator's Guide*.

This guide was prepared in conjunction with the initial release of the IRIS FailSafe DMF option and release 1.2 of the IRIS FailSafe software product. It describes IRIS FailSafe DMF software for DMF release 2.6.1 and higher.

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

Related Publications

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

In addition to this guide, other documentation for the IRIS FailSafe system includes the following:

- *IRIS FailSafe Administrator's Guide*
- *IRIS FailSafe Gauntlet Administrator's Guide* which documents the IRIS FailSafe Gauntlet option
- *IRIS FailSafe INFORMIX Administrator's Guide* which documents the IRIS FailSafe INFORMIX option
- *Addendum to IRIS FailSafe INFORMIX Administrator's Guide*
- *IRIS FailSafe Oracle Administrator's Guide* which documents the IRIS FailSafe Oracle option
- *IRIS FailSafe Programmer's Guide*
- *IRIS FailSafe Sybase Administrator's Guide* which documents the IRIS FailSafe Sybase option

- *IRIS FailSafe WebFORCE MediaBase Administrator's Guide* which documents the IRIS FailSafe WebFORCE MediaBase option
- *Origin200 FailSafe Network Server Important Note* which documents the Origin200 FailSafe Network Server
- *Origin200 FailSafe Network Server Quick Reference Guide* which documents the Origin200 FailSafe Network Server

Man pages (reference pages) also exist for IRIS FailSafe and its options.

Release notes are included with each IRIS FailSafe product.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
`command`	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
variable	Italic typeface denotes variable entries and words or concepts being defined.
`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.
[ ]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

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 part number of the document with your comments.

You can contact us in any of the following ways:

- Send electronic mail to the following address:

  techpubs@sgi.com

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

- Use the Suggestion Box form on the Technical Publications Library World Wide Web page:

  http://techpubs.sgi.com/library/

- Call the Technical Publications Group, through the Technical Assistance Center, using one of the following numbers:

  For Silicon Graphics IRIX based operating systems: 1 800 800 4SGI

  For UNICOS or UNICOS/mk based operating systems or CRAY Origin2000 systems: 1 800 950 2729 (toll free from the United States and Canada) or +1 651 683 5600

- Send mail to the following address:

  Technical Publications
  Silicon Graphics, Inc.
  2011 North Shoreline Boulevard, M/S 535
  Mountain View, California 94043–1389

We value your comments and will respond to them promptly.

Introduction [1]

This publication provides information about configuring the IRIS FailSafe product to run with the Data Migration Facility (DMF).

Note: IRIS DMF FailSafe does not support a dual-active configuration.

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

1.1 IRIS FailSafe DMF Daemon Monitoring

The IRIS Failsafe DMF option adds DMF monitoring and fail over scripts to the IRIS base software. The monitoring script monitors the availability and status of the DMF daemon. If the script detects that the daemon is not running or responding to requests, the script takes the following actions:

- Writes an error message to the monitoring script's log file (`/var/ha/logs/ha_dmf_lon.`*`$HOST`*`.log`). For example:

  ```
  Tue Sep 22 00:01:27 CDT 1998 DMF has STOPPED , check DMF logfiles
  ```

- Sends an error message to the system console indicating DMF has stopped. For example:

  ```
  Tue Sep 22 00:01:28 DMF STOPPED , check DMF logfiles
  ```

- Sends a mail message with error information to the mail address defined by the `mail-dest-addr` field. For example:

  ```
  Tue Sep 22 00:01:26 DMF STOPPED , check DMF logfiles
  ```

1.2 Overview of Configuring IRIS FailSafe for DMF

To configure an IRIS FailSafe cluster for fail over of DMF, complete the following steps:

- Install, configure, and test the version 1.2 base IRIS FailSafe software as described in the *IRIS FailSafe Administrator's Guide*.

- Choose how to configure the DMF software, databases, and DMF file systems, as described in Section 2.2, page 3.
- Install DMF on each node in the cluster.
- Set up DMF configuration files as described in Section 2.4, page 5.
- Set up tape configuration files if DMF is running the tape media-specific process (MSP), as described in Section 2.3, page 4.
- Install the FailSafe DMF software.
- Add DMF information to the `/var/ha/ha.conf` configuration file, as described in Section 2.5, page 6.
- Install the new configuration files on each node, as described in Procedure 1, page 6.
- Test DMF fail over as described in Section 2.6, page 7.
- Start FailSafe.

Configuring IRIS FailSafe for DMF [2]

This chapter provides information about configuring the IRIS FailSafe DMF database option for use on the IRIS FailSafe system.

2.1 Required Software

The required software for DMF fail over is as follows:

- DMF software as described in the *Cray DMF Administrator's Guide for IRIX Systems*.
- Base IRIS FailSafe software (see the *IRIS FailSafe Administrator's Guide* for information on installing FailSafe)
- IRIS FailSafe DMF software (included in the FailSafe DMF package):
  - `ha_dmf.books`, which contains this guide
  - `ha_dmf.man`, which contains the `relnotes` file
  - `ha_dmf.sw`, which contains the following:
    - DMF daemon monitoring script
    - FailSafe DMF fail over functions (`takeover`, `takeback`, `giveback`, `giveaway`)
    - Template file `ha.conf.dmf` for `ha.conf` modifications
    - Tape dismount software for configurations attached to STK silos running the Automated Cartridge System Library Software (ACSLS)

2.2 DMF File System Configuration

DMF databases, log files, journal files, and user file systems must be XFS file systems or XLV logical volumes and must be located on a shared disk within the cluster. User file systems must be configured with the `dmi` mount option in the `ha.conf` file. The file systems are normally created as NFS file systems so that they can be mounted and accessed remotely. Procedure 1, page 6 describes how to make changes to the configuration file.

2.3 Tape Configuration

The only tape configuration that is currently supported with FailSafe DMF is a cluster connected to an STK silo running ACSLS library control software. Only the DMF tape autoloader service configuration is supported. FailSafe DMF does not support Open Vault and Tape Management Facility (TMF) tape system configurations. Consult the *Cray DMF Administrator's Guide for IRIX Systems* and the DMF release online files (`Readme` and `News`) for a description of how DMF is configured for each of these tape management systems.

Each host in the cluster is connected to a separate set of drives in the tape library. You must create a drive identification file on each host; the file defines the drives that DMF uses on the other host in the cluster. These files are required so that FailSafe DMF will know what drives were in use when the fail over occurred. FailSafe DMF will dismount any tapes that were in use at the time of a fail over. These files are created in `/etc/config` on each host. Section 2.3.1 describes how to create and name these files.

2.3.1 Tape Drive Identification Files

Tape drive identification files identify tape components that DMF uses on each host. There are two types of components identified in these files: a tape loader and tape drives. The information for each component is obtained from the DMF autoloader tape configuration file `/etc/config/al_api.rc`. The template for these files is as follows:

`loader:`*loader name*
`drive:`*drive_1*`[:`*drive_2*`][:`*drive_3*`]...`

The loader name and drive names are those specified in `/etc/config/al_api.rc`.

Create the file `/etc/config/ha_serv.drives` on the backup node. It contains the drives that are connected on the server node. Create the file `/etc/config/ha_back.drives` on the primary server node. It contains the drives connected to the backup node.

Example 1: Creating Drive Identification Files

This example lists the `/etc/config/ha_back.drives` and `/etc/config/ha_serv.drives` files for a cluster containing a primary server machine `cm1` and a backup node `cm2`.

File `ha_back.drives` on machine `cm1` indicates that drives `t2` and `t3` are being used by DMF on node `cm2` and are managed by the loader `wolfy`. It contains the following lines:

```
loader:wolfy
drive:t2:t3
```

File `ha_serv.drives` on machine `cm2` indicates the drives `t1` and `t4` are being used by DMF on node `cm1` and are managed by the loader `wolfy`. It contains the following lines:

```
loader:wolfy
drive:t1:t4
```

2.4 DMF Configuration Files

DMF must be installed on each host in the cluster; therefore, each host will have a `dmf_config` configuration file.

If DMF is migrating files using only the FTP MSP (that is, the media-specific process (MSP) that runs over the file transfer protocol (FTP)), the `dmf_config` files on each host will be identical. If you will be migrating files to tape using the tape MSP, each configuration file will differ only in the actual tape drives they specify for the tape MSP. The configuration file on each host will specify the drives attached to that host as defined in the `/etc/config/al_api.rc` file.

For example, suppose the `dmf_config` file for each host has the following information for the MSP named `msptim`:

```
define  msptim
        TYPE                    msp
        COMMAND                 dmatmsp
        TAPE_TYPE               tim_drives
        CACHE_SPACE             800m
        CHILD_MAXIMUM           2
        DISK_IO_SIZE            1024k
        MAX_PUT_CHILDREN        2
enddef
```

In the `dmf_config` configuration file on the first host, the drives defined for `tim_drives` are as follows:

```
define  tim_drives
        TYPE            device
        LOADER_NAME     wolfy
        TAPE_UNITS      t1 t4
enddef
```

On the second host, the drives defined for `tim_drives` are defined as follows:

```
define  tim_drives
        TYPE            device
        LOADER_NAME     wolfy
        TAPE_UNITS      t3 t5
enddef
```

Note: The drives defined by `TAPE_UNITS` are the **only** difference between the two `dmf_config` files.

2.5 Adding DMF Information to the FailSafe Configuration File

This section describes the procedure for creating the `ha.conf` configuration file that includes DMF configuration information. The procedure assumes that a configuration file that doesn't include DMF has been created, installed, and tested as described in the *IRIS FailSafe Administrator's Guide*. Using Procedure 1, add DMF information to the configuration file. Install the configuration file as `/var/ha/ha.conf` on both nodes as described in the *IRIS FailSafe Administrator's Guide*.

Procedure 1: Making Changes for DMF in the Configuration File

Complete the following steps:

1. Make a copy of the `/var/ha/ha.conf` file on one node.
2. Add all of the file systems and volumes that will be used for DMF to the copy of `ha.conf`. See the *IRIS FailSafe Administrator's Guide* and the *Cray DMF Administrator's Guide for IRIX Systems* for more information on volume and file system configuration.

   Note: When you are setting up a file system block for a DMF user file system, the `dmi` mount option must be specified as one of the mount options.

For example, if the file system `fs1` is a DMF user file system that contains files to migrate, the file system description block in `ha.conf` might look like the following:

```
filesystem fs1
{
     mount_point     /fs1
     mount_info
     {
          fs_type = xfs
          volume_name = fs1
          mode = dmi, rw, noauto, wsync
     }
}
```

3. Make a copy of the `/var/ha/templates/ha.conf.dmf` file. In the copy's `dmf` block, modify the definitions of the `server-node` and `backup-node` fields. (The server node is the node that normally would be running DMF and the backup node would serve as a backup platform for DMF within the cluster.) For more information, see Section 3.1, page 11.
4. Append the modified copy of `ha.conf.dmf` to the end of the `ha.conf` copy.
5. Define an NFS block in the copy of the `ha.conf` configuration file for each DMF user file system if the file systems will be accessed remotely.
6. Using the information in Section 3.2, page 11, prepare the `action dmf` and `action-timer dmf` blocks.
7. Use the information in section about creating the configuration file in the*IRIS FailSafe Administrator's Guide* to verify the `ha.conf` copy and then install it on each node. You can begin with the step involving the `ha_cfgverify` command.

2.6 Testing DMF Fail Over

The following procedure explains how to test the DMF configuration and fail over.

Procedure 2: Testing the Fail Over

Complete the following steps:

1. Install DMF on each node in the cluster.

2. Stop FailSafe if it is running by issuing the following command:

   ```
   /etc/init.d/failsafe stop
   ```

3. Make a backup copy of the `ha.conf` file on each node in the cluster. Install the `ha.conf` file created in Procedure 1, page 6 on each node in the cluster.

4. Bring up FailSafe by issuing the following command:

   ```
   /etc/init.d/failsafe start
   ```

5. Verify that all the DMF file systems defined in `ha.conf` are mounted and that the DMF daemon is running by issuing the command:

   ```
   # /etc/dmf/dmbase/etc/dmdstat -v
   Daemon status OK; '1' responses received.
   ```

6. Stop the DMF daemon by issuing the following command:

   ```
   /etc/init.d/dmf stop
   ```

7. Verify the following events:

   - An error message is sent to the system console indicating that DMF has stopped.
   - An error message is issued to the DMF monitor log in `/var/ha/logs/ha_dmf_lmon.`*$HOST*`.log`.
   - A mail message with the error information is sent to the `fsafe_admin` alias.

8. Bring DMF back up by issuing the following command:

   ```
   /etc/init.d/dmf start
   ```

9. Issue the `ha_admin -fs` command to put the host running DMF into standby mode. Verify the DMF file systems, the DMF daemon, and the DMF MSPs are failed over to the other node in the cluster.

10. Issue the following command on the host that was put in standby mode in step 9 in order to make the node rejoin the cluster:

    ```
    ha_admin -fr
    ```

    If this step was successful, the following will be true:

    - DMF is running on the reactivated node

- DMF user, log, and database file systems are remounted on the reactivated node
- The host is running in `normal` mode. You can determine the status by issuing the following command:

  ```
  ha_admin -a
  ```

Configuration File Blocks for DMF [3]

Configuration parameters for FailSafe DMF must be specified in the configuration file `/var/ha/ha.conf`. The sections in this chapter describe each DMF-specific block defined in the FailSafe DMF template file `/var/templates/ha.conf.dmf`. This file is modified and appended to `/var/ha/ha.conf` as described in Section 2.5, page 6.

The examples in this chapter show the DMF configuration file blocks discussed in Section 2.5, page 6.

3.1 DMF Application-Class Block

The following example shows the application-class block in a DMF configuration. The cluster contains the machines `cm1` as the server (primary) node and `cm2` as the backup node.

Example 2: Application-Class Block

```
application-class dmf
{
        server-node = cm1
        backup-node = cm2
}
```

3.2 DMF Action and Action-Timer Blocks

The following example shows the action and action-timer blocks for DMF.

Example 3: Action and Action-Timer Blocks

The action block specifies the path names of the local monitoring script and the action-timer block specifies default monitoring timing and timeout values for the monitoring of the DMF daemon.

```
action dmf
{
        local-monitor = /var/ha/actions/ha_dmf_lmon
}
action-timer dmf
{
        start-monitor-time = 60
        lmon-probe-time = 120
        lmon-timeout = 60
        retry-count = 1
}
```

The parameters used in action and action-timer blocks for DMF are as follows:

`local-monitor`	The path name of the local monitoring script for DMF. Do not change this value.
`start-monitor-time`	Specifies the amount of time that the application monitor waits before it starts using the local monitoring script to monitor the DMF daemon.
`lmon-probe-time`	Specifies (in seconds) how often local monitoring of the DMF daemon is done.
`lmon-timeout`	Specifies (in seconds) how often local monitoring of the DMF daemon is done if no response is received.
`retry-count`	Specifies the number of times the local monitoring script retries its probes of the DMF daemon. This values does not affect the `lmon-timeout` value.