IRIS FailSafe™ Version 2
Samba Administrator's Guide

Document Number 007-3988-002

CONTRIBUTORS

Written by Anita Manders

Illustrated by Dany Galgani

Edited by Rick Thompson

Production by Glen Traefald

Engineering contributions by Jeremy Allison, Herb Lewis, Paddy Sreenivasan, Saadia Khan

New Features in This Guide

This rewrite of the *IRIS FailSafe Version 2 Samba Administrator's Guide* supports the IRIS FailSafe 2.1 release.

New Features Documented

The following change has been made to this guide:

- Added support for multiple Samba servers in a FailSafe cluster.
- Streamlined the Samba resource configuration process.
- Added support for the `monitor_level` resource attribute.
- Added support for multiple `IP_address` resource dependencies.
- Added support for multiple `filesystem` resource dependencies.
- Removed support for the `probe` script.
- Updated several examples.

Miscellaneous editing and formatting changes were also made.

Record of Revision

Version	Description
001	November 1999 Incorporates information for the IRIS FailSafe 2.0 release.
002	March 2001 Incorporates information for the IRIS FailSafe 2.1 release.

Contents

Figures

Tables

Examples

About This Guide

This guide provides information about how to configure IRIS FailSafe 2.1 systems with the IRIS FailSafe 2.1 Samba option. This option enables Samba resources to be failed over from one node to another if a component fails. 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 Samba option. It describes the IRIS FailSafe Samba software for the Samba for IRIX release 2.0.7. 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 Samba software. These system administrators must be familiar with the Samba for IRIX configuration.

Structure of This Document

This guide contains the following chapters:

- Chapter 1, "IRIS FailSafe Samba Overview," describes the IRIS FailSafe components that are included in the IRIS FailSafe Samba option. It also provides an overview on the steps that are required to add Samba as a highly available service to an IRIS FailSafe cluster.

- Chapter 2, "Configuring IRIS FailSafe for Samba," explains the procedure for adding Samba as a highly available service.

Related Documentation

For more information about Samba, see the *Samba for IRIX 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 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*

The IRIS FailSafe man 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_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

Release Note	Product
`cluster_admin`	Cluster administration services
`cluster_control`	Cluster node control services
`cluster_services`	Cluster services
`failsafe2`	IRIS FailSafe 2.1 release
`failsafe2_informix`	FailSafe /INFORMIX
`failsafe2_nfs`	FailSafe /NFS
`failsafe2_oracle`	FailSafe /Oracle
`failsafe2_samba`	FailSafe /Samba
`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:

Italics   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

Bold fixed-width type
> Commands and text that you are to type literally in response to shell and command prompts

#
> IRIX shell prompt for the super user (`root`)

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`

- 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 Samba 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 Samba 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 Samba to the highly available services that are failed over on an IRIS FailSafe cluster.

The major sections in this chapter are as follows:

- "IRIS FailSafe Samba File System Monitoring" on page 1
- "Overview of Configuring IRIS FailSafe for Samba" on page 3

IRIS FailSafe Samba File System Monitoring

The IRIS FailSafe Samba option adds NetBIOS name aliasing support to the IRIS FailSafe base software. The name aliasing monitoring scripts monitor the Samba NetBIOS names. 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 Samba resource action scripts do not interact with any FailSafe agent directly, except for the `ha_srmd` process.
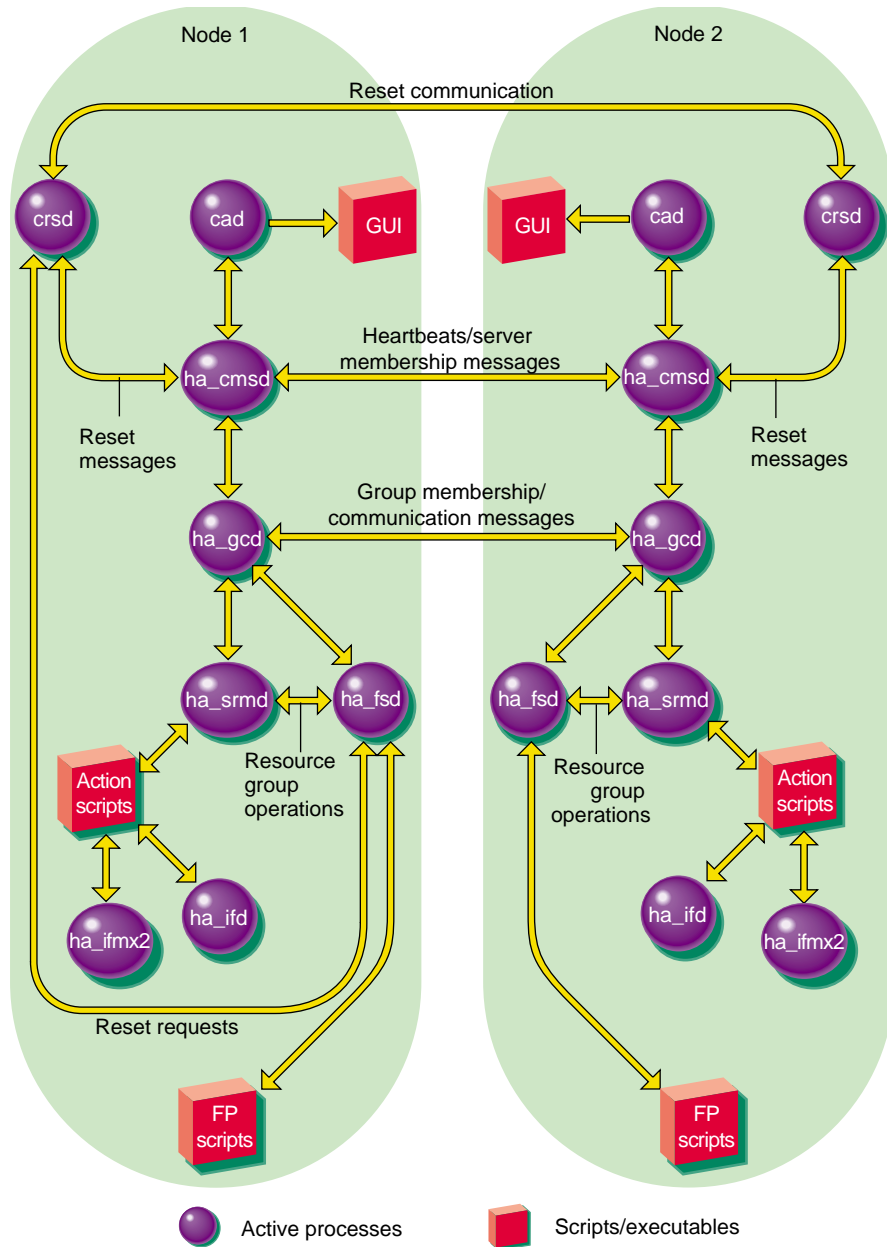
Figure 1-1   Message Paths in IRIS FailSafe in a two-node cluster

Overview of Configuring IRIS FailSafe for Samba

To configure an IRIS FailSafe cluster for failover of Samba, 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 the Samba for IRIX software on the cluster nodes and any additional software if required. See the section "Required Software" in Chapter 2.

3. Configure the Samba resources. See the section, "Adding Samba Resource Information to the Configuration Database (CDB)" in Chapter 2.

4. Add the Samba resource type, resource attribute, and resource information to the configuration database (CDB) file. See the section "Adding Samba Resource Information to the Configuration Database (CDB)" in Chapter 2.

5. Test the Samba failover. See the section "Testing the Samba Resource" in Chapter 2.

Configuring IRIS FailSafe for Samba

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

- "Required Software" on page 5

- "How does Samba failover work?" on page 7

- "What does the FailSafe Administrator need to set up?" on page 9

- "Adding Samba Resource Information to the Configuration Database (CDB)" on page 11

- "Samba Action Scripts" on page 15

- "Testing the Samba Resource" on page 16

Required Software

The required software for Samba failover is as follows:

- Samba software

  See the *Samba for IRIX Installation and Administration Guide* (part number 007-3965-004) for more information about Samba.

- 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.

- IRIS FailSafe Samba software

  - The *IRIS FailSafe Version 2 Samba Administrator's Guide* subsystem is `failsafe2_samba.books.book_AG`.

  - The man page subsystem is `failsafe2_samba.man.man`.

     –   The release notes subsystem is `failsafe2_samba.man.relnotes`.

     –   The base software subsystem is `failsafe2_samba.sw.base`.

Planning Samba Resources

The IRIS FailSafe 2.1 Samba option enables IRIS FailSafe 2.1 to provide failover protection for Samba resources designated by the NetBIOS name.

IRIS FailSafe operates on the concept of resources. A resource can be a filesystem, IP address or any entity that can be moved from one node to another when a problem (or scheduled downtime) occurs. A Samba resource is based on the NetBIOS name of a machine. Each node in a FailSafe cluster that runs Samba has a distinct NetBIOS name. (The NetBIOS name is usually the same as the host name of the machine, but in uppercase, and it must conform to the NetBIOS naming restrictions, which require that the name be in uppercase and contain 15 characters or less.) The ability of a Samba server to respond to service requests on this name is the resource that can be failed over between nodes.

In an IRIS FailSafe 2.1 cluster, one or more nodes can export Samba resources. If a node that exports Samba resources fails, another node provides backup service.

A Samba resource is defined by a NetBIOS name, and it depends on the following configuration parameters that must be defined within the FailSafe Samba resource definition:

Configuration Parameters for the Samba Resource

Table 2-1 lists the resource names and configuration parameters for the Samba resource.

Table 2-1   Samba Configuration Parameters

Resource Parameters	Comments
resource-name	Defines the name of the Samba resource; the NetBIOS name used to access the resource.
resource attribute	
monitor-level	Defines the monitoring action performed by the monitoring script. Monitor level 1 checks for the existence of the smbd and nmbd processes. Monitor level 2 checks for the correct responses from smbd and nmbd. Monitor level 3 checks for the existence of nmbd and for the correct responses from smbd.
resource dependencies	
filesystem	Name of the High Availability filesystem that this particular Samba resource is exporting. No checks are made that the paths exported in smb.conf are contained in this filesystem, but this lets the Samba resource be dependent on a FailSafe filesystem resource and be failed over with the filesystem.
IP_address	The High Availability IP (HA-IP) address to which the NetBIOS name belongs.

The procedure for configuring Samba resources for IRIS FailSafe 2.1 is described in the following section, "How does Samba failover work?"

How does Samba failover work?

In order to fail over the Samba resources, Samba must be able to register a secondary NetBIOS name alias. When FailSafe starts up, there must be no normal Samba /usr/samba/lib/smb.conf configuration file. Instead, the smb.conf file is constructed dynamically from a template file, as follows:

/usr/samba/lib/smb.conf.template

This template file must contain the standard Samba (`global`) setup for the workgroup, but it must not contain any of the `smb.conf options`, `netbios name`, `netbios aliases`, or `interfaces`. These lines in the `smb.conf` file are automatically added or modified as FailSafe starts the Samba resources and corresponding lines in the `smb.conf` template file are ignored. You should not define any nonglobal (`share`) names in this template file.

When Samba does not use FailSafe, it is given its primary NetBIOS name by adding the following line in `smb.conf`:

```
netbios name = name
```

NetBIOS aliases are added by using the following command in `smb.conf`:

```
netbios aliases = alias1 alias2
```

Multiple IP interfaces are supported in Samba by adding them to the following command in `smb.conf`:

```
interfaces = IP1/netmask1 IP2/netmask2
```

(where *IPx* is a dotted decimal IP address, and *netmaskx* is a dotted decimal netmask such as 255.255.255.0).

When Samba is run under FailSafe, the FailSafe daemons are responsible for maintaining these lines in the generated `smb.conf` file.

A Samba NetBIOS name must be defined for each Samba resource. For example, a Samba resource name of `server1` can be defined for a server that has an IP_address resource dependency with an HA-IP of 123.45.6.78 and a netmask of 255.255.255.0. FailSafe then generates the `smb.conf` file when the `server1` resource is brought online by first copying the contents of the `smb.conf.template` file and then appending the following lines to the generated `smb.conf` file:

```
netbios name = server1
interfaces = 123.45.6.78/255.255.255.0
include=/usr/samba/lib/smb.conf.%L
```

The `include=/usr/samba/lib/smb.conf.%L` line refers to a file that describes the *shares* that Samba serves out for a specific NetBIOS name. At run time, the `%L` is expanded into the NetBIOS name that the client is using to refer to the server The FailSafe administrator must create one of these files for each NetBIOS name. See the following examples.

This tells Samba to respond to connection requests on the given IP address under the name server1, and to register that name via broadcast on the given interface. A second FailSafe node with an HA-IP address of 123.45.6.90 and the same netmask, and a Samba resource name of server2 would have the following lines appended to its generated smb.conf file:

```
netbios name = server2
interfaces = 123.45.6.90/255.255.255.0
include=/usr/samba/lib/smb.conf.%L
```

In the previous example, Windows clients requesting a connection to server1 would access the resources on the first machine, and those clients requesting a connection to server2 would access the resources on the second.

Now, consider what happens when server server2 fails, and server server1 must take over the resource. Once server2 has failed and the HA-IP address 123.45.6.90 has been failed over to server1, the smb.conf file on server1 will be rewritten to contain the following lines:

```
netbios name = server1
netbios aliases = server2
interfaces = 123.45.6.78/255.255.255.0 123.45.6.90/255.255.255.0
include=/usr/samba/lib/smb.conf.%L
```

Following this, the nmbd daemon on server1 will be restarted (so it can register the additional NetBIOS name) and the smbd daemons on server1 will receive a HUP signal to cause them to reread the smb.conf file and offer the additional resources.

The Windows clients that now connect to the NetBIOS name server2 will be accessing server1 instead, although the change in physical servers will be transparent to them. (The change in physical servers will not be transparent to clients that had active Server Message Block (SMB) resources on server2. They will have to drop these resource handles and reconnect. The success of this procedure depends on what operation the Windows clients were using).

What does the FailSafe Administrator need to set up?

The FailSafe administrator has to set up the contents of the /usr/samba/lib/smb.conf.template file for each server in the cluster. (An example is provided with the plugin.) This file should be edited to contain the correct

workgroup for the Samba server and any site-specific (`global`) section changes that are needed. Copy this file to all nodes of the cluster that will be serving the Samba resource.

Note: The options `netbios name`, `netbios aliases`, and `interfaces` should not be added to this file, and there should not be any share points defined in this file.

Then, create the `/usr/samba/lib/smb.conf.`*servername* files (where *servername* is the NetBIOS name of the resources that are exported by the FailSafe cluster). Each of these files should contain a list of share definitions that each NetBIOS name will offer to clients. The path names in the share definitions should exist under the high-availability filesystem resource that the Samba resources depend on. In other words, the area of the filesystem being shared by Samba should be able to be failed over from one node to another node.

Copy these `/usr/samba/lib/smb.conf.`*servername* files to all the nodes of the cluster that will be serving the Samba resources, so that all nodes are able to share out all the resources.

For example, if there is a two-node cluster with the HA filesystem resource `/fs1` on NetBIOS name `server1`, and `/fs2` on NetBIOS name `server2`, the two files might look like the following:

```
Contents of /usr/samba/lib/smb.conf.server1:
[fs1]
        comment = failsafe filesystem1
        path = /fs1
        read only = no
        guest ok = yes
Contents of /usr/samba/lib/smb.conf.server2:
[fs2]
        comment = failsafe filesystem2
        path = /fs2
        read only = no
        guest ok = yes
```

If resource `server2` is failed over to `server1`, then as filesystem `/fs2` becomes mounted on `server1`, the Samba server on `server1` will start serving out the SMB share `fs2` under the NetBIOS name `server2`.

Next, create the FailSafe Samba resource. The following `cmgr` script creates a Samba resource:

```
define resource <<NetBIOS-name>> of resource_type Samba in cluster
<<cluster name>>
     set monitor-level to <<monitor level: either 1, 2, or 3>>
          add dependency <<IP_address resource name>> of type IP_address
     add dependency <<filesystem resource name>> of type filesystem
done
```

Finally, start the Samba FailSafe resources. The Windows clients should be able to see all the resources that you have defined under the server names.

Adding Samba Resource Information to the Configuration Database (CDB)

This section describes procedures that show you how to create a Samba resource type, resource, resource group, and how to test and monitor the Samba resource. These procedures assume that a CDB that does not include Samba has already been created, installed, and tested as described in the *IRIS FailSafe Version 2 Administrator's Guide*.

Creating a Samba Resource Type

To create a Samba resource type, this subsection assumes that you are already familiar with the concepts of resource types. The Samba resource type defines the following resource attribute. See Table 2-1 on page 7 for the definition of this resource attribute.

- *resource-name*, which defines the NetBIOS name of the resource
- *monitor-level*, which defines the type of monitoring performed

The Samba resource type is created at cluster creation time. If this automatic resource creation fails, the administrator must create the resource type before a Samba resource is created. The Samba resource type must be installed if you want to add a Samba resource to a cluster that was created before the Samba software was installed.

Example 2-1   Samba Installation Process

```
Inst> l *
Reading product descriptions ..  13%
Reading /var/inst/hist
Reading product descriptions ..  25%
```

```
 Setting distributions to .
 Reading product descriptions .. 100% Done.
   View:      distribution
   Status:    N=new, U=upgrade, S=same, D=downgrade
   Selection: i=install, r=remove, k=keep

   Subsystem Types [bdrpoc]: b=reBoot needed, d=Default, r=Required,
   p=patch, o=overlay, c=Client only

i N failsafe2_samba.books.FailSafe2Samba_AG [d] 444+ IRIS FailSafe 2.1
                                    Samba Administrator's Guide
i N failsafe2_samba.man.relnotes [d] 16+ FailSafe Samba Agent 2.1
                                    release notes
i N failsafe2_samba.sw.base [d]      84+ FailSafe Samba Agent 2.1 base
                                    software

Disk space summary (Kbytes):              /       /fs1

Current free space               2263692   4437968
- Selections net change              548+        0
- Temporary inst overhead           1400+        0
= Minimum free during install    2261736   4437968

Final projected free space       2263144   4437968

Inst> g
Pre-installation check ..    8%
Checking space requirements ..   16%
Pre-installation check completed
Installing/removing files ..  16%
Installing new versions of selected failsafe2_samba.books subsystems
Installing/removing files ..  62%
Installing new versions of selected failsafe2_samba.man subsystems
Installing/removing files ..  69%
Installing new versions of selected failsafe2_samba.sw subsystems
Installing/removing files ..  94%
Removing orphaned directories
Running exit-commands ..  94%
failsafe2_samba.books.FailSafe2Samba_AG: ( (chroot $rbase
/usr/sbin/insightAdmin -avfn C
/usr/share/Insight/library/SGI_bookshelves/SGI_Admin>>$rbase/var/share/
Insight/install.insight.log)2>>$rbase/var/share/Insight/install.insight
.log )
Running exit commands ..  99%
Checking dependencies .. 100% Done.
```

```
Installations and removals were successful.
You may continue with installations or quit now.

Inst> quit
```

Creating the `Samba` **Resource Type**

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

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

- Run cluster manager (`cmgr`) and install the resource type, as follows:

  ```
  cmgr> show resource_types installed

  NFS
  template
  Netscape_web
  statd
  statd_unlimited
  Oracle_DB
  MAC_address
  IP_address
  INFORMIX_DB
  filesystem
  volume

  cmgr> install resource_type Samba in cluster eagan

  cmgr>
  ```

- Instead of using cluster manager (`cmgr`) to create the Samba resource type in cluster `eagan`, you can use the IRIS FailSafe 2.1 **Load Resource Type** GUI task to load the Samba resource type in the cluster `eagan`.

Creating a Samba Resource

After you have defined the resource type, the administrator must define the Samba resources based on the resource type. Each resource requires a unique resource name (for example, the Samba resource type is the Samba instance name). Then, the administrator must supply the resource parameters. To create the resource, either use the cluster manager (`cmgr`), the `cmgr-create-resource-Samba` scripts, or the GUI.

Example 2-2   Creating a Samba Resource Using `cmgr`

```
root@hans1:~ # /usr/cluster/bin/cluster_mgr
Welcome to IRIS FailSafe Cluster Manager Command-Line Interface

cmgr> define resource server1 of resource_type Samba in cluster
interop-ha
Enter commands, when finished enter either "done" or "cancel"

Type specific attributes to create with set command:

Type Specific Attributes - 1: monitor-level

Resource type dependencies to add:

Resource Dependency Type - 1: IP_address
Resource Dependency Type - 2: filesystem

resource /server1 ? set monitor-level to 2
resource /server1 ? add dependency /fs1 of type filesystem
resource /server1 ? add dependency 163.154.48.34 of type IP_address
resource /server1 ? done
Successfully modified resource server1

cmgr> show resource server1 of resource_type Samba

monitor-level: 2

Resource dependencies
IP_address 163.154.48.34
filesystem /fs1

cmgr> exit
```

Creating a Samba 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`).

To create an effective resource group, you must include all of the resources that the Samba resource is dependent on, such as filesystems, volumes, and IP addresses. The following example shows the creation of a typical resource group:

Example 2-3   Creating a Samba Resource Group

```
root@hans1:~ # /usr/cluster/bin/cluster_mgr
Welcome to IRIS FailSafe Cluster Manager Command-Line Interface

cmgr> define resource_group hans1 in cluster interop-ha
Enter commands, when finished enter either "done" or "cancel"

resource_group hans1 ? set failover_policy to hans1-primary
resource_group hans1 ? add resource 163.154.48.34 of resource_type
IP_address
resource_group hans1 ? add resource /fs1 of resource_type filesystem
resource_group hans1 ? add resource vol1 of resource_type volume
resource_group hans1 ? add resource server1 of resource_type Samba
resource_group hans1 ? done
Successfully defined resource group hans1

cmgr> show resource_group hans1 in cluster interop-ha

Resource Group: hans1
        Cluster: interop-ha
        Failover Policy: hans1-primary

Resources:
        server1 (type: Samba)
        163.154.48.34 (type: IP_address)
        /fs1 (type: filesystem)
        vol1 (type: volume)

cmgr> exit
```

Samba Action Scripts

There are several scripts that FailSafe uses to start, stop, and monitor the Samba servers on a node. These are found in the directory /var/cluster/ha/resource_types/Samba and are the following scripts:

start   Starts all Samba resources active on this node

stop   Stops all Samba resources active on this node

monitor   Extended test to see if Samba is running (see Table 2-1 on page 7 for a description of the *monitor-level* resource attribute)

`exclusive`	Tests to see if a resource list exists on this node
`restart`	Attempts to restart crashed Samba daemons

In addition to these scripts, the following scripts are used internally by the Samba FailSafe option:

`add_name`	Adds a NetBIOS name to a generated `smb.conf` file
`remove_name`	Removes a NetBIOS name from a generated `smb.conf` file
`restart_samba`	
	Restarts `smbd` and `nmbd`

Testing the Samba Resource

To ensure that the Samba resource has been correctly configured, you can test individual actions by executing the scripts. This assumes that all of the dependent resources have already been started. The easiest way to test the new Samba resources is to create a resource group that contains all of the dependent resouces except the new Samba resource. You can then start this resource group and test the individual Samba action scripts.

Each script, located at `/var/cluster/ha/resource_types/Samba`, requires two arguments, an input file and an output file. The content of the input file is the resource name. 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*.

In the following example, you can test the `start` script by starting the Samba resource with the resource name `samba8`.

```
$ cd /var/cluster/ha/resource_types/samba
$ echo "samba8" > /tmp/ipfile
$ ./start /tmp/ipfile /tmp/opfile
```

This should start the `Samba8` instance, named by the Samba resource `samba8`.

To view the individual script actions, you must edit the script and add "`set -x`" to the `action` function.

In the following example, to see what the start_samba() function does, edit the start script and add "set -x" in the start function.

```
#
# Start the resource on the local machine.
# Return HA_SUCCESS if the resource has been successfully started on
# the local machine and HA_CMD_FAILED otherwise. The resource name is
# the samba instance id.
#
start_samba()
{
set -x
# for all samba resources configured
...
```

You are now ready to run the script. In this example, /dev/null is the name of the output file.

```
cm2> echo "/disk1" > /tmp/ipfile
cm2> /var/cluster/ha/resource_types/Samba/start /tmp/ipfile /dev/null
cm2> echo $?                                          (assumption: shell is ksh)
0
```

Testing the start Script

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

1. To test the start script, you must start the Samba resource on the node, as in the following example:

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

   where *resource-name* is the name of the Samba resource that will be started.

2. Check to ensure that the resource has started by entering the following command:

   ```
   # /usr/samba/bin/smbclient -L server -U%
   ```

where *server* is the NetBIOS name. If the script is working correctly, you will receive a list of shares that the NetBIOS name is offering.

Testing the `stop` Script

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

1. To test the `stop` script, you must start the Samba resource on the node (if it is not already running), as in the following example:

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

   where *resource-name* is the name of the Samba resource that will be started.

2. To stop the Samba resource, enter the following commands:

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

   where *resource-name* is the name of the Samba resource that will be stopped.

3. Check to ensure that the resource has shut down by entering the following command:

   ```
   # ps -ef | grep mbd
   ```

   The processes `smbd` or `nmbd` should not be running.

Testing the `monitor` Script

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

1. To monitor the `Samba` resource, enter the following commands:

   ```
   # echo "resource-name" > /tmp/ipfile

   # /var/cluster/ha/resource_types/Samba/monitor /tmp/ipfile /dev/null
   ```

2. The output is logged to the script file located at `/var/cluster/ha/log/script_nodename`. Check this file for the script log messages.

Testing the `restart` Script

You can use the following procedure to test the restart script.

1. To restart the Samba resource, enter the following commands:

    # **echo "***resource-name***" > /tmp/ipfile**

    # **/var/cluster/ha/resource_types/Samba/restart /tmp/ipfile /dev/null**

2. Check to ensure that the resource has restarted by entering the following command:

    # /usr/samba/bin/smbclient -L *server* -U%

    where *server* is the NetBIOS name. If the script is working correctly, you will receive a list of shares that the NetBIOS name is offering.

Testing the `exclusive` Script

You can use the following procedure to test the exclusive script.

1. To test the exclusive script, enter the following commands:

    # **echo "***resource-name***" > /tmp/ipfile**

    # **/var/cluster/ha/resource_types/Samba/exclusive /tmp/ipfile /dev/null**

2. The output is logged to the script file located at /var/cluster/ha/log/script_*nodename*. Check this file for the script log messages.

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 online resource_group Samba in cluster eagan

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

cmgr> admin offline resource_group Samba in cluster eagan

Index

S

Samba action scripts, 15
Samba failover
  how it works, 7
  required software, 5
Samba FailSafe
  file system monitoring, 1
Samba FailSafe overview, 1
Samba installation process, 11
Samba resource
  configuration parameters, 7
Samba resource configuration parameters, 7
Samba resource group
  creating, 14
Samba resource type
  creating, 11
Samba resources
  planning, 6
smb.conf file, 7

T

testing the exclusive script, 19
testing the failover policy, 19
testing the restart script, 19
testing the start script, 17