

Configuring the SFS Environment with the Menu System [3]

This chapter describes how to configure Cray Research systems to use the SFS feature and how to create shared file systems by using the UNICOS installation and configuration menu system. If you want to configure your SFS software manually, see Chapter 4, page 25.

The menu system does not support GigaRing-based I/O configurations.

For this chapter, it is assumed that your Cray Research systems are up and running, and that they support the menu system.

When initializing the SFS environment with the menu system, you perform the following tasks:

1. Configure the SFS device nodes on one system in the SFS cluster
2. Indicate which hosts systems are connected to the SFS device nodes
3. Create the SFS device nodes and the SFS configuration file you have defined
4. Define the disk device nodes for the SFS cluster
5. Define the file systems for the SFS cluster
6. Copy the configuration information from the system where it has been defined to each of the systems you want to add to your SFS cluster
7. Correct machine-specific information for each of the systems that receives a copy of the configuration information

3.1 Before You Begin

Before you run the UNICOS installation and configuration menu system to configure your SFS environment, you should be sure that you have all the information at hand that the menu system will request.

For each SFS semaphore device that you configure, you will need to provide the following information:

/dev/hdd/smp I/O cluster (IOC) number	___
/dev/hdd/smp IOP number	___
/dev/hdd/smp I/O channel number	___
/dev/hdd/smp Ifield	___
/dev/smp port number	___
/dev/dsk/slr minor device number	___
/dev/hdd/slr minor device number	___
/dev/hdd/slr HIPPI disk device type	___
/dev/hdd/slr I/O cluster (IOC) number	___
/dev/hdd/slr IOP number	___
/dev/hdd/slr IOP Channel number	___
/dev/hdd/slr starting sector number	___
/dev/hdd/slr number of sectors	___
/dev/hdd/slr HIPPI facility address	___
/dev/hdd/slr HIPPI Ifield	___
SFS Arbiter	___

You will also need to provide an associated host name for each SFS semaphore device that you define, on a screen in the following format:

Host name	___ Connected Arbiter
-----------	-----------------------

For each disk that you will use in an SFS configuration, you will be asked for the following information:

Device name	___
Device type	___
I/O cluster (IOC) number	___
IOP number	___
IOP Channel number	___
Alternate path?	___
Alternate I/O cluster (IOC) number	___
Alternate IOP number	___
Alternate IOP channel number	___
Unit	___
Size in blocks	___
HIPPI Facility (Unit bits 0-8)	___
HIPPI Raid Partition (Unit bits 9-15)	___
HIPPI disk Ifield	___
SSD spare memory chip configuration file	___

The device names are limited to 14 characters. ND devices can be configured as HD-16, HD-32, and HD-64 devices. HD-16 is used for HIPPI disk devices with a 16 kbyte sector size, HD-32 is used for HIPPI disk devices with a 32 kbyte sector size, and HD-64 is used for HIPPI disk devices with a 64 kbyte sector size.

For each physical slice on a disk, you will be asked for the following information:

Slice name	_____
Physical device name	_____
Minor device number	_____
Starting number	_____
Size of slice	_____
Slice unit of measure is depicted in	_____

For each logical device you are configuring, you will be asked for the following information:

Logical device name	_____
Logical minor number	_____
Member device type	_____
Member device name(s)	_____

For each SFS file system you are defining, you will be asked for the following information:

Logical disk device name	_____
File system mount point	_____
File system type	SFS
Backup frequency (in days)	_____
File system check pass number	
Read only option	
Quota control file	
User defined option	
Mount at multi-user startup?	
Mount point owner	_____
Mount point group	_____
Mount point mode	_____

Note that you need to specify SFS as the file system type.

3.2 Initializing the SFS environment on the first system

To initialize the SFS environment on the first system in the SFS cluster, follow this procedure:

1. Verify that the existing disk and file system configuration information in the menu system is valid and complete. The SFS configuration information is appended to the existing system information.
2. Start the UNICOS installation and configuration menu on the first system in the cluster from the `/etc/install` directory using the special `./configsfs` program rather than the standard `./install` program.

Note: You cannot run the `configsfs` program and the standard `install` program at the same time, unless you are running one of the programs in read-only mode. If you are running one of these utilities to configure the system, and you want to run the other to configure the system, you must first exit the utility you are currently running.

When you run the `configsfs` program, the following screen appears. The default values of the SFS configuration screens are the values entered at the last invocation of the menu system.

```
Shared File System (SFS) Configuration

A-> Configure SFS semaphore device nodes ==>
    SFS hosts ==>
    Create SFS device nodes ...

    Disk configuration ==>
    File system (fstab) configuration ==>
    Encapsulate SFS device information ...

    Merge SFS devices to Disk Configuration ...
    Remove SFS devices from Disk Configuration ...

    Remove SFS device nodes ...
    Import root mount point

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI
```

3.2.1 Configuring SFS device nodes

To configure the SFS semaphore device nodes, perform the following steps:

1. Go to the Configure SFS semaphore device nodes menu.

```
Shared File System (SFS) Configuration
->Configure SFS semaphore Device Nodes
```

A form screen appears indicating which SFS arbiters have been defined for this system.

```

                                Configure SFS Semaphore Device Nodes

hdd/smp IOC  hdd/smp IOP  hdd/smp channel  hdd/smp Ifield  smp port  dsk/s
-----
Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

File "cfdb/sfsnode.cfg" is empty. Use the 'n' key to add a record.
```

2. To add a new arbiter, press the n Return keys. The following screen appears, showing a new record which you must edit.

```

                                Configure SFS Semaphore Device Nodes

S-> /dev/hdd/smp I/O cluster (IOC) number          0
    /dev/hdd/smp IOP number                        0
    /dev/hdd/smp I/O channel number                030
    /dev/hdd/smp Ifield                            0
    /dev/smp port number                           0
    /dev/dsk/slr minor device number                1
    /dev/hdd/slr minor device number                1
    /dev/hdd/slr HIPPI disk device type            HD16
    /dev/hdd/slr I/O cluster (IOC) number          0
    /dev/hdd/slr IOP number                        0
    /dev/hdd/slr IOP Channel number                030
    /dev/hdd/slr starting sector number            0
    /dev/hdd/slr number of sectors                 0
    /dev/hdd/slr HIPPI facility address            2
    /dev/hdd/slr HIPPI Ifield                      0
    SFS Arbiter                                    __change__

Keys:  ^? Commands  H Help    Q Quit    V ViewDoc  W WhereAmI

                                Record 1 of 1
    
```

This record view requests machine-specific information that is used by the system to define the `/dev/hdd/smp`, `/dev/smp`, and the `/dev/hdd/slr` device nodes. The screen first appears with default values, which you must modify for your system.

The `/dev/hdd/smp` node is the HIPPI-disk node that contains the low-level I/O path and the HIPPI I-field that are used for physically accessing the SFS arbiter.

The `/dev/smp` node is the low-level interface to the semaphore device.

The `/dev/hdd/slr` node defines the area of the shared media pool that you will use for the Shared Lock Region, an area that is used by UNICOS SFS to record common information such as semaphore assignment and system state.

SFS Arbiter indicates the user-provided name of the SFS arbitration device you are defining.

Note: The SFS Arbiter entry must be a number. For the first arbitration device, set this field to 0; for the second arbitration device, set this field to 1, and so forth.

3. Enter the requested information for the device node you are configuring.

Note: You must modify any field that appears as `__change__` to create a functional SFS cluster.

4. Go to the Shared File System Hosts record screen.

```
Shared File System (SFS) Configuration
->SFS Hosts
```

The Shared File System Hosts form appears as follows:

```

                                     Shared File System Hosts

host name  arbiter name
-----
-----

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

File "cfdb/sfshost.cfg" is empty. Use the 'n' key to add a record.
```

5. To enter the host name, press the n Return keys, and the Shared File System Hosts record appears.

```

                                     Shared File System Hosts

S-> Host name                               __change__
     Connected Arbiter                       __change__

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

                                     Record 1 of 1
```

6. Enter the requested information. This information will be used to create the `/etc/config/sfs` file.

Note: You must modify any field that appears as `__change__` to create a functional SFS cluster.

For each SFS arbiter selection that you defined with the `Configure SFS Semaphore Device Nodes` menu, you must define at least one associated host with the `Shared File System Hosts` menu.

The name of the `Connected Arbiter` must match one of the numbers you defined as the `SFS Arbiter` in the `Configure SFS Semaphore Device Nodes` menu. If the host is connected to more than one SFS arbiter, then the list of names or numbers should be comma-separated.

7. Return to the main `Shared File System (SFS) Configuration` menu and execute `Create SFS device nodes`. This creates the SFS nodes that you have configured and the `/etc/config/sfs` file that contains the names of the SFS hosts and arbiters.

The following is a sample of the output that selecting this action yields.

```
Corrected output of action SFS device nodes actions:

Preparing to create the needed SFS device nodes

Creating /etc/config/sfs hostfile
Done creating the host section of /etc/config/sfs hostfile
Preparing to create arbiter #0 SFS device nodes

Verifying /dev/hdd/smp-0's existence
Node name /dev/dsk/smp-0 being created
Verifying /dev/smp-0's existence
Node name /dev/smp-0 being created
Verifying /dev/hdd/slr-0's existence
Node name /dev/dsk/slr-0 being created
Verifying /dev/dsk/slr-0's existence
Node name /dev/dsk/slr-0 being created
Verifying /dev/sfs-0's existence
Node name /dev/sfs-0 being created
Verifying /dev/smnt-0's existence
Node name /dev/smnt-0 being created

Done making the SFS device nodes
Finished creating /etc/config/sfs
```


8. To verify that you have configured your device nodes correctly, you can run the `sfs_start` command from a different window. This will yield errors if there are problems with the device nodes.

For information on the `sfs_start` command, see Chapter 5, page 33.

3.2.2 Configuring disk devices

To configure the disk devices for an SFS cluster, perform the following steps. This procedure is identical to the procedure for configuring disk devices for a standard file system.

1. Go to the Disk Configuration menu.

```
Shared File System (SFS) Configuration
->Disk Configuration
```

The Disk Configuration menu appears as follows:

```

                                     Disk Configuration
M-> Physical devices ==>
    Physical device slices ==>
    Logical devices (/dev/dsk entries) ==>
    Mirrored devices (/dev/mdd entries) ==>
    Striped devices (/dev/sdd entries) ==>
    Logical device cache ==>
    Verify the disk configuration ...
    Review the disk configuration verification ...
    Dry run the disk configuration ...
    Review the disk configuration dry run ...
    Update disk device nodes on activation?      YES
    Allow active root node updates?            NO
    Import the disk configuration ...
    Activate the disk configuration ...

Keys:  ^? Commands  H Help    Q Quit    V ViewDoc  W WhereAmI
```

The process of defining disk configuration for shared file systems is identical to the process of defining disk configuration for file systems that will not be shared.

First define the physical devices, then the physical device slices, then the logical devices.



Caution: The device names you define for shared file systems must be the same for all of the systems in the SFS cluster.

2. To define physical devices, go to the Physical Devices menu.

```
Shared File System (SFS) Configuration
->Disk Configuration
  ->Physical Devices
```

3. To enter a new physical device, press the n Return keys and the following screen appears.

```
Physical Devices

S-> Device name                0130
   Device type                 DD60
   I/O cluster (IOC) number    0
   IOP number                  0
   IOP Channel number         030
   Alternate path?            NO
   Alternate I/O cluster (IOC) number
   Alternate IOP number
   Alternate IOP channel number
   Unit                       0
   Size in blocks
   HIPPI Facility (Unit bits 0-8)
   HIPPI Raid Partition (Unit bits 9-15)
   HIPPI disk Ifield
   SSD spare memory chip configuration file
   Note: The device names are limited to 14 characters.

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

Record 1 of 1
```

4. Enter the information for each device you are configuring.

When defining physical devices on CRAY EL series systems or CRAY J90 series systems, you do not need to define the I/O cluster number (IOC) or I/O processor number (IOP).

5. To define the physical device slices, go to the Physical Device Slices menu.

```
Shared File System (SFS) Configuration
->Disk Configuration
  ->Physical Device Slices
```

6. To enter a new physical device slice, press the n Return keys and the following screen appears.

```
Physical Device Slices

S-> Slice name           slicel
    Physical device name 0130
    Minor device number  1
    Starting number      0
    Size of slice        0
    Slice unit of measure is depicted in sectors

Notes: Any SSD or RAM slices must have blocks as the unit.
       HIPPI disks are recommended to use blocks.
       The slice names are limited to 14 characters.

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

Record 1 of 1
```

7. Enter the information for each device slice you are configuring.

When defining physical device slices for /dev/hdd devices, you must ensure that the minor numbers are unique.

8. To define the logical devices, go to the Logical Device (/dev/dsk) Configuration menu.

```
Shared File System (SFS) Configuration
->Disk Configuration
  ->Logical Device (/dev/dsk) Configuration
```

9. To enter a new logical device, press the n Return keys and the following screen appears.

```
Logical Device (/dev/dsk) Configuration

S-> Logical device name          ldd1
    Logical minor number        1
    Member device type          hdd
    Member device name          slice1

Note:  The logical device names are limited to 14 characters.

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

Record 1 of 1
```

10. Enter the information for each device you are configuring.

Select the Mirrored devices (/dev/mdd entries), Striped devices (/dev/sdd entries), and Logical device cache entries, as necessary, and provide the requested information.

11. Before leaving the Disk Configuration menu, make sure that the Update disk device nodes on activation is set to YES; you will not be able to complete the process of defining your SFS configuration otherwise.

12. Select the menu options to perform the following tasks, reviewing them as desired:

- Verify the disk configuration
- Dry-run the disk configuration
- Activate the disk configuration.

When you perform any of these tasks, you may receive the warning gap/overlap <before slice sfs>. This could indicate that the menu system has reserved space on the disk because the SFS feature requires a portion of the HIPPI device for the Shared Lock Region. If you receive additional gap/overlap warnings, however, you may need to check your configuration.

13. After configuring the disk devices, you can run the ddstat(8) command to verify that the configuration was successful.

3.2.3 Defining an SFS file system

You define an SFS file system just as you define a file system that is not to be shared, using the same Standard File Systems Configuration menu.

1. Select the File System (fstab) Configuration menu from the main Shared File System (SFS) Configuration menu.

```
Shared File System (SFS) Configuration
->File System (fstab) Configuration
```

The File System (fstab) Configuration menu appears as follows:

```

                                File System (fstab) Configuration

M-> Standard file systems ==>
    NFS file systems ==>
    Shared file system ==>
    Import fstab configuration ...
    Activate fstab configuration ...

Keys:  ^? Commands  H Help    Q Quit    V ViewDoc  W WhereAmI
```

2. From this menu, select the Standard file systems entry and the following screen appears.

```

                                Standard File System Configuration

Device Name          Mount Point          FsType  Freq  Pas
-----
Keys:  ^? Commands  H Help    Q Quit    V ViewDoc  W WhereAmI
```

3. To add a new file system, press the n Return keys and the following screen appears.

```
Standard File System Configuration

S-> Logical disk device name           /dev/dsk/fs
    File system mount point           /mnt
    File system type                   SFS
    Backup frequency (in days)        1
    File system check pass number      2
    Read only option
    Quota control file
    User defined option
    Mount at multi-user startup?      NO
    Mount point owner                  root
    Mount point group                  root
    Mount point mode                   0755

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

Record 1 of 1
```

4. Enter your file system information just as you would for a file system that is not to be shared, with the exception that you use SFS as the File system type instead of NC1FS. A file system name must be the same on all the systems in the SFS cluster.

5. Go to the File System (fstab) Configuration menu:

```
Shared File System (SFS) Configuration
->File System (fstab) Configuration
```

6. To create the /etc/fstab file, select Activate fstab configuration.

3.2.4 Checking the SFS configuration

To check the SFS configuration, move to another window or exit the menu system and perform the following procedure:

1. Run the `mkfs(8)` command on the file systems you have created, being sure to use `-q` and the `-s #` options. The `-s #` option indicates that this is an SFS file system and specifies the number of semaphores to assign to the file system. For information on using the `mkfs` command to create shared file systems, see Section 4.2.3, page 30.



Warning: Do not run the `mkfs(8)` command on a mounted file system.

2. Bring up the SFS system to see if the environment has been successfully initialized by executing the `sfs_start` command. For information on the `sfs_start` command, see Chapter 5, page 33.

3.2.5 Encapsulating SFS configuration information

In order to copy the SFS configuration information to the other systems in the cluster, you must encapsulate the configuration information into a file. Use the following procedure:

1. Return to the menu system, go to the main Shared File System (SFS) Configuration menu, and select Encapsulate SFS device information. This creates the `sfs_encap` file in the `/etc/install` directory. This file contains the encapsulated device information that the other systems in the SFS cluster will use.
2. Exit the menu system.

3.3 Initializing the SFS environment on additional systems

For each additional system in the SFS cluster, perform the following steps.

1. Copy the `sfs_encap` file from the system on which you created it to the `/etc/install` directory of the other systems in the SFS cluster with the `ftp` or `cp` command. Note that you must change the permissions on the `sfs_encap` file before copying it.
2. Log in to a different system in the cluster and start the `configsfs` program from the `/etc/install` directory, which contains `sfs_encap`.
3. From the main Shared File System (SFS) Configuration menu, perform the action Merge SFS devices to Disk Configuration.
4. Go to the Configure SFS Semaphore Device Nodes menu.

```
Shared File System (SFS) Configuration
->Configure SFS semaphore Device Nodes
```

5. Provide the machine-specific information for the `/dev/hdd/smp`, `/dev/smp`, and `/dev/hdd/slr` device nodes that need to be changed for the current system.



Caution: When defining the physical device driver port for /dev/smp, take care to ensure that the port number is unique for every machine in the cluster. The valid port number range is 0 to 63. The menu system itself does not ensure that each /dev/smp port number is unique for each machine in the cluster.

6. Go to the Physical Devices menu.

```
Shared File System (SFS) Configuration
->Disk Configuration
  ->Physical Devices
```

The menu appears as follows:

```

                                Physical Devices

Name          Type  Ioc          Iop          Chn          Alt?  AltIoc  AltI
-----
0230.0        DD62  0            2            030          NO
0334.16.0.13  HD16  0            3            034          NO
0334.16.0.21  HD64  0            3            034          NO
E-> 0334.18.0.13  HD16  __change__  __change__  __change__  NO

Keys:  ^? Commands  H Help  Q Quit  V ViewDoc  W WhereAmI

Use the + and - keys to access other pages within this menu.
Record 4 of 4 Page 1 of 1
    
```

7. Change the information in the categories indicated with the text string __change__. The following fields will need to be changed:

Column numbers 3, 4, 5, 12, 13, and 14 must have the correct information for each system's access to the SFS device. The specifications to be changed for SFS devices are:

- I/O cluster (IOC) number
- IOP number
- IOP Channel number
- HIPPI Facility (Unit bits 0-8)
- HIPPI Raid Partition (Unit bits 9-15)

(If you will not be using this partition, leave the default values as they first appear or the system will not allow activation.)

- HIPPI disk switch address

8. Go back to the Disk Configuration menu.

```
Shared File System (SFS) Configuration
-> Disk Configuration
```

9. Make sure that the Update disk device nodes on activation is set to YES.

10. Select the menu options to perform the following tasks:

- Verify the disk configuration
- Dry-run the disk configuration
- Activate the disk configuration.

When you perform any of these tasks, you may again receive the warning gap/overlap <before slice sfs>. This indicates the space that the menu system has reserved on the HIPPI device for the Shared Lock Region.

11. After configuring the disk devices, you can run the `ddstat(8)` command to verify that the configuration was successful.

12. Go to the File System (fstab) Configuration menu.

```
Shared File System (SFS) Configuration
->File System (fstab) Configuration
```

13. Select Activate fstab configuration. This creates the `/etc/fstab` file.

Remember to repeat steps 1 through 13 for each additional system in the SFS cluster.

