

SFS Startup with the `sfs_start` Command [5]

After defining the UNICOS SFS environment and creating the file systems that the Cray Research systems will share, you can run the `sfs_start` command to initialize the UNICOS SFS environment. The `sfs_start` command runs a series of tests to ensure that the system is properly defined, and initiates the `sfsd` and `codeblue` system daemons. The `sfs_start` command also runs file system checks on the SFS file systems in the `fstab` file and mounts the shared file systems.

5.1 `sfs_start` program logic

The `sfs_start` daemon performs the following operations:

1. Ensures that the Cray Research system can read the `/etc/config/sfs` file, which indicates which systems are part of the SFS cluster.
2. Ensures that the current host is listed in the `/etc/config/sfs` file.
3. Ensures that the `/etc/services` file includes an entry for `sfs_config`.
4. Ensures that the various SFS commands and daemons are executable.
5. Ensures that the `/dev/smp` and `/dev/sfs` nodes have been defined as character special devices.
6. Ensures that the `slr` and `smnt` slices have been defined as a block special or character special device.
7. Compares the local configuration (`/etc/config/sfs`) with those of all other hosts in the SFS cluster.
8. Compares the device nodes of the local host with those of all other hosts in the SFS cluster.
9. Executes the `sema` command to test semaphore 0.
10. Sets semaphore 5, the SFS start semaphore, to ensure that only one system at a time is running `sfs_start`.
11. Executes `sfsping` for each member of the SFS cluster identified in `/etc/config/sfs` to determine which systems in the cluster are up (live) and which are not up (dead).

12. Prints a list of live systems.
13. Checks whether `/etc/sfsd` is currently running. The `/etc/sfsd` command initializes the UNICOS SFS environment. For information about this command, see Section 5.2, page 35.
14. If the `/etc/sfsd` command is not running, checks whether this is the first system up in the SFS cluster. If this is the first system up, then `sfs_start` executes `sfsd` with the `-F` option to force semaphore initialization. If this is not the first system up, it executes `sfsd` without the `-F` option, which would clear any semaphores already in use.
15. Checks whether recovery is currently taking place; if so, waits until recovery is finished before proceeding.
16. If other systems in the SFS cluster are up, checks whether the `/dev/smp`, `/dev/sfs`, `/dev/dsk/slr`, and `/dev/smnt` devices are compatible. For each of these devices, the major and minor device numbers must be identical on every system.
17. Retrieves the list of SFS file systems from the `/etc/fstab` file.
18. If this is the only live system, clears the recovery semaphore (semaphore 2), if it is set, and clears the Ports Needing Recovery (PNR) mask by executing `/etc/sfspnr -c` (clear PNR mask), if it is not zero.
19. Performs the following steps for each SFS file system:
 - Compares the SFS file system node on this host with the SFS file system node on all systems that are up. If the nodes do not match, the file system will not be mounted and an error is reported.
 - Checks whether the SFS file system is currently mounted on other up systems in the SFS cluster.
 - If the SFS file system is currently mounted, executes `/etc/sfsck -rv` on the file system. This command checks the file system without requiring that the file system be unmounted.
 - If the SFS file system is not currently mounted, executes `/etc/sfsck -uS` on the file system.
 - Mounts the SFS file system, creating the mount point for the file system if necessary.
20. Starts `/etc/codeblue -d /etc/sfsprefix`. This monitors the status of other systems in the SFS cluster.
21. Clears semaphore 5, the SFS start semaphore.
22. Prints that SFS initialization is complete.

5.2 The `sfsd` command

For every system in the UNICOS SFS environment, the `sfs_start` program executes the `sfsd(8)` command to initialize the SFS environment. The first time `sfsd` is run, it uses the `-F` option to force the initialization of the semaphores. Using this option after the first Cray Research system has been brought up, however, will clear any semaphores in use, which will likely cause systems currently using the UNICOS SFS feature to panic.

The `sfsd` command provides multiple services as a system daemon. It validates the site license, then initializes the SFS devices. It also listens for queries from other SFS client programs such as `sfsping`, `sfsstat`, and `sfsddstat`, described on Section 5.3, page 36.

After `sfs_start` executes `sfsd`, messages appear on the system console indicating the status of the UNICOS SFS environment. The following example shows a typical message display.

```
00:55:43(GMT) uts/cl/io/smp-ipi3.c: INFO          IPI-3 SMP device initialized (1024 semas,
        64 ports)
00:55:43(GMT) uts/cl/io/smp.c: INFO          smpopen: SMP device supports 1024 semas, 1024 will
        be used
00:55:43(GMT) uts/cl/io/smp.c: INFO          smpopen: smp[0] initialized:
        type=04 SMP-IPI3 phys=1024 start=0 len=1024 port=2
00:55:45 uts/cl/io/slr.c-06: INFO          slr_open: lock region area is </dev/dsk/slr>
00:55:45 uts/cl/io/slr.c-36: INFO          slr_open: SLR is not mirrored
slr_open: slr_size_size = 128
00:55:45 uts/cl/io/slr.c-68: WARNING       slr_init: SLR area 'init' semaphore not set
esd_set_slr_mirror_mode: DEBUG: called with newmask 07
00:55:45 uts/cl/io/slr.c-73: WARNING       slr_init: SLR area initialized
00:55:45 uts/cl/io/slr.c-24: INFO          slr_open: SLR open completed
00:55:45 uts/cl/io/slr.c-86: INFO          slr_init: port 2 setting Validity sema 15
00:55:45 uts/cl/io/esd_hb.c-25: INFO       ESD-port 2, name <ice>, num <2> logged on
SMP Heartbeat started.
```

If the system detects an error, a message indicating the nature of the error appears on the window in which `sfsd` was executed. For example, if you are not licensed to run the UNICOS SFS feature, the following message appears:

```
craylm: checkout of "sfs" failed.
lm_checkout: no such feature exists
This system is not licensed to run SFS.
Please contact your system administrator, or
Cray Research Software Technical Support
to obtain a Shared File System license.
```

Once `sfsd` is running, `sfs_start` can check and mount SFS file systems.

5.3 Additional SFS commands

In addition to `sfsd`, there are three SFS commands the `sfs_start` program uses when a Cray Research system is being booted and brought up to multiuser mode. These commands are as follows:

<code>/etc/sfsping</code>	Used to ping another system, to determine if it is running UNICOS SFS software.
<code>/etc/sfsstat</code>	Used to obtain <code>stat</code> information about a device node on another Cray Research system in an SFS cluster.
<code>/etc/sfsddstat</code>	Used to obtain <code>ddstat</code> information about a disk configuration on another Cray Research system in an SFS cluster.

5.4 Sample `sfs_start` output

The following shows typical output from the execution of the `sfs_start` command.

```
# /etc/sfs_start
sfs_start - INFO: Beginning Shared File System startup
sfs_start - INFO: Locking SFS init sema 5
sfs_start - INFO: SFS init sema 5 Locked.
sfs_start - INFO: Checking to see if cluster member hosta is running SFS...
sfs_start - INFO: Checking IPI-3 SMP port numbers... hostb
sfs_start - INFO: Clearing SMP device during SFS initialization
sfs_start - INFO: SFS daemon (sfsd) started.
sfs_start - INFO: Mounting Shared File Systems...
sfs_start - INFO: SFS Cluster monitor started
sfs_start - INFO: Clearing SFS init sema 5
sfs_start - INFO: SFS initialization complete.
```