After you have run the `sfs_start` command and your SFS cluster is up and running, you can still create and mount additional SFS file systems. In order to do this, you must manually check and mount the file system instead of relying on the `sfs_start` command to do this automatically.

7.1 Checking an SFS file system

After creating a file system with the `mkfs`(8) command, use the standard `fsck`(8) command to verify the file system structure. The `fsck` command needs to be executed on only one system in the SFS cluster before the file system is mounted anywhere in the cluster.

Caution: Once a file system is mounted on any system in the SFS cluster, you must not execute `fsck` on the file system. Mounted file systems may be checked with the `sfsck` command.

Following are examples of the output from `fsck` on a shared file system.

```
# /etc/fsck /dev/dsk/sfs_test0
/sfs_test0: file system opened
/sfs_test0: super block fname sfs_test, fpack sfsp_tes, (SFS enabled)
/sfs_test0: Phase 1 - Check Blocks and Sizes
/sfs_test0: Phase 2 - Visit Directories
/sfs_test0: Phase 3 - Checking Directories
/sfs_test0: Phase 4 - Checking Non-Directories and Link Counts
/sfs_test0:
/sfs_test0: i-node - 22080, name - ./hostb/lioAAAa00815/lio_815
/sfs_test0:   Clearing SFS (I) locks <0100000000000000000000, 040, 0, 040, 0>
/sfs_test0:         - 19 SFS locks cleared
/sfs_test0: Phase 5 - Verify Dynamic Information
/sfs_test0: dynamic block flags indicate mounted or not checked
/sfs_test0:             reflagged
/sfs_test0: Phase 6S - Rebuilding SFS Block Information
/sfs_test0: file system summary
/sfs_test0:     75776 total i-nodes (75448 free i-nodes)
/sfs_test0:     303104 total blocks  (296412 free blocks)
/sfs_test0:  ***** FILE SYSTEM WAS MODIFIED *****

# /etc/fsck /dev/dsk/sfs_test0
/sfs_test0: file system opened
```

```
/sfs_test0: super block fname sfs_test, fpack sfsp_tes, (SFS enabled)
/sfs_test0: clean exit for clean file system

# /etc/fsck /dev/dsk/sfs_test1
/sfs_test1: file system opened
/sfs_test1: super block fname sfs_test, fpack sfsp_tes, (SFS enabled)
/sfs_test1: Phase 1 - Check Blocks and Sizes
/sfs_test1: Phase 2 - Visit Directories
/sfs_test1: Phase 3 - Checking Directories
/sfs_test1: Phase 4 - Checking Non-Directories and Link Counts
/sfs_test1: i-node - 2816, name - ./dct/data/rtstest.gdb
/sfs_test1:   Clearing SFS (R) locks <020000000000000000000, 0, 042, 042, 02>
/sfs_test1:         - 2 SFS locks cleared
/sfs_test1: Phase 5 - Verify Dynamic Information
/sfs_test1: dynamic block flags indicate mounted or not checked
/sfs_test1:             reflagged
/sfs_test1: Phase 6S - Rebuilding SFS Block Information
/sfs_test1: file system summary
/sfs_test1:     75776 total i-nodes (75366 free i-nodes)
/sfs_test1:     303104 total blocks  (271452 free blocks)
/sfs_test1:  ***** FILE SYSTEM WAS MODIFIED *****

# /etc/fsck /dev/dsk/sfs_test1
/sfs_test1: file system opened
/sfs_test1: super block fname sfs_test, fpack sfsp_tes, (SFS enabled)
/sfs_test1: clean exit for clean file system
```

7.2 Mounting SFS file systems

After successful execution of the `sfs_start` command, you can mount shared file systems. Attempting to mount a shared file system if you have not initialized the UNICOS SFS environment yields an error.

When a shared file system is mounted, the UNICOS kernel recognizes that the file system was made as a shared file system and an assignment of the specified number of semaphores is attempted. If the assignment was successful, then the assignment is recorded in the Shared Lock Region for all other systems to see.

You may mount a shared file system on every Cray Research system in the SFS cluster. To view the Shared Mount Table, which shows all current mounted SFS file systems, use the `mount`(8) command with the `-s` option.

Caution: After a file system is mounted, you should not execute `mkfs` on that file system again until it is unmounted from all Cray Research systems in the SFS cluster.

When you execute the `mount` command with the `-s` option, the display will indicate whether the file system is currently mounted on some other Cray Research system. The following example shows a display that results from executing `mount -s`.

```
# /etc/mount /dev/dsk/sfs_tst1 /sfs/tst1
# mount -s
/sfs/tst0 on /dev/dsk/sfs_tst0 SFS ports=--m-----,dev=34/205, semas=29-156,
  validity=0:2:303104:16:24576:10240,mode=rw
/sfs/tst1 on /dev/dsk/sfs_tst1 SFS ports=--m-----,dev=34/206,semas=29-156,
  validity=0:2:303104:16:34816:10240,mode=rw
```

The `ports=` display indicates which ports the file systems are currently mounted on, with one dash serving as a place holder for each port. The leftmost port is port 0. In this example, `ports=--m-----` indicates that the file system is mounted on port 2, because the place holders for the first two ports, port 0 and port 1, are indicated by dashes.

The `dev=` display indicates the major and minor number for the file system.

The `semas=` display indicates which semaphores have been assigned to the file system.

The `validity=` display is a "magic cookie" indicator to ensure that all the systems have compatible configurations.

The maximum number of Cray Research systems that can share a file system is 64. This is also the maximum number of systems that you can have in one SFS cluster.