

# Introduction to Shared File Systems [1]

---

The UNICOS Shared File System (SFS) is a software feature available on Cray Research computer systems that allows multiple Cray Research systems to perform both read and write operations on one or more shared file systems. A shared file system provides most of the features present in the NC1FS file system, such as device striping, mirrored file systems, and asynchronous I/O.

**Note:** Mirroring is not supported on CRAY T3D or CRAY T3E systems.

To users, a UNICOS file system defined as a shared file system (file system type SFS) appears the same as an NC1FS file system, except that users can access the file system from any Cray Research system that is configured to use that shared file system. The Cray Research systems that share file systems and the media that contains the shared file systems constitute an *SFS cluster*.

Cray Research systems in an SFS cluster can mount any of the file systems defined as shared that are configured on the shared media accessible to the machines in that cluster.

All Cray Research systems that use a shared file system are peers. In an SFS configuration, there is no server system that accesses file systems on behalf of client systems. File system integrity is maintained through an external arbitration device that manages access to file data and file system structures.

## 1.1 SFS configurations

Different system configurations can take advantage of the UNICOS Shared File System feature. The following list gives some examples:

- **Fault tolerant configurations** of multiple Cray Research systems can be built where the tasks being performed by a system that fails (and the tasks are stored in data and command files on a shared file system) can be taken over and recovered by another machine in the SFS cluster when the system failure is detected.
- **Work Sharing SFS clusters** can be built where a queue of work to be done is maintained (either by Cray Research's Network Queuing Environment or by custom software) and the data and result files needed to perform the work reside on one or more shared file systems. Any machine in the SFS cluster can then work on any of the tasks because all the data and other programs needed to complete the items in the work queue are available to all the machines in the SFS cluster.

- **Division of Labor:** some machines in an SFS cluster may be producers of data while other systems in the SFS cluster may be consumers of that data. This would be typical of a configuration involving, for example, a CRAY C90 compute server which writes data files requiring a good deal of computation to produce to a shared file system also accessed by smaller CRAY J90 or CRAY Y-MP series systems, which might be well suited to further refining, rendering, archiving, or otherwise post-processing the data.

In general, large file transfers are best suited for SFS operations.

## 1.2 Basic hardware requirements

SFS file systems are supported on CRAY T3E series, CRAY T3D series, CRAY T90 series, CRAY C90 series, CRAY EL series, CRAY J90 series, and CRAY Y-MP E series systems.

The primary hardware requirements for an SFS cluster are the following:

- *Shareable media*, such as Model E SSD storage devices (SSD-Es) or ND-12/14/30/40 High Performance Parallel Interface (HIPPI) disks. SSD storage devices may be shared, providing that the mainframes and SSD storage devices can be positioned close enough that VHISP channels may be used to interconnect the clustered mainframes and shared SSD storage devices.

When ND-12/14/30/40 HIPPI disks are used as the shareable media, a set of switched HIPPI channels are used to allow the computers to access the SFS arbiter and the HIPPI disks.

- A *Shared File System arbiter*, used by UNICOS SFS to control access to the shared media. A HIPPI-SMP (H-SMP) can be used as the SFS arbiter. The H-SMP is a moderate performance, low-cost semaphore device that can coexist on a HIPPI switch with HIPPI disks and other devices. It is accessed by the Cray Research systems in the SFS cluster by using an extension to the IPI-3/HIPPI standard that provides semaphore functionality.

The UNICOS operating system supports SFS configurations that contain more than one arbitration device associated with a particular Cray Research system. For example, the UNICOS operating system can support a configuration in which one arbitration device is required for a set of file systems shared by systems A and B, and a second arbitration device is required for a different set of file systems shared by systems B and C. In this configuration, system B accesses two different arbitration devices, each controlling a unique pool of file systems.

## 1.3 Licensing

To use the UNICOS SFS feature you must obtain a software license. Licenses for the UNICOS SFS feature are maintained and administered through the Flexible License Manager (FLEXlm) product. The license keys and installation instructions are sent to you by electronic mail.

