

CentralStor™

Frequently Asked Questions

Numerical Applications, Inc.
601 Williams, Suite 4W
Richland, WA 99352
(509)-943-0861

www.numerical.com

How big a CentralStor™ system do I need?

Every customer's needs are different, of course. You can estimate by looking at your current environment to see how much file storage is online and how much data you are currently writing to tapes and other off-line storage like CD-ROM.

How many gigabytes per year do you expect to archive? This is usually a much smaller number than how many gigabytes you need to backup, since only some files will be archived for long-term storage. The size of tape robot subsystem that you need is determined by the total online storage you will require. The smallest system we sell will hold 250 gigabytes. The largest holds multiple terabytes.

What if I outgrow my initial CentralStor™ system?

The main physical components of CentralStor are the server computer system, the disk caches, and the tape robotic subsystem. Each of these components can be upgraded. The server computer system and the disk cache size affect the response time of the system. A faster server computer and more disk space help the system respond faster to user requests.

The total storage capacity of the system is essentially unlimited because tapes can be stored offline and mounted by an operator as needed, but convenient access to files requires that all the tape cartridges in use fit in robotic mechanisms. You can upgrade by purchasing replacement or additional tape robot subsystems.

For example, you might start with a system that has 12 gigabytes of disk cache and a 1x7 (one drive, seven storage slots) robot with a standalone second drive using DLT-40 tapes. After two years go by, you could attach a second 1x7 robot and add more disk space. The next upgrade might be a new 2x20 robot using a new type of tape (DLT-80 for instance). CentralStor can support multiple media types at the same time.

As your organization grows and the number of CentralStor users increases, you could choose to upgrade the server computer with a faster CPU, a faster network connection, more RAM, etc.

What about media failures?

All of the tapes in the Tape Management System (TMS) are written in duplicate, an A tape and a B tape. All normal reading from tape is done with the A tape only. The B tape is a backup in case of media failure. So long as the A tape is being actively appended to, the B tape remains on-line also as it is automatically

updated. Once an A tape is filled however, the corresponding B tape can be removed from the robot for permanent storage elsewhere.

In addition, there are administrator commands to make extra copies of the tapes at any time for offsite storage. For example, one large customer we have brings two spare tapes up to date each evening. Another smaller site updates a single spare tape once a month.

Is there a limit to the size of a file? To the number of files? To the number of versions of a file?

The current absolute limit on the size of an individual file is two gigabytes, but this is adjustable downwards by the customer. Most organizations will define a much smaller maximum size. There is no hard limit to the number of files that can be stored in CentralStor. There is a limit of two billion directories; however limits imposed by the cache file systems will likely be encountered first.

The number of versions of a file is also limited to two billion. A more practical limit is around a thousand because the interactive response time will gradually get worse as the number of filenames in a directory increases (and each version is represented by a separate filename).

What if a new storage technology comes along?

The Tape Management System (TMS) can work with multiple media types simultaneously, so you can just add new storage hardware as it becomes available. When the old tapes fill, just don't initialize any more of them. All new files will go onto the new media.

If you wish to completely migrate from one storage technology to another, you can copy the older media onto the newer type and then disconnect the old hardware. This migration can be done over a period of time while operating in a mixed environment.

How much does CentralStor cost? What do I get?

CentralStor systems are available in a variety of configurations, including custom solutions. Prices range from \$30,000 to \$170,000 for complete solutions including hardware and software. The hardware consists of a UNIX or Linux server computer system with an attached DLT tape robotic system. The specific models of computer and tape robot will vary between configurations.

How many users does CentralStor support?

This depends completely on the type and number of files that your organization produces and stores. There is no "typical" installation, but standard installations are available for 25 to 250 users.

I already have an HSM. Can I use it with CentralStor software?

The CentralStor DMS software is compatible with an HSM. You would purchase a software license only. The cost is negotiated on a case by case basis and depends on the size of the HSM system and the number of supported users.

What platforms are supported?

The CentralStor server software runs on Hewlett-Packard computers running HP-UX 10.20 or on RedHat Linux systems. The CentralStor client software is available for HP-UX 9.x or 10.x, Linux, Windows NT, Windows 95/98, and Macintosh. Other UNIX clients can be made available at customer request.

What about a major catastrophe like fire or earthquake?

All of the tapes in the Tape Management System (TMS) are written in duplicate, an A tape and a B tape. Once an A tape is filled, the corresponding B tape can be removed from the robot for storage elsewhere. Since filling a tape may take several months (some kinds of tapes have capacities over 50 gigabytes), you may want to create additional backups of the A tape before it is full. CentralStor includes a separate program for making extra tape copies. You can use the tape copier to copy the partially filled A tape on a regular schedule. The copier program has a resume option so you can bring the copy up-to-date without having to re-copy the entire contents.

A recommended procedure is to keep two extra copies of each A tape and store them offsite. Alternate these extra tapes daily or weekly, bringing one offsite copy back each time to bring it up to date with recent changes.

In addition, the database files kept on disk to index the tapes need to be backed up. These should be part of the regular backup schedule as well as having special offsite copies made.

Another way to protect against major catastrophe is to purchase special hardware connections that support long distances from computer to peripherals.

The tape drives used by CentralStor use SCSI controllers. There are special devices that use optical connectors to allow SCSI peripherals to be placed hundreds or thousands of meters from the computer. You might place the B tape in a drive that is physically located in a separate facility, and remotely mirror the disk drives holding the database files.