

Enstore Glossary

accounting server

This server maintains statistical information on a running system. It is Enstore's interface to a database of transfer-related data.

active file

Any file in Enstore that is not deleted.

alarm server (AS)

The Alarm Server maintains a record of alarms raised by other servers, and creates a report that's available online.

bfid

Bit file id; an Enstore-assigned, unique identifier for a data file.

cern wrapper

A file family wrapper that accommodates data files up to $(2^{64} - 1)$ bytes. It matches an extension to the ANSI standard, as proposed by CERN, and allows data files written at Fermilab to be readable by CERN, and vice-versa. See **file family wrapper**.

configuration server (CS)

The Configuration Server maintains and distributes the information about Enstore system configuration, such as the location and parameters of each Enstore component and/or server.

cpio_odc wrapper

A file family wrapper which allows the file to be dumpable via cpio. This wrapper has a file length limit of $(8G - 1)$ bytes. See **file family wrapper**.

crc (Cyclic Redundancy Check)

Used to verify that data has been stored properly; it's used like a checksum, but is less prone to multiple-bit errors. During a transfer, both sides calculate the crc and compare the values, unless the `--nocrc` option is specified. Enstore uses an Adler 32 crc.

cwd

current working directory

dCache

DCache is a data file caching system which acts as an intelligent manager between the user and the data storage facilities. It optimizes the location of staged copies according to an access profile. It decouples the (potentially slow) network transfer rate from the (fast) storage media I/O rate in order to keep the mass storage system from bogging down.

dCap

DCap is a dCache-native C-API access protocol.

DESY

Deutsches Elektronen-SYNchrotron; a laboratory in Hamburg, Germany that conducts particle physics research.

direct I/O

Direct I/O differs from normal disk read/writes in that it by-passes the file system's buffer cache. This is achieved by skipping the (normally done) copy that goes from the application memory space to the kernel buffer's memory space. Direct I/O is an SGI/Linux extension. Compare to **memory-mapped I/O** and **POSIX I/O**.

door (for dCache)

A door is a protocol converter (e.g., for FTP, dCap) between clients and internal dCache systems. Each door is associated with a particular port on the dCache server, and has its own access profile.

drivestat server

The drivestat server maintains statistical information of the drives.

ecrc

Stands for Enstore crc (cyclic redundancy check); see **erc**. The **ecrc** program calculates the crc of a local file.

encp

Encp as an end-user command is considered to be deprecated. It was designed to be used with Enstore, and used to copy data files from disk to storage media and vice-versa. This command is distributed as part of the **encp** product, available from kits under
`ftp://ftp.fnal.gov/products/encp/` or
`ftp://ftp.fnal.gov/KITS/<OS>/encp/`, e.g.,
`ftp://ftp.fnal.gov/KITS/Linux/encp/`

Enstore

Enstore is the mass storage system implemented at Fermilab as the primary data store for experiments' data sets. It provides distributed access to data on tape or other storage media both locally and over networks.

enstore.conf

A configuration file to allow for multiple network interface cards dedicated to Enstore, and to map the interfaces to routers. The default location for the file is `/etc/enstore.conf`. The location of the file can be overridden with the environment variable `ENSTORE_CONF`. See **ENSTORE_CONF**.

ENSTORE_CONF

Environment variable that can be used to override the location of the `enstore.conf` file. See **enstore.conf**.

ENSTORE_CONFIG_HOST

An environment variable which points to the Enstore server that is running the configuration server

ENSTORE_CONFIG_PORT

An environment variable which sets the port number; the value is (by convention) 7500 for all installations at Fermilab.

ensync

A wrapper for encp that allows you to copy the contents of an entire directory structure to Enstore via a single command.

event relay (ER)

The Event Relay is a server that forwards messages based on subscription. All the Enstore servers send messages to the ER. Any server may “subscribe” to the ER in order to have messages of particular types forwarded to it.

fairshare

A mechanism used in Enstore’s queue management that helps to keep any one storage group (experiment or group) from monopolizing tape drives. Fairshare determines which storage groups currently have jobs in progress (at a mover) and which ones don’t, then gives preference to requests associated with those that don’t.

file clerk (FC)

The File Clerk is a server that tracks files in the system. It manages a database of metadata for each data file in the Enstore system.

file family

A file family is a grouping of data; it defines a category, or family, of data files. Each experiment defines a set of file families for its data. A given storage volume may only contain files belonging to one file family.

file family width

File family width is a value used to limit write-accessibility on data storage volumes. At any given time, Enstore limits the number of volumes associated with a given file family that are open for writing to the value of the file family width.

file family wrapper

A file family wrapper consists of information that gets added to the front and back of data files as they’re written to media, and defines the files’ format on the storage volume. The format of the wrapper depends on the type of wrapper used. (See **cern wrapper** and **cpio_odc wrapper**.)

filemark

A filemark is a physical mark on tape indicating end of file. Tape drives recognize it and can do high speed searches over it.

ftp

File transfer protocol.

gridftp

See **GSI ftp**.

GSI ftp

An implementation of ftp that uses Grid Proxies for authentication and authorization and is compatible with popular tools such as globus-url-copy (from the globus toolkit).

information server

A read-only server that maintains detailed file and volume information.

inquisitor

The Inquisitor monitors the Enstore servers, obtains information from them, and creates reports at regular intervals that can be viewed on the web.

job

In Enstore terminology, a job is what a user submits to **encp**. See **request** for comparison.

Kerberized ftp client

A Kerberized ftp client is an ftp client that implements Kerberos v5 authentication.

layer

Pnfs stores metadata about each file in “layers”, each layer containing a specific type of metadata. Each stored data file has its own set of these layers. Currently, only layers 1 and 4 are used.

library

A library in Enstore is comprised of both the physical data storage media, robotic devices and drives. An Enstore library is typically called a robot.

library manager (LM)

A Library Manager is a server which controls a virtual library. LMs receive requests for file copies from users via **encp** and they distribute the requests to the Movers.

log server (LS)

The Log Server (LS) receives messages from other processes and logs them into formatted log files that are available online.

media changer (MC)

The Media Changer mounts and dismounts the media into and out of drives according to requests from the Movers.

memory-mapped I/O (mmapped I/O)

With this type of I/O, part of the CPU’s address space is interpreted not as accesses to memory, but as accesses to a device; once you map a file to memory, changes made to the memory map are propagated back to the file. Mmapped I/O strives to avoid memory copies of the data between the application memory space and the kernel memory space (also see **direct I/O** and **POSIX I/O**). Mmapped I/O is in the POSIX standard.

monitor server (MS)

The Monitor Server is available for investigating network-related problems. It attempts to mimic the communication between an **encp** request, the corresponding library manager, and the mover.

mover (MV)

A Mover is a process responsible for efficient data transfer between the **encp** process and a single, assigned media drive in a library (robot). The Mover receives instructions from a Library Manager (LM) on how to sat-

isfy the users' requests. The Mover sends instructions to the Media Changer (MC) that services the Mover's assigned drive in order to get the proper volume mounted.

pnfs layer

See "layer".

pnfs namespace

Pnfs is an independent namespace package, written at DESY. It presents a collection of library database entries as a UNIX-like file system, and thus allows users to browse stored files as though they reside in this file system. Pnfs is mounted like NFS, but it is a virtual file system only. It maintains file grouping and structure information via a set of tags in each directory.

pnfs tags

See tags.

POSIX I/O

POSIX is a name applied to a widely used family of open system standards based on UNIX. POSIX I/O refers to the POSIX standards for I/O.

request

In Enstore terminology, after a user submits a job to encp, encp sends a request to Enstore to process the job. See **job** for comparison.

resubmit

Encp has functionality to retry and resubmit requests, where we distinguish between these two terms. Encp will *resubmit* a request if it has been waiting for a mover for over 15 minutes, but not due to an error condition. See **retry**.

retry

Encp has functionality to retry and resubmit requests, where we distinguish between these two terms. Encp will *retry* (i.e., resend) a request after an error occurs. See **resubmit**.

SRM

SRM (Storage Resource Management) is the middleware for managing storage resources for the grid.

storage group

A storage group is an identifier corresponding to an experiment that Enstore uses as it controls and balances assignment of resources such as tape drives and media. Each storage group (i.e., each experiment) is assigned an area in PNFS.

storage volume

A unit of mass storage, e.g., a tape.

streaming (of files on tape)

Streaming refers to the sequential access of adjacent files on tape at the maximum tape read/write speeds.

striping (of files on tape)

Striping refers to single files (usually large ones) being split onto two or more volumes, each writing simultaneously, in order to expedite the writing process. (Striping is not supported under enstore.)

suspect volume

A volume becomes suspect when a mover communicates to the appropriate library manager that it had a problem with the volume. It is not yet established that the volume is faulty.

tags

Pnfs uses tag files (usually just called tags) in the `/pnfs` namespace to specify file-specific configuration information, and **enclp** transfers this information to Enstore. Tags are associated with directories in the `/pnfs` namespace, not with any specific file, and thus apply to all files within a given directory.

virtual library

A Virtual Library (VL) is a subset of an Enstore data storage library. It can contain one and only one type of media. It is paired with its own library manager which controls it.

volume

See storage volume.

volume assert

A job in which a volume (tape) gets mounted and certain attributes are read in order to check, and thus assert, that the volume is “ok” (without actually checking the entire contents). Requesting volume asserts is an administrative task, and these requests are assigned the lowest priority.

volume clerk

The Volume Clerk (VC) is a server that stores and administers storage volume (tape) information.

volume family

The triplet “storage group + file family + file family wrapper” is called a volume family. In order for different data files to be stored on the same volume, all three of these pnfs tags for the files must match.

wrapper

See file family wrapper.