

MONITORING A PETABYTE SCALE STORAGE SYSTEM

Jon Bakken, Eileen Berman[#], Chih-Hao Huang,
Alexander Moibenko, Don Petravick, Michael Zalokar
FNAL, Batavia, IL 60510, USA

Abstract

Fermilab operates a petabyte scale storage system, Enstore, which is the primary data store for experiments' large data sets. The Enstore system regularly transfers greater than 15 Terabytes of data each day. It is designed using a client-server architecture providing sufficient modularity to allow easy addition and replacement of hardware and software components. Monitoring of this system is essential to insure the integrity of the data that is stored in it and to maintain the high volume access that this system supports.

The monitoring of this distributed system is accomplished using a variety of tools and techniques that present information for use by a variety of roles (operator, storage system administrator, storage software developer, user). Essential elements of the system are monitored: performance, hardware, firmware, software, network, data integrity. We will present details of the deployed monitoring tools with an emphasis on the different techniques that have proved useful to each role. Experience with the monitoring tools and techniques, what worked and what did not will be presented.

INTRODUCTION

The Enstore distributed systems installed at Fermilab, consist of more than 100 commodity computers, 5 robotic tape libraries, and more than 90 tape drives. Many different types of processes run on these nodes providing high throughput reliable access to modular and scalable data storage. [1] Monitoring of these systems is done in order to -

- insure data integrity
- enable problems to be quickly and efficiently located and solved
- enable users to monitor progress of their requests
- provide information to be used to predict future needs of users and of the storage system itself
- make available current state and activity

The monitoring is done in real time by a separate server, and periodically by many cron jobs. Mount and transfer information is recorded in a database for easy generation

of plots. The monitoring uses three types of notification to facilitate problem solving -

1. Automatic generation of pages to storage system administrators upon problem occurrence. This is done for serious problems that may result in loss of service.
2. Automatic generation of an alarm indicating a problem that can be fixed at a later time.
3. Automatic updating of web pages.

What is Monitored

To accomplish the goals of the monitoring stated above, a wide variety of components of the Enstore systems are monitored including Enstore Software, Transfers, Volumes, Automated Tape Libraries, Hardware and Firmware, Security, Networks, and Metrics.

Monitoring Roles

Monitoring information is supplied to benefit and support the needs of several different roles -

- Users
- Storage system administrators
- Storage system developers
- Site system operators
- Managers

MONITORING DETAILS

Enstore Software

Monitoring of the Enstore software is done on many different levels. A top view of the system as a whole and its status gives an immediate grasp of the system state (Illustration 1). The color of the ball associated with each piece of the system indicates its health: green-operational, yellow-minor problem, and red-major problem. The color of the ball next to the 'enstore' label is a function of the colors of each individual server. Generation of this information is connected to the Fermilab helpdesk and automatic paging system. Users, administrators, and site operators all find this page useful. It is the only page monitored by the site operators and is structured so that problems can be noticed from across a room.

[#]berman@fnal.gov

Overall Status

- enstore
- ADIC AML/2
- alarms
- D0 Mezzanine Powderhorn

Illustration
1: Enstore
Status-At-A-
Glance
Monitoring

Enstore Individual Server Status

- | | | |
|---|--|--|
| <p>Servers</p> <ul style="list-style-type: none"> ● Alarm Server ● Event Relay ● Inquisitor ● Volume Clerk | <p>Library Managers</p> <ul style="list-style-type: none"> ● Configuration Server ● File Clerk ● Logger ● D0-9940B.library_manager ● mezsilo.library_manager ● samlto.library_manager ● samlto2.library_manager ● samm2.library_manager | <p>Media Changers</p> <ul style="list-style-type: none"> ● aml2r1.media_changer ● stk.media_changer |
| <p>Movers</p> | | |
| <ul style="list-style-type: none"> ● 994004.mover ● 994012.mover ● 994024.mover | <ul style="list-style-type: none"> ● 994007.mover ● 994019.mover ● 994025.mover | <ul style="list-style-type: none"> ● 994009.mover ● 994020.mover ● 9940B26.mover |
| <ul style="list-style-type: none"> ● 994011.mover ● 994023.mover ● 9940B27.mover | | |

Another level of detail is shown by the monitoring of the request queues and the status of the Enstore servers and movers. Request queue information includes the list of active and pending jobs and hints as to why a job is pending. This is helpful to users to determine when their job will complete. Storage administrators often monitor the request queues to obtain a more detailed view of a problem. Problems with a server's status will automatically generate an alarm.

Completed transfers are listed separately (Figure 2).

alarms is also useful to developers as it provides a quick overview of problems with the system.

Additional monitoring of the Enstore software includes the monitoring of -

- Current Enstore configuration
- Log File size – alarm if too large
- Files being transferred or in queue
- Size of databases - alarm if too large

Enep History										
Time	Node	User/Storage Group	Mover Interface	Bytes	Volume	Network Rate (MB/S)	Transfer Rate (MB/S)	Drive Rate (MB/S)	Disk Rate (MB/S)	Overall Rate (MB/S)
200-08-15 10:12:20	dlc0132	enepDC	@Barracuda	99260912 (4)	enep7803014	5.37	5.37	21.0	5.27	3.87
200-08-15 10:12:20	dlc0278	enepDC	@Barracuda	62793168 (3)	enep7807014	4.57	4.57	41	4.47	3.81
200-08-15 10:11:20	dlc01000	enepDC	@Barracuda	1431960 (3)	enep78014	1.1	5.96	11.9	13.1	3.263
200-08-15 10:11:17	dlc01017	enepDC	@Barracuda	42466092 (4)	enep78042	0	0	40.9	19	1.84
200-08-15 10:11:08	dlc01030	enepDC	@Barracuda	11289488 (5)	enep78042	0.0	0.0	40.9	1.86	0.42
200-08-15 10:11:06	dlc01000	enepDC	@Barracuda	40419408 (8)	enep78010	0.53	0.53	21.2	1.62	0.71

Figure 2: Completed Transfers

Transfer information includes the time the transfer completed, where the transfer request came from, which volume was accessed, and several measurements of the rate of data flow during the transfer. Users, administrators, and developers all find this page useful to identify transfer, network, and drive problems. The variety of rate measurements has proven useful in detecting network problems on both the user and the Enstore end.

Alarms generated in response to problems are gathered on a single web page. This list serves as a work list for the storage system administrators. Alarm information includes first and last time the alarm was seen, number of times the alarm was seen, where it came from, information about the error, and whether or not an automatic page was generated for the alarm. This list of

Networks

Monitoring of the networks connecting the Enstore systems to the users is mostly done automatically. Simulated transfers test access to and from the Enstore nodes. Additional monitoring of UDP rates and general network rates are used similarly. Drops in network rates are reported as alarms. This information is mainly used by developers and has proved crucial on occasion when diagnosing transfer problems.

Security

Security monitoring is done to prevent unauthorized access to Enstore machines and to help protect the data stored on tape. Open socket connections are monitored and alarmed if unrecognized. All Enstore machines are fully kerberized and access to them requires a legal kerberos ticket. This access can be denied easily by turning off someone's ability to obtain a legal ticket.

All tapes marked as full in the tape libraries have their write protect tabs flipped. This protects against software error, user error, admin error, and malicious intent. Security monitoring is especially useful to managers.

Hardware and Firmware

Monitoring of the hardware and the firmware has been mostly automated to generate alarms. Monitored elements include -

- Node configuration
- Status of raid disks
- CPU & Memory Information
- System event log
- Baseboard and Processor Voltages
- Processor Fan Speeds
- Baseboard and Processor Temperatures
- Console logs

This automation has taken a considerable amount of effort and it is unclear what benefit has been derived from some of it. CPU load monitoring proved to be useful in debugging the Enstore servers.

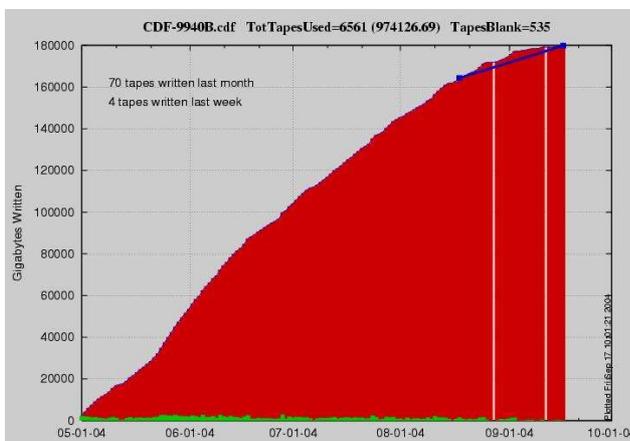


Figure 4: Tapes in Library

Volumes and Tape Libraries

Many aspects of tape and library usage are monitored. This monitoring has become a good tool to predict trends and usage patterns and directly feeds into planning activities for future purchases by users and storage administrators. For example, plots of the total used slots in a library indicate when space is getting scarce. Plots of tape usage by an experiment indicates when new tapes need to be purchased (Figure 4).

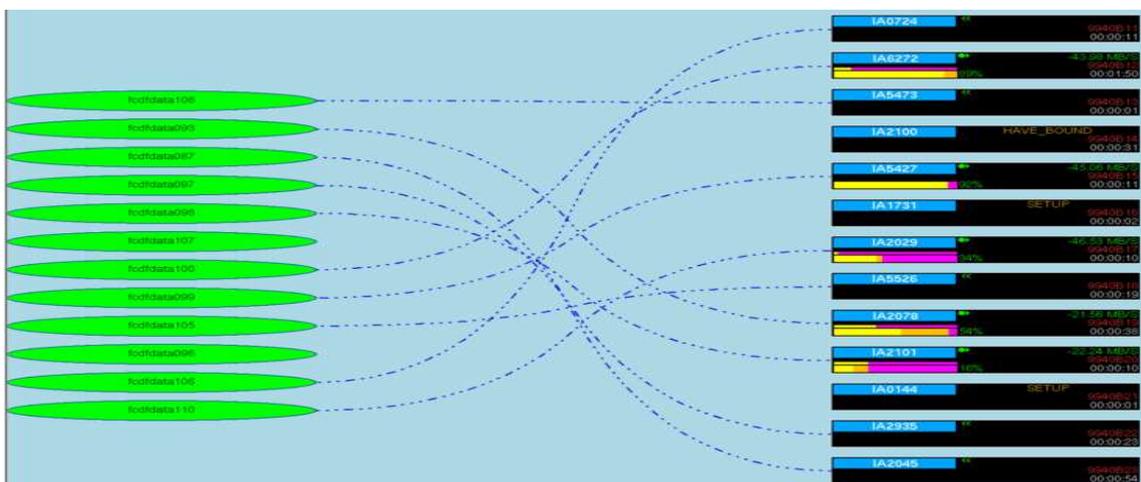
Additional monitoring which has been found useful includes -

- Available library quota – indicates when additional quota should be allocated.
- Cleaning tape status – are new tapes needed
- List of failed tape drive transfers – useful to quickly find problem drives or tapes.
- Tape library logs
- Firmware version, serial # of tape drives – useful to help debug drive problems
- Volume Information – useful to users to have available online
- Volumes with too many mounts – generates alarms
- Volumes with errors
- Volume Inventory
- Random read of files – part of data integrity checks, problems can indicate volume or system problems.

ENTV

Enstore TV (entv) is a realtime monitoring program that graphically represents the movement of data into and out of the Enstore mass storage system (Figure 5). It is the single most useful monitoring tool available to Enstore. It is used by users, administrators, developers, and managers. The display shows ellipses (on the left) which represent user nodes transferring data to/from Enstore. On the right side are rectangles which represent individual tape drives and the mover software which interfaces to them to perform the reads and writes. Animated dashed/dotted lines connect user nodes with tape drives representing active connections and data that is flowing. The movement of these lines represent the direction of data flow and the speed of the movement the rate of transfer. Information within the mover/tape drive display areas give tape information, transfer state including amount completed, and transfer rate. This tool presents an easily understandable and interpreted view of the Enstore system state. There are several locations on

Figure 5: ENTV



site where these displays are viewable and users make considerable use of them.

Metrics

The Enstore system collects and records several different types of metrics. These metrics are used by users (for information), administrators (to monitor the use and performance of the system), developers (to debug problems), and managers (to report on system use). Gathered metrics include -

- Total user data on tape (currently 2.2 petabytes)
- Bytes transferred/day (individually and total for all systems) (Figure 6)
- Separate experiment use of Enstore (Figure 7)
- Transfer activity & Instantaneous transfer rates
- Drive utilization snapshots
- Number of mounts/day
- Backup times
- Status of the suite of cronjobs the system supports

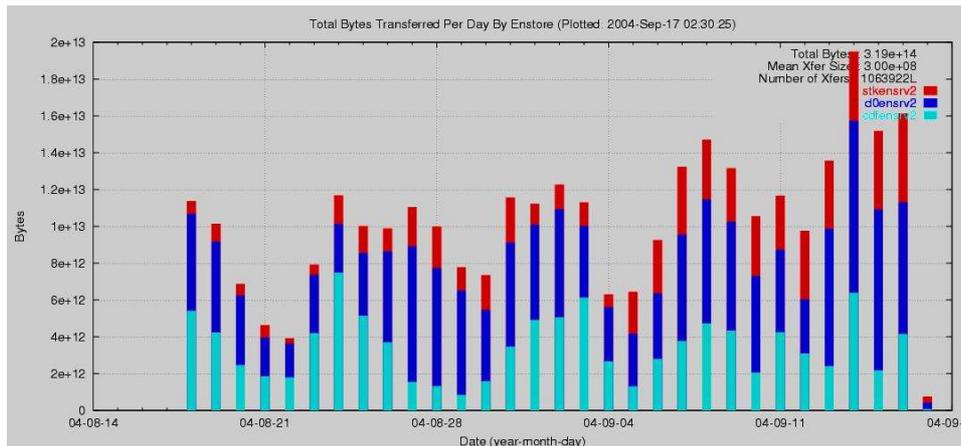


Figure 6: Total Bytes Transferred (all systems)

LESSONS LEARNED

What Worked Well

The most useful tool is ENTV. Originally ENTV was developed as an experiment in monitoring. It has proved highly useful and is continually being enhanced. Additionally, the metrics have proved useful for daily monitoring and reporting. The list of alarms which functions as an administrator work list, provides a simple method for keeping track of what needs to be done. In general, it has been very useful to be able to automate the monitoring as much as possible. Automated paging of administrators insures a quick response to problems without needing to constantly poll the system.

What Did Not

Automated monitoring of PC firmware took more effort than was originally thought and has supplied very little useful results. It was difficult to read out the firmware consistently and many false alarms were generated. Effort was used to generate an online web page that provided a list of user files that were involved in active or pending transfers. This was thought to be of use to the user who was new to Enstore and only knew a file name. Practice shows that users rarely used this feature and would go straight to the request queue pages and were able to find information there.

CONCLUSIONS

The monitoring of the mass storage system at Fermilab has been a great success.

Users, administrators, operators, developers, and managers make use of this monitoring to do problem resolution, maintain data integrity, follow file requests, and plan for future needs, both in the long term and the short term.

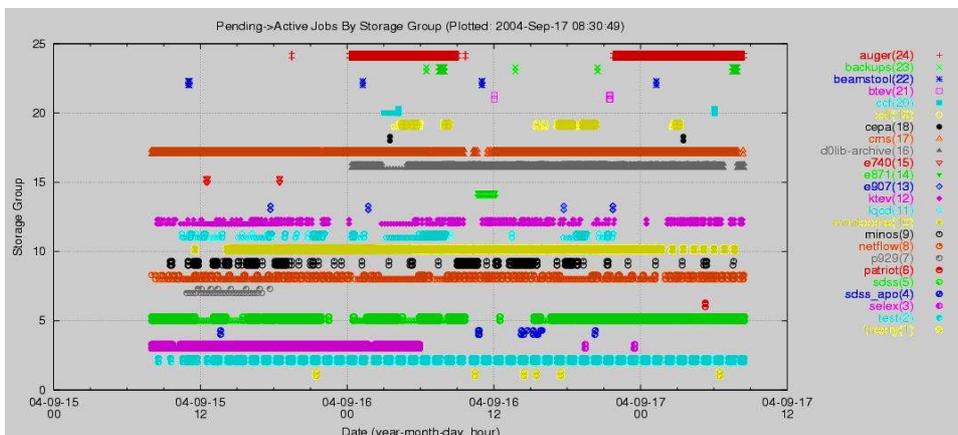


Figure 7: Experiment Use of Enstore

REFERENCES

[1] J. Bakken, et al, "The Status of the Fermilab Data Storage System", CHEP 04, Sep 2004.