

HSM and backup services at INFN-CNAF

Alessandro Cavalli¹, Daniele Cesini¹, Enrico Fattibene^{1,}, Andrea Prosperini¹, Vladimir Sapunenko¹*

¹INFN-CNAF Bologna, Italy

Abstract. IBM Spectrum Protect (ISP) software, one of the leader solutions in data protection, contributes to the data management infrastructure operated at CNAF, the central computing and storage facility of INFN (Istituto Nazionale di Fisica Nucleare – Italian National Institute for Nuclear Physics). It is used to manage about 55 Petabytes of scientific data produced by LHC (Large Hadron Collider at CERN) and other experiments in which INFN is involved, stored on tape resources as the highest latency storage tier within HSM (Hierarchical Space Management) environment. To accomplish this task, ISP works together with IBM Spectrum Scale (formerly GPFS - General Parallel File System) and GEMSS (Grid Enabled Mass Storage System), an in-house developed software layer that manages migration and recall queues. Moreover, we perform backup/archive operation of main IT services running at CNAF, such as mail servers, configurations, repositories, documents, logs, etc. In this paper we present the current configuration of the HSM infrastructure and the backup and recovery service, with particular attention to issues related to the increasing amount of scientific data to manage, expected for the next years.

1 Introduction

The computing and data centre hosted at INFN-CNAF is one of the 13 Tier 1s of the WLCG (Worldwide LHC Computing Grid), that receive data produced by the LHC experiments. CNAF provides also computing and storage facilities for 30 other experiments in which INFN is involved, belonging to Astrophysics, Astro-particle Physics and High Energy Physics domains. Data are stored on both disk and tape resources. At the time of writing, ~20 PB of data reside on disk and ~55 PB on tape.

Moreover, it is responsibility of CNAF to implement, manage, maintain and support many services which are fundamental for the entire INFN, such as general ICT services (top level domain management; coordination of academic and INFN Wi-Fi roaming infrastructure; backup of INFN Certification Authority data; centralized management of INFN Web sites; creation of software development tools and collaborative tools) and INFN Information System services.

* Corresponding author: enrico.fattibene@cnafe.infn.it

2 HSM service

Tape-based storage is the highest latency tier within a HSM environment. The infrastructure is based on a tape library Oracle-StorageTek SL8500 equipped with 16 T10000D tape drives (other 9 T10000C drives are used only for the backup and recovery service). The overall capacity of the SL8500 library is 10000 slots, so ~85 PB could be stored with the existing technology. Data stored on tape are organised into Spectrum Scale file systems. Disk storage spaces are used as buffer, to allow writing and reading operations on/from tape.

ISP [1] offers a HSM extension to manage migrations from disk to tape and recalls from tape to disk of data hosted on Spectrum Scale [2] file systems. Figure 1 shows an high level schema of the HSM system operated at INFN-CNAF. The system is managed by an ISP server, connected to the tape drives with 16GFC (16 Gb Fibre Channel) technology. This server mounts a remote file system containing configuration, database, active and archive logs. Another ISP server, configured in the same way, is in stand-by and can be turned on in case of unavailability of the first one, after having mounted the remote file system. Each file system containing datasets that have to be stored on tape is managed by an active HSM server, equipped with Fiber Channel connections to both disk storage and tape drives. Each active server shares the configuration with a stand-by server. At the time of writing, five active HSM servers (one for each LHC experiment and one for all the others) run at CNAF and other five are in stand-by mode. Three servers (CMS, LHCb and no-LHC) can exploit a 16GFC card and the other two (ALICE and ATLAS) are equipped with an 8GFC (8 Gb Fibre Channel) connection, that will be upgraded soon to 16GFC.

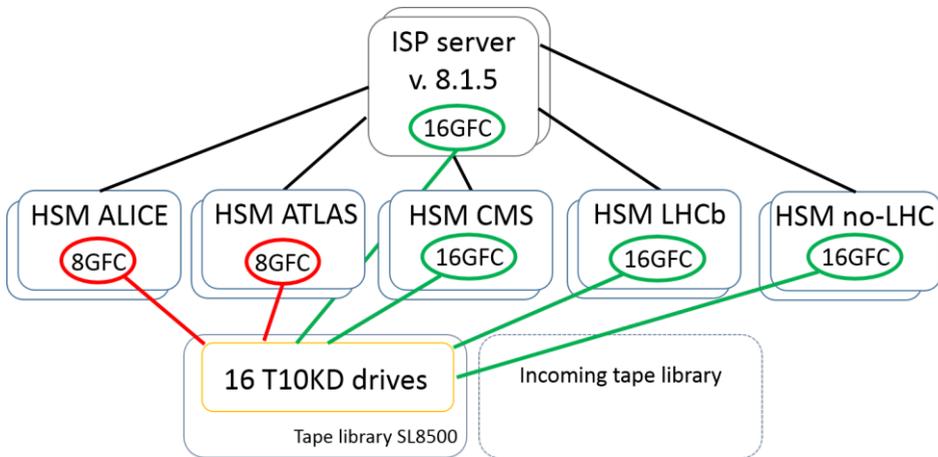


Fig. 1. HSM system schema at INFN-CNAF.

Data are accessed through StoRM [3] (Storage Resource Manager), a software developed by INFN and based on SRM (Storage Resource Management) interface to access storage resources. Migrations and recalls are managed through the HSM servers thanks to GEMSS [4], another INFN software that provides a full HSM integration of Spectrum Scale, ISP and StoRM. GEMSS has been designed to optimize migration and recall operations. Recalls can be triggered by a periodic scan of StoRM bring-online table (the table in the StoRM database that contains the files to copy on the disk buffer), can be requested through GEMSS command line or can be caused by a direct file access. GEMSS passes the list of files to recall to the ISP-HSM service, that groups the files by cartridge, ordered on the basis of their position on the tape. At this point, a recall process can start for

each cartridge, corresponding to a drive dedicated to that process. Among other things, GEMSS offers the possibility to customize, for each handled file system, the maximum number of recall threads running and performs periodic regeneration of tape ordered file lists to include new requests in already existing lists.

2.1 HSM input and output rate

Currently the HSM servers equipped with a single 8GFC connection to the Tape Area Network are capable of handling 800 MB/s simultaneously for inbound and outbound traffic. This limit is 1.6 GB/s for the HSM servers with a 16GFC card. However, the real rate depends on the number of drives in use (250MB/s native rate per each T10000D tape drive). Anyway, the observed mean daily transfer rate from disk to tape for HSM servers with 8GFC hits 650 MB/s, as shown in figure 2, because of a bit of inefficiency, that is caused by the time spent in file system scan to select the candidates. We experienced inefficiency also for recalls. In this case, it could depend by the fact that files to read are not put sequentially on the tape or by the small size of files to recall. Moreover, the data transfer rate is expected to grow in the next years, both for writing and reading operations. Data produced by the scientific experiments will fulfil the SL8500 tape library in production during 2019, so that we will purchase a new one. By the end of 2023, data stored on tape at CNAF should amount to 220 PB. On the other hand, our customers, such as the scientific communities, disclosed the trend to use tapes as near-line (or “slow”) disk, thereby increasing the reading traffic rate. For all these reasons, we recently double the Fibre Channel connectivity for three HSM servers and we plan to update the other two next year, by purchasing new servers equipped by 16GFC connection to the Tape Area Network.



Fig. 2. Reading (yellow line) and writing (green line) data rate from/to tape of an HSM server equipped with 8GFC connection, over a period of 24 hours.

2.2 Tape drive sharing

The 16 T10000D tape drives are shared among the file systems handling the scientific data. At this moment, there is no way to allocate dynamically more or less drives to recall or migration activities on the different file systems. In fact, the HSM system administrators can only set manually the maximum number of migration or recall threads for each file system by modifying the GEMSS configuration file. As effect of this static setup, we experience that frequently some drives are idle and, at the same time, we notice a certain number of pending recall threads that could become running by using those free drives. In order to overcome this inefficiency, we designed a software solution [5], that will be a GEMSS extension, to dynamically change the number of tape drives usable by recalls or migrations on the file systems on the basis of the number of pending requests and of the recent usage of the drives. The algorithm gives a greater priority to those file systems that

have less used the system in the recent time (for example in the last 24 hours) or that have a smaller number of threads running at that moment.

3 Backup and recovery service

CNAF backup and recovery service is running to protect different kinds of IT services data (mail servers, repositories, service configurations, logs, documents, etc.). This service was re-designed during 2016 after a couple of episodes of data loss that needed restore of backed-up data from the system. In those cases, data were recovered successfully, but that experience moved the system administrators to make the service more efficient and secure. The knowledge of ISP gained with the HSM system administration is being reused for backup and recovery service management.

The system is composed by:

- 1 ISP server version 8.1.5 on a host equipped with an Intel Xeon E5-2640-v2 8-core processor and 32GB RAM
- A disk area of 64TB (4 LUN, 16TB each) on DDN storage
- 9 tape drives Oracle-StorageTek T10000C
 - 7 for data and 2 for ISP server database backup
- 60 T2 tapes, 5.4 TB each
 - 55 for data (300 TB) and 5 for ISP server database backup (27 TB)
- 16 client nodes (ISP version 6, 7 or 8) on bare-metal or virtual machines (on virtualization environment oVirt or VMWare)
 - 12 standard clients on Linux (CentOS7 or SL6)
 - 2 standard clients on Windows Server 2012
 - 2 Oracle databases (Tivoli Data Protection for Oracle Database - TDPO) on RedHat Enterprise Linux

Since we had on different machines a not so great amount of data to backup, we thought to make these data visible to some servers hosting ISP client, that became collectors of data to backup. In this way we are able to minimize ISP license costs. Moreover, we installed ILMT (IBM License Metric Tool) software in order to use the sub-capacity licensing for Virtual Machines, i.e. counting the virtual cores instead of the physical cores on the hypervisor.

The service allows to exploit both *Incremental backup* and *Archive* features available with ISP. In the scope of *Incremental backup* the most recent version of an object saved on the system is defined *Active*. This version stays on the system for an indeterminate period of time. When an object is modified or deleted from its original location, at the subsequent backup the new version becomes *Active* and the modified version (or the deleted one) becomes *Inactive*. *Inactive* data are stored on the system for a period of time specified as retention for each client node by the ISP server administrators.

Figure 3 shows the data workflow in the backup system. Every day data are sent from clients to ISP server. The *Active* version is saved on both tape and disk, as we configured an *Active* data pool on disk. *Inactive* data are stored on tape, with different retentions depending on the client, from 1 month to 1 year. For some critical clients, we also save another copy of *Active* and *Inactive* data on tape. We also do *Archive* of some data that do not need to be changed.

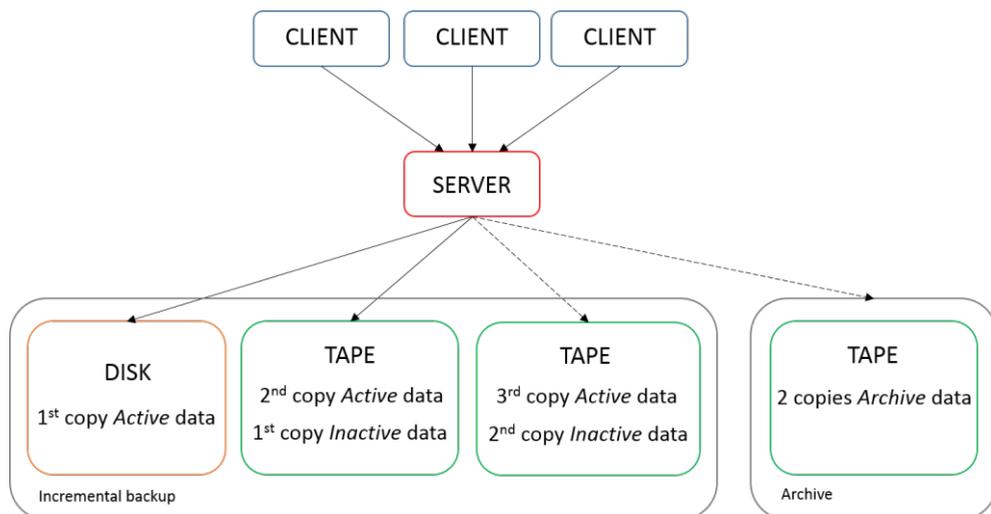


Fig. 3. CNAF backup workflow. Full lines represent data flow for every client nodes; dotted lines are depicted to show additional data flow for selected clients (second copy of Inactive data, third copy of Active data or Archive copies).

At the time of this paper, the backup system is holding roughly the following amount of data:

- 23 TB / 66 millions of *Active* objects
- 22 TB / 10 millions of *Inactive* objects
- 2.7 TB / 2.2 millions of *Archive* objects

The used space on disk is 23 TB (only *Active*) and 75 TB are stored on tape (*Active*, *Inactive* and *Archive* data).

In order to simplify the service operations and to minimize errors and warnings, we set up a notification system based on ISP Operation Center, a monitoring and administration tool released with the Spectrum Protect suite. Each CNAF department having at least a client node under backup receives via email a daily report on the nodes managed by the department itself. The report is composed by the results of 3 queries to the server database:

- scheduled backups in the last 24 hours, with schedule time, real start time, end time, final status (Completed, Failed, Future, Missed, Started, Restarted) and the output that gives information about the presence of warnings or errors;
- used space and number of files per client on each storage pool (disk or tape);
- used space and number of files per filespace on all pools.

Each client administrator is in charge to analyze the report and take the suitable actions to solve the problems, possibly with the support of the backup system administrators.

To test the correct operation of the backup system and to minimize the restore time, we consider to run a couple of data restores for each client node on a yearly basis. This test consists of a restore of all *Active* data from the disk storage pool and a sample of *Inactive* data from tape. These tests are planned together by backup system administrators and client operators.

Figure 4 shows the duration of a set of backup jobs of different clients on a period of 30 days during 2018, depending on the number of copied objects. Three size intervals are represented by different colors of the dots. On a logarithmic scale we have the number of objects backed up against the duration of backup jobs in seconds. We notice two groups of dots (circled in red) with similar numbers of backed up objects (~10000-15000) and same

size range (100-350 GB). For these groups there is a clear difference in the backup duration, less than 1 hour against ~8-10 hours.

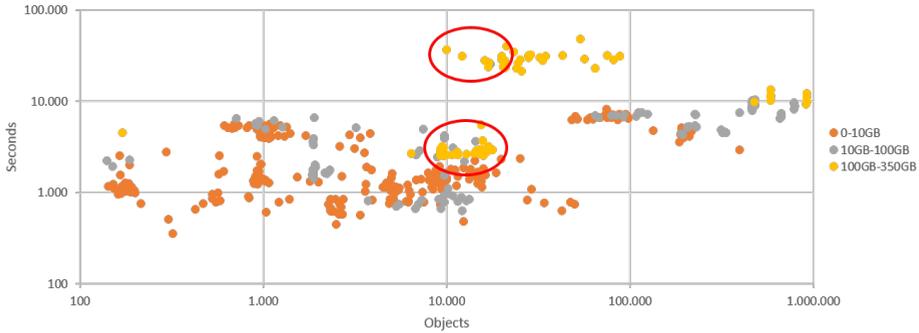


Fig. 4. Duration of backup jobs on a period of 30 days. Three size intervals are represented by different colors.

Figure 5 shows the same data plotted in Figure 4, but the yellow dots of the previous plot are divided here in two groups: blue dots represent backup jobs from clients with a number of overall objects greater than 2 millions. If the number of objects to scan in the client is over this value (2 millions) the backup job duration becomes of the order of several hours.

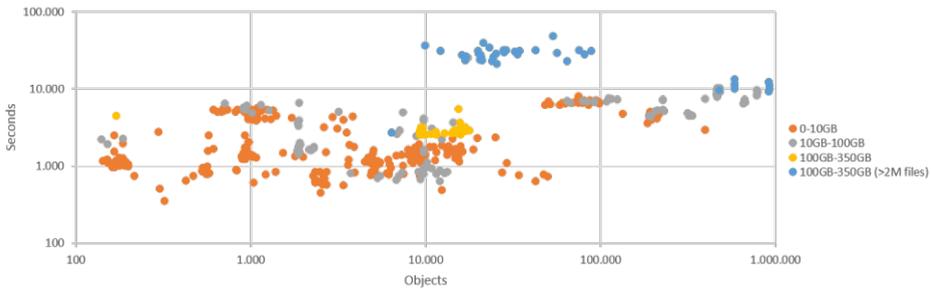


Fig 5. Duration of backup jobs on a period of 30 days. Three size intervals are represented by different colors. For the greatest size interval (100-350 GB), backup jobs from clients with a number of overall objects greater than 2 millions are distinguished from the others.

4 Conclusions

CNAF exploits ISP for both HSM service and backup and recovery service. Through the HSM service, about 55 PB of scientific data are handled on shared resources. The prevision of important growing of this amount of data, together with the future increase of recall activities, move us to plan the upgrade of the Fiber Channel connection for HSM servers and to develop a software solution to dynamically allocate additional drives to file systems and to manage concurrent requests. The backup and recovery service is running to protect different kinds of IT services data (mail servers, repositories, service configurations, logs, documents, etc.). This service was recently re-designed to become more secure and

efficient. The backup job duration is acceptable unless the client nodes has a large number of files to scan (> 2 millions).

References

1. IBM Spectrum Protect web page: www.ibm.com/systems/storage/spectrum/protect [accessed at February 27, 2019]
2. IBM Spectrum Scale web page: www.ibm.com/systems/storage/spectrum/scale [accessed at February 27, 2019]
3. R. Zappi, E. Ronchieri, A. Forti, A. Ghiselli, *An efficient Grid data access with StoRM*, Lin S., Yen E. (eds) *Data Driven e-Science*, Springer, New York, NY, 239 (2011) DOI: 10.1007/978-1-4419-8014-4_19
4. D. Bonacorsi, A. Cavalli, L. Dell’Agnello, D. Gregori, A. Prosperini, P. P. Ricci, L. Rinaldi, V. Sapunenko, V. Vagnoni, *The Grid Enabled Mass Storage System (GEMSS): the Storage and Data management system used at the INFN Tier1 at CNAF*, J. Phys. Conf. Ser., **396**, 042051 (2012) DOI: 10.1088/1742-6596/396/4/042051
5. A. Cavalli, D. Cesini, E. Fattibene, L. Morganti, A. Prosperini, P. P. Ricci, V. Sapunenko, *Dynamic sharing of tape drives accessing scientific data*, J. Phys. Conf. Ser., **1085**, 042039 (2018) DOI: 10.1088/1742-6596/1085/4/042039