SCIENTIFIC DATA STORAGE AND ACCESS

CAPACITY PLAN

Version 1.0

Effective Date: 10/01/2013

Expiry Date: 09/30/2014

Fermilab

VERSION HISTORY

Version	Implemented By	Revision Date	Approved By	Approval Date	Reason
1.0	*Gene Oleynik*	*06/28/2013*	*Robert M Harris*	*07/01/2013*	Draft based on v1.5 Plan Template with SDSA content.

TABLE OF CONTENTS

1 INTRODUCTION

PURPOSE OF CAPACITY PLAN

This capacity plan provides an analysis of existing service capabilities and a description of the service plans with respect to capacity.. The analysis is based on trends of the existing support resources, current and expected SLA requirements, and currently understood plans for improvement from the service owners and Service Level Management.

1.1 SERVICE OVERVIEW

The Scientific Data Storage and Access (SDSA) services provide direct tape and disk backed tape services for storing and accessing scientific data. The direct tape access storage service offering is called Enstore and the disk back tape storage service offering is called dCache. dCache can be run with Enstore as the tape back-end, or as a volatile cache without tape backing. Both services share a common file namespace called Chimera. These service offerings are described in more detail in the SLA for these services in CS-doc-5032.

Both of these service offerings involve two types of capacity: storage and access bandwidth. There are two components to bandwidth – bandwidth to the hardware: tape drive and disk drive/raid rates, and transfers to and from these systems by the user. User access for both services is over Ethernet networks. The tape storage access is restricted to the Fermilab LAN and is accessed through software called encp. dCache can be accessed through a number of protocols some of which are restricted to the Fermilab LAN, while others can be used over the WAN.

For the Enstore service offering, storage capacity is expanded by adding more tapes and, if necessary, adding more tape libraries to hold them. Storage capacity can also be increased by moving to newer tape technology, that has a higher data density per tape cartridge. Moving to a significantly higher density media combined with migrating data from existing media to the higher density media can free up slots in the tape libraries. This can lead to an excess in unoccupied tape slots and opens up the possibility of retiring tape libraries. Planning for tape storage capacity needs to strike a balance between high slot occupancy while maintaining enough margin to accommodate unexpected data ingests or new customer opportunities. We would like to operate between 50% and 75% occupancy. This is also complicated by the fact that there are three library complexes in two buildings and there are inefficiencies in storage because of this.

Tape storage bandwidth capacity is added by adding more tape drives to the libraries. Tape drives are relatively expensive and hence are limited resources. It takes significant time to mount a tape, seek a file to it, read or write the file, and dismount the tape. For these reasons bandwidth capacity doesn't necessarily scale with the number of tape drives – scaling is strongly dependent on usage patterns. Enstore has tape drives that are dedicated to specific customers who have paid for them (e.g. CMS or Run II), but also drives that are shared by many users (so called "Public" Enstore). Sharing the drives can lead to efficiencies, but also can result in problems if some users' access patterns are not optimal or efficient. For example, the rate of reading or writing files from a tape to local disk can run much more slowly than the full tape drive bandwidth – thus tying up the tape drive for longer than is necessary. Another example is that many transfers from a single host can saturate the host’s NIC bandwidth and thus slow down the tape drive (though Enstore has a protection mechanism designed to prevent this). For these reasons, it is difficult to predict how many tape drives are needed and what the total bandwidth needs may be. Typically we monitor tape drive-hour utilization plots to determine whether more drives are required and are flexible in that we can temporarily move underutilized drives owned by one group to a group that is having a peak demand.

Enstore has a feature called Small File Aggregation (SFA). The feature collects small files on internal disk storage, and when enough are accumulated, packages them up into a single tar file and writes them to tape. For reads, the package is read from tape to disk and the individual files are unpacked to disk. This provides quicker access to other files in the package. SFA is an add-on feature that is available by request. SFA disk and bandwidth capacities are resources that need to scale with the number of users using the feature and the amount of data and files that are pushed through the system.

Enstore provides a multitude of plots and data for monitoring the capacity, bandwidth and historical usage and patterns and all of this information is recorded in databases, which are used for usage and trends plots and for making purchasing decisions.

dCache is a disk cache system. It can be tape backed or a tapeless “volatile” cache[1]. The disk space is organized into pools that are distributed across a number of file servers. Pools can be read pools, write pools, or read+write pools. Since it is a cache, the read pools never really fill – older files, unless they are “pinned”, are evicted to make room for newer files. Write or read-write pools can fill if there are enough pending writes to tape since these files are not eligible for eviction until they are written to tape.

dCache stores whole files – there is no striping to disk; it is optimized for aggregate throughput rather than per-file performance. dCache provides load balancing across storage pools in a pool-group, which consists of a number of data pools which usually are separate file servers. Load balancing is achieved by calculating a cost function to determine which pool in a group that a file should be read or written to. Pool-group selection is based on a “storage group” (usually an experiment) specified in the user’s request or it can be based on host-name or network. Having multiple customers and users share a pool-group can provide for more efficient usage of disk resources, but it can be overwhelmed - effectively a denial of service - by a single misbehaving user. There is no mechanism to protect against this and the response is administrative intervention to fence off the user and to manually remove the misbehaving user’s files. The impact to other users is they may encounter write failures (write pools) or their cached files have been evicted (read or read+write pools) and they would have to read them back in from tape.

Because dCache is a caching file system, it is not as simple to know, as it is with a normal file system, whether or not more capacity is needed. For write pools there should be enough space margin to buffer for peak demand, tape drive congestion, and tape library service windows. For read+write pools these margins may need to be larger to accommodate this without affecting unpinned cached files intended for reading. For read pools the margin has to be large enough such that the lifetime of the files in the cache is sufficiently long to be efficient for the users.

[1] Currently there is only one customer of the Fermilab volatile cache – Fermigrid. All other users use tape backed cache.

dCache provides several mechanisms (some localized customizations) to monitor and manage storage capacity, partition it for multiple users, and monitor and plan for capacity. As already mentioned, dCache can be partitioned into storage pool-groups and dCache requests can be directed to them based on a storage group or by hostname or network. This provides a mechanism for administratively balancing loads if a given pool-group is overloaded. In addition, the dCache administrative interface provides tools to migrate files between pools in the background. Plots for dCache include file lifetime histograms and current usage (cached space, unused space, pending write space, pinned files, etc.) for each pool.

1.1.1 Technology

Fermilab tape storage uses automated tape libraries. These libraries have robotic arms that move tape cartridges from their storage slots to their tape drives for reading and writing, or from a tape I/O port to/from tape slots in the library. These slots are "universal" and as such can hold a variety of types of tape cartridges (LTOx, 9940x, T10000x). The libraries include software and firmware that tracks cartridge and drive locations and for mounting and dismounting tapes to drives. Tapes are entered or removed from the library by command through the library' I/O cap. In general, automated libraries can be expanded for more tape slots up to a physical limit and accommodate additional tape drives up to a physical limit. Automated tape libraries can be strung together with pass-thru ports to form a "complex". Tapes can automatically be exchanged between tape libraries in a complex through their pass-thru ports.

As of fiscal year 2013, the Fermilab automated tape libraries have a mix of commodity LTO4 (800 GB cartridges) and enterprise T100000T2 (5400 GB cartridges) and is in the process of migrating data off of LTO4 to T10000T2 media. Potential capacity is 375 Petabytes (T10000T2) and up to 15 GB/s (T10000C drives).

The goal for dCache performance is high aggregate throughput. Commodity disk arrays can be used to achieve this. dCache disk at Fermilab is implemented with commodity SATA disk arrays capable of RAID6. Each array has on the order of several tens of terabytes of usable storage.

1.1.2 Customer Base and Organization

There are several instances of the SDSA Enstore and dCache services that serve different customers:

- Cdfen Enstore, which serves the CDF experiment
- D0en Enstore, which serves the D0 experiment
- Stken (AKA "Public") Enstore, which serves CMS and all other experiments and customers, in particular this includes the Intensity Frontier experiments
- Cdfen dCache which serves CDF
- Stken (AKA Public) dCache that serves all other experiments and customers besides CMS[2]

[2] CMS has its own dCache that the CMS Facilities Department manages and is not part of the SDSA services.

The plots[3] below are a snapshot of the distribution of tape storage capacity for these customers (June 2013). Not all customers are shown, very small datasets are omitted (the blue "other") in order not to clutter the plots. The first two plots show data on tape, which double counts a large amount of data that has been migrated to new media and has not had the old media ejected. The second pie chart counts only active data, which is not double counted. The last plot shows the amount of data transferred to/from tape, which includes migrating data from LTO4 to T10000T2.

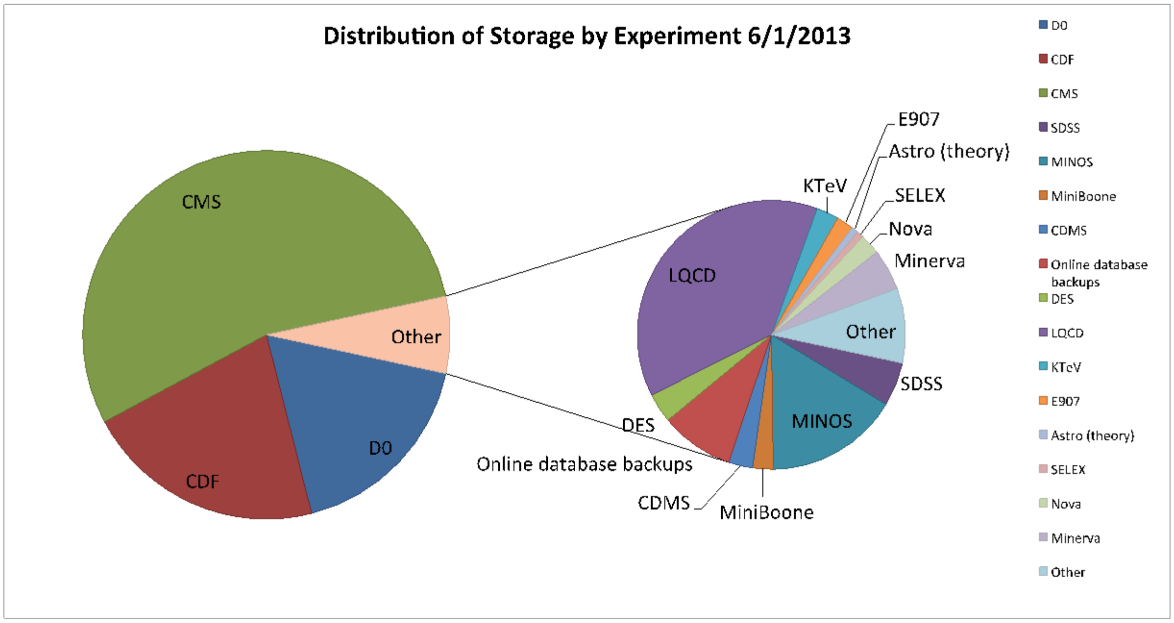

Figure 1. Distribution of tape storage in use (double counts a large amount of migrated data)

[3] Current plots can be found in CS-DOC-2478 in the latest tape_plots Excel spreadsheet

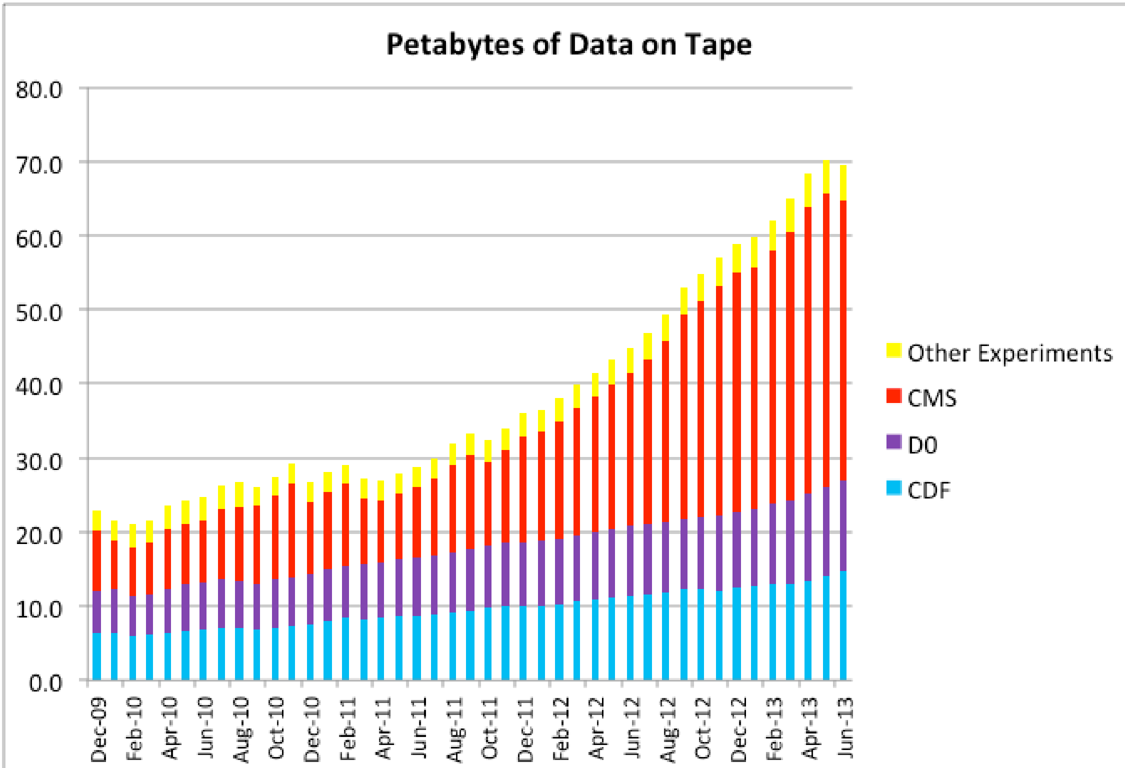

Figure 2. Historical Data on Tape (double counts a larger amount of migrated data)

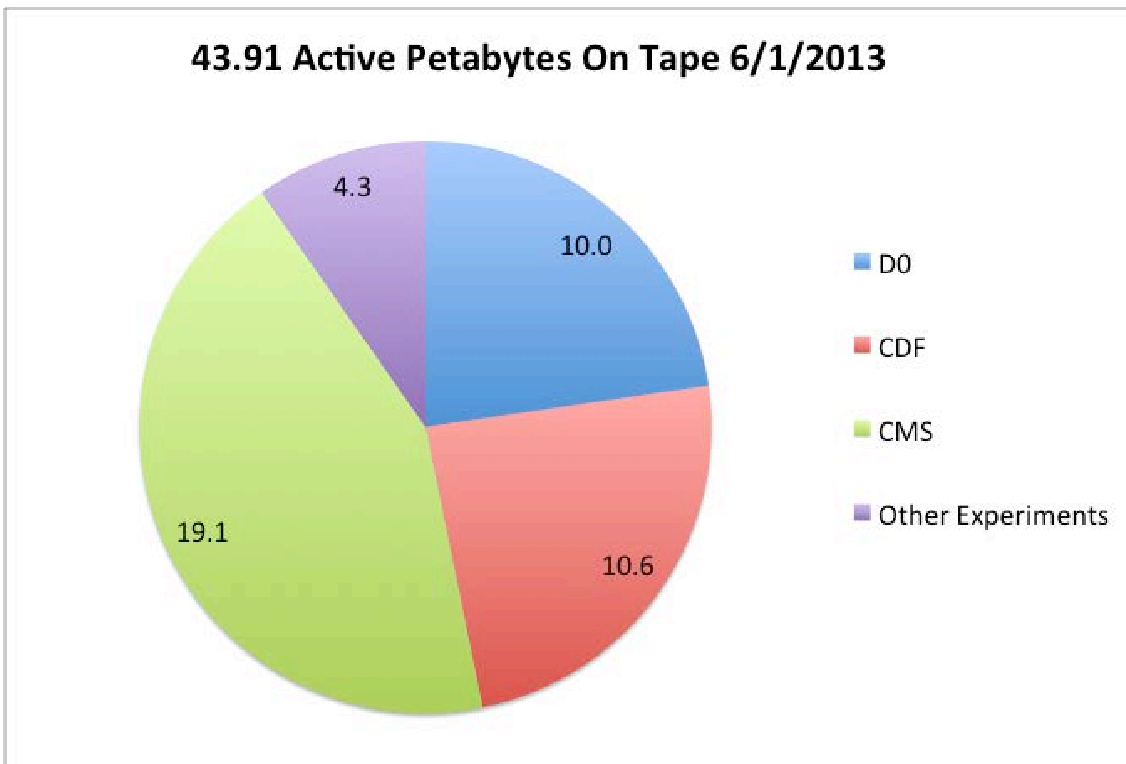

Figure 3. Distribution of active data (no double counting) on tape

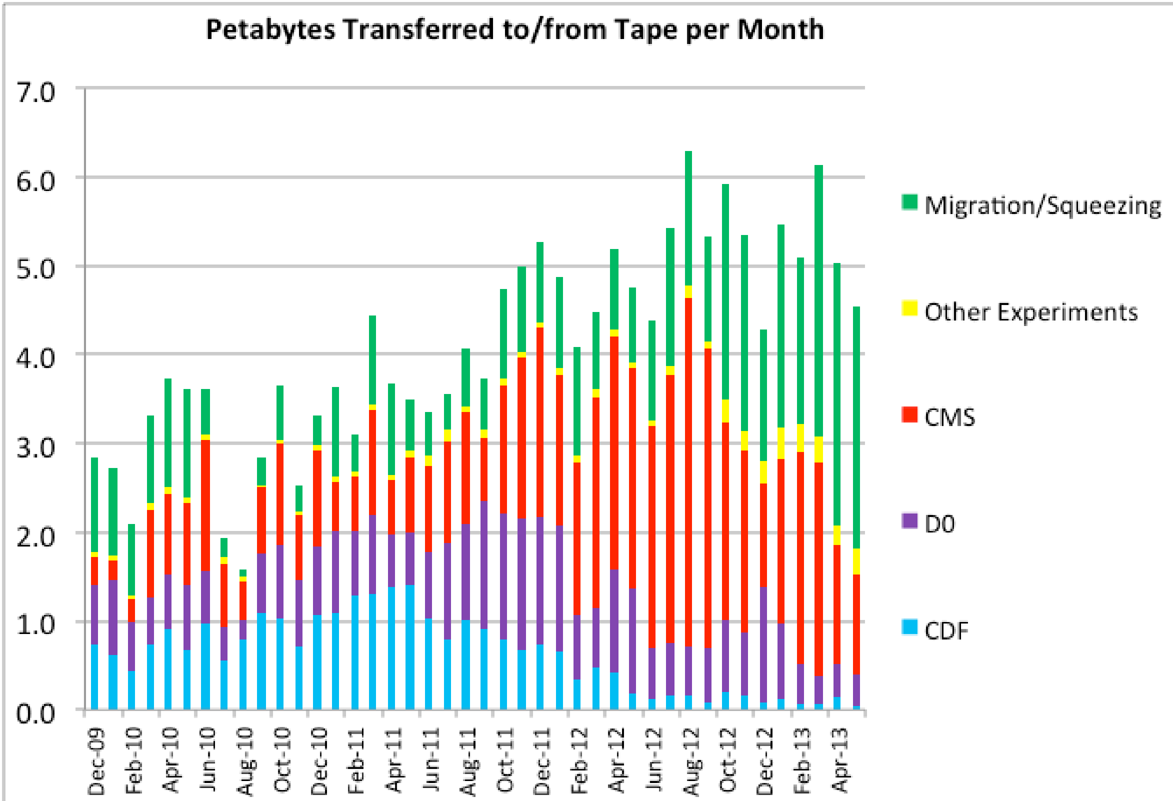

Figure 4. Petabytes of data transferred to and from tape per month

DES and MINOS have dedicated Public dCache disk. The rest of Public dCache is currently shared space.

1.1.3 Customer Chargeback

See the Cost and Budget section below

1.1.4 Snapshot of Current Tape and Slot Capacity

The pie-charts in Figures 5-7 below are a snapshot of the current capacity in each of the three library complexes:

- GS is the 10,000 slot stand-alone at GCC (Run II RAW and "Public" data)
- G1 is the 30,000 slot complex at GCC (CMS data)
- F1 is the 30,000 slot complex at FCC (Run II non-RAW and Public data)

In these plots

- "Empty" = slots without tape cartridges

- "Blank T2" and "Blank LTO4" are tapes with no data yet that represents usable storage capacity; 5.4 TB/slot for a T2 and 800 GB/slot for an LTO4.
- "Migrated LTO4" are LTO4 tapes whose files have been migrated to T2 tapes. These files are user accessible from the T2 and are marked deleted on the LTO4, but still occupy slots. These are potential for slot capacity (if ejected) or can be reused for LTO4 blank tape capacity.
- "Active(Used) LTO4" tapes have user accessible data and will eventually be migrated to T2 tapes.
- "Active(Used) T2" tapes have user accessible data.

In the pie charts below, shades of red represent slots with cartridges that have active data on them, shades of green represent storage capacity in the form of empty slots or slots with blank tapes, and what is blue represent potential slot or storage capacity in slots or migrated media that can be reclaimed.

	Blank tapes	Empty slots (assume T2=5.4 TB/slot)	Migrated (assume T2=5.4 TB slot)
F1	11.8 PB	11 PB	15 PB
G1	4.9 PB	31.4 PB	84 PB
GS	0.6 PB	8.3 PB	32.5 PB

Table 1. Current Estimated Capacity and Potential Capacity (6/21/2013)

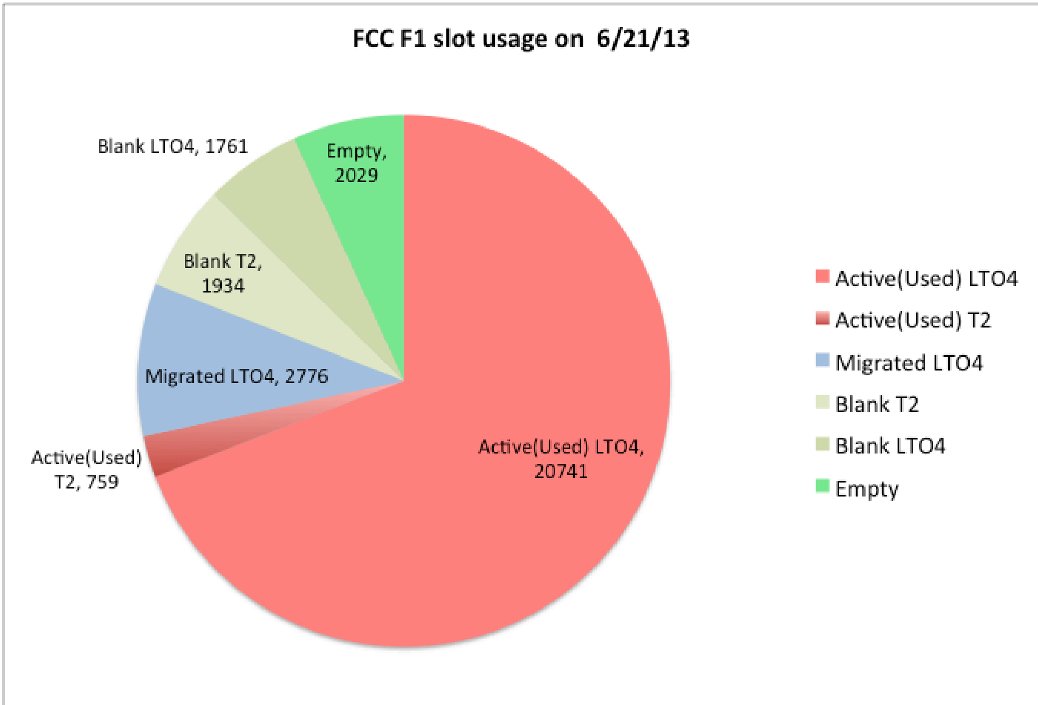

Figure 5. FCC F1 tape library complex slot usage

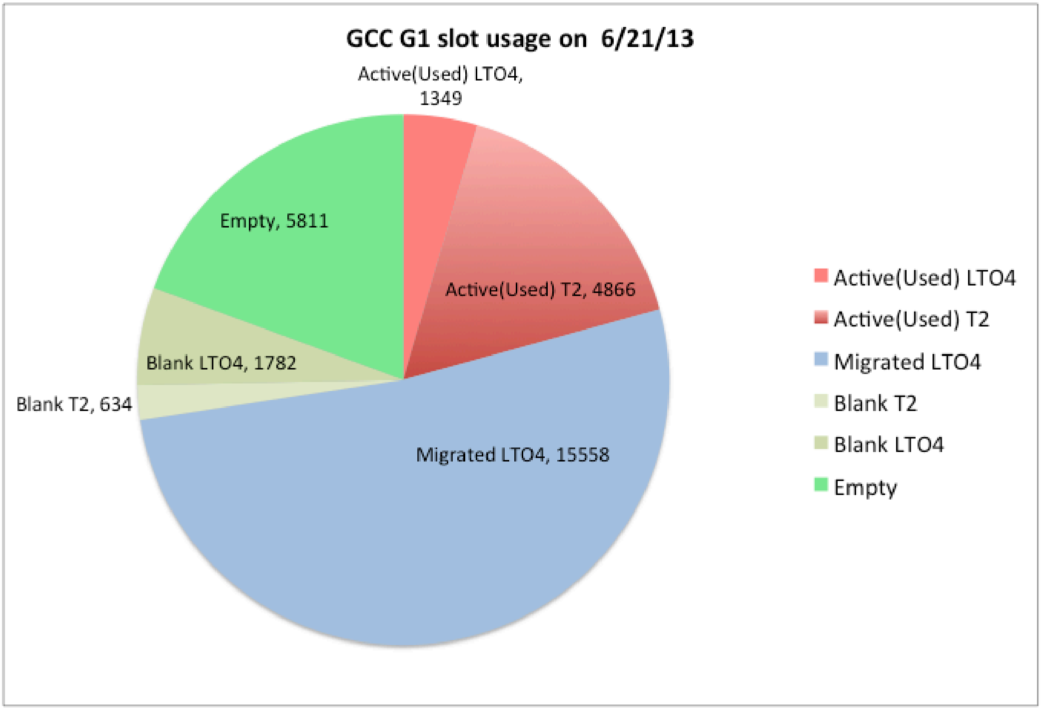

Figure 6. GCC G1 tape library complex slot usage

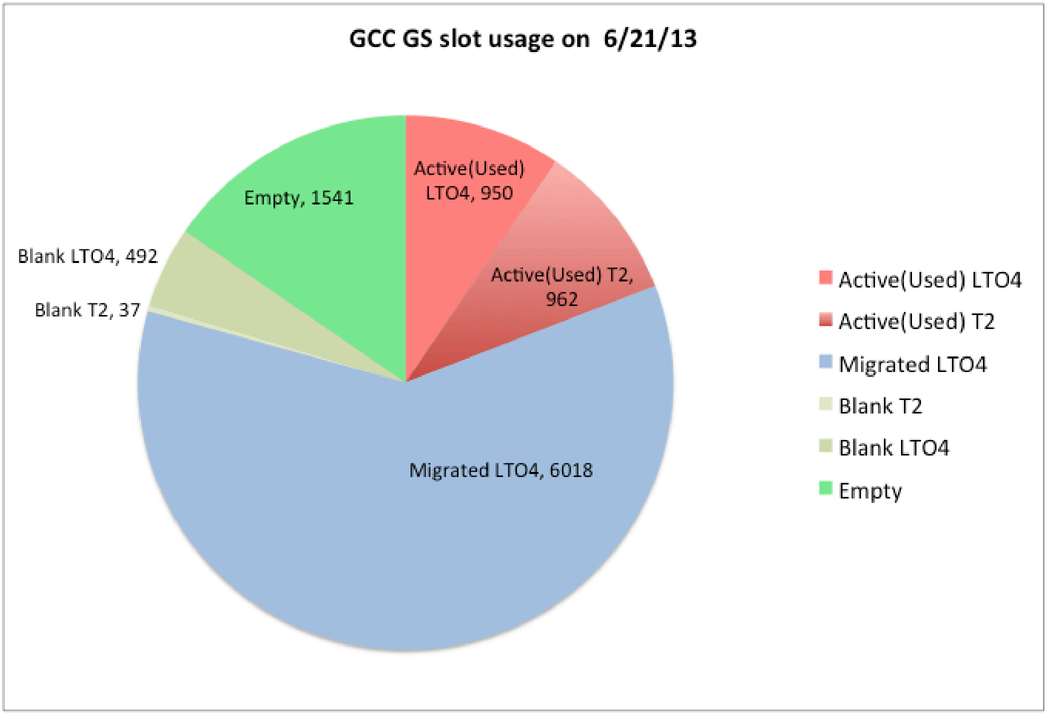

Figure 7. GCC GS tape library complex slot usage

1.2 ASSUMPTIONS/CONSTRAINTS

This Capacity Plan assumes that the service will be implemented and operate as is for the remainder of the FY13 fiscal year, except where noted below.

1. There will be increased demand as the accelerator comes online towards the end of the fiscal year. There is enough capacity on site to handle this.
2. We currently have one SFA file server in production and will soon be adding a second to handle demand.
3. We are instrumenting some additional Public dCache and upgrading a number of the pools NIC from 1 GE to 10 GE to more efficiently utilize the T10000C drives.
4. We will be using LTO4 tapes retired from CMS and Run II migration to satisfy Intensity Frontier demands.
5. We are switching NOvA to T10000C and later in FY13, we will do the same for Minerva.
6. We will continue to use drive resources to aggressively migrate LTO4 to T10000T2.

This Capacity Plan assumes there are sufficient infrastructure resources, such as Facilities and networking.

This Capacity Plan treats marginal costs, the costs to add or remove capacity to an existing service, unless otherwise noted. Neither maintenance and operations costs for the service, nor the costs to research potential changes to the service, are considered in this plan.

2 PLAN SUMMARY

2.1 CAPACITY MANAGEMENT SUMMARY

The following table summarizes the capacity plan for this service:

Capacity Metric	Capacity Requirement [4]	Existing Capacity	Predicted Growth + Timescale	Capacity Threshold	Threshold Response Strategy/Tuning (Action to Be Take Upon Reaching Threshold(s), includes any tuning or demand management strategies)
Storage Slot Capacity	+1800 53,000 Total by end FY13	9381 empty 70,000 Total	*30%/yr*	10% unused slots	The free slots for the tape libraries monitored from the Enstore Volume Audit page and Quota pages linked from it. Unused slots includes blank tapes and unfilled slots. There are currently 3 complexes that hold different data-sets: • GS stand alone library at GCC • G1 Complex at GCC • F1 Complex at FCC The threshold is applied to each independently. There are several strategies that we have available if the threshold is exceeded (can be a combination of these): 1. Request recycling of tapes from the recyclable tape candidate list. If the media is not of the correct type, they may be ejected to make space. 2. Migrate data off of tapes with many deleted files thus recovering unused space on tape media. This takes time however. 3. Purchase more libraries. 4. Shelve infrequently accessed tapes. The last is a last resort as we do not as part of operations shelve data. For determining capacity needs in slots, this dynamic information is used to supplement usage predictions in the tape planning document (CS-doc-2478) “tape storage planning” spreadsheet. This spreadsheet takes into account migration, duplication of migrated data, recycling of migrated volumes, and user deletions and inefficiencies. We don’t expect slot capacity to be an issue for quite some years since we just made a big jump in media density (a factor of over 6x). If we were totally migrated to the new media we would only be at around 15% capacity.
Storage Blank Tape Capacity	Two month blank supply @ current write rates +16 PB, 65 PB total	17+ PB in blanks (with more ordered)	*30%/yr*	2 months supply	Experiments budget for enough tape to last past the end of the FY and they are purchased throughout the year. The blank tape burn rate for customers is monitored from the Enstore audit and quota pages. When the threshold is reached the customer is contacted. Strategies for providing blank tapes include: 1. Purchase additional tape media.

[4] Slot and PB capacity requirement estimates are based off of the predictions in the tape-planning document.

	end FY13				**2.** Request recycling of tapes from the recyclable candidate list (this may include migration source volumes). **3.** Migrate data off of tapes with many deleted files thus recovering unused space on tape media. This takes time however. **4.** Purge obsolete files from tape and reuse them. If an unexpected increase in demand is such that getting new blanks delivered in time is risky, we may borrow from other customers, with their permission, if they have surplus.
Network Capacity	> 75 Gbps		*10%/yr*	75%	Generally determined for each experiment in the budget planning process determined for each experiment in the budget process. This should be less than the Tape drive bandwidth capacity. Enstore transfer page, ENTV and network MRTG plots are used to monitor and diagnose network bandwidth issues.
Tape drive capacity	> 100 Gbps (60+ drives[5])	61 T10000C, 104 LTO4 Libraries can hold 64 x 7 = 448 tape drives	*10%/yr*	75%	Generally determined for each experiment in the budget planning process There are several strategies if the threshold is exceeded (all of which have been exercised at one time or the other): 1. Purchase additional tape drives. Each library can host 64 tape drives (a total of 448 for 7 libraries. Currently we have about 170 drives) 2. Move underutilized drives from one customer to the customer that is having a higher demand than usual. 3. Slow down competing migration processing freeing up drive resources This happens when experiments decide, for example, to reprocess select datasets.
Disk Cache	125 TB	125 TB	*20%/yr*	Need threshold	Generally determined for each experiment in the budget planning process. There is a large (approx.. 4 PB) purchase in FY13. File lifetime and utilization plots are used to determine if more dCache disk is needed. For the Enstore Small File Aggregation ZFS cache, rate and queue depth are used to determine if expansion is required.
Server Memory	64 GB		*5GB/yr*	90%	Request additional funding to increase memory in affected servers
Server CPU Utilization	50%		*5%/yr*	80%	We use host monitoring for system load. Usually a load of 5 or less is tolerable, but this depends on the server function. Actions we can take for high load threshold are: 1. Rearrange service components across the servers. 2. Purchase more servers 3. Address source of high loads via software changes

[5] Based on T10000C drives. We have over 100 LTO4 drives that are being phased out.

The table contents are described in detail in Section 3 Capacity Analysis.

2.2 RISKS

Risks associated with Capacity

1. We have moved to a single source tape library and drive vendor. Storage Tek is a large vendor that has been around for a long time and has done well through a couple of acquisitions. The risk is thought low.
2. We have a single source of tape media, Fujifilm. This is also a well established manufacturer. The impact of Fujifilm not being able to deliver blank tapes is high, though the risk is thought to be low. There was no impact to tape on Fermilab supplies from the recent Earthquake/Tsunami in Japan.
3. There is some risk that we have underestimated (or overestimated) tape drive needs, in particular since new experiments are starting up, and we do not fully understand the demand that Run II will place on the system over the next several years (their drive capacity may be an overestimate).

2.3 RECOMMENDATIONS

Metrics for determining tape drive utilization currently only exist on a per tape library basis. In order to understand this at the customer level to minimize the risk of too little drive capacity, Enstore software is being updated to track tape utilization metrics per customer (per storage group to be exact).

[This is filled out by the Service Owner and Capacity Manager.]

[Following the example above, matching recommendation numbers to risks numbers:

1.1 Migrate Minos DB to virtual server hosting service. There is sufficient excess capacity to mitigate this risk

OR

1.2 Increase CPU resources for Minos DB by X to mitigate this risk]

2.4 DECISIONS

[This is filled out by the Service Owner, Capacity Manager, and Financial Manager.]

[Enter all your priority 1 capacity related budget request items here. Update if you are granted more than your priority 1 capacity budget requests. Example: if recommendation 2 above was the direction we were going…

1. *Purchase new server ABC for Minos DB]*

2.5 NEXT REVIEW DATE

[Enter approximate next review date or timeframe. Expect annual reviews.]

3 CAPACITY ANALYSIS

3.1 CAPACITY METRICS AND PROCESSES

In addition to the capacity requirements obtained from customers in the budget input planning, this service has distinct technical infrastructure which is managed via the following metrics:

1. Enstore total and per storage-group (aka experiment) data on tape per month
2. Enstore tape burn-rate plots and Enstore free slots plots from the Enstore Volume Audit page. The burn-rate plots are used to predict near-term utilization of blanks (see 3) and the free slots the remaining capacity in the libraries .
3. Tape planning document CS-doc-2478. This incorporates current usage and predicted usage for the year either based on current burn-rate or estimates from the customers.
4. Enstore total and per Storage group data transferred per month and per library tape and slot utilization charts (CS-doc-2478) and from queries to the enstore databases.
5. Enstore recyclable tape candidate list. This list contains tapes that contain nothing but (user) deleted files and are candidates for reuse or ejection (explicit customer permission required).
6. Enstore Quota pages (linked to from the Volume Audit Page). Each customer is given a quota of tapes. Once this quota is exceeded, they cannot draw any new blank tapes. Staff and software monitor these quotas, and alarms are generated when a customer is approaching their quota. In response to the alarm, staff normally raises the quota and may contact the customer.
7. Enstore tape drive-hour utilization plots (soon per-storage group)
8. Networking MRTG plots for Enstore tape drive mover and dCache pool computers
9. Enstore encp transfer pages and Enstore EnTV monitoring are used to monitor network, user host, and Enstore bandwidth issues.
10. Enstore Small File Aggregation (SFA) queue depth and data rate plots
11. dCache file lifetime plots
12. dCache pool utilization plots
13. Zabbix (or equivalent) monitoring is used to monitor server load and memory and CPU utilization.

In addition, staffing resources need to be considered for the capacity planning for this service. Staffing levels will be reviewed, reported, and updated yearly in the *Tactical Plan For Data Movement and Storage*, available at ***https://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=4836***.

3.2 CAPACITY AND PERFORMANCE REQUIREMENTS

As part of the budget planning process each year, estimated capacity requirements are determined for each customer of the SDSA services.

3.3 TRENDING AND PREDICTIVE ANALYSIS

Analysis on how much capacity is needed is performed by each customer as part of the budget process for that fiscal year. Typically the budget estimation includes enough capacity for 3 months past the end of the fiscal year. This is done in order to buffer against potential purchasing difficulties at the beginning of the next fiscal year due to continuing resolution, budget difficulties, etc.. There is always some uncertainty in these estimates, so the usage is monitored throughout the year. In addition, new experiment's estimates have a much larger uncertainty and so estimates must be conservative. If a customer's current utilization indicates that they may outstrip their budgeted capacity for the fiscal year, they are sent a courtesy warning mail, at which point, they will decide on a mitigation if needed – either purchase more resources, curtail demand, or recycle resources.

When experiments are taking data, their ingest rate is fairly constant as is production analysis output. There are peaks in the output of production,. We have some limited flexibility for moving resources between customers to meet peak demands.

Historical trends in Metrics 1-6 are used to estimate how many tapes to budget and purchase to satisfy all experiments to the end of the fiscal year + 3 months (end of December). They are used to determine and predict slot capacity needs.

Historical trends in Metrics 7-9 are used to determine if additional tape drives or networking resources are needed.

Historical trends in Metric 10 is used to determine if more SFA storage resources are required.

Metrics 8, 11 and 12 are used to determine if dCache capacity upgrades are required

3.4 IMPACT OF NEW TECHNOLOGY/TECHNIQUES/UPGRADES

Tape is still cheaper than disk. Providing our own tape storage is still less expensive than cloud storage and it is not clear cloud storage could sustain the I/O rates we need (even at a significant cost).

New denser tape technologies are expected soon. Moving to such a technology can reduce library space. In general, we refresh tape technology every other generation because of new media and tape drives costs. This is usually more cost effective than every generation if new media needs to be purchased. However, the T10000D tape drive, expected to be available by 2014, can use the same tape media as what we are currently using for T10000C, but write it at a higher density. No new media is needed in this particular case. However, the tape drives are expensive and we currently have an abundance of potential capacity in the libraries.

Migration is an Enstore process where data is migrated off of old media onto newer, denser media. The files are staged off tape to disk on a "migration station", then copied from that disk onto new tape. Typically the disk on the migration system is a fairly large RAID array of spinning disks, and reading/writing disk contention on these arrays decrease performance. We are moving away from this to single large enterprise Solid State disk drives for this process.

We have need for a robust globally accessible file system for the Small Files Aggregation (SFA) Component of Enstore. For this we use a ZFS appliance. We have purchased a second ZFS server and are deploying it in FY13 in order to scale SFA up to more users.

dCache is the de facto cache file system in use at Fermilab. There are numerous global file systems that are available; some are not mature enough for production use, while others are not well suited to the type of use in High Energy Physics. Fermilab US-CMS T1 have evaluated dCache along with Lustre, HFS and EOS and found dCache best met their storage needs. NAS storage has been (and is being) used in by several Intensity Frontier experiments and found not to scale well. For the near term, dCache appears to be the storage solution to replace this. In FY13 we are purchasing between 0.5 and 1.5 PB of disk for dCache based storage for the IF experiments.

3.5 THRESHOLDS AND RESPONSES

The capacity thresholds are conservative and based on past experience. For tape, a percentage of the total capacity is probably not the best measure for a threshold. A better measure is probably based on currently used capacity. Burn rate is monitored and projected into the future to determine if adjustments to the initial estimates may be needed.

For dCache, file lifetimes are monitored and, if they are too short, more disk space may be acquired or disk allocation redistributed (for instance for peak demand by an experiment).

3.6 EXTERNALLY DRIVEN MANDATES

There are no externally driven mandates for this service at this time.

3.7 COST AND BUDGET

Budget to address the risks and perform the recommendations for this service are reflected in the tactical plans of this service and any of those services listed in Section 3.1.

For tape storage, “high” consumption internal customers budget each fiscal year for the incremental slots they expect to use, the tapes they need to purchase, any additional tape drives, and the maintenance for the total number of slots they expect to have used by the end of the fiscal year. The tapes, drives and maintenance funds are applied directly to these items. Since the libraries have already been purchased, the incremental slot costs are considered “contributions” and are used to pay for other storage costs (servers, etc.). Since estimates are often not perfect, the previous years expenditures are compared to the actual usage in that year and the customer may be credited or charged more. Customers are expected to pay for migration costs when technology nears its end of life or for programmatic migration by the Scientific Computing Division. Migration costs include drives and media.

The Scientific Computing Division will, at their discretion, pay the storage costs for internal customers with small datasets.

A pilot is underway for external customer archiving. It is planned to use an amortized cost model with prepayments. This cost model is documented in CS-doc-2478 document “Amortized Cost Model Spreadsheet”.

For dCache storage, the Scientific Computing Division, at their discretion, will provide shared tape backed storage for low-use customers. This dCache space is shared with other customers, and the amount of space is limited. The Scientific Computing Division will purchase with customer funds, install and operate dedicated disk for internal customers.

This disk is dedicated to the internal customer and is not shared. The internal customer is expected to pay the cost of refreshing the disk and server hardware when it has run past its warranty period (typically 4-5 years).

Typical costs are:

- Tape drives $15,000 - $22,000 + $2,500 for mover computer
- Tape library slot $45
- Tape media: $25-$40 / TB
- Disk media: < $400 / Raw TB

Enstore and dCache servers are now purchased with 4-5 year extended warranties and RAID arrays with 5 yr.

3.8 PLANNING INTEGRATION

The tactical plan includes activities and budget line item to execute against (procurement and planning).

3.9 MONITORING AND REPORTING

Metrics 1,2,5, 6 and 7 are automatically generated by Enstore. Metrics 11 and 12 are automatically generated by dCache.

Enstore Historical usage data is stored in Enstore databases on a per transfer and per tape mount granularity. At the beginning of each month the service owner queries the database and makes various historical plots of capacity used with a breakdown by customers and per month. This spreadsheet is uploaded to CS-DocDB monthly and the plots are used for planning and weekly operations reports. These are metric 4.

A spreadsheet predicting how much capacity is needed and how many slots and tapes is based upon the customer’s estimates for the fiscal year. This is periodically updated in DocDB. This is the source for metric 3 and is updated at least once per year.

Appendix A: Capacity Plan Approval

The undersigned acknowledge they have reviewed the **Fermilab Capacity Plan** and agree with the approach it presents. Changes to this **Capacity Plan** will be coordinated with and approved by the undersigned or their designated representatives.

Signature: Robert M Harris   Date: July 1, 2013

Print Name: Robert M. Harris

Title: Scientist II

Role: IT Capacity Manager

Signature: ____________________   Date: ____________

Print Name: ____________________

Title: ____________________

Role: ____________________

Signature: ____________________   Date: ____________

Print Name: ____________________

Title: ____________________

Role: ____________________

APPENDIX B: Capacity Reports

[Many of you in CCD have chosen to do Quarterly Reports. For each metric you listed above, you should report on current usage and any trends you want to call out because it may be an issue. These reports are usually placed into DocDB. You can simply reference those reports in this section with a DocDB number.

Do not list reports that you are not in fact creating. Document the as-is state of reporting.

Add rows to the table as necessary.]

The following table summarizes the documents referenced in this document.

Document Name and Version	Description	Location
<Document Name and Version Number>	*[Provide description of the document]*	*<URL or Network path where document is located>*