

GENERAL			
This document is under the Change Management Control Policy.			
Description	Service Level Agreement for Scientific Data Storage and Access		
	Service Level agreement for the services offered under Scientific Data Storage and Access		
Purpose	The purpose of this document is to publish agreed service level commitments between the service owner and the service customers.		
Applicable to	<i>All processes</i>		
Supersedes	<i>N/A</i>		
Document Owner	Gene Oleynik	Owner Org	<i>Scientific Computing Division</i>
Effective Date	November 2013	Review Date	Annually

VERSION HISTORY				
Version	Date	Author(s)	Approved by (if needed)	Change Summary
1.1	11-15-2013	Gene Oleynik	Gene Oleynik	Third Revision
1.2	11-07-2014	Gene Oleynik	Gene Oleynik	Section on data integrity considerations
1.3	2-27-2015	Gene Oleynik	Tammy Whited Panagiotis Spentzouris	New service offering: Active Archive Facility

TABLE OF CONTENTS

INTRODUCTION	4
1.1 EXECUTIVE SUMMARY	4
2 SERVICE AREA OVERVIEW	4
2.1 SERVICE AREA DESCRIPTION.....	4
2.2 SERVICE OFFERINGS	5
2.3 LIFECYCLE MANAGEMENT CONTEXT	10
3 RESPONSIBILITIES.....	11
USER RESPONSIBILITIES	13
4 DATA INTEGRITY CONSIDERATIONS	14
5 COMPUTER SECURITY CONSIDERATIONS	15
6 SERVICE SUPPORT PROCEDURE.....	15
6.1 REQUESTING CD SERVICE SUPPORT	15
6.2 SPECIAL SUPPORT COVERAGE	16
6.3 SERVICE BREACH PROCEDURES	16
7 SERVICE COMMITMENTS.....	16
7.1 INCIDENTS – SOMETHING IS BROKEN	16
7.2 REQUESTS – SOMETHING IS DESIRED.....	16
7.3 STANDARD REQUESTS	17
7.4 SERVICE AVAILABILITY– THE SYSTEM IS USABLE WHEN EXPECTED.....	17
7.5 TICKET PRIORITIES	18
7.6 COMMITMENT AND TARGETS	19
7.7 CRITICAL INCIDENT HANDLING.....	20
8 CUSTOMER REQUESTS FOR SERVICE ENHANCEMENT	20
9 SERVICE CHARGING POLICY.....	20
10 SERVICE MEASURES AND REPORTING	21
APPENDIX A: SUPPORTED HARDWARE AND SOFTWARE.....	22
APPENDIX B: SLA REVIEW PROCEDURE.....	22

APPENDIX C: OPERATIONAL LEVEL AGREEMENT (OLA) CROSS-REFERENCE	23
APPENDIX D: UNDERPINNING CONTRACT (UC) CROSS-REFERENCE	23
APPENDIX E: TERMS AND CONDITIONS BY CUSTOMER.....	23
APPENDIX F: ESCALATION PATH.....	23
APPENDIX G: ITIL PROCESSES ACROSS SERVICE BOUNDARIES	23
G.1 INCIDENT MANAGEMENT	24
G.2 PROBLEM MANAGEMENT	24
G.3 CHANGE MANAGEMENT	24
G.4 RELEASE MANAGEMENT	24
G.5 CONFIGURATION MANAGEMENT	24
G.6 CAPACITY MANAGMENT	24
G.7 AVAILABILITY MANAGEMENT	24
G.8 SERVICE LEVEL MANAGEMENT	24
G.9 SUPPLIER MANAGEMENT	24
G.10 SERVICE CONTINUITY MANAGEMEN	24

INTRODUCTION

1.1 EXECUTIVE SUMMARY

This Service Level Agreement (“SLA”) for the Scientific Data Storage and Access services (SDSA) with the Fermilab community documents:

- The service levels and support availability provided for the offerings under Scientific Data Storage and Access services
- Service Availability
- Support Availability
- The responsibilities of the Service Owner, Customer(s), and Users
- Specific terms and conditions relative to the standard Service Offering.

NOTE: For the purposes of this document, **Customer** refers to the **organization** which requests and receives the service; **User** refers to those **individuals** within the customer organization who access the service on a regular basis.

The service levels defined in this agreement are in effect during normal operations, in the case of a continuity situation they may change.

2 SERVICE AREA OVERVIEW

Service Area Owner: Gene Oleynik

Service Overview

Scientific Data Storage and Access service provides direct file based tape storage, called Enstore, and disk cache storage, called dCache (both tape backed and non-tape backed). These offerings have tiers of service that accommodate different needs of customers.

2.1 SERVICE AREA DESCRIPTION

Enstore Tape Storage

Enstore tape storage provides for direct archival storage of files on tape media over the Fermilab Ethernet LAN. Enstore provides a client API to the end user for transferring their local files to tape. Customers can have a mixture of Tiers for storing their data.

Enstore Service Pre-requisites

Clients must be running an officially supported Fermilab Linux OS Systems. Customer computers directly accessing Enstore storage are required to be located on the Fermilab LAN. It is recommended that computers that will frequently be reading and writing to

tape have 10 GE network connectivity since the tape drive technology transfer speeds are over twice what 1 GE Ethernet can handle. Transfer to/from Linux systems is supported. Windows or MACs are not supported for direct transfer to/from tape. Client nodes will typically need to be registered to be able to directly access files on tape.

dCache Disk Cache Storage

dCache storage provides high-performance, low latency file based storage on disk. The storage has a finite lifetime as dCache is a caching file system. dCache can automatically back up files to tape using Enstore (any Enstore Tier is possible), and retrieve them from tape when they are not resident on disk. Files can be transferred to/from dCache using a number of protocols.

For full client support, clients must be running an officially supported Fermilab OS and APIs supported by dCache and Fermilab. Transferring large amounts of data over the WAN will require consultation with Fermilab networking.

Active Archive Facility Storage

We provide an active archive facility for external customers that are not part of the Fermilab program. These services are provided as a combination of Enstore and dCache, and transfers in and out of Fermilab use either one of the supported dCache WAN transfer protocols. This service is negotiated with customers and includes a formal statement of work and a cost model that is different from that used for internal customers.

The Enstore Small File Aggregation (the Enstore “+ tiers”) is currently only available on-site and is not part of the Archive facility offering.

2.2 SERVICE OFFERINGS

The following Service Offerings are available:

2.2.1 STANDARD SERVICE OFFERINGS

2.2.2 Enstore Tape Storage

Offering Owner: Gene Oleynik

Offering Overview: Enstore tape storage provides archival storage of files on tape media over the Fermilab Ethernet LAN. Performance level is high latency.

Offering Description: Enstore provides a Linux client API, called encp (ENstore CoPy), to the end user for transferring their files to and from tape. The

files are stored using a namespace called Chimera, which is shared with the dCache service described later in this document.

Tape storage has its own peculiar challenges: inefficiency for small files due to its serial access, the rare potential for mechanical damage to the delicate tape, and resource usage inefficiencies due to limited client bandwidth while they have exclusive access to the drives. Enstore addresses these challenges with tiered feature sets. Customers can have a mixture of Tiers for storing their data. Enstore provides organization abstractions - file-families, storage groups, and a hierarchical name space for this purpose.

Tier 2: Basic tape storage

- Description: This storage tier provides a single copy of user's data. Data can be read from and written to tape by users over the Local Area Ethernet using the Enstore encp program. Transfer resources are provided from a pool of shared drives. Support is 24x7 on the tape libraries, 8x5 on the tape drives, and 8x5 on tapes.
- Criteria: If the user's data is statistical in nature and spans many tapes, or the customer provides a second copy of their data independently, then the single copy storage of customer's data with Tier 1 may be appropriate.

Tier 2+: Basic Tape Storage with Small File Aggregation

- Description: This is the same as Tier 2 except that all or a subset of the customer's files will be aggregated on disk and packaged into larger files before being written to tape. Support is 24x7 on the tape libraries, 8x5 on the tape drives, 8x5 on tapes, and 8x5 on files in the Small Files Cache. Space and resources for SFA is limited and this feature may not be available for all customers. It is therefore good idea to ask for this service far in advance.
- Criteria: If a customer's data or subset of their data consists of small files (less than several GB in size) then it may be appropriate to aggregate them before writing them to tape. This is more efficient and results in less wear and tear on tape drives. It does require more central disk resources though and needs to be negotiated with DMS.

Tier 1: Tape Storage with Multiple copies

- Description: N copies of a user's files are made, with N typically set to 2. The cost to the customer is N times the cost of Tier 2 storage. Copies can be made in tape libraries in the Feynman Computing Center, or between FCC and the tape libraries in the Grid Computing Center. Support is 24x7 on the tape libraries, 8x5 on the tape drives, and 8x5 on tapes.
- Criteria: This Tier is meant for data that may not be statistical in nature, is critical single copy data, or is statistical in nature, but doesn't span many tapes (such that one tape may hold a significant portion of the customer's data).

Tier 1+: Tape Storage with multiple copies and Small File Aggregation

- Description: This is the same as Tier 1 except that all or a portion of the customers files will be aggregated on disk and packaged into larger files before being written to tape. If multiple copies are specified the packaged files will be duplicated. Support is 24x7 on the tape libraries, 8x5 on the tape drives, and 8x5 on tapes, and 8x5 on files on the Small Files Cache. Criteria: This Tier is meant for customers who have both small files and need data redundancy,

Some data may have multiple copies and some not, some subset of files that are small may be aggregated and some not.

Enstore Service Pre-requisites:

Clients must be running an officially supported Fermilab Linux OS Systems. Customer computers directly accessing Enstore storage are required to be located on the Fermilab LAN. It is recommended that computers that will frequently be reading and writing to tape have 10 GE network connectivity since the tape drive technology transfer speeds are over twice what 1 GE Ethernet can handle. Transfer to/from Linux systems is supported. Windows or MACs are not supported for direct transfer to/from tape. Client nodes will typically need to be registered to be able to directly access files on tape.

Offering Support Availability: 8:00AM to 5:00PM Monday to Friday

24x7 support for servers, libraries and other critical components for all tier 1 and 2 services or if more than 50% of the drives in a logical tape library go offline.

Offering Cost: The customer bears the full cost for storage, license fees, software maintenance and hardware maintenance. An engagement between the customer or a designee and the service owner is necessary. The goal of the meeting is to define the customer requirements with respect to storage space and performance. These requirements will be mapped onto storage tiers with specific storage configurations. The combination of storage tiers, configurations and ultimately storage resource utilization will determine the cost of the service to the customer.

2.2.3 dCache Disk Storage Service

Offering Owner: Gene Oleynik

Offering Overview: dCache storage provides high-performance, low latency file based storage on disk. The storage has a finite lifetime as dCache is a caching file system. dCache can automatically back up files to tape using Enstore (any Enstore Tier is possible), and retrieve them from tape when they are not resident on disk. Files can be transferred to and from dCache using a number of protocols.

Offering Description:

Tier 1: Customer dedicated tape backed storage pools.

- Description: This Tier provides storage space dedicated to the customer. With this tier, dedicated storage hardware is purchased and managed on the customer's behalf with the appropriate transfer of money. This Tier of dCache requires the selection of a Tier of Enstore Service.
- Criteria: If a customer needs a guaranteed amount of cache storage or needs to keep a data set on low latency disk for a prolonged period or has high throughput needs, then they should consider this tier.

Tier 2: Shared tape backed storage pools:

- Description: This tier provides cache storage space shared with other users. This Tier of dCache services also requires the selection of a Tier of Enstore Service.
- Criteria: If a customer is writing files infrequently or does not need low latency read access to a set of files for prolonged periods then Tier 2 can provide them with tape backed low latency file access.

Tier 3: Volatile cache storage.

- Description: As the name implies, this tier is not tape backed.
- Criteria: This tier is typically used to temporarily store files. volatile storage, which is used. An example is Fermigrid, which uses a volatile pool in the Public dCache system for input and temporary output for a batch jobs.

dCache Service Offering pre-requisites:

For full client support, clients must be running an officially supported Fermilab OS and APIs supported by dCache and Fermilab. Transferring large amounts of data over the WAN will require consultation with Fermilab networking.

Supported Protocols:

- WAN access: SRM, GSI (Grid) FTP, WEBDAV, kerberized FTP
- LAN access: NFS V4.1, dcap (a Posix-like protocol), weak FTP

WAN based protocols can be used to transfer files to/from tape from offsite by going through dCache.

Offering Support Availability: 8:00AM to 5:00PM Monday to Friday

24x7 support for servers, libraries and other critical components for all tier 1 and 2 services or if more than 50% of the drives in a logical tape library go offline.

Offering Cost: The customer bears the full cost for storage, license fees, software maintenance and hardware maintenance. An engagement between the customer or a designee and the service owner is necessary. The goal of the meeting is to define the customer requirements with respect to storage space and performance. These requirements will be mapped onto storage tiers with specific storage configurations. The combination of storage tiers, configurations and ultimately storage resource utilization will determine the cost of the service to the customer.

24x7 support is provided for the following (configuration) items, which are required for these service offerings to be available:

- *Public Dcache*
- *CDF Dcache*
- *D0 Dcache*
- *CDFen*
- *STKen*
- *D0en*
- *stkensrv2n*
- *cdfensrv2n*
- *d0ensrv2n*

2.2.4 Active Archive Facility Service

Offering Owner: Gene Oleynik

Offering description: This service is a combination of dCache and Enstore that provides archival storage to off-site customers. The customer transfers their data to/from tape-backed Fermilab dCache through a dCache WAN protocol. Enstore Small File Aggregation is not available for this offering.

Active Archive Facility is part of the Public Enstore and dCache storage instances so scheduled Public downtimes apply.

Offering Support Availability: 8:00AM to 5:00PM Monday to Friday

Offering Cost: A cost model is provided to the customer and the customer makes yearly payments based on the amount of data they have stored, the amount of media they will require, and any incidental costs they might incur. External customers are charged for the media they will use in the year, and a per Terabyte cost (costs besides incidental and media) that is applied to the total amount of data stored at the end of the year. Details of the media, per TB storage cost, incidental cost, cost transfer, and requirements are covered in a Statement of Work (SOW) between the external customer and the Scientific Computing Division.

8x5 support is provided. Incidents and requests are provided through The Fermilab Service Now service desk (see the Service Support Procedure section below).

2.2.5 ENHANCED OFFERINGS

N/A

2.3 LIFECYCLE MANAGEMENT CONTEXT

Plan

The Service Owner, along with the customer, will help plan and requisition the proper storage required to meet the customer's needs. The equipment, including tape media, will be fully managed and maintained by the Service Owner and through its OLAs and underpinning contracts.

DMS/SDSA is the dCache, Enstore and Active Archive Facility Service Owner and provides dCache service application support, including capacity and availability planning. SDSA has an OLA with the ECF (Experiment Computing Facilities) department to provide for dCache hardware procurement, deployment, management, OS and system support. FEF is the Service Owner for providing dCache system and hardware support. In the sections below, Service Owner can refer to either ECF (dCache hardware and system), or DMS (Enstore, Enstore hardware and dCache application)

Purchase

The Service Owner will create purchase requisition orders along with the required documentation. S/He will coordinate with the Building Facilities Managers to ensure that adequate floor space, power and cooling are available for the equipment. S/He will coordinate with procurement, receiving, PREP and the vendor to ensure the proper installation of the equipment into the Fermilab Datacenter(s). S/He will create purchase requisitions for tape media as needed for capacity or migration.

Deploy

Storage resources will be deployed in accordance to the Plan developed initially between the Service Owner and the customer.

Manage

The Service Owner will manage and maintain the operational integrity of the hardware and software required to maintain the storage service to the customer. This includes implementing/coordinating repairs, upgrades and replacements as necessary.

Retire/Replace

The Service Owner coordinates storage disposal with PREP. PREP ensures the destruction of data on storage devices prior to disposal in accordance with lab security policies.

Equipment that is no longer supportable by the vendor (End-of-Support) must be replaced. The customer agrees to pay for the replacement cost for the storage that s/he is using.

Tape drive technology increases in capacity per cartridge on a cycle of every two or three years. The Service Owner generally refreshes the media technology every one to three cycles and migrates the data from the older technology to the new technology. This reduces footprint in the library and library cost at the expense of new drives and (usually) new media. The customer agrees to pay for the differential cost migrating to the new technology.

The customer agrees to pay for yearly maintenance costs. For the tape libraries, this includes per slot/year maintenance and may include tape drive maintenance as well. For external archive customers, this is in general included in their rate as specified in a Statement Of Work.

3 RESPONSIBILITIES

The Customer, Users, and Service Providers are expected to abide by applicable Fermilab policies, including but not limited to:

- [Fermilab Policy on Computing](#)

- [Guidelines for Incidental Computer Usage](#)
- [Fermilab Human Rights Policy](#)
- [Fermilab Director's Policy Manual](#)

[In particular, from these documents,](#) it is up to the customer to assure that the data their users store at Fermilab:

1. Does not support legally prohibited activities
2. Is not offensive and will not result in public embarrassment to Fermilab
3. Is not sexually explicit
4. Does not violate a license or other computer related contract provisions
5. Is not Personally Identifiable Information (according to the Fermilab PII policy)

Fermilab has the right to remove any data not meeting these criteria.

Customer Responsibilities

The Customer agrees to:

- Convey pertinent information to the users about the content of this service agreement.
- Participate in customer specific SLA reviews
- Provide representation for Continual Service Improvement (CSIP) activities. CSIP activities can be triggered in the event of an SLA breach or as part of normal Service Owner/Customer meetings. During this time, the customer and Service Owner can discuss what services are working well, which are not, and come up suggestions as to what areas need improvements. During this time, the Service Owner may also discuss with the customer upcoming Service improvements/changes/additions and poll the Customer for an opinion regarding these topics.
- The customer is the de-facto Data Owner and Data Manager for the storage, and the "Fermilab Policy on Computing" policies on these responsibilities apply. The customer may designate any one task or both tasks to someone else. If the customer elects to transfer this responsibility, the customer should also notify the Service Owner of this change.

Data Owner responsibilities include:

- Handling transfer of ownership requests for files/folders/directories
- Specifying which users have access to files/folders/directories
- Quota increase requests
- Receiving announcements of service disruptions and communicating them to their users.
- Working with the service provider to provision and establish the service for the customer
- Authorizing NFS exports to new customer nodes (internal users)
- Requesting new user authorization (for dCache)

- Authorizing of the physical deletion of files on tape (tapes that contain only files marked deleted can be recycled)
- Insuring that the users have obtained or know how to obtain proper security credentials and other authentication/authorization items needed to access the storage services and that they are registered with the proper Fermilab authentication services. Examples of the former are X.509 certificates and Fermilab uid/gid and the latter the Fermilab Gums service and dCache gPlazma service.
- Ensuring that the SDSA services performance, availability and reliability will meet their requirements.

Data Manager responsibilities include:

- Changing permissions on files/folders/directories
 - Changing ownership of files/folders/directories
 - Creating folder/directory hierarchy/structure
 - Deleting files/folders/directories (for housekeeping)
- Programmatically assure that their files have been successfully written to tape before deleting any local copies. There are several ways to do this and the exact means can be worked out with the service provider.
 - If the data is critical, ensuring that there is enough local storage to be able to buffer at least a 24 hour outage.
 - Provide and maintain their own catalog of files with cross-reference to the names of files in the Enstore/dCache namespace. The Enstore/dCache namespace is not a catalog and treating it as such for searches can lead to DoS to other users.
 - Coordinate standard maintenance downtimes requiring a service outage to occur every 3rd Thursday of the month between the hours of 6:00AM – 4:00PM. Notification of a service outage will be provided to the customer via email and Operations meeting at least 2 weeks in advance of an outage.

USER RESPONSIBILITIES

The users agree to:

- Read user documentation and consult with the Service provider on how to use APIs to access the storage. It is very important to handle errors on writes to the storage system to make sure the file is successfully written to tape..
- In case of issues, read documentation provided on the web (<http://www-ccf.fnal.gov/enstore/documentation.html>)
- To not create a denial-of-service (DoS) through improper controls for dCache or the shared dCache/Enstore namespace client access. Examples are flooding dCache with thousands of small files in a short period of time and using “ls” on the namespace to find files.
- Report incidents and service requests via the service desk (<http://servicedesk.fnal.gov>)
- When reporting an incident, be available for support (call back via support staff)

Scientific Data Storage and Access services manager agrees to:

General responsibilities:

- Create new shares/NFS exports for access to the storage as requested by the Data Owner and authorized by the customer contact.
- Service requests for authorization of new dCache users
- Maintain appropriately trained staff.
- Coordinate standard, non-emergency, maintenance downtimes requiring a service outage to occur once per month on the 3rd Thursday¹ of the month between the hours of 6:00AM – 4:00PM. Notification of a service outage will be provided to the customer via email and Operations meeting at least 2 weeks in advance of an outage (unless deemed and emergency).
- Provide usage statistics, plots and service metrics on a daily basis which include:
 - Total and historical amount of data stored
 - Tape drive utilization
 - dCache storage and bandwidth usage
 - Request and incident ticket status and response and resolution statistics through Service Now
- Provide Monthly usage and billing information for Active Archive Facility customers

Provide a 99% availability for scheduled downtimes with at least a 2 week notice for service outages. 99% availability corresponds to up to 7 hours of downtime/month.

4 DATA INTEGRITY CONSIDERATIONS

Tape technology is mechanical in nature and sometimes, rarely, a file or a set of files may get damaged and become unreadable or corrupted. If this happens, SDSA administrators will make a best effort attempt to recover all files possible and copy recovered files to new, undamaged media. In this process some file may not be recoverable. In this case the customer will be notified as soon as possible of the unrecoverable files so that they may recover them from an alternate source, if available. The procedures SDSA perform for recovery are documented in CS-doc-5468.

It also may be possible, though it is unlikely, that a vendor that specializes in recovery (e.g. Oracle for Oracle tapes), may be able to recover some files. This may or may not have a cost associated with the attempted recovery. In the case that SDSA cannot recover files, we will consult with the customer on whether or not we should request a vendor to make further attempts at recovery.

The table below shows the result of an analysis of data integrity over the two year period of January 2013-January 2015 at the two computing centers with tape libraries at Fermilab.

¹ D0 and CDF instances of the storage service are on the first Tuesday and second Wednesday of the month, respectively.

Active Archive and Intensity Frontier files are stored in the Feynman Computing Center (FCC) tape libraries, with some duplicate copies kept in the Grid Computing Center (GCC) Tape Robot Room. The Grid Computing Center library complexes holds all CMS data, Run II RAW data, RMAN backup data, and as already mentioned, copies of FCC data.

Most losses occurred in the GCC TRR. There were two sources of data loss in that location: environmental and a mechanical issue. There has been no environment issue in the Feynman Computing center, and over the last year, a lot has been done to improve the environment in GCC TRR to prevent such occurrences in the future. There were no losses in either center due to software or operator error.

Library	Files Read	Unique Files Read	Unrecoverable file loss	Loss per million unique files read	Reliability
FCC	20,964,183	19,132,543	14*	0.73	99.999927
GCC	13,802,540	11,301,288	775**	68.58	99.9931
Total	34,766,723	30,433,831	789	25	99.9975

* 13 of these were from one LTO4 tape, which was sent out for data recovery. None of the 13 could be recovered. The cause for the file loss was undetermined.

** On one tape with environmental contamination, we stopped reading after 46 errors and with 3275 files remaining to be read and sent the tape to Oracle (most of these were recovered from a CMS copy and the tape was destroyed for analysis of the material on it). Only 46 files are counted for this tape, but there may have been more. Another tape was damaged by a tape drive. This tape had 15,707 files on it, and all but 537 of these were recovered. A firmware update, which prevents drives from damaging tapes in this manner, has been applied to all of our tape drives.

5 COMPUTER SECURITY CONSIDERATIONS

Computer Security incidents are to be reported to the FNAL Service Desk as defined by the Fermi Incident Response (FIR) Procedure located at <http://security.fnal.gov>.

6 SERVICE SUPPORT PROCEDURE

6.1 REQUESTING CD SERVICE SUPPORT

Access to all Computing Sector IT services should be requested through the Service Desk, via the [ServiceNow](#) application, or by phone (630-840-2345). More information about requesting service can be found in the Self Service section of ServiceNow.

External customers can use mail for incident reporting and service requests. Incident reports can be sent to aaf-incident@fnal.gov and service requests can be placed via aaf-request@fnal.gov. These incidents and requests will get assigned to the Storage

Services group by Service Now and assigned a medium priority. An alternative is to call 630-840-2345 during work hours (Central time) and specify that you are an external customer and that the ticket should be assigned to the “Storage Service” group for the “Active Archive Facility” service offering. Please include a detailed description of the incident or request.

6.2 SPECIAL SUPPORT COVERAGE

Requests for changes in support coverage should be made by opening a request with the Service Desk a minimum of 7 days before the coverage change is needed.

These requests must be negotiated and are subject to approval based on the staff available at the time and the nature of the additional support.

6.3 SERVICE BREACH PROCEDURES

Breaches in service are defined as not meeting agreed to commitments over a month’s time. Breaches are recorded, classified and reviewed on a monthly basis utilizing the Service Level Management process. Breaches and opportunities for improvement are available on the Monthly Service Performance Report.

7 SERVICE COMMITMENTS

7.1 INCIDENTS – SOMETHING IS BROKEN

- Incident Response: Ticket acknowledgement, which is measured during expected support availability window. Critical Incident response is measured 24 x 7.
- Incident Resolution: Ticket resolution, which is measured during expected support availability window. Critical Incident resolution is measured 24 x 7.

7.2 REQUESTS – SOMETHING IS DESIRED

Standard Request

Standard Requests are defined and have expected outcomes. Standard Request items are ordered in the Service Catalog.

- Request Fulfillment: Amount of time to deliver the item requested, which is measured during core business hours, 8a.m. to 5p.m Monday through Friday.

General and Non-Standard Requests

Request that may require negotiation or are outside of the standard defined items offered by the service.

- Request Response: Ticket Acknowledgement, which is measured during core business hours, 8a.m. to 5p.m. Monday through Friday.

There is no resolution time target for requests due the wide variety of requests being handled.

7.3 STANDARD REQUESTS

Enstore Tape Storage

- Enstore Storage Request
 - 8 business hour response target
- Add NFS Export Request
 - 8 business hour response target

dCache Disk Storage

- dCache Storage Request
 - 8 business hour response target
- Add dCache Users
 - 8 business hour response target

7.4 SERVICE AVAILABILITY– THE SYSTEM IS USABLE WHEN EXPECTED

System availability is measured as an uptime percentage during the expected service availability window. An Outage implies system unavailability and negatively impacts availability measurements. *An Outage during an 'agreed to maintenance window' does not impact the availability measurement.*

Enstore Tape Storage

- Maintenance Window – The maintenance window for Enstore Tape Storage depends on the storage type.
 - Public Storage: 3rd Thursday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
 - CDF Storage: 2nd Tuesday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
 - D0 Storage: 1st Tuesday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
- Outage – Enstore is stopped. No storage i/o requests are being taken and no requests are being processed.
- Degradation – Enstore is running but a particular library manager may not be running.

dCache Disk Cache StorageTape Storage

- Maintenance Window – The maintenance window for dCache storage depends on the storage type.
 - Public Storage: 3rd Thursday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
 - CDF Storage: 2nd Tuesday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
 - D0 Storage: 1st Tuesday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
 -
- Outage – dCache is stopped. No request are being taken and no requests are being processed.
- Degradation – A door (transfer protocol) or a pool is unavailable.

Active Archive Facility Service

- Maintenance Window –
 - 3rd Thursday of the month, lasting between 4 and 8 hours, beginning at 8:00a.m. CT.
- Outage –
 - The services are stopped, no requests are being processed.
- Degradation –
 - The service is degraded and may operate more slowly. Enstore may be unavailable, but dCache available. In that case transfer to the archive may proceed to dCache disk, but will be delayed in copying to tape until after Enstore is available.

7.5 TICKET PRIORITIES

- INCIDENTS – SOMETHING IS BROKEN
 - Incident **Priority** determines the target resolution commitment
 - **Priority** is a combination of Impact and Urgency
 - **Impact** is determined by the number of users experiencing the issue
 - **Urgency** is determined by how quickly the issue needs to be resolved for the lab to resume normal operations
- REQUESTS – SOMETHING IS DESIRED
 - Request **Priority** is pre-defined bases on what is being asked for
 - Currently, all Requests have the same target response time.

Impact \ Urgency	Extensive Service is out for Enterprise	Significant Service is out for many users or degraded for Enterprise	Moderate Service is out for 1 user or degraded for many	Localized Service is degraded for 1 user
Critical <i>Based on event</i>	Priority - Critical Respond – 1 H Resolve – 5 H	Priority - High Respond – 4 H Resolve – 35 H (1.5 D)	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)
High <i>Required</i>	Priority - High Respond – 4 H Resolve – 35 H (1.5 D)	Priority - High Respond – 4 H Resolve – 35 H (1.5 D)	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)	Priority - Low Respond – 8 H Resolve – 172 H (7 D)
Medium <i>Important</i>	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)	Priority - Low Respond – 8 H Resolve – 172 H (7 D)
Low <i>Desirable</i>	Priority - Medium Respond – 8 H Resolve – 97 H (4 D)	Priority - Low Respond – 8 H Resolve – 172 H (7 D)	Priority - Low Respond – 8 H Resolve – 172 H (7 D)	Priority - Low Respond – 8 H Resolve – 172 H (7 D)

Table 1. Response and Resolution Time Table

7.6 COMMITMENT AND TARGETS

Unless defined in the exception list, all Service Offerings in this agreement operate and manage to the commitments defined below. *See Foundation SLA for commitments and Targets.*

- Incidents
 - Response
 - 90% of Critical Incidents responded to within 1 clock hour
 - 90% of High Priority Incidents responded to within 4 hours, during the support availability window
 - 90% of Medium Priority Incidents responded to within 8 hours, during the support availability window.
 - 90% of Low priority tickets responded to within 8 hours, during the support availability window.
 - Resolution
 - 90% of Critical Incidents resolved within 5 clock hours
 - 90% of all High Priority Incidents resolved within 2 days, during the support availability window
 - 90% of all Medium Priority Incidents resolved within 4 days, during the support availability window
 - 90% of all Low Priority Incidents resolved within 7 days, during the support availability window

- Requests
 - General and Non-Standard Request Response

- 90% of Critical Requests responded to in 8 hours, 8a.m. to 5p.m. Monday through Friday.
 - 90% of High Requests responded to in 8 hours, 8a.m. to 5p.m. Monday through Friday.
 - 90% of Medium Requests responded to in 8 hours, 8a.m. to 5p.m. Monday through Friday.
 - 90% of Low Requests responded to in 8 hours, 8a.m. to 5p.m. Monday through Friday.
- Availability
 - Enstore Tape Storage: 95% Availability 24 x 7
 - dCache Disk Cache Storage: 95% Availability 24 x 7 hours by days
 - Active Archive Facility: 95% Availability 8x5

7.7 CRITICAL INCIDENT HANDLING

A critical incident is the highest priority incident, one in which a highly visible and important service depended upon by many users is no longer operable and there is no acceptable work-around. In addition to the faster response expectations listed in the priority/response table, critical incidents move to the front of the incident report queue and may be handled by a distinct Critical Incident Management process.

Critical incident response requires a phone call to the Service Desk (630-840-2345).

8 CUSTOMER REQUESTS FOR SERVICE ENHANCEMENT

Service enhancements are customer requests for planned changes in service, for example, providing for Small File Aggregation or changing storage pool affinity for dCache customers. It is required that the customer and Service Owner meet to fully understand the requirements and expectations from the enhancement. The customer will use these requirements to officially request a service enhancement via the service desk.

The Service Owner will respond to requests for service enhancements received with appropriate advance notice within 7 business days. This time is needed to discuss issues regarding power, cooling, support and budget to determine if the enhancement request is possible.

9 SERVICE CHARGING POLICY

The customer should work with Service provider to develop a budget for estimated costs of hardware required to provide the appropriate storage. Once established, this budget will be input into the Computing Sector Budget entry system.

10 SERVICE MEASURES AND REPORTING

The Service Offering dashboard is available in the service desk application under the report section. The dashboard measures each offering for each service against the incident response and resolution times and request response times defined in section 6.5 of this document. The dashboard shows performance trending for the Service Offerings on a weekly/monthly/yearly basis.

The Service Offering dashboard is available to Service Owners and Providers, Business Analysts, Process Owners and Senior IT Management.

Service Level breaches are identified in the service offering dashboard and are monitored by the Service Owners, Incident Manager and Service Level Manager.

Customer Reports are available in ServiceNow in the Service Management Reports section.

APPENDIX A: SUPPORTED HARDWARE AND SOFTWARE

- Sun Sparc Servers for ACSLS
- X86 Platforms running Scientific Linux from a variety of vendors
- Qlogic Fibre-Channel Adapters
- Nexsan Satabeast, Sataboy and E18 storage arrays
- Oracle Storage Tek SL8500 tape libraries
- Storage Tek/IBM LTO-4 tape drives
- Storage Tek T10000C tape drives
- Advanced HPC ZFS appliances for Small File Aggregation disk
- Newport Electronics wSeries Wireless Sensor System (temp and humidity monitoring)
- Officially supported Fermilab Scientific Linux OS
- Solaris OS for ACLS computers
- OpenSolaris for ZFS appliances
- NexentaStor ZFS management software for ZFS appliances
- Oracle ACSLS tape management software
- Oracle Tape Analytics software
- Oracle SDP (remote monitoring) software
- Newport Electronics Virtual Coordinator Software

APPENDIX B: SLA REVIEW PROCEDURE

The series of steps will be conducted on at least an annual basis. Either the Customer, Service Owner or Service Level Manager may request additional reviews as necessary.

1. At least one month prior to the expiration of this agreement, the customer will be sent notification via email requesting that a face-to-face review be conducted between the Service Owner and the Customer regarding the Service.
2. During the review, customer may negotiate changes to the Service Level Agreement with the Service Owner. Requests for changes are subject to approval based on the limitations of resources from the Service, a supporting organization, funding and effort available.
3. If additional meetings are required, those meetings will be held as necessary in order to renew the SLA prior to expiration.
4. Should an agreement not be reached prior to the expiration date, service will continue on a month to month basis using the existing SLA agreement requirements.

APPENDIX C: OPERATIONAL LEVEL AGREEMENT (OLA) CROSS-REFERENCE

This service depends on the following IT Services to operate within their respective SLAs / OLAs:

- Networking [CS-doc-4312](#)
- Authentication Services [CS-doc-4314](#)
- Service Desk [CS-doc-3716](#)
- Facilities [CS-doc-4594](#)
- OLA with the ECFdepartment [CS-doc-5249](#)
- OLA with DELL Managed Services [CS-doc-5032](#)

APPENDIX D: UNDERPINNING CONTRACT (UC) CROSS-REFERENCE

Supplier Lists, including contact information can be found [here](#).

Key vendor contracts supporting this service are:

- Oracle Maintenance Contract for Oracle Tape Storage hardware and software.
- Nexsan Maintenance contract for Nexsan storage arrays.
- Various white box server maintenance contracts.
- NexentaStor/ZFS appliance maintenance agreement with Advanced HPC

APPENDIX E: TERMS AND CONDITIONS BY CUSTOMER

N/A

APPENDIX F: ESCALATION PATH

This section outlines the Escalation Path to be followed if there is a specific workflow for this service. The default escalation path is defined in the Foundation SLA as hierarchic (line management).

APPENDIX G: ITIL PROCESSES ACROSS SERVICE BOUNDARIES

This section defines how the services will work together to execute ITIL processes across the service boundaries, if and only if there are variations or unexpected responsibilities compared to the overall ITIL process definitions. Perhaps there is some

historical variation to capture, or an explicit agreement to blur service boundaries to make more efficient use of resources or talents. The list below is a suggestion of topics to consider. Include only those areas needed.

- G.1 INCIDENT MANAGEMENT**
- G.2 PROBLEM MANAGEMENT**
- G.3 CHANGE MANAGEMENT**
- G.4 RELEASE MANAGEMENT**
- G.5 CONFIGURATION MANAGEMENT**
- G.6 CAPACITY MANAGEMENT**
- G.7 AVAILABILITY MANAGEMENT**
- G.8 SERVICE LEVEL MANAGEMENT**
- G.9 SUPPLIER MANAGEMENT**
- G.10 SERVICE CONTINUITY MANAGEMENT**