

General			
This document is under the Change Management Control Policy.			
Description	Service Catalog Document for High Performance Computing. Describes the services offered under High Performance Computing.		
Purpose	The purpose of this document is to publish agreed service level commitments between the service owner and the service customers.		
Supersedes	<i>Service Level Agreement, Capacity Plan and Availability Plan for High Performance Computing</i>		
Document Owner	Amitoj Singh	Owner Org	<i>Scientific Computing Sector</i>
Effective Date	04-01-2016	Review Date	<i>At minimum annually</i>

Version History				
Version	Date	Author(s)	Approved by Change Number	Change Summary
0.1	03-01-2016	Amitoj Singh		Initial document

TABLE OF CONTENTS

EXECUTIVE SUMMARY	4
1 SERVICE AREA OVERVIEW	4
2 SERVICE OFFERINGS	6
2.1 Parallel and Tightly Coupled Batch Computing	6
2.2 Application Support	8
2.3 User Accounts	8
2.4 File-System Support	8
3 SERVICE CAPACITY	11
3.1 Business Capacity Management	11
3.2 Service Capacity Management	12
3.3 Component Capacity Management	13
4 BUSINESS REQUIREMENTS, SERVICE ENTITLEMENTS AND COST	13
4.2 Business Requirements	13
4.3 Service Entitlements	13
4.4 Service Charging Policy	13
5 SERVICE REQUESTS	14
5.2 Standard Requests	14
6 SERVICE COMMITMENTS	15
6.2 Service Availability	15
6.2.1 Parallel and Tightly Coupled Batch Computing	16
6.2.2 Application Support	16
6.2.3 User Accounts	16
6.2.4 File-System Support	16
6.3 Other Service Levels	17
7 SERVICE SUPPORT	17
7.2 Requesting Service Support	17
7.2.1 Special Support Coverage	18
7.3 Customer requests for Service Enhancements	18
8 SERVICE LIFECYCLE	18
9 RESPONSIBILITIES	19
9.2 General Responsibilities	19
9.3 Service Specific Responsibilities	19

9.3.1	CUSTOMER RESPONSIBILITIES	19
9.3.2	USER RESPONSIBILITIES	19
9.3.3	SERVICE OWNER	20
10	SERVICE CONTINUITY	20
10.1	Recovery Strategy	21
10.2	Strategy for initial recovery	21
10.3	Overall recovery strategies	21
10.4	Recovery Scenarios	22
10.4.1	Building not accessible (Data Center Available)	22
10.4.2	Data Center Failure (Building Accessible)	22
10.4.3	Building not accessible and Data Center Failure	23
10.4.4	Critical recovery team not available	23
10.4.5	Government Mandated Shutdown of Services	24
10.5.4	Return to Operations	24
11	SERVICE MEASURES AND REPORTING	24
11.2	Standard Service Measures and Reports	24
11.3	Service specific Measures and Reports	25
	APPENDIX A: SUPPORTED HARDWARE AND SOFTWARE	26
	APPENDIX B: SLA AND OLA CROSS-REFERENCE	26
	APPENDIX C: SERVICE DEPENDENCY CROSS-REFERENCE	28
	APPENDIX D: UNDERPINNING CONTRACT (UC) CROSS-REFERENCE	28
	APPENDIX E: TERMS AND CONDITIONS BY CUSTOMER	28

EXECUTIVE SUMMARY

This document provides details and commitments of the High Performance Computing Service Area and Service Offerings.

The descriptions of the Service Area and Service Offerings together with their service commitments and targets, owner, support organization and the type of Foundation Service Level Agreement that they conform to, are maintained and controlled in the CMDB under change control. This document contains the approved service parameters extracted from the CMDB at the time of approval of the document. Future versions of this document will contain URL's to reports from the CMDB (Service Now) rather than embedded tables of data extracted under change control.

In addition to those parameters, common to all Services, this document contains specific terms and conditions of the services for this Service area

This document, together with the applicable Foundation Service Level or Operational Level Agreement, forms the Service Level Agreement "SLA" or Operational Level Agreement "OLA" (for internal service offerings) for these services with the Fermilab community. Taken together they fully describe the responsibilities of the Service Owner, Customer(s) and Users, the Service Levels, Service Commitments, Service Support and Service breach procedures, computer security responsibilities, and specific terms and conditions for the services described below.

1 SERVICE AREA OVERVIEW

Service Area:	High Performance Computing
Service Area Owner:	Amitoj Singh
	Design, procure and operate state-of-the-art tightly coupled parallel computing resources to meet the computing needs of scientists from the USQCD collaboration, the Astrophysics and Accelerator Modeling communities.
	Not ISO20K certified

The High Performance and Parallel Computing (HPC) department designs and operates state-of-the-art tightly coupled and parallel computing resources to help scientists and engineers from Fermilab and University communities meet their computing needs most effectively and in a highly cost effective manner. HPC department currently operates two primary facilities as listed below:

USQCD facility: The HPC department operates this world-class facility in support of the US DOE Office of Science (SC) program to build and operate computational hardware infrastructure for the lattice field theory. The USQCD facility at Fermilab provides computing to the USQCD national collaboration. HPC department designs, procures and operates several Infiniband-

based clusters that are highly performance-to-cost optimized for the tightly coupled and parallel computations of lattice field theory. The USQCD facility at Fermilab also operates GPU-accelerated clusters which support many-GPU parallel computing. Cache disk storage is provided through the Lustre parallel and distributed file system. Long term archival storage is provided through the Fermilab Enstore tape backed file system.

Wilson Facility: The Wilson facility is primarily used for simulations and engineering studies of cosmic rays, particle beams and accelerators. Stakeholders in this facility include the Accelerator Physics Center, the Theoretical Astrophysics Group, the Accelerator Simulation Group and members of the Fermilab Technical Division. A modest sized Intel Phi-based and NVIDIA GPU-based cluster for software development is primarily used by the various Muon and Neutrino experiments, current and next generation accelerator simulations and toolkit designers. Cache disk storage is provided through a modest sized Lustre parallel and distributed file system.

Service Offering	Short Description	Offered	Owner
USQCD Facility Parallel and Tightly Coupled Batch Computing	This consists of the batch queuing system that provides rule-based access to the CPU-only and the accelerated HPC compute clusters for the USQCD Computing Facility.	Customer-facing	Amitoj Singh
USQCD Facility Application Support	Assisting users with the installation of a collaboration or experiment specific application on the facility clusters. This also includes diagnosing application failures when running jobs on the facility clusters.	Customer-facing	Amitoj Singh
USQCD Facility User Accounts	Creating new user accounts on the various facility clusters for a user with a valid Kerberos principal. Also assisting users with any remote access issues especially with Kerberos and Secure SHell. Creating project specific groups and modifying directory and file permissions as requested.	Customer-facing	Amitoj Singh
USQCD Facility File-System Support	Diagnosing performance and access related issues with the Parallel & Distributed file-systems (Lustre) or the Networked File Systems (NFS). Also includes first order diagnosis of failures related to file transfers especially to remote sites using Globus Online.	Customer-facing	Amitoj Singh

Wilson Facility Parallel and Tightly Coupled Batch Computing	This consists of the batch queuing system that provides rule-based access to the CPU-only and the accelerated HPC compute clusters for the Wilson Computing Facility.	Customer-facing	Amitoj Singh
Wilson Facility Application Support	Assisting users with the installation of a collaboration or experiment specific application on the facility clusters. This also includes diagnosing application failures when running jobs on the various facility clusters.	Customer-facing	Amitoj Singh
Wilson Facility User Accounts	Creating new user accounts on the various facility clusters for a user with a valid Kerberos principal. Also assisting users with any remote access issues especially with Kerberos and Secure SHell. Creating project specific groups and modifying directory and file permissions as requested.	Customer-facing	Amitoj Singh
Wilson Facility File-System Support	Diagnosing performance and access related issues with the Parallel & Distributed file-systems (Lustre) or the Networked File Systems (NFS). Also includes first order diagnosis of failures related to file transfers especially to remote sites using Globus Online.	Customer-facing	Amitoj Singh

2 SERVICE OFFERINGS

2.1 USQCD Facility Parallel and Tightly Coupled Batch Computing

Parallel and tightly coupled batch computing consists primarily of the batch queuing system for the various facility clusters. Each cluster in the USQCD facility has a dedicated login and server head node. The batch queuing system consists of the batch server (Torque) and scheduler (Maui) that run on the dedicated server head node. Users access the cluster resources by submitting jobs to the batch queuing system by logging in to the corresponding cluster login head node. The cluster compute nodes are available as a pool of resources to the users through the batch compute server.

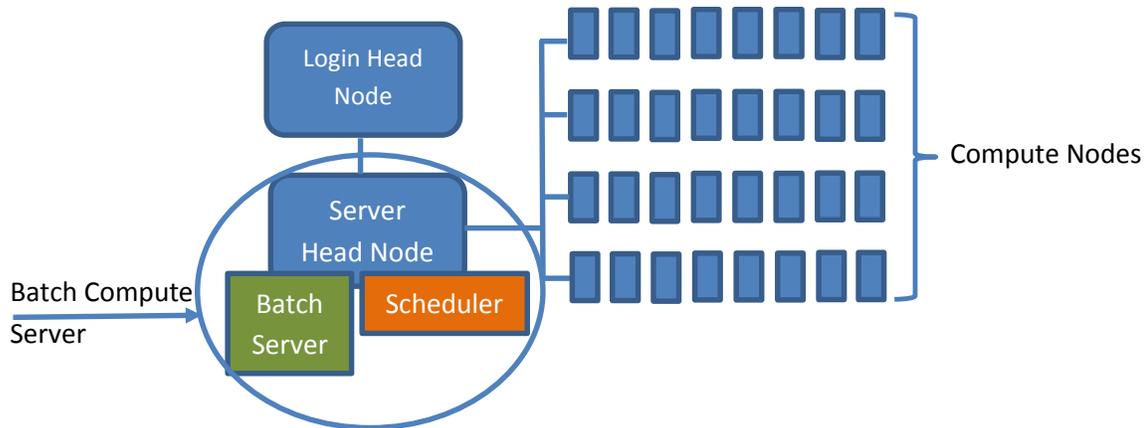


Figure 1: Batch Computing System for the USQCD Facility

This service offering is further tiered as follows:

- 2.1 USQCD Facility Parallel and Tightly Coupled Batch Computing
 - 2.1.1 CPU Computing
 - 2.1.2 Accelerated Computing

USQCD Facility Parallel and Tightly Coupled Batch Computing	
Other Information	The customer is the de-facto owner and information manager for the outputs of jobs that are submitted to completion through the batch queuing system. Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 2 - somewhat critical
Supported by	High Performance Computing 8to17by5

2.2 USQCD Facility Application Support

Assisting users with the installation of a collaboration or experiment specific application on the facility clusters. This also includes diagnosing application failures when running jobs on the various facility clusters.

USQCD Facility Application Support	
Other Information	The customer is the de-facto owner and information manager for the outputs of their applications. Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 3 - less critical
Supported by	High Performance Computing 8to17by5

2.3 USQCD Facility User Accounts

Creating new user accounts on the various facility clusters for a user with an approved Fermilab ID and Kerberos principal. Also assisting users with any remote access issues especially with Kerberos and Secure SHell. Creating project specific groups and modifying directory and file permissions as requested.

USQCD Facility User Accounts	
Other Information	Users contact Fermilab Service Desk to fulfill their request for a valid Fermilab ID and Kerberos principal. The customer is solely responsible to keep their Fermilab ID and Kerberos principal up to date. Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 2 - somewhat critical
Supported by	High Performance Computing 8to17by5

2.4 USQCD Facility File-System Support

Diagnosing performance and access related issues with the Parallel & Distributed file-systems (Lustre) or the Networked File Systems (NFS). First order diagnosis of failures related to file transfers especially to remote sites using Globus Online.

This service offering is further tiered as follows:

2.4 USQCD Facility File-System Support

2.4.1 Parallel and Distributed File-System

2.4.2 Network File-System

2.4.3 File Transfer

2.4.3.1 Globus Online

USQCD Facility File-System Support	
Other Information	The customer is the de-facto owner and information manager for data stored on the various HPC managed file-systems. Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 2 - somewhat critical
Supported by	High Performance Computing 8to17by5

2.5 Wilson Facility Parallel and Tightly Coupled Batch Computing

Parallel and tightly coupled batch computing consists primarily of the batch queuing system for the various facility clusters. Each cluster in the Wilson facility has a dedicated login and server head node. The batch queuing system consists of the batch server (Torque) and scheduler (Maui) that run on the dedicated server head node. Users access the cluster resources by submitting jobs to the batch queuing system by logging in to the corresponding cluster login head node. The cluster compute nodes are available as a pool of resources to the users through the batch compute server.

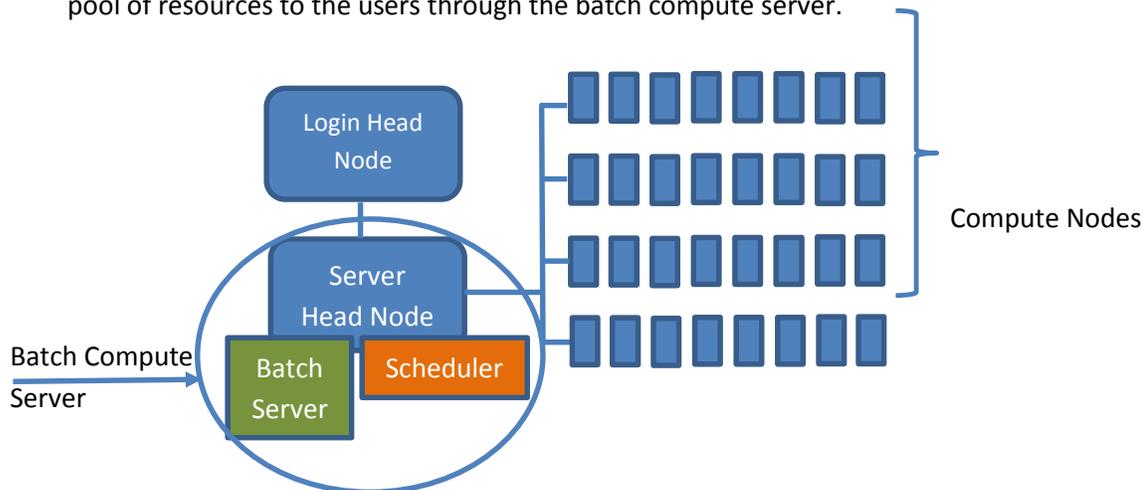


Figure 2: Batch Computing System for the Wilson Facility

This service offering is further tiered as follows:

2.1 Wilson Facility Parallel and Tightly Coupled Batch Computing

2.1.1 CPU Computing

2.1.2 Accelerated Computing

Wilson Facility Parallel and Tightly Coupled Batch Computing	
Other Information	The customer is the de-facto owner and information manager for the outputs of jobs that are submitted to completion through the batch queuing system. Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 2 - somewhat critical
Supported by	High Performance Computing 8to17by5

2.6 Wilson Facility Application Support

Assisting users with the installation of a collaboration or experiment specific application on the facility clusters. This also includes diagnosing application failures when running jobs on the various facility clusters.

Wilson Facility Application Support	
Other Information	The customer is the de-facto owner and information manager for the outputs of their applications. Does not qualify for Critical for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 3 - less critical
Supported by	High Performance Computing 8to17by5

2.7 Wilson Facility User Accounts

Creating new user accounts on the various facility clusters for a user with an approved Fermilab ID and Kerberos principal. Also assisting users with any remote access issues especially with Kerberos and Secure SHell. Creating project specific groups and modifying directory and file permissions as requested.

Wilson Facility User Accounts	
Other Information	Users contact Fermilab Service Desk to fulfill their request for a valid Fermilab ID and Kerberos principal. The customer is solely responsible to keep their Fermilab ID and Kerberos principal up to date.

	Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 2 - somewhat critical
Supported by	High Performance Computing 8to17by5

2.8 Wilson Facility File-System Support

Diagnosing performance and access related issues with the Parallel & Distributed file-systems (Lustre) or the Networked File Systems (NFS). First order diagnosis of failures related to file transfers especially to remote sites using Globus Online.

This service offering is further tiered as follows:

2.4 Wilson Facility File-System Support

2.4.1 Parallel and Distributed File-System

2.4.2 Network File-System

2.4.3 File Transfer

2.4.3.1 Globus Online

Wilson Facility File-System Support	
Other Information	The customer is the de-facto owner and information manager for data stored on the various HPC managed file-systems. Does not qualify for Critical Incident Handling
Type of SLA	Based on Foundation Agreement Business Criticality - 2 - somewhat critical
Supported by	High Performance Computing 8to17by5

3 SERVICE CAPACITY

3.1 Business Capacity Management

The objective is to translate business needs and plans into capacity and performance requirements for Computing services and infrastructure, and to ensure that future capacity and performance needs can be fulfilled.

The Lattice Quantum ChromoDynamics (LQCD) computing project funded by the DOE Office of Science provides funds for the acquisition and operation of computational systems that will serve as a principal computational resource for the national LQCD user community (USQCD). The current LQCD computing project titled "LQCD-ext II" is funded from FY15 through FY19. Using these computational resources, the LQCD theorists can better provide theoretical insight and guidance to the HEP community.

The *LQCD-ext II Project Execution Plan* (see docdb#5565) document describes the plan and related methodologies to be followed while executing the Lattice Quantum Chromodynamics (LQCD) computing facility project for the period FY2015 through FY2019.

Allocation of the USQCD Facility resources amongst various USQCD projects is done by the USQCD Software Program Committee. An annual call for proposals for awards of time on the USQCD computer resources dedicated to lattice QCD and other lattice field theories is sent out to all USQCD projects. The awards are for calculations that further the scientific goals of the collaboration, as laid out in recent white papers and USQCD proposals noting that an important reason for funding is relevance to the DOE experimental program.

For the Wilson Cluster Facility a model somewhat similar to what the Scientific Portfolio Management Team follows is used by the Wilson Cluster Facility oversight committee. The oversight committee members are representatives of key stakeholders in the Wilson Cluster Facility hardware. The key stakeholders are the Accelerator Physics Center, the Theoretical Astrophysics Group, the Accelerator Simulation Group and members of the Fermilab Technical Division. The oversight committee entertains requirements from all parts of the lab's scientific program and reviews the requests and needs and prioritizes the resources to be provided in the coming one or two years in each area, based on Fermilab scientific strategies and priorities.

3.2 Service Capacity Management

The objective is to manage, control and predict the performance and capacity of operational services. This includes initiating proactive and reactive action to ensure that the performances and capacities of services meet their agreed targets.

The *LQCD-ext II Project Execution Plan* (see docdb#5565) document outlines staffing needs to match hardware acquisitions and performance delivery milestones. These staffing needs are based on a 5-year acquisition and operational plan which manages, controls and predicts the performance and capacity of operational services.

For the Wilson Cluster Facility staffing levels will be reviewed, reported, and updated yearly in the *Tactical Plan*.

3.3 Component Capacity Management

The objective is to manage, control and predict the performance, utilization and capacity of IT resources and individual IT components.

The Networking Group records transfer rates over public network links connected to the compute clusters that make up the USQCD and Wilson Facilities. On an annual basis, the Networking Group either upgrades capacity on saturated (for very long periods of time) network links or adds additional lanes to existing network backbones.

For the USQCD Facility the *LQCD-ext II Project Execution Plan* (see docdb#5565) provisions for compute and storage capacity growth over the 5-year life-time of the project. For the Wilson Facility we record detailed job statistics to track queue depths (i.e. average user job wait times) and available storage capacity on the various file-systems. Longer queue wait times and close to full capacity on storage targets are strong indicators of an overprovisioned compute and storage facility.

4 BUSINESS REQUIREMENTS, SERVICE ENTITLEMENTS AND COST

4.2 Business Requirements

In the annual budget process the business requirements are reviewed and aggregated so that the High Performance Computing area owner may plan adequate technical resources to meet the business needs. Refer to:

- Tactical Plan and Budget process described in Financial Management Policy and Procedures (see docdb#4112)
- Capacity Plans (see docdb#4047)
- Business Impact Assessment (see docdb#4571)
- Continuity of Operations Plans (see docdb#5097,4969 and #4571)

4.3 Service Entitlements

Service Entitlements are defined in the applicable Foundation Service Level Agreement. Exceptions to those entitlements (if any) are listed below.

4.4 Service Charging Policy

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

USQCD Facility Parallel and Tightly Coupled Batch Computing
no charge but showback in normalized core hours for both compute and storage
USQCD Facility Application Support
no charge
USQCD Facility User accounts
no charge
USQCD Facility File-System Support
no charge
Wilson Facility Parallel and Tightly Coupled Batch Computing
no charge but showback in normalized hours for compute
Wilson Facility Application Support
no charge
Wilson Facility User accounts
no charge
Wilson Facility File-System Support
no charge

5 SERVICE REQUESTS

5.2 Standard Requests

Service Catalog Items	
Service Offering	Catalog Item
USQCD Facility Parallel and Tightly Coupled Batch Computing	None
USQCD Facility Application Support	None
USQCD Facility User Accounts	None
USQCD Facility File-System Support	None
Wilson Facility Parallel and Tightly Coupled Batch Computing	None
Wilson Facility Application Support	None
Wilson Facility User Accounts	None
Wilson Facility File-System Support	None

6 SERVICE COMMITMENTS

Except as otherwise stated below the Availability commitments and targets and the Service Level commitments and targets for both response and resolution of Incident (something is broken) and Request tickets is as described in the applicable Foundation Service Level or Operational Level Agreement.

SERVICE COMMITMENTS AND TARGETS			INCIDENT							REQUEST		
			Response			Resolution				Response		
Service offering	Target	Ticket Priority	1h	4h	8h	2d	4d	5d	7d	1h	4h	8h
USQCD Facility Parallel and Tightly Coupled Batch Computing	95%	Priority 2		X		X					X	
USQCD Facility Application Support	95%	Priority 2		X		X					X	
USQCD Facility User Accounts	95%	Priority 2		X		X					X	
USQCD Facility File-System Support	95%	Priority 2		X		X					X	
Wilson Facility Parallel and Tightly Coupled Batch Computing	90%	Priority 2		X		X					X	
Wilson Facility Application Support	90%	Priority 2		X		X					X	
Wilson Facility User Accounts	90%	Priority 2		X		X					X	
Wilson Facility File-System Support	90%	Priority 2		X		X					X	

6.2 Service Availability

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

Service Availability and Targets	
Service offering/Component	Availability
USQCD Facility Parallel and Tightly Coupled Batch Computing	95% Availability 24X7
USQCD Facility File-System Access	95% Availability 8to17by5
Wilson Facility Parallel and Tightly Coupled Batch Computing	95% Availability 24X7
Wilson Facility File-System Access	95% Availability 8to17by5

6.2.1 USQCD Facility Parallel and Tightly Coupled Batch Computing

- Maintenance Window – As needed. The outage usually lasts a few hours to a few days. All planned outages will be communicated ahead of time.
- Outage – All users are unable to access or use any of the compute facility resources.
- Degradation – Only a part of the compute facility resources are available.

6.2.2 USQCD Facility Application Support

- Maintenance Window – Not applicable.
- Outage – Not applicable.
- Degradation – Not applicable.

6.2.3 USQCD Facility User Accounts

- Maintenance Window – Not applicable.
- Outage – Kerberos authentication service is down.
- Degradation – Kerberos authentication service is slow to respond.

6.2.4 USQCD Facility File-System Support

- Maintenance Window – As needed. The outage usually lasts a few hours to a few days. All planned outages will be communicated ahead of time.
- Outage – All users are unable to access their data on the various file-systems managed by the HPC department.
- Degradation – Only a part of data is accessible on the various file-systems managed by the HPC department.

6.2.5 Wilson Facility Parallel and Tightly Coupled Batch Computing

- Maintenance Window – As needed. The outage usually lasts a few hours to a few days. All planned outages will be communicated ahead of time.
- Outage – All users are unable to access or use any of the compute facility resources.
- Degradation – Only a part of the compute facility resources are available.

6.2.6 Wilson Facility Application Support

- Maintenance Window – Not applicable.
- Outage – Not applicable.
- Degradation – Not applicable.

6.2.7 Wilson Facility User Accounts

- Maintenance Window – Not applicable.
- Outage – Kerberos authentication service is down.
- Degradation – Kerberos authentication service is slow to respond.

6.2.8 Wilson Facility File-System Support

- Maintenance Window – As needed. The outage usually lasts a few hours to a few days. All planned outages will be communicated ahead of time.
- Outage – All users are unable to access their data on the various file-systems managed by the HPC department.
- Degradation – Only a part of data is accessible on the various file-systems managed by the HPC department.

6.3 Other Service Levels

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

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

7 SERVICE SUPPORT

7.2 Requesting Service Support

Access to all USQCD Facility Computing IT services should be requested by emailing lqcd-admin@fnal.gov.

Access to all Wilson Facility Computing IT services should be requested by emailing tev-admin@fnal.gov.

The service provider for these services will enter a ticket in the Fermilab Service Management Application on-behalf of the user.

Fermilab ID, Kerberos principals or renewing an account password should be requested through the Service Desk, via the [ServiceNow](#) application, or by phone (630-840-2345). More

information about requesting these services can be found in the Self Service section of ServiceNow.

Unless otherwise noted Support Availability is 8:00AM to 5:00PM Monday to Friday excluding holidays

7.2.1 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.

7.3 Customer requests for Service Enhancements

Customers can request Service Enhancements by contacting their Computing Sector Liaison. The Computing Sector Liaison will work with the Service Owner and the Customer to prioritize and track progress.

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.

8 SERVICE LIFECYCLE

Plan

The Service Owner, along with the customer, will help plan and requisition the proper software, compute or storage required to meet the customer's needs. The Service Owner will also help coordinate requirements with the networking, storage and database backup providers.

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).

Deploy

Software, compute or 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 agreed upon service to the customer. This includes implementing or coordinating repairs, upgrades and replacements as necessary.

Retire/Replace

Application software will be upgraded regularly in accordance with vendor offerings as part of a normal software maintenance lifecycle. The Service owner will coordinate with the customer and underpinning service providers to perform these types of upgrades. The Service Owner will coordinate with the customer the retirement of any compute or storage hardware.

9 RESPONSIBILITIES

9.2 General Responsibilities

The applicable Foundation Service Level agreement defines the general responsibilities of the User, Customer and Service Owner including Computer Security responsibilities. It describes how to report incidents and the responsibilities with respect to service tickets.

9.3 Service Specific Responsibilities

9.3.1 CUSTOMER RESPONSIBILITIES

The applicable Foundation Service Level agreement defines the general responsibilities of the Customer.

9.3.2 USER RESPONSIBILITIES

The users agree to:

- Read the user documentation provided and consult with the Service Owner on how to use the batch queuing system to access the available resources.
- Obtain the appropriate Fermi credentials (Kerberos Principal)
- In case of issues, read documentation provided on the web (<http://www.usqcd.org/fnal> and <http://tev.fnal.gov>)
- To not create a denial-of-service (DoS) through improper web queries or information downloads.
- Report incidents and service requests via the appropriate email list (lqcd-admin@fnal.gov, tev-admin@fnal.gov) or service desk (<http://servicedesk.fnal.gov>).
- When reporting an incident, be available for support (call back via support staff).
- To be responsible for data content.

9.3.3 SERVICE OWNER

High Performance Computing service offering owners agree to:

- Provide the services as outlined in this agreement.
- Meet service targets as outlined in this agreement and invoke improvements as needed.
- Maintain appropriately trained staff.
- Coordinate standard, non-emergency, maintenance downtimes. Downtimes will be scheduled in coordination with underlying services. Notification of a service outage will be provided to the customer via email and/or announced in the Operations meeting at least 5 business days in advance of an outage (unless deemed an emergency).

10 SERVICE CONTINUITY

Recovery Time Objective (RTO) is defined as the length of time processes could be unavailable before the downtime adversely impacts business operations.

Recovery Point Objective (RPO) is defined as the maximum interval of data loss since the last backup that can be tolerated and still resume the business process.

Recovery Objectives	RTO	RPO
Service offering		
USQCD Facility Parallel and Tightly Coupled Batch Computing	12 hours	up to 2 business days
USQCD Facility Application Support	24 Hours	up to 2 business days
USQCD Facility User Accounts	12 hours	up to 2 business days
USQCD Facility File-System Support	12 Hours	up to 2 business days
Wilson Facility Parallel and Tightly Coupled Batch Computing	12 hours	up to 2 business days
Wilson Facility Application Support	24 hours	up to 2 business days
Wilson Facility User Accounts	12 Hours	up to 2 business days
Wilson Facility File-System Support	12 hours	up to 2 business days

The Service Continuity plan for this service (if it has a unique plan) is stored in the docdb entry associated with this document

The plan works in conjunction with the Continuity of Operations and Disaster recovery plans for Core or Scientific IT Services.

10.1 Recovery Strategy

- Current strategy is to handle on a case by case basis:
 - Communicate and cooperate with Service Desk, Service Manager and senior level management.
 - Communicate and cooperate with OLA partners to get infrastructure ready for compute, applications and filesystems recovery.
 - Recover/restore compute, applications and filesystem from backups.
 - Verify recovery.
 - Release recovered areas to application service owners.

10.2 Strategy for initial recovery

Current strategy typically involves:

- Assessing the situation and stabilizing compute, application and file-system resources to extent possible.
- Informing the Service Desk and Service Owners.
- Informing senior management of the situation and status.
- Contacting dependent OLA partners to ascertain recovery status of their services, as needed.
- Communicate and cooperate with all interested parties to develop and execute a plan of action to restore services as soon as possible.

10.3 Overall recovery strategies

- Choose a recovery coordinator.
- Perform a damage assessment if any.
- Notify customers of service status.
- Coordinate recovery with underlying services if needed.
- Implement recovery plan.
- Verify recovery is completed.
- After communicating with Service Desk and clearing any possible issues with Facilities, Networking and Authentication make sure all the services are stable. Communicate with the end users and ask to verify that the service has been re-established and is functioning properly.

10.4 Recovery Scenarios

The high level plan for recovery of the High Performance Computing services is:

- Data: Recover data from backups on disk (kept up to 1 year) or on tapes.
- Compute: Recover OS images from backups on disk (kept up to 1 year) or on tapes.

10.4.1 Building not accessible (Data Center Available)

Completed	Action
	Contact the Critical Incident Command Center. The Computing Sector Continuity plan lists the location of these centers as well as alternate locations.
	Coordinate with the Critical Incident Command Center to execute the overall strategy for recovery.
	When authorized by the Critical Incident Command Center, restore the hardware and software. If restoration from tape is required that may require physical intervention when the building and the tape libraries become accessible.

10.4.2 Data Center Failure (Building Accessible)

Completed	Action
	Contact the Critical Incident Command Center. The Computing Sector Continuity plan lists the location of these centers as well as alternate locations.
	Coordinate with the Critical Incident Command Center to execute the overall strategy for recovery.
	When authorized by the Critical Incident Command Center, restore the hardware and software. If restoration from tape is required that may require physical intervention when the building and the tape libraries become accessible.

10.4.3 Building not accessible and Data Center Failure

Completed	Action
	Contact the Critical Incident Command Center. The Computing Sector Continuity plan lists the location of these centers as well as alternate locations.
	Coordinate with the Critical Incident Command Center to execute the overall strategy for recovery.
	When authorized by the Critical Incident Command Center, restore the hardware and software. If restoration from tape is required that may require physical intervention when the building and the tape libraries become accessible.

10.4.4 Critical recovery team not available

Completed	Action
	Contact the Critical Incident Command Center. The Computing Sector Continuity plan lists the location of these centers as well as alternate locations.
	Coordinate with the Critical Incident Command Center to execute the overall strategy for recovery.
	When authorized by the Critical Incident Command Center, restore the hardware and software. If restoration from tape is required that may require physical intervention when the building and the tape libraries become accessible.
	Perform work that can be done remotely. Example - Server power may be cycled remotely, but some actions may require on-site presence (if restoration from Tape requires a physical loading of media, for example).
	If restoration from tape is required that may require physical intervention when personnel become available.

10.4.5 Government Mandated Shutdown of Services

Completed	Action
	Contact the Critical Incident Command Center. The Computing Sector Continuity plan lists the location of these centers as well as alternate locations.
	Coordinate with the Critical Incident Command Center to execute the overall strategy for the shutdown of services.
	Ensure that all services/servers are in a safe and secure state for recovery at a later time.
	Maintain periodic contact with the Critical Incident Command Center
	Coordinate with the Critical Incident Command Center to execute the overall strategy for recovery.
	When authorized by the Critical Incident Command Center, restore the hardware and software. If restoration from tape is required that may require physical intervention when the building and the tape libraries become accessible.

10.5.4 Return to Operations

- Coordinate with underlying services owners and the service desk to communicate regarding the services being restored.
- Proceed with the recovery plan.

11 SERVICE MEASURES AND REPORTING

11.2 Standard Service Measures and Reports

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 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.

11.3 Service specific Measures and Reports

There is no need for service specific measures or reports.

APPENDIX A: SUPPORTED HARDWARE AND SOFTWARE

N/A

APPENDIX B: SLA and OLA CROSS-REFERENCE

The services in this Service Area depend on the following IT Services to operate within their respective SLAs / OLAs.

Critically depends on usually means that the Service Offering will be unavailable (or at minimum degraded) if the depends on Service Offering is unavailable.

Depends on means that there is a dependency for Availability and Continuity but the extent of the dependency can vary.

Service Offering	Depends on /Uses			Availability/Continuity
	Service Area	Service Offering	through	Comments
Parallel and Tightly Coupled Batch Computing	Critically depends on	Kerberos Certificate Authority (KCA)	SLA	
		Data Center Networking	SLA	
		Network Facilities	SLA	
		Lattice Computing Center (LCC)	SLA	Data center, power, environment
		Grid Computing Center (GCC) A	SLA	Data center, power, environment
		Grid Computing Center (GCC) C	SLA	Data center, power, environment
	Uses	Managed Scientific Workstation	SLA	
	Uses	Scientific Linux Engineering	SLA	
Application Support	Uses	Managed Scientific Workstation	SLA	
	Uses	Scientific Linux Engineering	SLA	
User accounts	Uses	ServiceNow Self Service	SLA	

	Critically depends on	Kerberos Certificate Authority (KCA)	SLA	
File-System Support	Critically depends on	Kerberos Certificate Authority (KCA)	SLA	
		Data Center Networking	SLA	Network/Storage cabling
		Network Facilities	SLA	
		Lattice Computing Center (LCC)	SLA	Data center, power, environment
		Grid Computing Center (GCC) A	SLA	Data center, power, environment
		Grid Computing Center (GCC) C	SLA	Data center, power, environment
	Uses	Managed Scientific Workstation	SLA	
		Scientific Linux Engineering	SLA	

APPENDIX C: SERVICE DEPENDENCY CROSS-REFERENCE

The following Services/Service Areas depend on Services described in this document

Table to be inserted when dependencies are all loaded into SNOW

APPENDIX D: UNDERPINNING CONTRACT (UC) CROSS-REFERENCE

Supplier Lists, including contact information can be found here.

Key vendor contracts supporting this service area are:

- KOI Computers
- Seneca Data
- Imation

APPENDIX E: TERMS AND CONDITIONS BY CUSTOMER

N/A