



FY11 Tactical Plan for SCF/System Administration

Prepared by: Jason Allen

Date: 23 April 2010

Relevant Strategic Plans - Scientific Facilities, Computing Division, Grids, Run II Computing, Intensity Frontier Program

System Administration Goal

Provide a stable Linux computing environment for Fermilab experiments.

Executive Summary of Objectives for FY11

Maintenance/Compliance Drivers

1. Keep operating system and SLF software patched per policy.
2. Proactively monitor systems for failures.
3. Provide incident and problem management.
4. Purchase worker nodes, storage, and servers.
5. Fulfill service requests.
6. Provide consulting and technical support.

Upgrades and Enhancement Drivers

1. Refresh aging hardware (console servers, servers, etc)
2. Improve collection and reporting of key operational metrics.
3. Implement change control.

Strategic Drivers

1. Virtualization
2. Place a stronger emphasis on "pre-production" system management.

Activities and Work Definition

System Administration/Short Term Projects

Activity type: New Project

Description: D0 CAB status webpage enhancements. Improve <http://fefweb.fnal.gov/cab> by implementing features requested by D0 users.

Timescale: Start Q4 CY2010; Complete Q1 CY2011

Metrics: N/A

Service Activities

System Administration/Online Systems Management

Activity type: Service

Description: Maintain computer systems, storage, and other hardware related to CDF and D0 Online operations.

Timescale: Continuous

Metrics: Number of systems supported and service desk tickets.

System Administration/Compute Node Management

Activity type: Service

Description: Provision, maintain, and monitor batch system computers (worker nodes) and related infrastructure.

Timescale: Continuous

Metrics: Number of nodes, availability, server/sysadmin ratio

System Administration/Server Management

Activity type: Service

Description: Provision, maintain, and monitor hardware and operating system.

Timescale: Continuous

Metrics: Number of nodes, availability, server/sysadmin ratio

System Administration/Storage Management

Activity type: Service

Description: Provision, maintain, and monitor disk based storage hardware: direct attached RAID, NAS, and SAN.

Timescale: Continuous

Metrics: Volume of storage managed. Number of storage units.

System Administration/Batch System Management

Activity type: Service

Description: Maintain and support Torque batch system software. The scope of this activity currently only includes the D0 CAB batch system.

Timescale: Continuous

Metrics: Number of tickets, job slots

System Administration/Event and Incident Management

Activity type: Service

Description: Resolve interruptions, or potential interruptions, in service.

Timescale: Continuous

Metrics: Number of Service Desk tickets, availability, incidents reporting by monitoring software.

System Administration/Problem Management

Activity type: Service

Description: Perform incident root cause analysis with the goal of reducing service interruptions.

Timescale: Continuous

Metrics: Number of problems resolved.

System Administration/Operational Planning and Consulting Support

Activity type: Service

Description: Provide scientific computing design and maintenance related technical assistance to Fermilab experiments.

Timescale: Continuous

Metrics: ---

System Administration/Procurement Support

Activity type: Service

Description: Execute budget for scientific computing hardware, services, and support contracts. Assist customers with annual budget and capacity planning.

Timescale: Continuous

Metrics: Total dollar amount of approved requisitions. Number of reqs approved.

System Administration/Professional Development

Activity type: Service

Description: Misc activities dedicated to the improvement of employee technical skill and knowledge.

Timescale: Continuous

Metrics: Number of classes taken, certifications.

System Administration/System Administration Management

Activity type: Service

Description: Line management, change management, service level management

Timescale: Continuous

Metrics: Staff levels

Detailed Tactical Plan Objectives and Priorities

Maintenance and Compliance Drivers

Objectives:

1. Maintain systems in accordance with Fermilab computing security baselines. In particular, ensure system software is kept current per the "Patching Timeline Policy" (<http://security.fnal.gov/internal/patchingtimeline.html>)
2. Proactively monitor systems for significant events and failures. To achieve the goal of providing a stable computing environment, it's important that we monitor systems so that we can identify issues are currently causing, or may cause, an interruption in service.
3. Closely related to monitoring is incident and problem management. Once an interuption in service has been identified it's important that we resolve the issue as quickly as possible. Equally important is problem management where we determine the root cause of incidents or provide work arounds.
4. Purchase good value hardware for the experiments and assist with budget planning and execution.
5. Provide quick and accurate resolution to customer service requests.
6. Be a resource to Fermilab experiments and CD if technical assistance or guidance is needed.

Assumptions and Risks:

1. Inadequate manpower to provide quality services to the Fermilab experiments. Insufficient effort is the universal and most obvious risk to IT operations.
2. As a consequence of a reduced budget hardware may not be refreshed at regular intervals. This in turn means that hardware is more prone to failures and more effort is required to maintain systems.

3. Lack of well defined computing policies. Specifically, guidance to experiments regarding the age and types of systems CD is willing to support. The consequence is a increased demand on a system administration effort and a feeling of frustration from users.
4. Low morale, decreased productivity, and an increase in resignations due a perceived winding down of CDF and D0 experiments.
5. Risk of security incidents if systems are not patched per policy.
6. Reduced availability if machines are not monitored and incident/problem management services are not properly provided.

Upgrade and Enhancement Drivers

Objectives:

1. Refresh aging hardware. Older hardware tends to be more expensive to maintain in terms of effort and maintenance costs so it's important that we regularly replace systems including system management infrastructure such as console servers and PDUs.
2. Collection of and reporting of key metrics. Relevant metrics allow us to measure measure the level and quality of service we provide to customers. We first need to carefully identify key metric and then automate the collection and reporting process.
3. Implement change control. We expect to be "on-boarded" into formal CD change control in FY11.

Assumptions and Risks:

1. Inadequate budget to replace out of maintenance hardware.
2. Change control implemented in a way that slows down resolution of service requests.

Strategic Drivers

Objectives:

1. Virtualization. Continue to implement systems in accordance with the CD strategy of leveraging virtualization to consolidate servers and increase reliability. In particular, the roll-out of GPCF [tactical plan] will provide flexible computing resources for the upcoming IF experiments.
2. Place a stronger emphasis on "pre-production" system management. Ideally our system management infrastructure would make it easier to completely rebuild a system then to troubleshoot it. We're almost to that point with worker nodes but servers are more challenging due to their complex hardware and system configurations. In FY11 we want to spend more effort on system configuration management techniques and packaging standard Linux builds. The goal is to making provisioning so easy a sysadmin can simply click a button a webpage will completely reinstall a server from scratch. The obvious benefits to this model are greater availability and faster provisioning. Perhaps more importantly, it would give us a tiered, scalable approach to system management. Junior system administrators would be responsible for daily operations while senior staff would focus on pre-production issues and problem management.

Assumptions and Risks:

1. Large scale virtualization at Fermilab is an unproven technology and there are a number of risks. First, virtualization by its nature is extremely software dependent, so we are vulnerable to buggy implementations or application availability issues. In FY09, the virtualization management software selected to run D0 and CDF Online virtual clusters was purchased by large company and then immediately discontinued. Second, expectations need to be managed. It's tempting to view the consolidation features provided by virtualization as a panacea to budget shortfalls, but not all systems are good candidates for virtualization. In particular heavily loaded machines or systems with large amounts of local disk in many cases are better suited to running "bare metal", non-virtualized machines.

Staffing Issues

Staffing levels for system administration (FEF Dept) for FY11 should be sufficient if we remain fat the FY10 level of 10 FTE.