



---

FermiGrid

Keith Chadwick



# Overall Deployment Summary

---

## 5 Racks in FCC:

- 3 Dell Racks on FCC1
  - Can be relocated to FCC2 in FY2009.
  - Would prefer a location on FCC2 that can be served by panels fed from 2 UPS for the FermiGrid-HA services (requires 4 x L5-30R).
- 2 Racks on FCC2
  - Will be retiring 1 “soon”.

## 1 Dell Rack in WH8 “falcon’s nest” systems area

## Several Racks in GCC:

- GCCA
- GCCB
- Expect to add additional racks in FY2009 and FY2010



# Physical Services

---

<http://fermigrid.fnal.gov/fermigrid-organization.html>

<http://fermigrid.fnal.gov/cdfgrid-organization.html>

<http://fermigrid.fnal.gov/d0grid-organization.html>

<http://fermigrid.fnal.gov/gpgrid-organization.html>

<http://fermigrid.fnal.gov/fgtest-organization.html>



# FermiGrid @ FCC1

---

## OSG [Gratia] Services Rack

- 8 Dell 2850 Systems, many near end of warranty.
- Support: 8x5 + Best Effort
- These systems are proposed to be upgraded/replaced this FY.
- New systems for OSG ReSS will also be added to this rack this FY.
- Old systems will be moved to WH and retained for development/integration/test.

## FermiGrid Gatekeeper Services Rack

- 15 systems, various vendors.
- Support: 8x5 + Best Effort
- Many systems have been recently upgraded.

## FermiGrid HA Services Rack

- 7 Dell 2950 Systems with dual power supplies.
- 4 power controllers spread across 2 panels.
- Support: 24x7 for GUMS and SAZ services, all others 8x5 + Best Effort.



## FermiGrid @ FCC2

---

### OSG Resource Selection Services Rack

- ReSS systems overdue for hardware refresh.
- Replacement systems (2 x Dell 2950) will be installed in OSG Services Rack (on FCC1).
- OSG Bestman Test nodes will be installed in this FCC2 rack.

### FermiGrid ITB Worker Node Rack

- In the process of being replaced/retired.
- Replacement systems already in GCC-B.



## FermiGrid @ Wilson Hall 8

---

### FermiGrid Test System Dell Rack

- 4 x Dell 2850 (old production fermigrid[1-4]).
- 1 x Dell 2950 (old production fermigrid5).
- 3 x 1U systems.



## FermiGrid @ GCC

---

These systems are mostly managed by FEF.

GCC-A:

GCC-B:



# FermiGrid-HA - Highly Available Grid Services

---

The majority of the services listed in the FermiGrid service catalog are deployed in high availability (HA) configuration that is collectively known as “FermiGrid-HA”.

- The goal for FermiGrid-HA is > 99.999% service availability.
  - Not including Building or Network failures.
  - These will be addressed by FermiGrid-RS (redundant services) in FY2010/11.
- For the period of 01-Dec-2007 through 30-Jun-2008, FermiGrid achieved 99.997% service availability.

FermiGrid-HA utilizes three key technologies:

- Linux Virtual Server (LVS).
- Xen Hypervisor.
- MySQL Circular Replication.






# HA Services Deployment

---

## FermiGrid employs several strategies to deploy HA services:

- Trivial monitoring or information services (such as Ganglia and Zabbix) are deployed on two independent virtual machines.
- Services that natively support HA operation (Condor Information Gatherer, FermiGrid internal ReSS deployment) are deployed in the standard service HA configuration on two independent virtual machines.
- Services that maintain intermediate routing information (Linux Virtual Server) are deployed in an active/passive configuration on two independent virtual machines. A periodic heartbeat process is used to perform any necessary service failover.
- Services that do not maintain intermediate context (i.e. are pure request/response services such as GUMS and SAZ) are deployed using a Linux Virtual Server (LVS) front end to active/active servers on two independent virtual machines.
- Services that support active-active database functions (circularly replicating MySQL servers) are deployed on two independent virtual machines.

# HA Services Communication





# Non-HA Services

---

The following services are not currently implemented as HA services:

- Globus gatekeeper services (such as the CDF and D0 experiment globus gatekeeper services) are deployed in segmented pools.
  - Loss of any single pool will reduce the available resources by approximately 50%.
- MyProxy
- OSG Gratia Accounting service [Gratia]
  - not currently implemented as an HA service.
  - If the service fails, then the service will not be available until appropriate manual intervention is performed to restart the service.
- OSG Resource Selection Service [ReSS]
  - not currently implemented as an HA service.
  - If the service fails, then the service will not be available until appropriate manual intervention is performed to restart the service.

We are working to address these services as part of the FermiGrid FY2009 activities.



# FermiGrid Service Level Agreement

---

## Authentication and Authorization Services:

- The service availability goal for the critical Grid authorization and authentication services provided by the FermiGrid Services Group shall be 99.9% (measured on a weekly basis) for the periods that any supported experiment is actively involved in data collection and 99% overall.

## Incident Response:

- FermiGrid has deployed an extensive automated service monitoring and verification infrastructure that is capable of automatically restarting failed (or about to fail) services as well as performing notification to a limited pager rotation.
- <http://fermigrid.fnal.gov/fermigrid-metrics.html>
- It is expected that the person that receives an incident notification shall attempt to respond to the incident within 15 minutes if the notification occurs during standard business hours (Monday through Friday 8:00 through 17:00), and within 1 (one) hour for all other times, providing that this response interval does not create a hazard.

## FermiGrid SLA Document:

- <http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=2903>



# Measured Service Availability

	Last Week	Last Month	Last Quarter	Since 10-Jul-2008
Hardware	100.000%	100.000%	100.000%	100.000%
Core Services	100.000%	99.992%	99.995%	99.976%
Gatekeepers	100.000%	99.864%	98.474%	99.082%
Batch Services	100.000%	99.801%	99.813%	99.844%
ReSS	100.000%	100.000%	99.995%	99.717%
Gratia	100.000%	100.000%	99.731%	99.812%

The above figures are through the week ending Saturday 03-Jan-2009 23:59:59.

The goal for FermiGrid-HA is > 99.999% service availability.

The SLA for GUMS and SAZ is 99.9% (during data taking) / 99.0% (otherwise).