

FermiGrid and FermiCloud: What Experimenters need to know (FIFE Workshop 6/4/2013)

Steven C. Timm
FermiGrid Services Group Lead
FermiCloud Project Lead
Grid & Cloud Computing Department

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

What is FermiGrid?

FermiGrid is:

The interface between the Open Science Grid and Fermilab.

A set of common services for the Fermilab site including:

- The site Globus gateway.
- The site Virtual Organization Membership Service (VOMS).
- The site Grid User Mapping Service (GUMS).
- The Site AuthoriZation Service (SAZ).
- The site MyProxy Service.
- The site Squid web proxy Service.

Collections of compute resources (clusters or worker nodes), aka Compute Elements (CEs).

Collections of storage resources, aka Storage Elements (SEs).

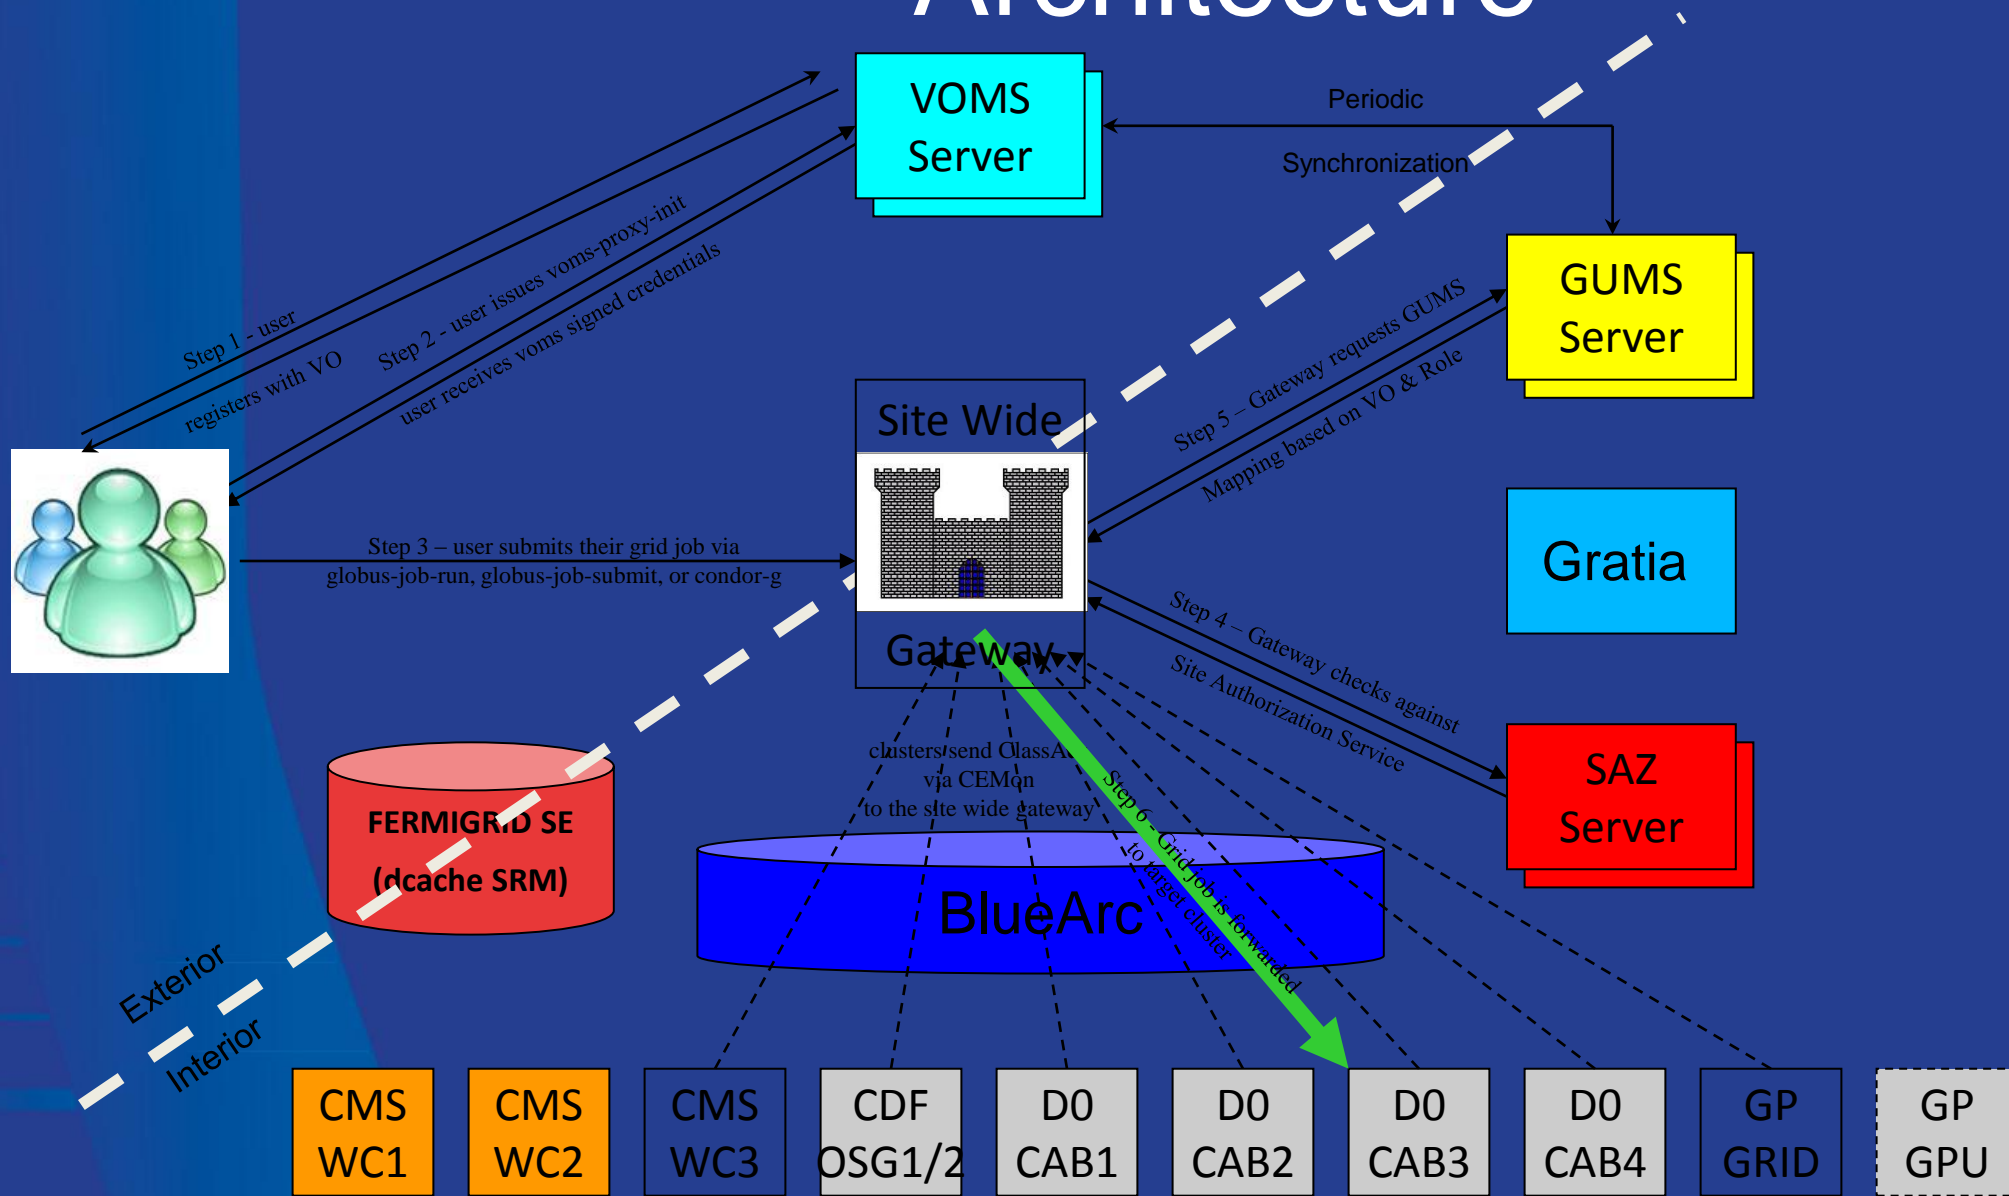
More information is available at <http://fermigrid.fnal.gov>

On November 10, 2004, Vicky White (then Fermilab CD Head) wrote the following:

In order to better serve the entire program of the laboratory the Computing Division will place all of its production resources in a Grid infrastructure called FermiGrid. This strategy will continue to allow the large experiments who currently have dedicated resources to have first priority usage of certain resources that are purchased on their behalf. It will allow access to these dedicated resources, as well as other shared Farm and Analysis resources, for opportunistic use by various Virtual Organizations (VOs) that participate in FermiGrid (i.e. all of our lab programs) and by certain VOs that use the Open Science Grid. The strategy will allow us:

- to optimize use of resources at Fermilab
- to make a coherent way of putting Fermilab on the Open Science Grid
- to save some effort and resources by implementing certain shared services and approaches
- to work together more coherently to move all of our applications and services to run on the Grid
- to better handle a transition from Run II to LHC (and eventually to BTeV) in a time of shrinking budgets and possibly shrinking resources for Run II worldwide
- to fully support Open Science Grid and the LHC Computing Grid and gain positive benefit from this emerging infrastructure in the US and Europe.

FermiGrid - Current Architecture



Who can use FermiGrid?

Any Fermilab employee, contractor, or user can run up to 25 jobs at once as member of “Fermilab” VO.

Usage above this level must be approved by Scientific Computing Division Management and the Computer Security Board.

Liaison should submit “New VO or Group Support on FermiGrid” request via ServiceNow.

Policy on new group/VO acceptance is in

<http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=3429>

Allocations (Quotas)

General Purpose Grid Cluster

(Previously known as “Farms”)

High priority for production work

Each experiment has a quota of “batch slots”

Quota is maximum number of slots you can use

Based on physics priorities of the lab.

Quotaed slots are not pre-emptable

Quotas are oversubscribed by ~200%

Rare that the cluster fills up with quota jobs.

Getting more quota

Your liaison should submit:

“Increased Job Slots or Disk Space on FermiGrid” request in ServiceNow.

Requests are processed by senior SCD management.

We will expect a presentation at the Computing Sector Liaisons meeting on what you need the extra slots for, and another presentation when you are done.

First question we will ask with any quota increase: Can you use opportunistic slots?

Opportunistic usage

Use as many slots as you want.

Quotaed usage has priority.

If cluster is full, opportunistic jobs will be sent a pre-empt signal and have 24 hours to finish before they get killed.

Balance of General Purpose Grid, CDF, D0, and CMS cluster all are available to Intensity Frontier users and opportunistic use.

Any Intensity Frontier groups using gpsn01 (and soon FIFE) have a separate entry point to submit opportunistic jobs.

FermiCloud Background

Infrastructure-as-a-service facility for Fermilab employees, users, and collaborators

- Project started in 2010.
- OpenNebula 2.0 cloud available to users since fall 2010.
- Condensed 7 racks of junk machines to 1.5 racks of good machines
- Provider of integration and test machines to the OSG Software team.
- OpenNebula 3.2 cloud up since June 2012

Who can use FermiCloud

- Any employee, user, or contractor of Fermilab with a current ID.
- Most OSG staff have been able to get Fermilab “Offsite Only” ID’s.
- With Fermilab ID in hand, request FermiCloud login via Service Desk form.
- Instructions on our new web page at <http://fclweb.fnal.gov>
- Note new web UI at <https://fermicloud.fnal.gov:8443/>
- **Doesn't work with Internet Explorer yet**

FermiCloud capabilities

Infiniband interconnect

Persistent live-migratable storage on SAN

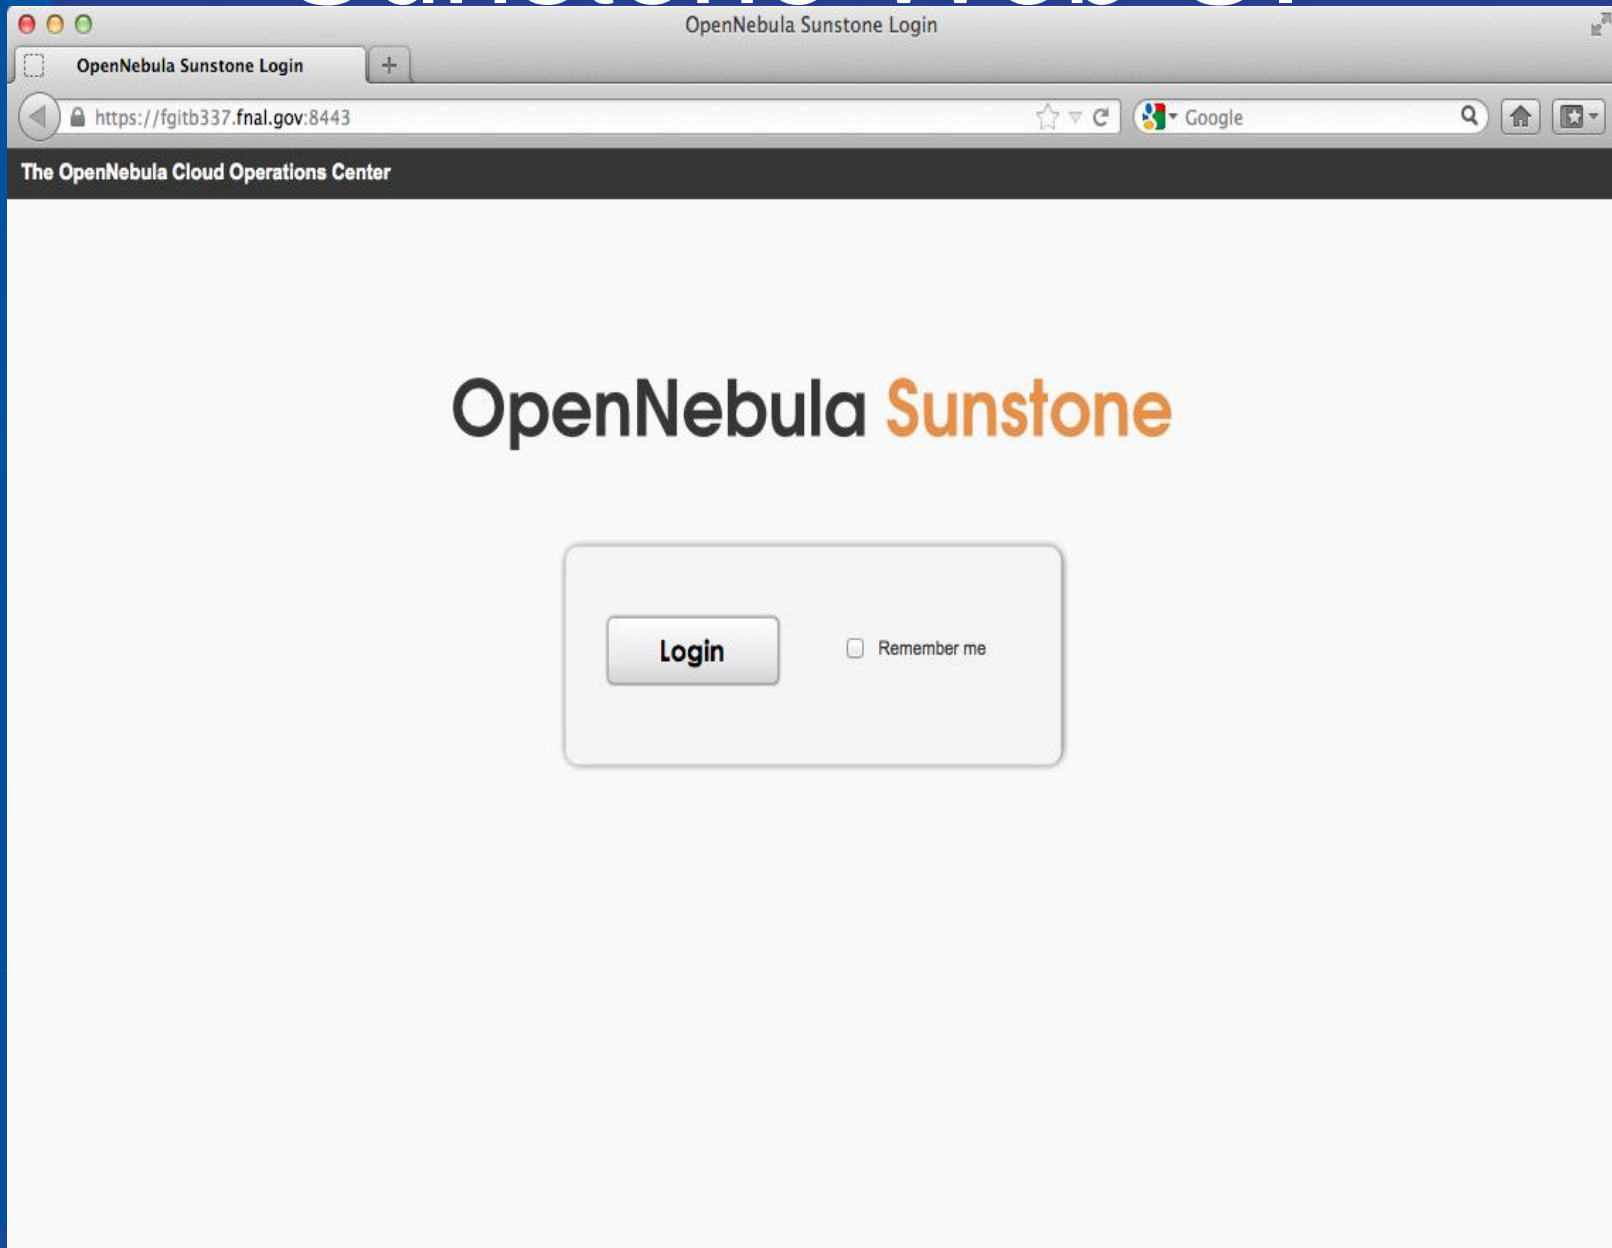
Public/private network clusters

Storage virtual machines

Simulate fault-tolerance behavior in multi-machine systems.

Coordinated launch of clients and servers.

Sunstone Web UI



Selecting a template

OpenNebula Sunstone Documentation | Support | Community Welcome timm2 | Sign out

Dashboard
Virtual Machines
Templates
Virtual Networks
Images
Configuration

+ New Update properties Instantiate Previous action Delete

Show 10 entries Search:

<input type="checkbox"/> All	ID	Owner	Group	Name	Registration time
<input type="checkbox"/>	50	oneadmin	oneadmin	SLF6 New Base	16:16:25 11/26/2012
<input type="checkbox"/>	53	oneadmin	oneadmin	SCT-slf58	22:31:08 11/28/2012

Showing 1 to 2 of 2 entries First Previous 1 Next Last

Copyright 2002-2012 © OpenNebula Project Leads (OpenNebula.org). All Rights Reserved. OpenNebula 3.2.1

Launching the Virtual Machine

The screenshot displays the OpenNebula Sunstone Cloud Operations Center interface. The browser address bar shows the URL `https://fgitb337.fnal.gov:8443`. The page title is "OpenNebula Sunstone: Cloud Operations Center". The navigation bar includes links for "Documentation", "Support", "Community", and "Welcome hyunwoo | Sign out". The left sidebar contains a menu with "Dashboard", "Virtual Machines" (highlighted), "Templates", "Virtual Networks", "Images", and "Configuration". The main content area shows a table of virtual machines with the following data:

<input type="checkbox"/>	ID	Owner	Group	Name	Status	CPU	Memory	Hostname	Start Time	VNC Access
<input type="checkbox"/>	119	hyunwoo	users	one-119	PROLOG	0	0K	fgitb338	14:04:03 10/16/2012	

Below the table, it indicates "Showing 1 to 1 of 1 entries" and provides navigation links: "First", "Previous", "1", "Next", "Last".

Copyright 2002-2012 © OpenNebula Project Leads (OpenNebula.org). All Rights Reserved. OpenNebula 3.2.1

Monitoring VM's

OpenNebula Sunstone Documentation | Support | Community **Welcome timm2 | Sign out**

Dashboard

- Virtual Machines
- Templates
- Virtual Networks
- Images
- Configuration

Summary of resources

VM Templates (total/public)	3 / 0
VM Instances (total/running/failed)	0 / 0 / 0
Virtual Networks (total/public)	1 / 0
Images (total/public)	4 / 0

Quickstart

New:

- VM Template
- VM Instance
- Virtual Network

Image

Historical monitoring information

Total VM count total active error

Total VM CPU cpu_usage

Total VM Memory mem_usage

VM Network stats net_tx net_rx

Copyright 2002-2012 © OpenNebula Project Leads (OpenNebula.org). All Rights Reserved. OpenNebula 3.2.1

Is your experiment using FermiCloud already?

GridFTP servers for Minerva, Nova, gm2, mu2e, LBNE, microboone, argoneut, MINOS, marsmu2e

Admin servers for Minerva, Nova, gm2, mu2e, LBNE, microboone, argoneut, MINOS.

Event display for MINOS, argoneut, microboone.

CVMFS test servers for D0.

SAMGrid forwarding nodes for D0.

dCache 4.1 testing for CDF

Application testing for DES

Coming soon: CVMFS Stratum 1 server

FermiCloud Development Goals

Goal: Make virtual machine-based workflows practical for scientific users:

- **Cloud bursting:** Send virtual machines from private cloud to commercial cloud if needed
- **Grid bursting:** Expand grid clusters to the cloud based on demand for batch jobs in the queue.
- **Federation:** Let a set of users operate between different clouds
- **Portability:** How to get virtual machines from desktop → FermiCloud → commercial cloud and back.
- **Fabric Studies:** enable access to hardware capabilities via virtualization (100G, Infiniband, ...)

FermiCloud Summary

FermiCloud Development Collaboration:

- Leveraging external work as much as possible,
- Contribution of our work back to external collaborations.
- Using (and if necessary extending) existing standards:
- AuthZ, OGF UR, Gratia, etc.

FermiCloud Facility

- Deploying 24by7 capabilities, redundancy and HA.
- Delivering support for science collaborations at Fermilab
- Making new types of computing work possible

The future is mostly cloudy.