Grid Interoperability via the FermiGrid Site Gateway

Steven C. Timm
Neha Sharma, Dan R. Yocum, Keith Chadwick,
Tanya Levshina, Gabriele Garzoglio,
Parag Mhashilkar

Grid Facilities Department,
Fermi National Accelerator Laboratory
Batavia, IL, USA

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

Outline

- Overview of the FermiGrid Campus Grid
- Why a site gateway?
- Authentication and authorization
- Information and matchmaking with Condor ClassAds
- Interoperability with WLCG/EGEE
- Interoperability with PRAGMA
- Metrics and Performance
- Conclusions and Future Plans

The FermiGrid Campus Grid

- In 2005 Fermilab had several large computing clusters, each owned by a large experimental stakeholder.
- The Computing Division created FermiGrid as a meta-facility with four major components:
  - Common Grid Services
    - Sitewide Globus Job Gateway, VOMS, VOMRS, GUMS, SAZ, SQUID
  - Stakeholder bi-lateral interoperability
    - CMS, D0, CDF, and General Purpose computing can all run jobs on each other's clusters.
  - Open Science Grid Interface
    - The OSG software stack used in external and internal gateways.
  - Permanent Storage System Interfaces
    - SRM/dCache interface to tape-backed Enstore storage element and to volatile intermediate storage accessible by SRM/dCache.
    - Bluearc enterprise NAS appliance for NFS-based intermediate data

Beginnings of FermiGrid

Ziggy in 2005

- Started with just 28 processors available to the Open Science Grid in early 2005.
- Site gateway used MyProxy-based forwarding system.
- For past 2 years all of our clusters have been open to Grid use.
- Most major reconstruction and analysis work done using Grid submission.

FermiGrid Campus Grid, continued

- Currently about 18000 job slots available
- 62 TB of Bluearc NAS-based storage.
- One site gateway, ten internal gatekeepers to receive the jobs
- Site authentication and gateway servers are Dell PowerEdge 2950 and similar machines.
- All server machines virtualized using Xen for easy movement of logical machines between different physical machines if necessary.
- All server machines also use service IP's so they are not tied to any given host name.
- All OSG VO's supported uniformly on all clusters, but cluster owner has priority to run when they need it.

Why A Site Gateway

- Security Control
  - Globus GSI authentication is point to point by nature.
  - Fermi CD management and many OSG VO's wanted central decision point.
  - We do this by
    - A single globus gatekeeper open to the OSG
    - Centralized Site AuthoriZation service SAZ
- Internal site flexibility
  - Don't have to notify OSG every time internal structure changes
- Access to all FermiGrid resources from one point.
- Internal Fermilab users submit to site gateway too, can find opportunistic cycles wherever they are available at Fermilab

FermiGrid - Current Architecture

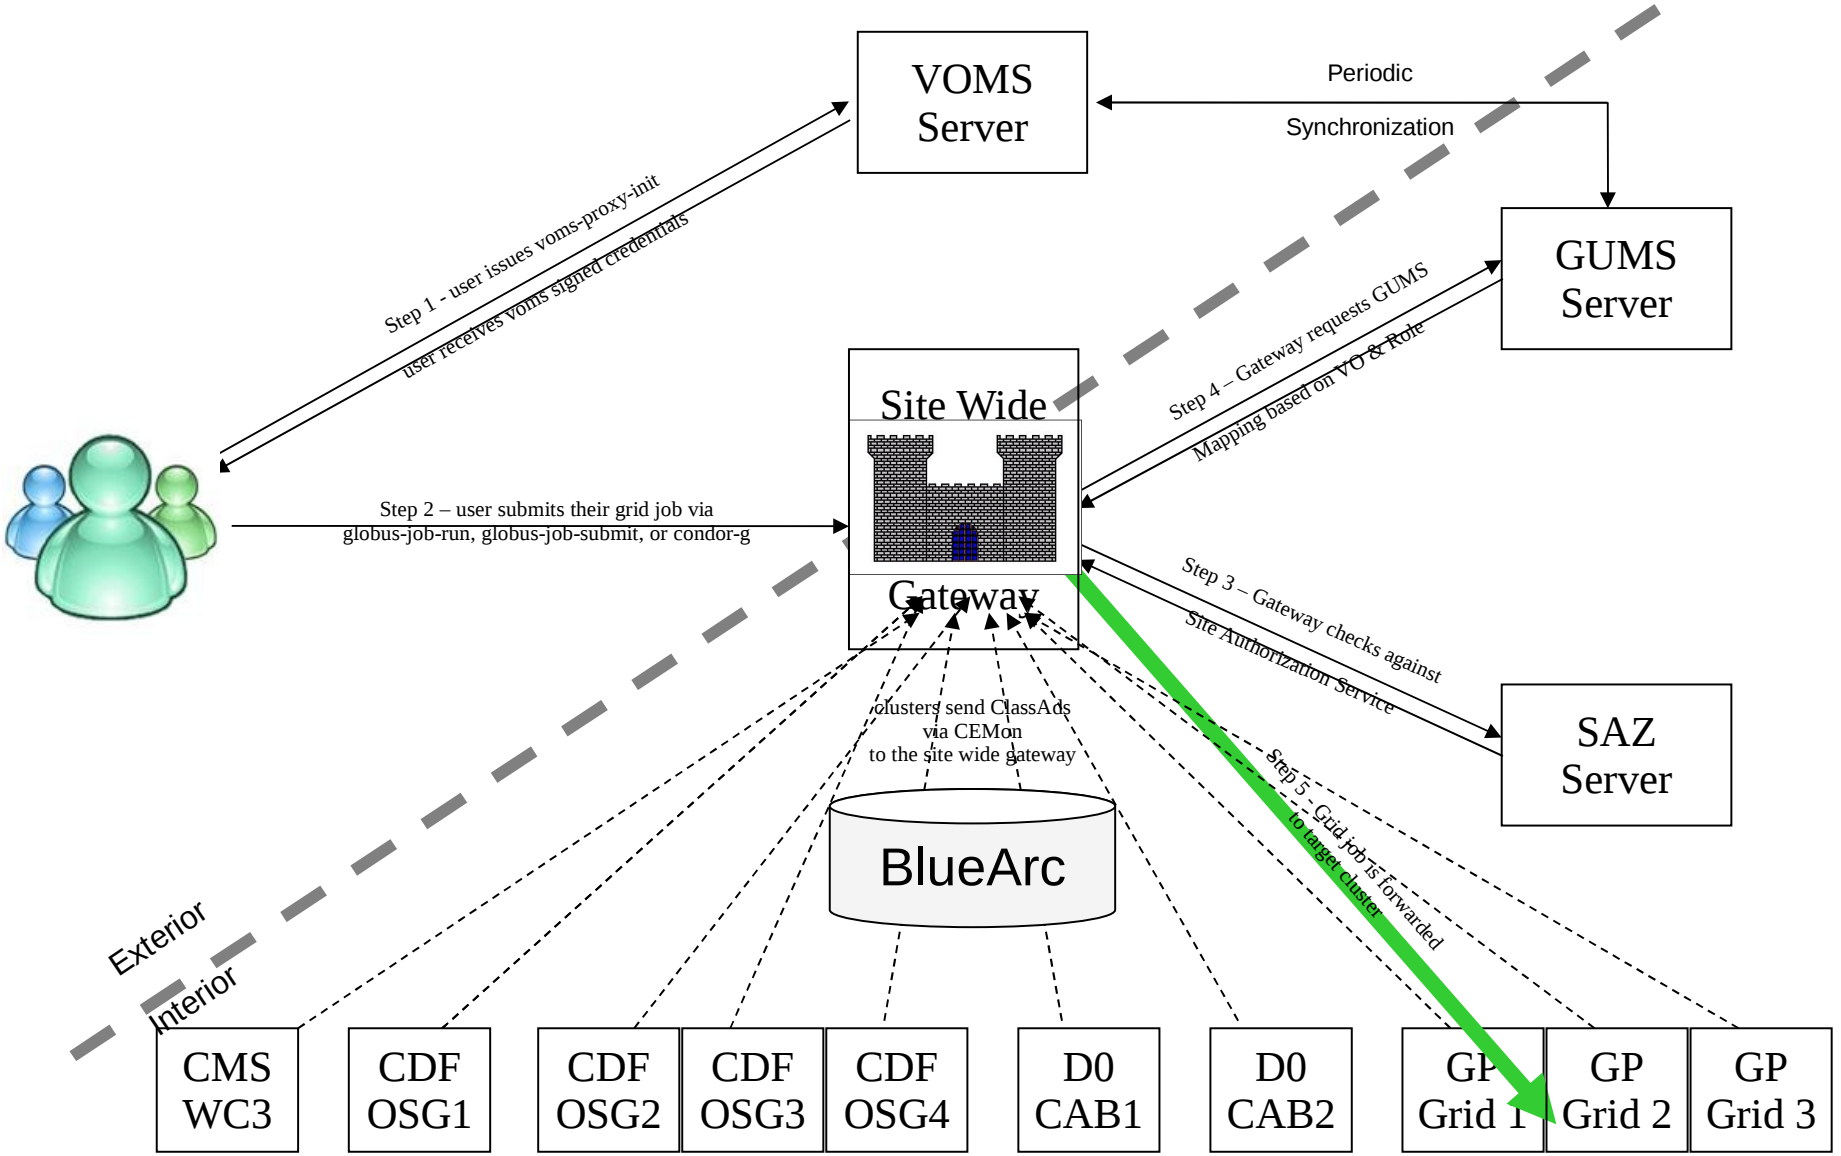

Central Authentication and Authorization

- VOMS: Virtual Organization Membership Service
  - Authenticates an individual's credentials as being part of a VO
  - Required for all those who run at Fermilab.
- GUMS: Grid User Mapping System
  - Given a Grid x509 credential, determines the username
  - Unified user mapping for the site, all potential usernames exist on all clusters.
- SAZ: Site AuthoriZation
  - Given a Grid x509 credential, returns YES or NO
  - Allows us to blacklist a person, a VO, a CA, or a VO Role.
- All of the above on highly available active-active servers, LVS frontend.

Job-forwarding services

- Our site gateway builds on work of
  - SAMGrid (D0 job submission and forwarding gateway)
  - GridX1—Canadian grid gateway
  - Condor team—which developed Condor-G matchmaking
  - Open Science Grid Resource Selection Service ReSS
    - (presented at this conference 2 years ago).
- All use Condor ClassAds to describe cluster.

Making classads with GIP, CeMON

- Generic Information Providers run on the compute and storage elements and gather data about the cluster, some static and some dynamic.
- An OSG-provided sensor type translates the GLUE Schema data into Condor ClassAds (sent to ReSS) and raw LDIF format, sent to the OSG BDII and from there to the LCG BDII.
- A web services layer is used for transport with both sender and receiver running under Tomcat 5.5.

Multiple ReSS classads per cluster

- One for every OSG VO accepted (currently 56)
- One for every different hardware subcluster type (2-6 varieties of hardware per cluster).
- One for every unique storage element available (most clusters have 3).
- Can lead to 56 x 6 x 3 = 1008 classads for one cluster
- Some of our bigger clusters can have >800 classads
  - CDF VO, 5 slots free, 50 machines, 2.66 GHz, 15 GB storage
  - CDF VO, 5 slots free, 10 machines, 1.8 GHz, 10 GB storage
  - D0 VO, 200 slots free, 300 machines, 1.8 GHz, 10 GB storage

GLUE Schema—sample cluster ClassAd

```
MyType = "Machine"
TargetType = "Job"
Name = "fnpcfg1.fnal.gov:2119/jobmanager-condor-cdf.1619067735"
GlueCEInfoContactString = "fnpcfg1.fnal.gov:2119/jobmanager-condor"
GlueCEStateFreeCPUs = 1
GlueCEPolicyAssignedJobSlots = 50
GlueCEInfoApplicationDir = "/grid/app"
GlueCEPolicyPriority = 1000
GlueCEInfoControlBaseRule = "VO:cdf"
GlueCEStateWorstResponseTime = 0
GlueCEStateEstimatedResponseTime = 14400
GlueCEUniqueID = "fnpcfg1.fnal.gov:2119/jobmanager-condor-cdf"
GlueCEInfoDataDir = "/grid/data"
GlueCEInfoTotalCPUs = 1088
FermiGridCEInfoWSGridType = "gt4"
FermiGridCEInfoWSGridResource = "gt4 https://fnpcfg1.fnal.gov:9443 Condor"
FermiGridCEInfoLRMSMaster = "fnpccm1.fnal.gov"
GlueClusterName = "fnpcfg1.fnal.gov"
GlueSiteUniqueID = "FNAL_GPGRID_2"
GlueSubClusterUniqueID = "fnpc195.fnal.gov"
lueSubClusterPhysicalCPUs = 320
FermiGridHostWNTmpDirSize = 57000000
GlueHostProcessorVendor = "GenuineIntel"
GlueHostMainMemoryRAMSize = 4025
GlueHostOperatingSystemRelease = "5.1"
GlueHostProcessorClockSpeed = 2333
GlueHostProcessorModel = "Intel(R) Xeon(R) CPU 5148 @2.33GHz"
FermiGridSubClusterRSL = "condorsubmit=(requirements 'TotalCpus==4')"
GlueHostNetworkAdapterInboundIP = TRUE
GlueHostNetworkAdapterOutboundIP = TRUE
FermiGridHostBenchmarkFermiCycles = 3317
GlueHostBenchmarkSI00 = 2265
GlueHostBenchmarkSF00 = 2138
```

FermiGrid use of ReSS

- We run our own ReSS collector
- Most of our classads only go to our private collector.
- Jobs submitted to our scheduler match against only classads from our cluster, not from all OSG
- Match criteria: Most slots free plus it supports the VO
- Users can add user-defined criteria via special field GlueRequirements
  - For instance: JobManager != PBS, GlueCEInfoContactString != cmsosgce.fnal.gov
- GlobusRSL string that is incoming from the job is forwarded on to the final destination gatekeeper

Watching the jobs

- Once a job is matched to a final destination cluster, it has two hours to start running.
- If it does not match by that time, the gateway recalls the job and attempts to rematch it to a different cluster.
- RANK statement causes gatekeeper to preferably match to a cluster that is different than the one it ran on previously.

Forwarding the job

- We receive inbound job from the Grid as any condor-based site on the Grid would.
- Instead of submitting a “vanilla” universe job to a local condor cluster, we submit a grid/gt2 job via Condor-G, using the inbound proxy that came with the job for credentials
- This requires that our internal grid/gt2 gatekeepers accept “limited proxies” for running a job. They take jobs only from on-site at Fermilab.

Web Services (GT4)

- We recently added support for GT4 web services jobs.
- No standard way to advertise them in GLUE schema so we added our own fields for now.
- We use the GT4 delegation web service to forward credentials to the final destination.

WLCG Interoperability

- FermiGrid site gateway not part of WLCG MOU.
- Nevertheless we accept several LCG and EGEE VO's and our site is advertised to the LCG resource brokers via the BDII
- http://is.grid.iu.edu/cgi-bin/status.cgi
- Frequently run jobs from ATLAS, Geant4, Auger, CMS
- LCG and OSG software stacks very similar but we do find differences in usage pattern. LCG jobs tend to bring bigger executable tarball along instead of pre-staging in applications area.

PRAGMA Interoperability work

- T. Ikegami of Japan—Used our globus/gt2 interface to claim multiple slots on same node, low-bandwidth parallel computing with master on his end on workers on ours.
- Steve Androulakis of Monash U., Australia—testing NIMROD as a grid client submission.
- Currently there is an interoperability project between PRAGMA grid people and the Engagement group of Open Science Grid.

Metrics and performance

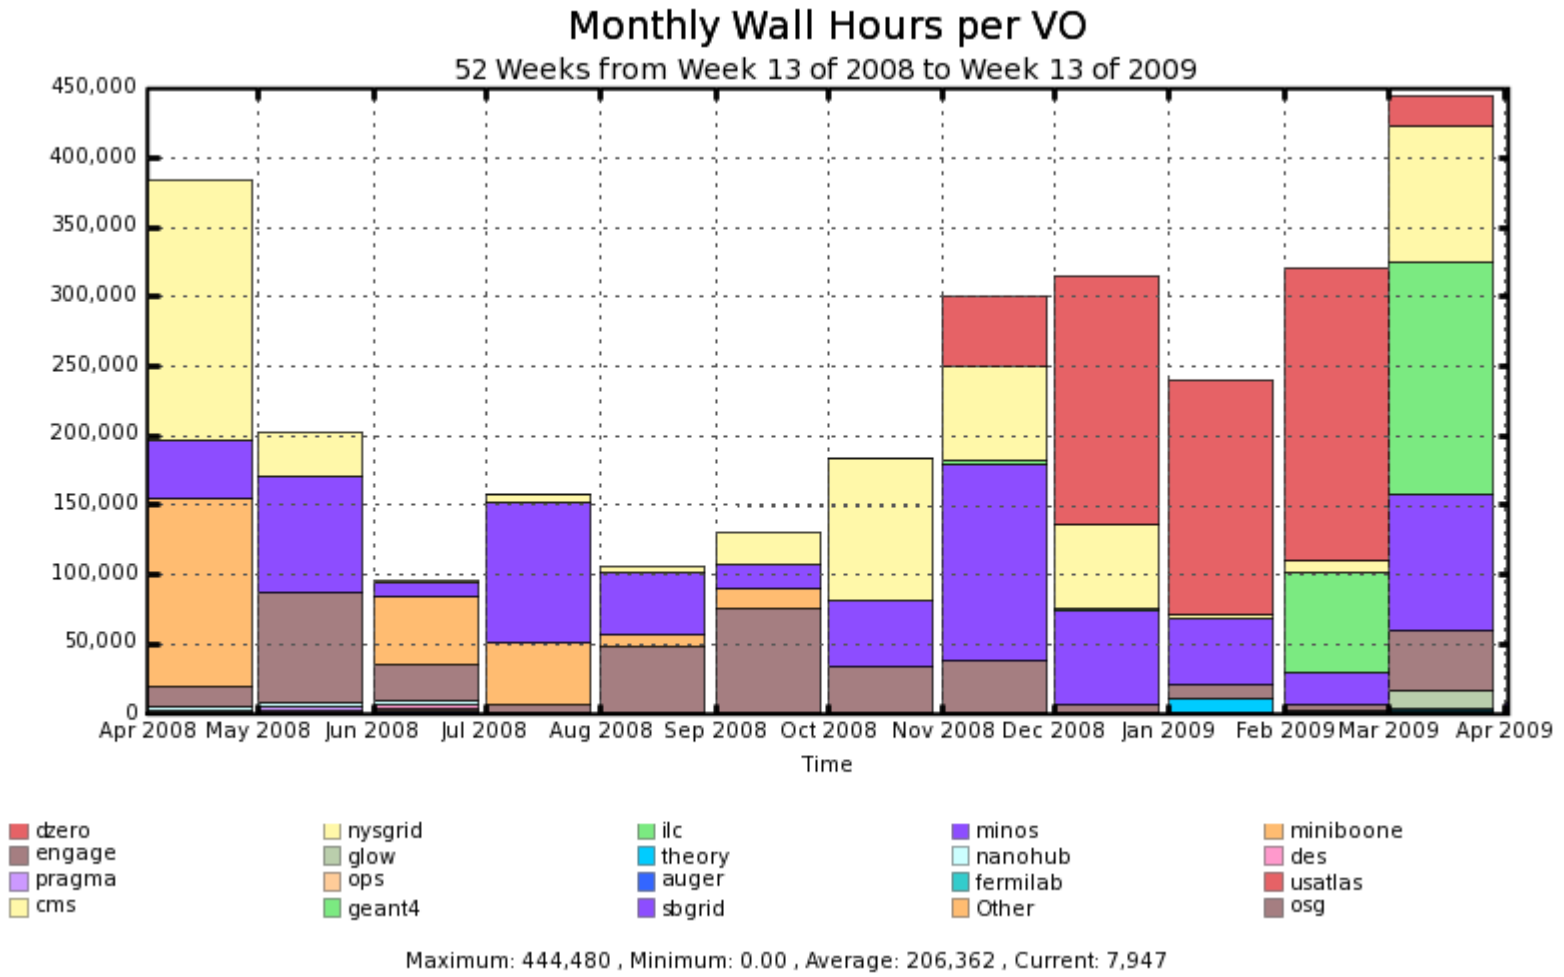

Current challenges

- Large numbers of jobs can build up one or two at a time on clusters where they have no chance to run.
  - Information system reports slots free but they are gone by the time the job gets there.
  - Need to tweak our matchmaking to include relative priorities
  - i.e. The chance that you will run on CDF's cluster when CDF has jobs waiting to run is 1 in 5*10^7.
- Large numbers of opportunistic jobs take over a cluster and the major stakeholder wants to get it back quickly.
  - Condor relies on pre-emption to do this. We give 48 hours for the job to finish. Sometimes major stakeholders want it to happen faster.
  - Currently tweaking timeouts so a job matched to one that is supposed to exit, can look for a different free slot when one comes free.

Conclusions

- The FermiGrid site gateway has allowed us to improve overall cluster utilization to > 80%
- We have saved hardware money by effectively using cycles that would otherwise have gone unused. Analyses that would have taken four months to complete on our small “general purpose” cluster can complete in a couple of weekends.
- We can provide burst capacity for large urgent projects
  - 10^6 cpu-hours delivered in 6 months for ILC 4th letter of intent.
  - Equal amount delivered for D0 Top matrix analysis.

Future Plans

- Make the gateway highly available.
  - Two condor collectors and failover for negotiator—Coming soon.
  - Two ReSS web services—Coming soon.
  - Have two globus-gatekeeper and condor_schedd machines so one will always be up—Easy enough
  - Active-passive failover for globus gatekeeper and condor_schedd. This is tricky, requires a HA file system that is very performant but still works right with Xen. NFS won't work.
- Improve monitoring and metrics
- Add our internal site BDII
  - Alternate method for site people to add resources
  - Use new feature of OSG software to show more of the structure of FermiGrid to the outside world.

Future Plans continued

- Working on making this software available to the rest of the Open Science Grid as part of the Virtual Data Toolkit. Other campus and organizational grids are interested.
- Continue to tune matchmaking algorithms to deal with corner cases in matchmaking
- Explore making the gateway less dependent on shared NFS volumes
- Leverage the technology to build other types of grid gateways. (See tomorrow's talk on OSG-TeraGrid gateway)