



Metrics and Monitoring on FermiGrid

Keith Chadwick

Fermilab

chadwick@fnal.gov



Outline

FermiGrid Introduction and Background
Metrics
Service Monitoring
Availability (Acceptance) Monitoring
Dashboard
Lessons Learned
Future Plans



Personnel

Eileen Berman, Fermilab, Batavia, IL 60510	berman@fnal.gov	
Philippe Canal, Fermilab, Batavia, IL 60510	pcanal@fnal.gov	
Keith Chadwick, Fermilab, Batavia, IL 60510	chadwick@fnal.gov	*
David Dykstra, Fermilab, Batavia, IL 60510	dwd@fnal.gov	
Ted Hesselroth, Fermilab, Batavia, IL, 60510	tdh@fnal.gov	
Gabriele Garzoglio, Fermilab, Batavia, IL 60510	garzogli@fnal.gov	
Chris Green, Fermilab, Batavia, IL 60510	greenc@fnal.gov	
Tanya Levshina, Fermilab, Batavia, IL 60510	tlevshin@fnal.gov	
Don Petravick, Fermilab, Batavia, IL 60510	petravick@fnal.gov	
Ruth Pordes, Fermilab, Batavia, IL 60510	ruth@fnal.gov	
Valery Sergeev, Fermilab, Batavia, IL 60510	sergeev@fnal.gov	*
Igor Sfiligoi, Fermilab, Batavia, IL 60510	sfiligoi@fnal.gov	
Neha Sharma Batavia, IL 60510	neha@fnal.gov	*
Steven Timm, Fermilab, Batavia, IL 60510	timmm@fnal.gov	*
D.R. Yocum, Fermilab, Batavia, IL 60510	yocum@fnal.gov	*



What is FermiGrid?

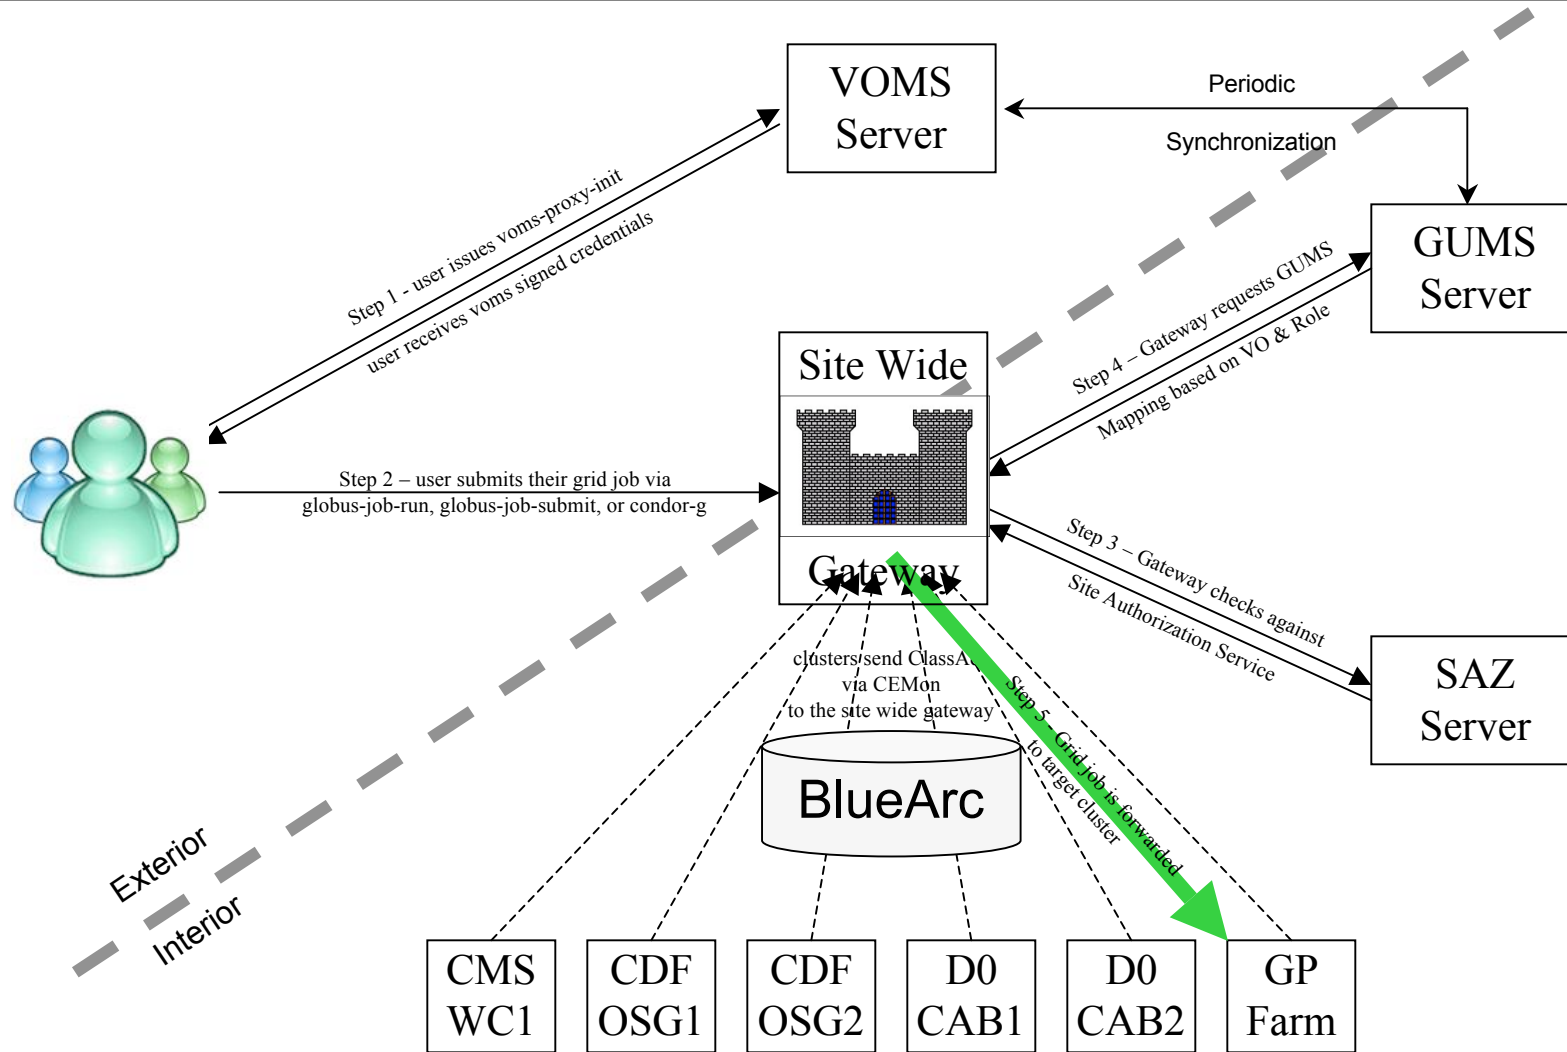
FermiGrid is:

- The Fermilab campus Grid and Grid portal.
 - The site globus gateway.
 - Accepts jobs from external (to Fermilab) sources and forwards the jobs onto internal clusters.
- A set of common services to support the campus Grid and interface to Open Science Grid (OSG) / LHC Computing Grid (LCG):
 - VOMS, VOMRS, GUMS, SAZ, MyProxy, Squid, Gratia Accounting, etc.
- A forum for promoting stakeholder interoperability and resource sharing within Fermilab:
 - CMS, CDF, D0;
 - ktev, miniboone, minos, mipp, etc.
- The Open Science Grid portal to Fermilab Compute and Storage Services.

FermiGrid Web Site & Additional Documentation:

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

FermiGrid - Current Architecture





Software Stack

Baseline:

- SL 3.0.x, SL 4.x, SL 5.0 (just released)
- OSG 0.6.0 (VDT 1.6.1, GT 4, WS-Gram, Pre-WS Gram)

Additional Components:

- VOMS (VO Management Service)
- VOMRS (VO Membership Registration Service)
- GUMS (Grid User Mapping Service)
- SAZ (Site AuthoriZation Service)
- jobmanager-cemon (job forwarding job manager)
- MyProxy (credential storage)
- Squid (web proxy cache)
- syslog-ng (auditing)
- Gratia (accounting)
- Xen (virtualization)
- Linux-HA (high availability)



Timeline

FermiGrid services were initially deployed in April 1, 2005.

The first formal metrics collection was commissioned in late August 2005.

- Initially a manual process.
- Automated during the fall of 2005.

Service monitoring was commissioned in June 2006.

VO Acceptance monitoring was commissioned in August 2006.

Availability monitoring was commissioned earlier this month.



Metrics vs. Monitoring

Metrics collection:

- Takes place once per day.

Service Monitoring:

- Takes place multiple times per day (typically once an hour).
- May have abilities to detect failed (or about to failed) services, notify administrators and (optionally) restart the service.
- Generates capacity planning information.

Acceptance Monitoring:

- Does a grid site accept “my” VO and pass a minimal set of tests.
- May not guarantee that a real application can run - just that it can get in the door.

Availability Monitoring:

- Very lightweight.
- Can be run very frequently (multiple times per hour).
- Optional automatic notification if results are “unexpected”.
- Feeds automatic “Dashboard” display.



Metrics Collection - Mechanics

Metrics collection is implemented on FermiGrid as follows:

- A central metrics collection system launches a central metrics collection process once per day.
 - `collect_grid_metrics.sh`
- The central metrics collection process in turn launches copies of itself (secondary metrics collection processes) via ssh across all systems (and the services) that are designated for metrics collection.
 - `collect_grid_metrics.sh <node> <service> <date> <...>`
- The secondary metrics collection processes identify the system, service and metrics to be collected, and then launch a script which has been custom written to collect the desired metrics from the specified service.
 - `collect-globus-metrics.sh <date> <...>`
 - `collect-voms-metrics.sh <date> <...>`



Metrics collected within FermiGrid

Globus Gatekeeper:

- # of authenticated, authorized, jobmanager, jobmanager-fork, jobmanager-managedfork
- batch (jobmanager-condor, jobmanager-pbs, etc.), jobmanager-condorg, jobmanager-cemon,
- jobmanager-mis, default.
- # of total IP connections, # of unique IP connections, # of unique IP connections from within Fermilab.

VOMS:

- # of voms-proxy-init's by VO.
- # of voms-proxy-init's by group within the fermilab VO.
- # of total IP connections, # of unique IP connections, # of unique IP connections from within Fermilab.

GUMS:

- # of successful GUMS mapping calls & # of failed GUMS mapping calls.
- # of total certificates, # of unique dn, # of unique mappings, # of unique Vos
- # of voms-proxy-inits, # of grid-proxy-inits.
- # of total IP connections, # of unique IP connections, # of unique IP connections from within Fermilab.

SAZ:

- # of successful SAZ calls & # of rejected SAZ calls.
- # of unique DN, # of unique VO, # of unique Role, # of unique CA.
- # of total IP connections, # of unique IP connections, # of unique IP connections from within Fermilab.


Metrics are stored using two mechanisms:

- First, they are appended to “.csv” files which contain a leading date followed by tag-value pairs. Example:
 - 22-Jun-2007,total=5721,success=5698,fails=53
 - total_ip=5721,unique_ip=231,fermilab_ip=12
- Second, the “.csv” files are processed and loaded in to round robin databases using rrdtool.


A set of “standard” png plots are automatically generated from the rrdtool databases.

All of these formats (.csv, .rrd and .png) are periodically uploaded from the metrics collection host to the central FermiGrid web server.


Globus Gatekeeper Metrics 1




Globus Gatekeeper Metrics 2



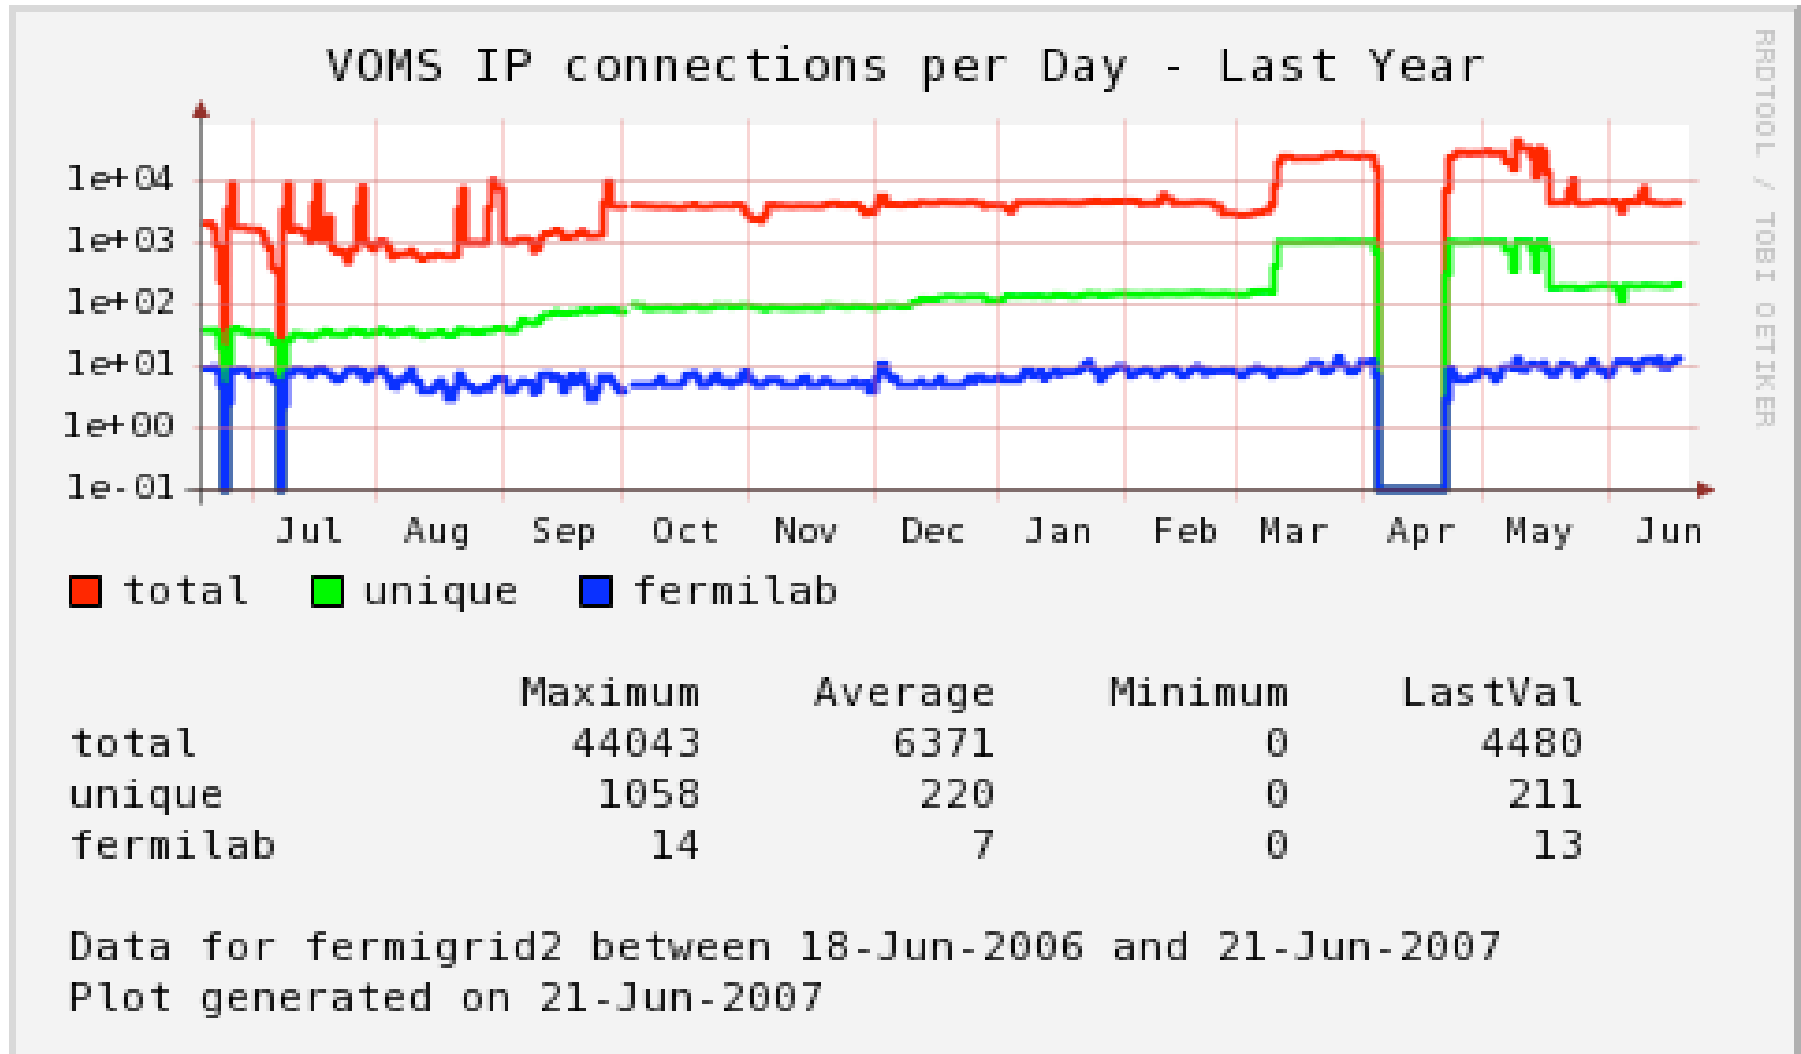
VOMS Metrics 1




VOMS Metrics 2




VOMS Metrics 3




GUMS Metrics 1




GUMS Metrics 2



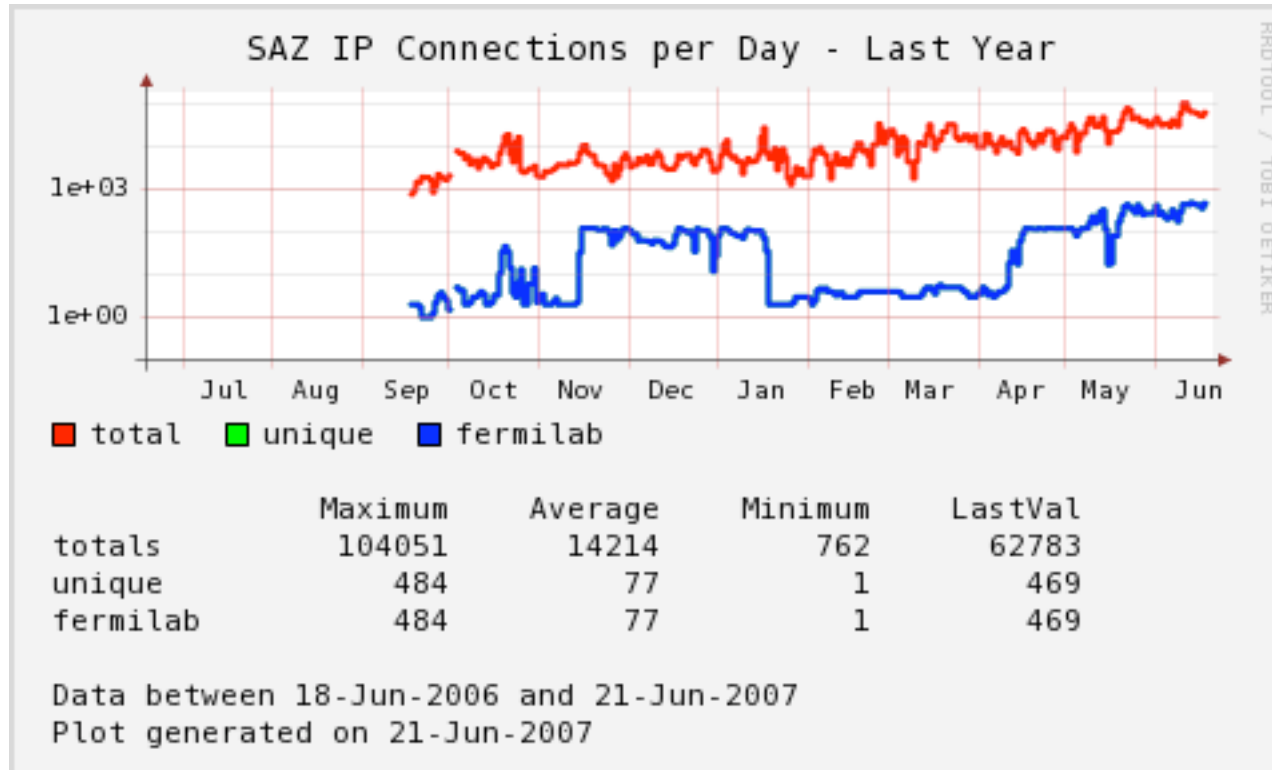
GUMS Metrics 3




SAZ Metrics 1



SAZ Metrics 2



SAZ Metrics 3





Service Monitoring - Mechanics

- A central service monitor system launches the central service monitor collection script once per hour.
 - `monitor_grid_script.sh`
- The central service monitor process in turn launches **background** copies of itself (secondary service monitor processes) across all systems (and the services) that are designated for service monitoring.
 - `monitor_grid_script.sh`
- The secondary service monitor processes identify the system, service to be monitored, and then launch a script which has been custom written to monitor the specified service.
 - `monitor_<service>_script.sh`
 - `monitor_gatekeeper_script.sh`
 - `monitor_voms_script.sh`
 - `monitor_gums_script.sh`
 - `monitor_saz_script.sh`



Service Monitor Configuration

Configuration of the service monitor system is via a central configuration file:

```
fermigrd0          fermigrd0.fnal.gov      master
fermigrd1          fermigrd1.fnal.gov      publish          var/www/html
#
fermigrd0          fermigrd0.fnal.gov      vo              fermilab
fermigrd1          fermigrd1.fnal.gov      gatekeeper
fermigrd2          fermigrd2.fnal.gov      voms            voms.fnal.gov
fermigrd3          fermigrd3.fnal.gov      gums            gums.fnal.gov
fermigrd3          fermigrd3.fnal.gov      mapping         cms
fermigrd3          fermigrd3.fnal.gov      mapping         dteam
fermigrd4          fermigrd4.fnal.gov      saz             saz.fnal.gov
fermigrd4          fermigrd4.fnal.gov      myproxy         myproxy.fnal.gov
fermigrd4          fermigrd4.fnal.gov      squid           squid.fnal.gov
#
fcdfosg1          fcdfosg1.fnal.gov      gatekeeper
fcdfosg2          fcdfosg2.fnal.gov      gatekeeper
d0cabosg1         d0cabosg1.fnal.gov      gatekeeper      ssh:/grid/login/chadwick
d0cabosg2         d0cabosg2.fnal.gov      gatekeeper      ssh:/grid/login/chadwick
###cmsosgce       cmsosgce.fnal.gov      gatekeeper      grid:/uscms/osg/app/fermilab/chadwick
###cmsosgce2     cmsosgce2.fnal.gov      gatekeeper      grid:/uscms/osg/app/fermilab/chadwick
```




Service Monitor - Information Collected

Globus Gatekeeper:

- # of authenticated, authorized, jobmanager, jobmanager-fork, jobmanager-managedfork, batch (condor, pbs, Isf, etc.), condorg/cemon, mis, default.
- The value of uptime, load1, load5 and load15.

VOMS:

- # of voms-proxy-init's
- # of apache and tomcat processes
- The rss and vmz of the Tomcat VOMS server process.
- The value of uptime, load1, load5 and load15.

GUMS:

- # of successful GUMS mapping calls & # of failed GUMS mapping calls.
- # of apache and tomcat processes
- The rss and vmz of the Tomcat GUMS server process.
- The value of uptime, load1, load5 and load15.

SAZ:

- # of successful SAZ calls & # of rejected SAZ calls.
- # of apache and tomcat processes
- The rss and vmz of the Tomcat SAZ server process.
- The value of uptime, load1, load5 and load15.



Service Monitor Storage and Publication

Results of the service monitors are stored using two mechanisms:


- First, they are appended to “.csv” files which contain a leading time (in seconds from the Unix epoch) followed by tag-value pairs. Example:
 - `time=1182466920,authenticated=42,authorized=26,jobmanager=26`


Second, the “.csv” files are processed and loaded in to round robin databases using rrdtool.

A set of “standard” png plots are automatically generated from the rrdtool databases.


All of these formats (.csv, .rrd and .png) are periodically uploaded from the metrics collection host to the central FermiGrid web server.

Globus Gatekeeper Monitor 1






VOMS Monitor 1





VOMS Monitor 2






GUMS Monitor 1







SAZ Monitor 1



Monitor the acceptance of a VO across a Grid in order to:

- Identify where the members of the VO can consider running jobs.
 - Not a guarantee that the job can actually run.
- Identify misconfigured sites that advertise that they “support” the VO but to not actually accept jobs from VO members.
- Log formal trouble tickets through the OSG GOC.
 - Ideally have the sites respond and fix their configuration.
 - Unfortunately some sites have not been very responsive.
 - And still other sites have responded by removing support for the VO.



VO Acceptance Monitoring Mechanics


How it is done:

- A cron script periodically launches kcroninit.
- kcroninit launches a script which does authentication:
 - `kx509`
 - `kxlist -p`
- Robot certificate “issued” by the Fermilab KCA:
 - `/DC=gov/DC=fnal/O=Fermilab/OU=Robots/CN=cron/CN=Keith Chadwick/UID=chadwick`
- Get VO signed credentials:
 - `voms-proxy-init -noregen -voms fermilab:/fermilab`
- Pulls the list of OSG sites from the OSG gridscan reports
 - http://scan.grid.iu.edu/cgi-bin/get_grid_sv?get=set1
- For each site in the report, the acceptance monitor tests:
 - Unix ping.
 - `globusrun -a -r` (authenticate).
 - `globus-job-run` (existing application - `typ /usr/bin/id`).
 - `globus-url-copy` (to and from).
- Periodically I review the list of failing sites and if appropriate, log trouble tickets.



VO Acceptance Monitor 1

fermigrid0 - fermilab VO Probe production Monitor - Last Year



RINTOOL / TOBI OETIKER

total
 vors
 success
 ping
 globus-job-run
 globus-url-copy
 catalog

	Maximum	Average	Minimum	LastVal
total	82	67	53	82
vors	43	30	0	37
success	32	24	5	29
ping	24	13	6	15
globus-job-run	48	26	18	33
globus-url-copy	18	5	0	6
catalog lookup	4	0	0	0

Data for fermigrid0 between 18-Jun-2006 and 21-Jun-2007
 Plot generated on 21-Jun-2007



Availability (Infrastructure) Monitoring

Designed to be very “lightweight”.

Currently running with the service monitor, but designed and implemented so that it can run much more frequently.


Monitors both the host system and the service which is running on the system.

Driven by the same configuration file as the service monitor.

<http://fermigrid.fnal.gov/monitor/fermigrid0-ping-monitor.html>



Base Infrastructure Monitor





Dashboard

Based on a secondary analysis of the infrastructure monitor data.

Design goal is to be a simple “health” dashboard:

<http://fermigrid.fnal.gov/monitor/fermigrid-dashboard.html>



Dashboard - Typical Display

FermiGrid Dashboard Summary - Updated at Thu Jun 21 18:01:50 CDT 2007

System	Ip Name	Ping Status	Service Name	Service Status	Service Alias
fermigrd0	fermigrd0.fnal.gov	Up	master	Up	
fermigrd0	fermigrd0.fnal.gov	Up	vo	Up	
fermigrd1	fermigrd1.fnal.gov	Up	gatekeeper	Up	
fermigrd1	fermigrd1.fnal.gov	Up	publish	Up	
fermigrd2	fermigrd2.fnal.gov	Up	voms	Up	
fermigrd3	fermigrd3.fnal.gov	Up	gums	Up	
fermigrd3	fermigrd3.fnal.gov	Up	mapping	Up	
fermigrd4	fermigrd4.fnal.gov	Up	myproxy	Up	myproxy.fnal.gov
fermigrd4	fermigrd4.fnal.gov	Up	suz	Up	
fermigrd4	fermigrd4.fnal.gov	Up	squid	Up	
d0cabosg1	d0cabosg1.fnal.gov	Up	gatekeeper	Up	
d0cabosg2	d0cabosg2.fnal.gov	Up	gatekeeper	Up	
fcdfosg1	fcdfosg1.fnal.gov	Up	gatekeeper	Up	
fcdfosg2	fcdfosg2.fnal.gov	Up	gatekeeper	Up	
fgitb-gk	fgitb-gk.fnal.gov	Up	gatekeeper	Up	
fnpp-osg	fnpp-osg.fnal.gov	Up	gatekeeper	Up	
fgtest1	fgtest1.fnal.gov	Up	gatekeeper	Up	
fgtest2	fgtest2.fnal.gov	Up	voms	Up	
fgtest3	fgtest3.fnal.gov	Up	gums	Up	
fgtest3	fgtest3.fnal.gov	Up	mapping	Up	
fgtest4	fgtest4.fnal.gov	Up	suz	Up	
fgtest5	fgtest5.fnal.gov	Up	gums	Up	
fgtest5	fgtest5.fnal.gov	Up	mapping	Up	
fgtest6	fgtest6.fnal.gov	Up	gatekeeper	Up	
fgtest6	fgtest6.fnal.gov	Up	suz	Up	

Metrics and Service Monitoring is difficult:

- Every service has it's own log file format (at least today).
 - find, grep, awk are your friends.
 - The format of the messages within the service log file will change as new versions of the services are deployed.
- Some services don't log all necessary and/or interesting information "out of the box", they need additional logging options enabled.
 - You may have to work with the service developers to insure that they log the necessary service information.
- Some services are extremely "talkative" and place lots of information (that I am certain is useful to the developers) in the log file along with the "golden nuggets" that is needed by the metrics collection and service monitoring.
 - You may have to work with the service developers to insure that they log the necessary service information.
- You may have to extract and correlate information from multiple logs.
- You must also monitor services that the monitored service depends on (especially apache and tomcat).

Out of band access and monitoring is quite useful and necessary.

- ssh, ksu as well as grid.

Using grid services to monitor other grid services may not correctly identify the problem:

- Did some local (non-grid) service fail?
 - `kx509`, `kxlist -p`
- Did the local grid service fail?
 - `voms-proxy-init`
- Did some intermediate service fail or timeout?
 - Network congestion
- Did the remote grid service fail or timeout?
 - Globus gatekeeper

Service monitoring with automatic service recovery can be very useful.

- Especially when responding to automated security probing,
- And also for getting a full nights rest...

Automatic service recovery will usually require some level of root access.

- Sites are understandably reluctant to grant “remote” root access (I know that I am...).

Robot certificates are extremely useful for automating grid service monitoring.



Plans for the Future

Continue with the development of additional metrics and monitor probes.

Continue with the development of automated reports & publication.

Integrate/incorporate the new OSG SAM probes to fermilab VO monitoring.

As part of the FermiGrid-HA deployment, enhance the metrics and monitoring infrastructure:

- Collect from all [voms,gums,saz] service instances.
- Collate a HA view of the services.

Work towards making this infrastructure more portable.

Any questions?