

Fermilab FermiGrid/FermiCloud

Keith Chadwick

Fermilab Grid and Cloud
Computing Department Head

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

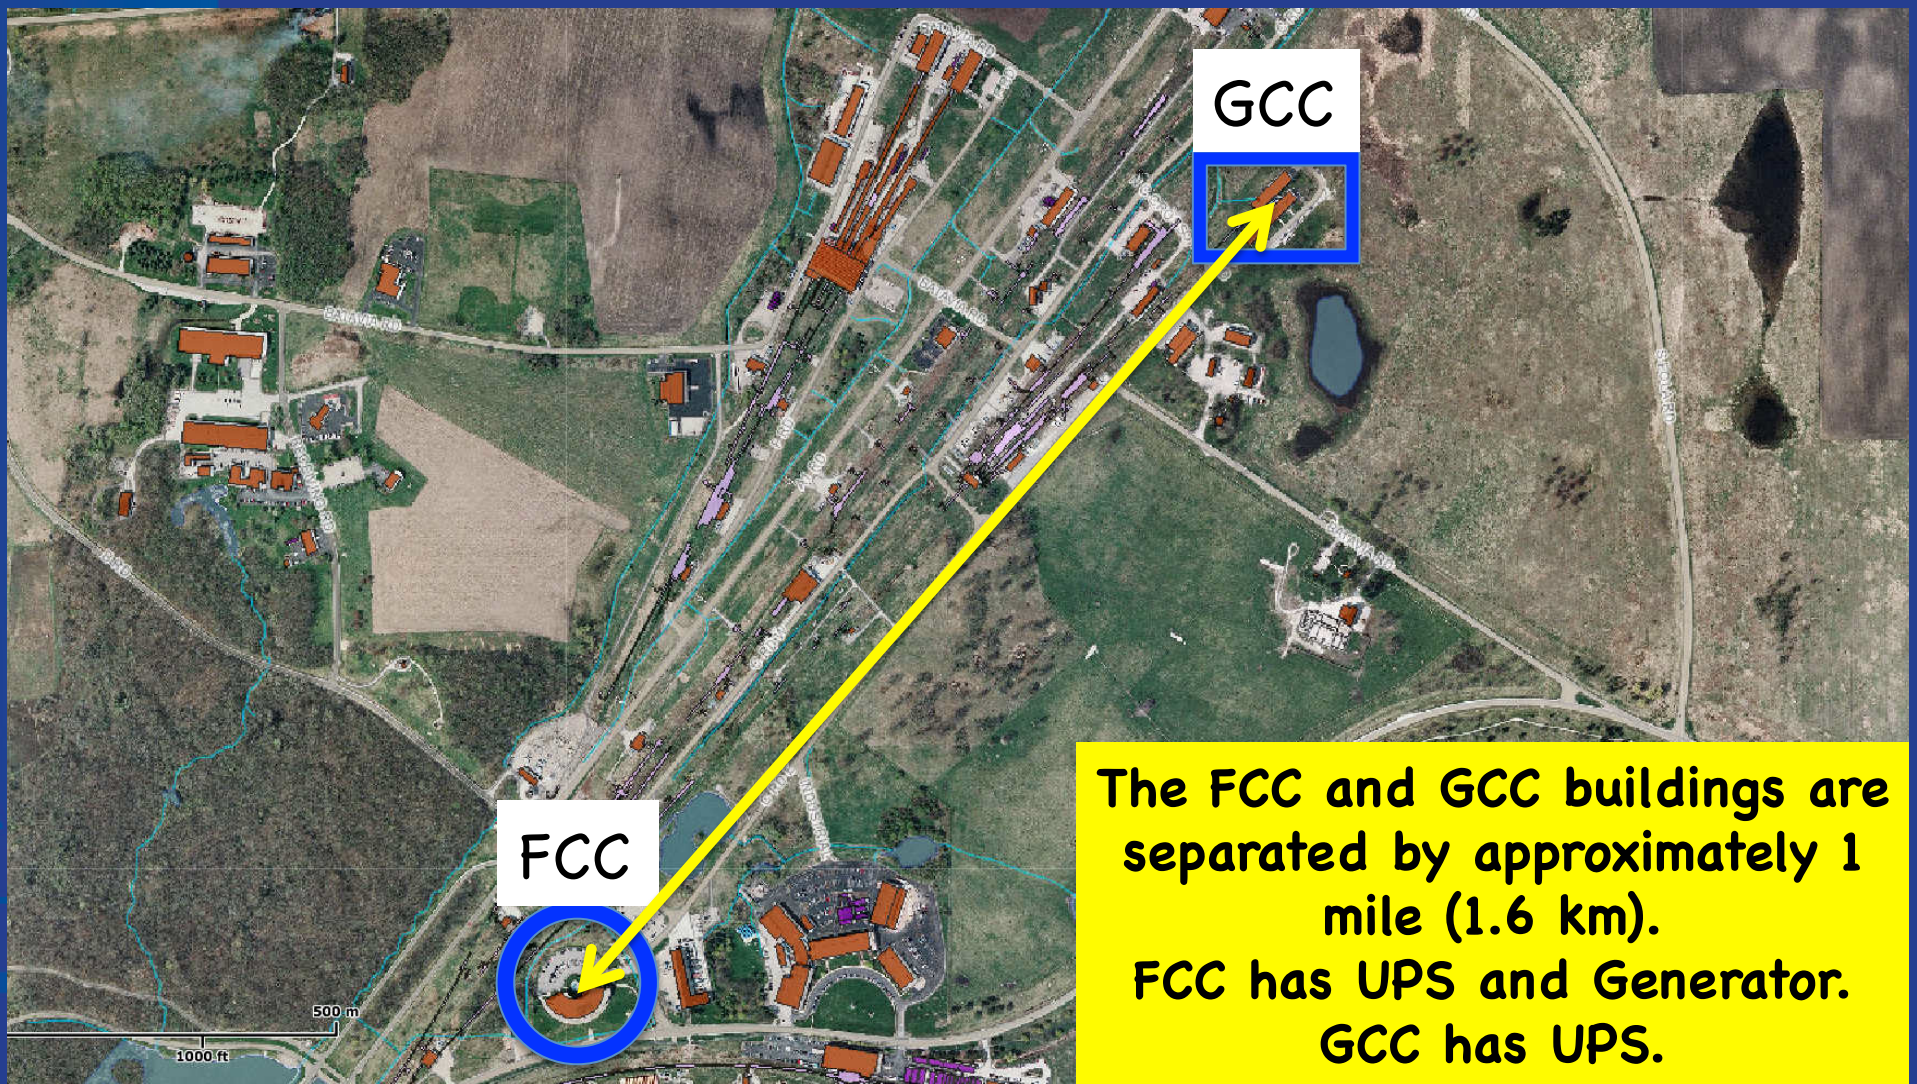
Background

- Fermilab is a stakeholder to the Gratia accounting project:
 - Provider of Gratia accounting software
 - Consumer of Gratia accounting reports
- The FermiGrid Services Group in the Fermilab Grid and Cloud Computing Department operates the Gratia collector/reporter infrastructure for both Fermilab and the Open Science Grid.
- Fermilab receives 0.25 FTE funding from the OSG to operate the OSG Gratia collector/reporter infrastructure and the OSG VOMS server. The majority (~80%) of this funding is charged against OSG Gratia operations. The last Gratia upgrade (completed in early 2012) required ~six months of elapsed time and almost entirely consumed the 0.25 FTE OSG provided funding for FY2012. Fortunately the other OSG services (VOMS, ReSS) operated at Fermilab did not require significant effort during FY2012.

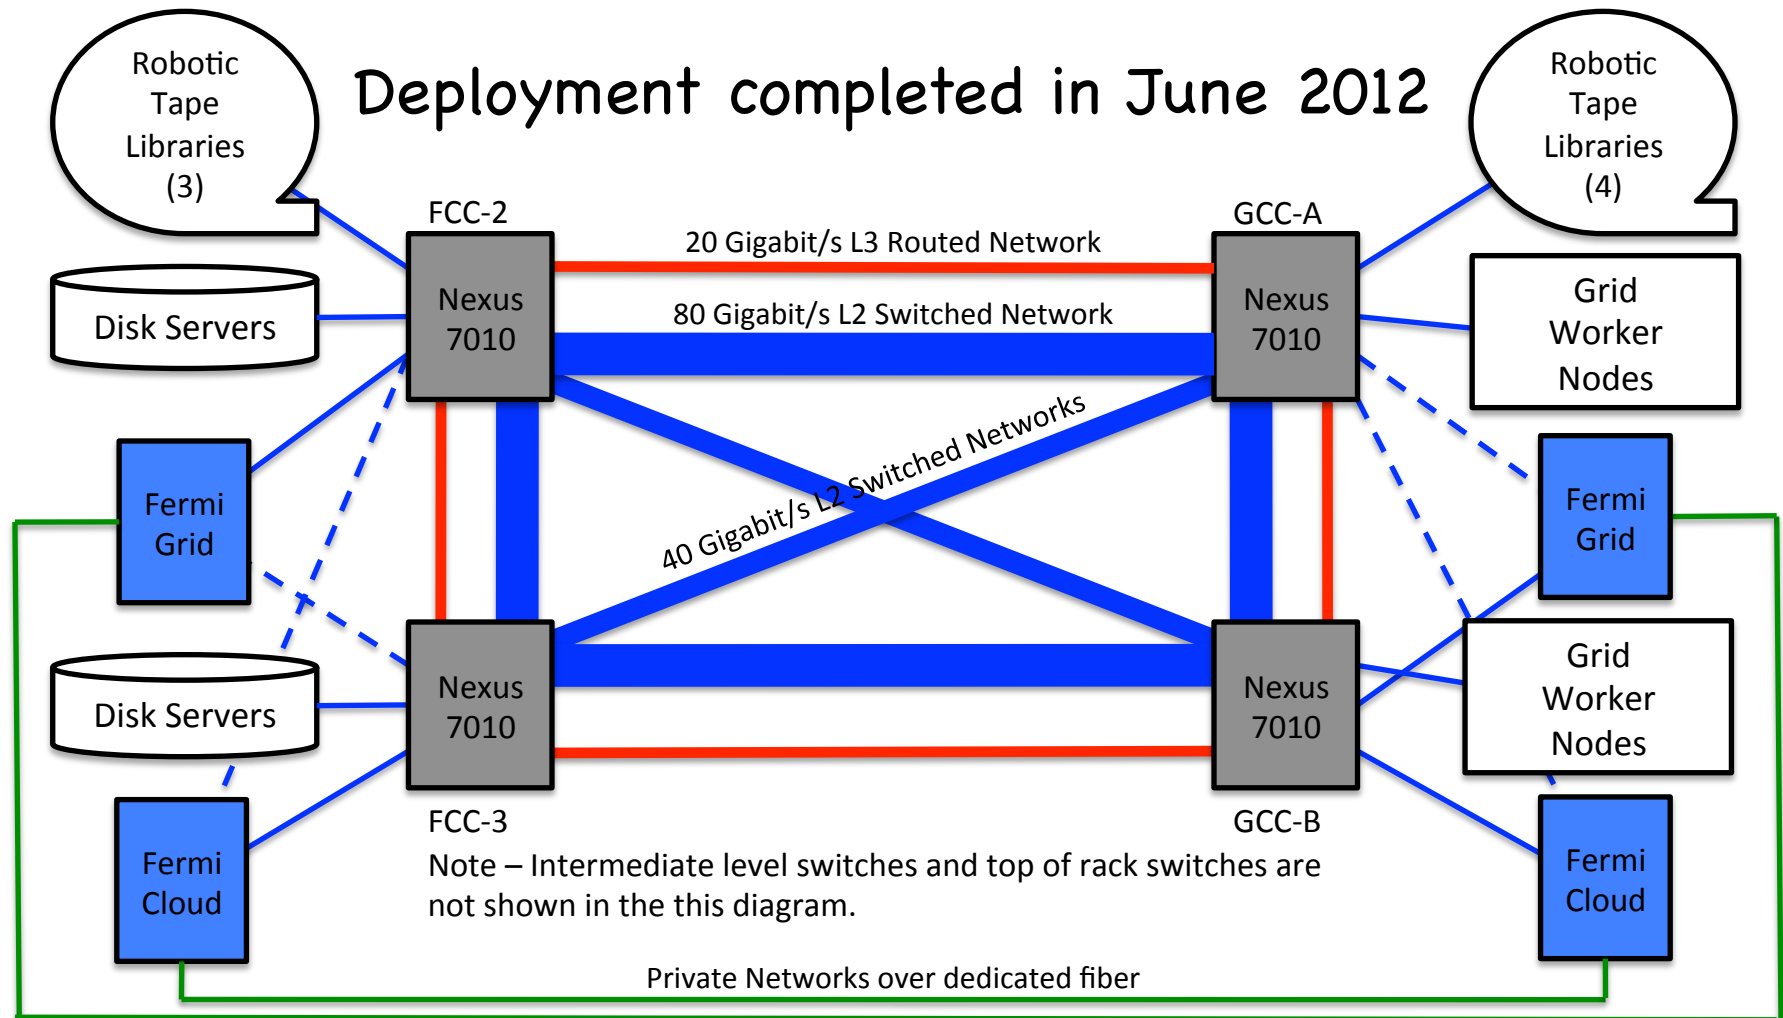
Gratia – Fault Tolerance

- The systems (and virtual machines) that host the Fermilab and OSG Gratia accounting services are:
 - Split across two buildings roughly 1 mile apart,
 - The buildings are connected via a fault tolerant network infrastructure.
- This is all part of the FermiGrid-HA2 infrastructure (completed in June 2011):
 - <http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=3739>

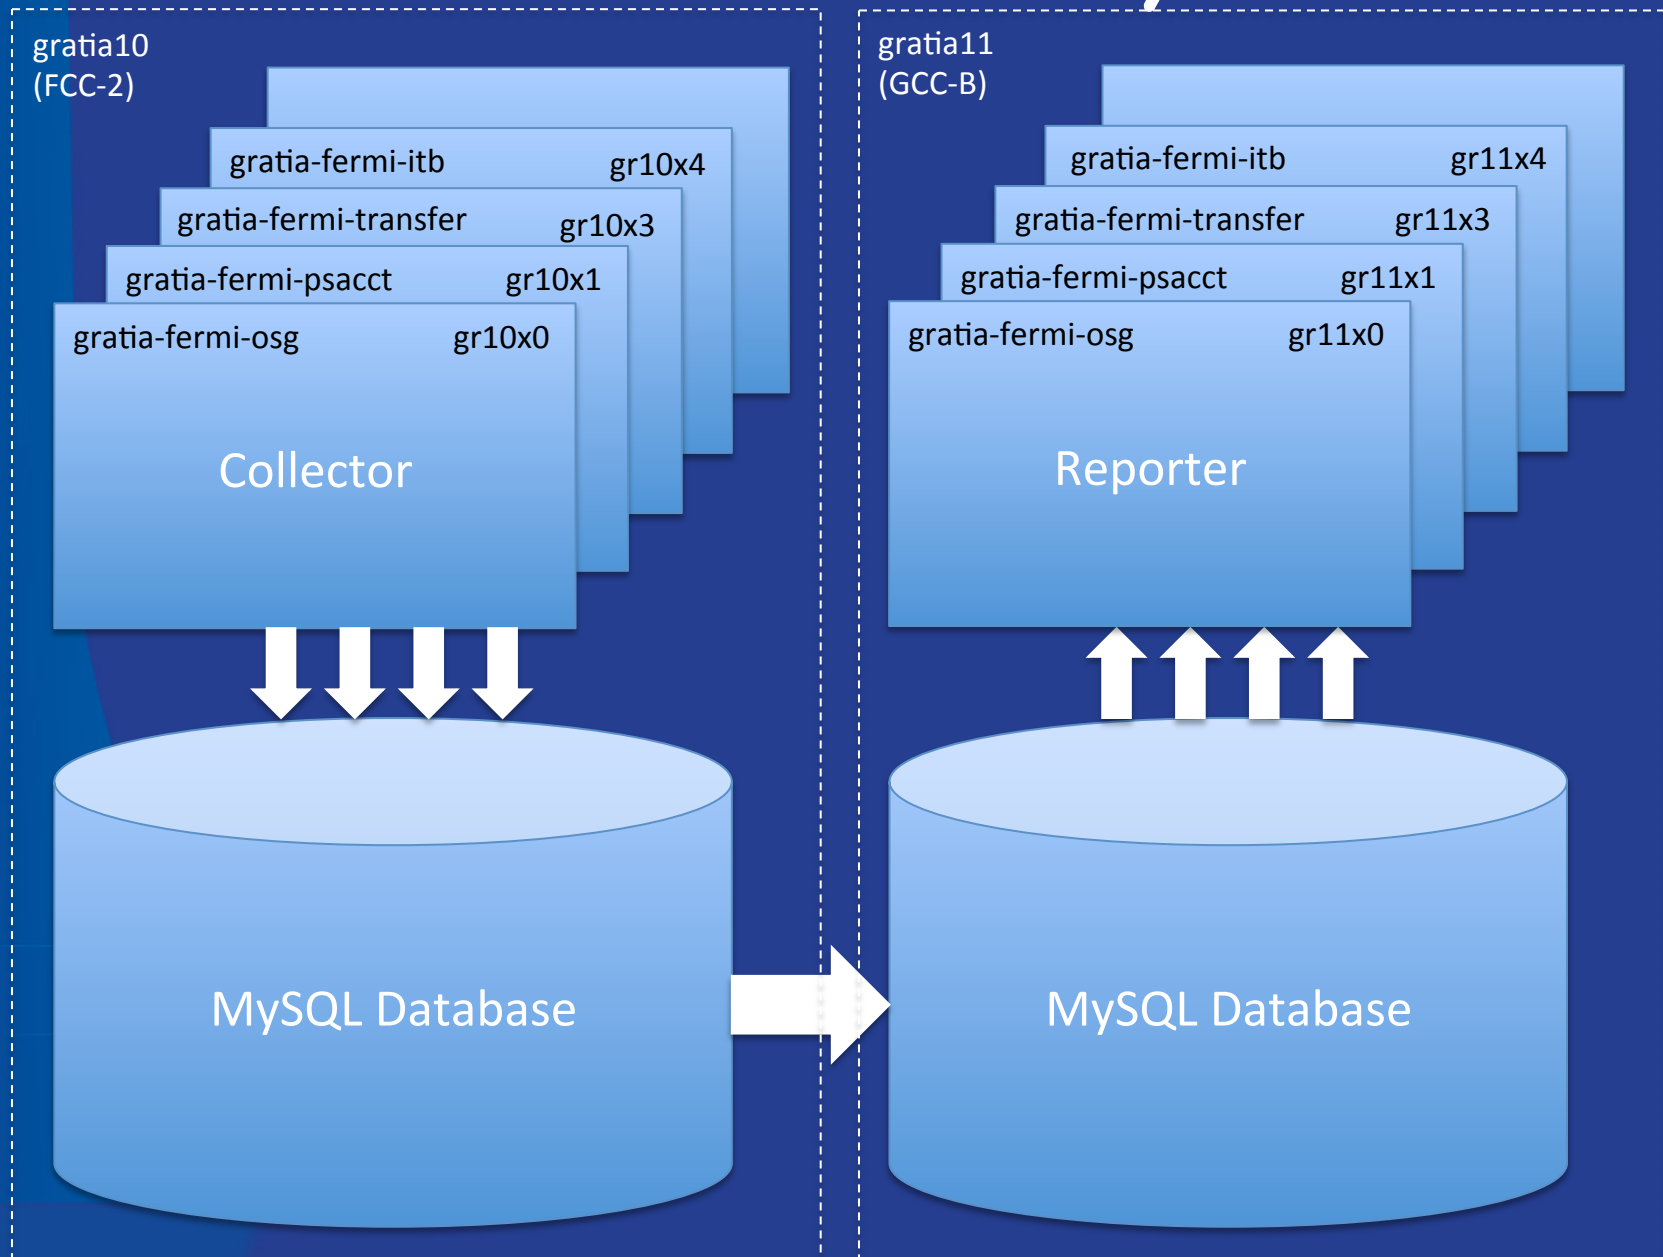
FCC and GCC



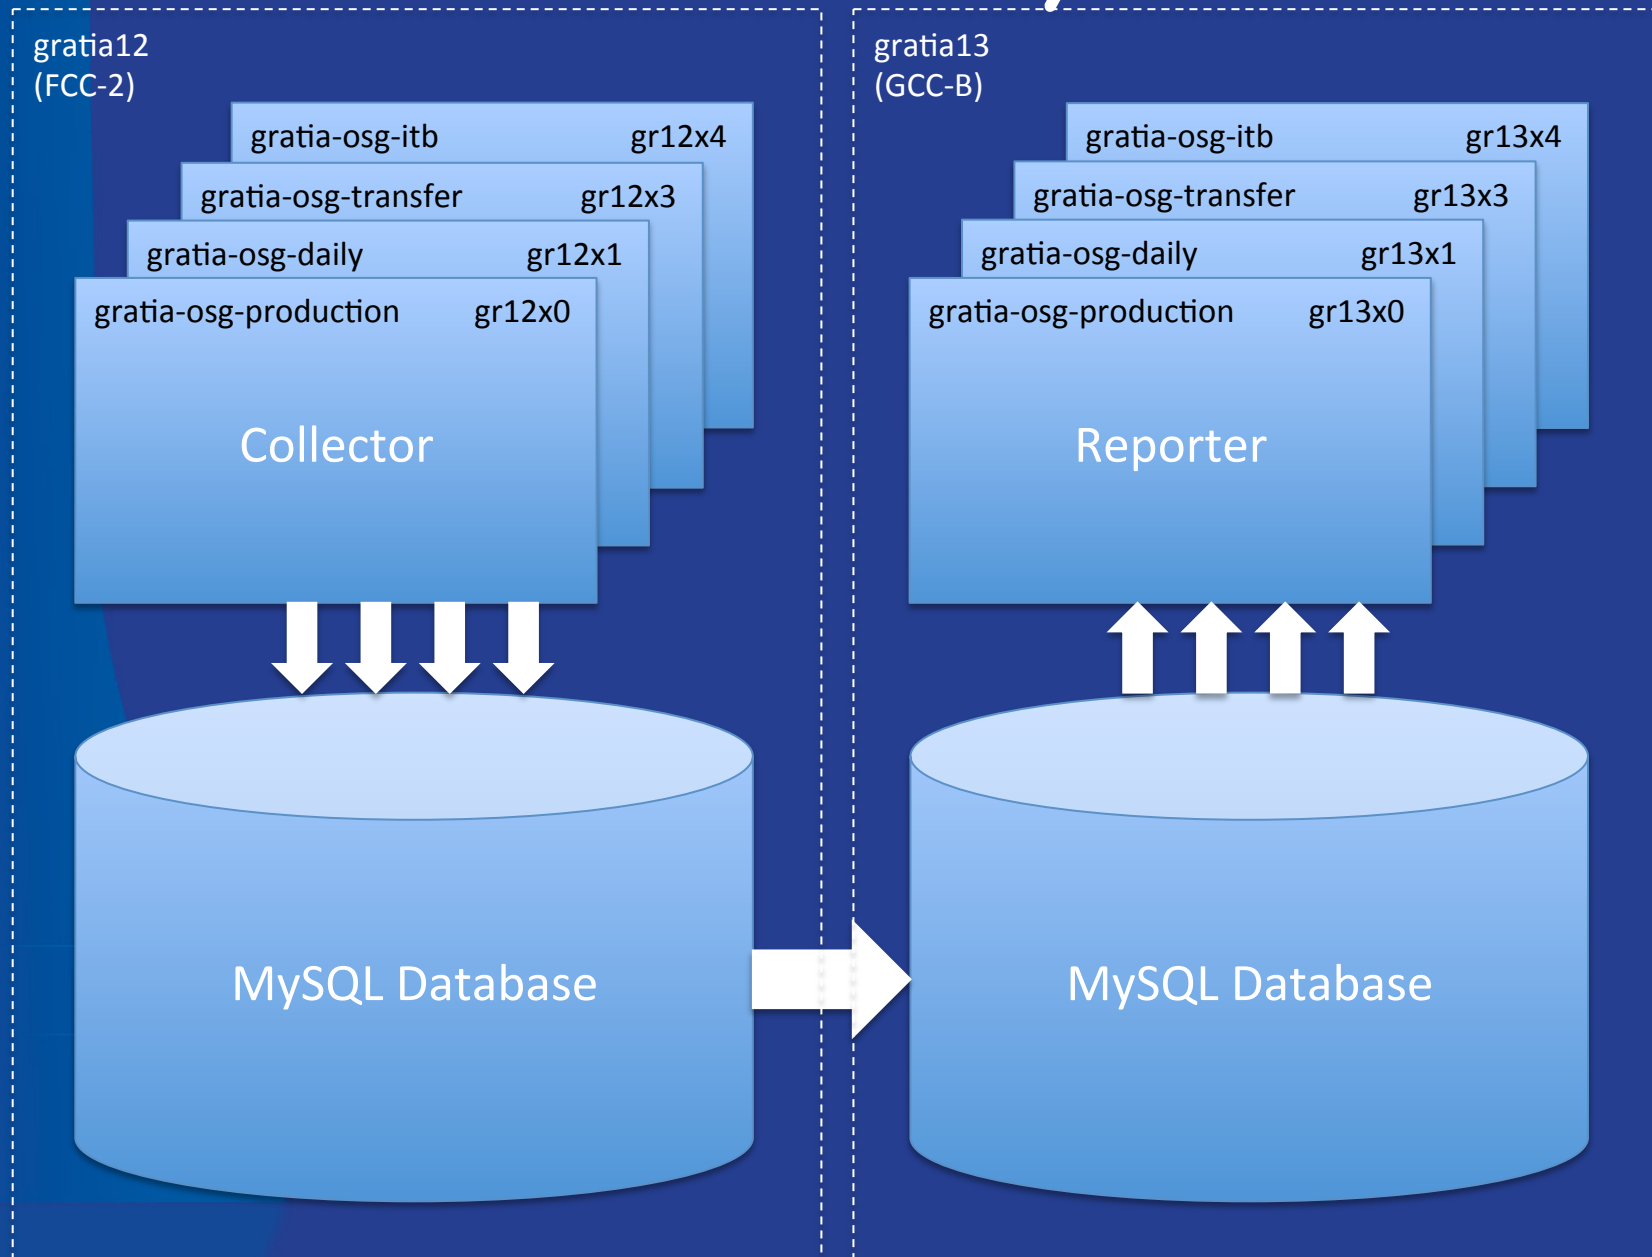
Distributed Network Core Provides Redundant Connectivity

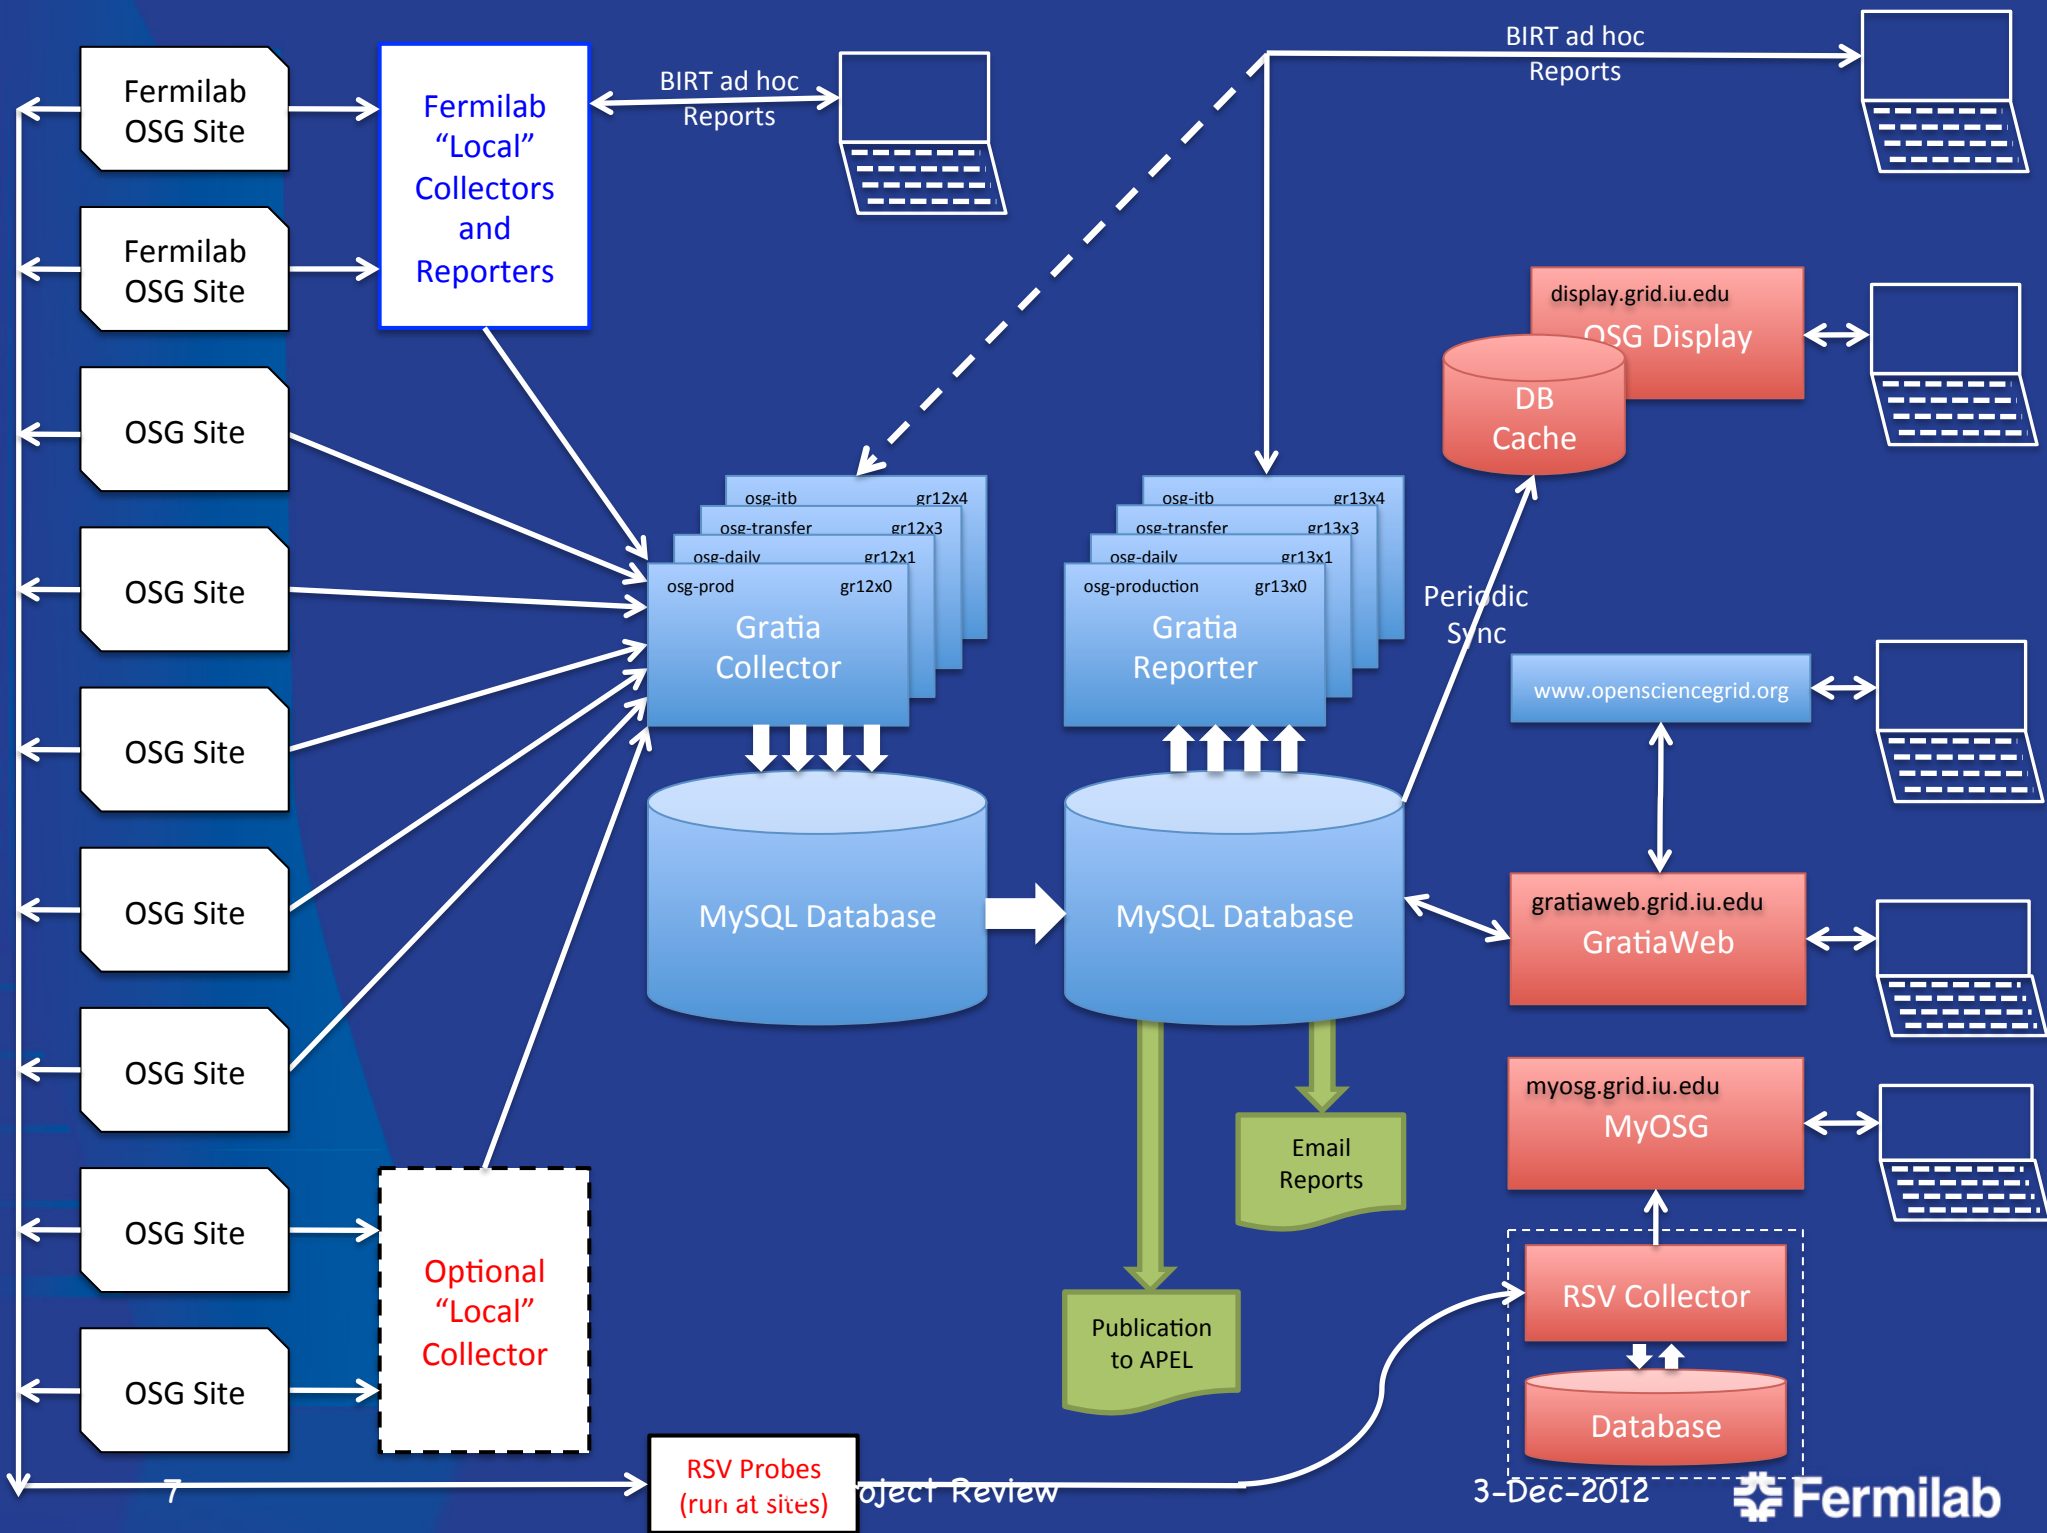


Fermilab Gratia Systems




OSG Gratia Systems






OSG Gratia Collector Database Size – Past Year



OSG Gratia Reporter Database Size – Past Year




Measured Service Availability

Gratia Service	Last Week	Last Month	Last Quarter	Last Year	Estimated Yearly Downtime (Minutes)
Fermilab Production	100%	100%	99.94%	99.98%	95
Fermilab Transfer	100%	100%	100%	100%	0
OSG Production	100%	100%	99.94%	99.98%	95
OSG Transfer	100%	100%	99.87%	99.96%	190
Fermilab Database	100%	100%	100%	100%	0
OSG Database	100%	100%	100%	100%	0

OSG Jobs/Month – Past Year



OSG Transfers/Month – Past Year



Past Issues – 1

- Development and support of the Gratia accounting system has been negatively impacted by lack of personnel and personnel turnover (especially in the area of project lead):
 - Matteo Melani (through 2006) -> Philippe Canal (2006 to 10/2011) -> Tanya Levshina (10/2011 onwards)...
 - M. Melani, P. Canal, C. Green, A. Garu, B. Bockelman, T. Levshina, H. Kim, etc.
- Operation of the Gratia accounting system is very complex:
 - Software issues – debugging software at scale only after has been deployed.
 - Maintenance issues – the last upgrade of the Gratia software on the central systems was a six month project.
 - (Remote) operational issues – Gratia is a complex system, and many (? most?) of the OSG site administrators tend to treat it as a black box.
 - OSG management issues – Unrealistic expectations of Gratia.

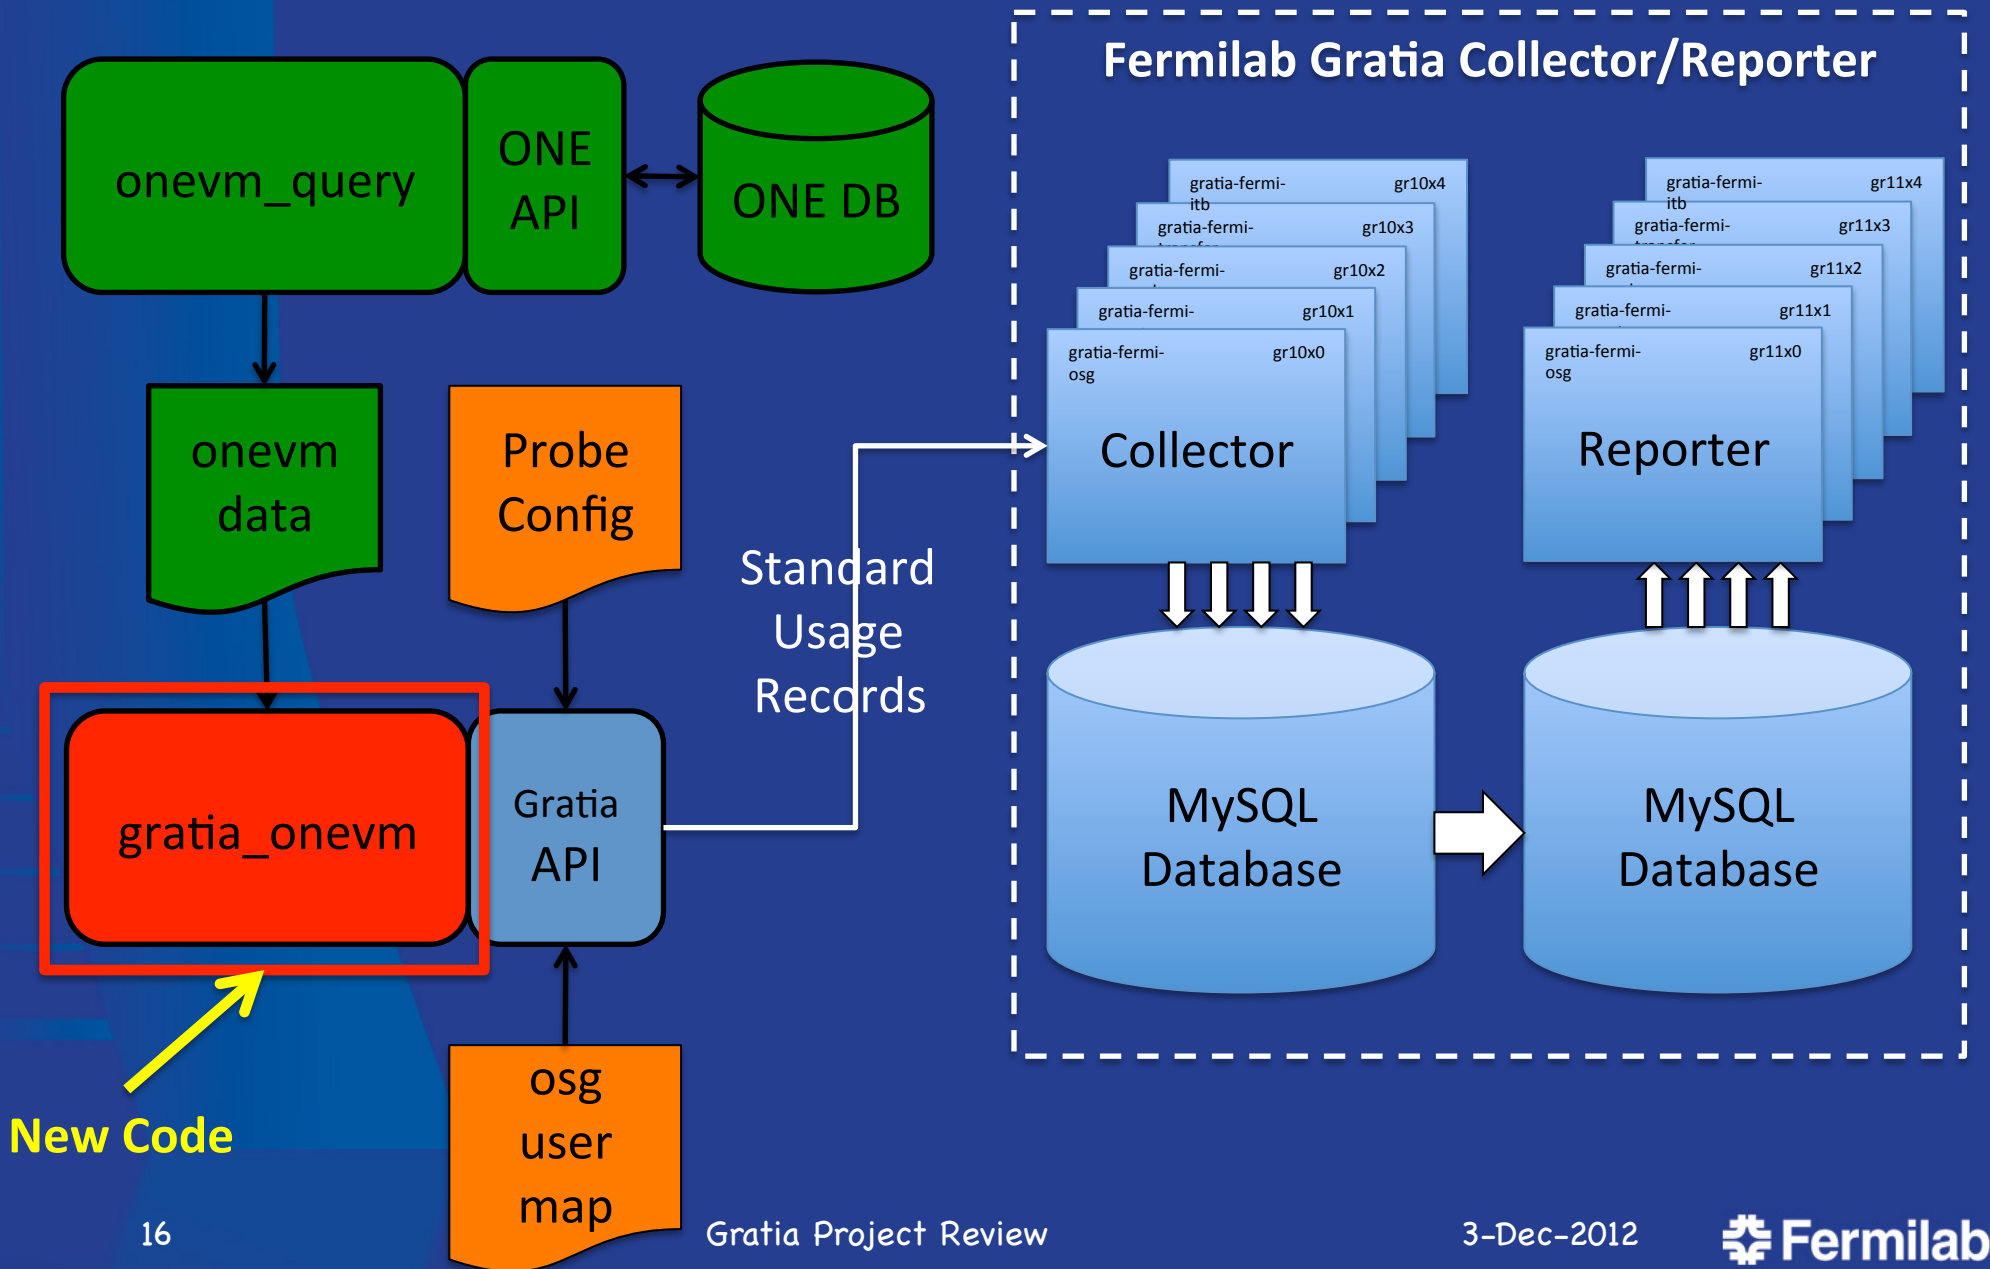
Past Issues – 2

- Gratia is an accounting service. Gratia is not:
 - A security service,
 - An audit service,
 - A monitoring service,
 - Or anything else other than an accounting service!
- The OSG has implemented new publically facing dependent services that consume Gratia data without any notification to or negotiation with the FGS operators prior to the design/commissioning/deployment:
 - OSG Display
 - GratiaWeb




FermiCloud Gratia Accounting

- Currently have two “probes” based on the Gratia accounting framework used by Fermilab and the Open Science Grid:
 - <https://twiki.grid.iu.edu/bin/view/Accounting/WebHome>
- Standard Process Accounting (“psacct”) Probe:
 - Installed and runs within the virtual machine image,
 - Reports to standard gratia-fermi-psacct.fnal.gov.
- Open Nebula Gratia Accounting Probe:
 - Runs on the OpenNebula management node and collects data from ONE logs, emits standard Gratia usage records,
 - Reports to the “virtualization” Gratia collector,
 - The “virtualization” Gratia collector runs existing standard Gratia collector software (no development was required),
 - The development of the Open Nebula Gratia accounting probe was performed by Tanya Levshina and Parag Mhashilkar.
- Additional Gratia accounting probes could be developed:
 - Commercial – OracleVM, VMware, Amazon EC2, Azure, ...
 - Open Source – Nimbus, Eucalyptus, OpenStack, ...

Open Nebula Gratia Accounting Probe



FermiCloud Gratia Reports



Feature Request from Fermilab HPC Department

- Enhance the Gratia storage probe to have the ability to distinguish GridFTP transfers attributable to Globus Online separately from other GridFTP transfers.

Future Issues / Questions – 1

- OSG Central Collector/Reporter:
 - Does it make sense for Fermilab to continue to operate them or should they be moved to the GOC at IU?
 - If these systems remain at Fermilab, then should the GOC run a mirror of the OSG Central Reporter database and run all the OSG Display/GratiaWeb/etc. queries against that copy?
 - Should the ability to submit arbitrary queries against the Gratia reporting interface be disabled and/or more strictly controlled?
- Effort required to install and upgrade the OSG Central Collector/Reporter needs to be significantly lowered.
 - Post deployment testing at scale is not a sustainable strategy.
- Better testing and installation tools / more installation automation is needed.

Future Issues / Questions – 2

- Plans for Gratia hardware upgrade/replacement?
 - The systems were last upgraded in 2010, they probably have another 2–3 years of life in them. Who funds and manages their replacement?
- Plans for Gratia software upgrade/replacement?
 - What happens if the Gratia accounting record rate increases 10x?
 - What happens if the Gratia accounting record rate increases 100x?
 - What happens if the Gratia accounting record rate increases 1000x?
- What about XSEDE?
 - What is the impact on Gratia?
- Is there a need for additional virtualization/cloud accounting probes?
 - Amazon EC2, Vmware, Azure, OpenStack, Eucalyptus, Nimbus, etc.
 - Who is going to fund their development?
 - The Gratia developers in the Grid and Cloud Computing department are willing to do this work, but the development must be funded by someone.