FermiCloud Update
ISGC 2013

Keith Chadwick
Grid & Cloud Computing Department
Fermilab

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

Fermilab

Outline

FermiCloud Introduction:

- Acknowledgements, Staffing, Previous Presentations,

FermiCloud Project:

- Mission, Project Phases,

FermiCloud Facility and Software:

- Fault Tolerance, Distributed Replicated File System, FermiCloud-HA,
- Current Capabilities, Typical Use Cases, Proposed Economic Model & SLA,
- Cloud Computing Environment, Idle VM Detection, vCluster
- Usage, Personnel Effort, Review

Summary & Conclusions

Acknowledgements

None of this work could have been accomplished without:

- The excellent support from other departments of the Fermilab Computing Sector – including Computing Facilities, Site Networking, and Logistics.
- The excellent collaboration with the open source communities – especially Scientific Linux and OpenNebula,
- As well as the excellent collaboration and contributions from the Korea Institute of Science and Technology Information (KISTI).

Fermilab Grid & Cloud Computing Department

Keith Chadwick (Department Head)
Gabriele Garzoglio (Associate Head)

Distributed Offline Computing Services	FermiGrid Services (operations)
Gabriele Garzoglio (Leader) David Dykstra Hyunwoo Kim Tanya Levshina Parag Mhashilkar Marko Slyz *2012 Summer students:* Siyuan Ma, Anupam Rajendran, Giovanni Franzini	**Steven C. Timm (Leader)** Gerard Bernabeu Altayo Kevin Hill (OSG Security) Neha Sharma Karen Shepelak

Previous Presentations @ ISGC

ISGC 2012:

- FermiGrid and FermiCloud Update
  - http://indico3.twgrid.org/indico/getFile.py/access?contribId=8&sessionId=26&resId=0&materialId=paper&confId=44

ISGC 2011:

- FermiCloud: Technology Evaluation and Pilot Service Deployment
  - http://event.twgrid.org/isgc2011/slides/GridsandClouds/1/FermiCloud-ISGC-2011.pdf
- FermiGrid Scalability and Reliability Improvements
  - http://event.twgrid.org/isgc2011/slides/OperationandManagement/2/FermiGrid-ISGC-2011.pdf

FermiCloud Mission

In late 2009 the Grid Facilities Department established the FermiCloud Project with the goal of developing and establishing Scientific Cloud capabilities for the Fermilab Scientific Program,

- Building on the very successful FermiGrid program that supports the full Fermilab user community and makes significant contributions as members of the Open Science Grid Consortium.
- Reuse High Availabilty, AuthZ/AuthN, Virtualization from Grid.

In a (very) broad brush, the mission of FermiCloud is:

- To deploy a production quality Infrastructure as a Service (IaaS) Cloud Computing capability in support of the Fermilab Scientific Program.
- To support additional IaaS, PaaS and SaaS Cloud Computing capabilities based on the FermiCloud infrastructure at Fermilab.

The FermiCloud project is split over several overlapping phases.

FermiCloud Project Delivered Over Overlapping Phases

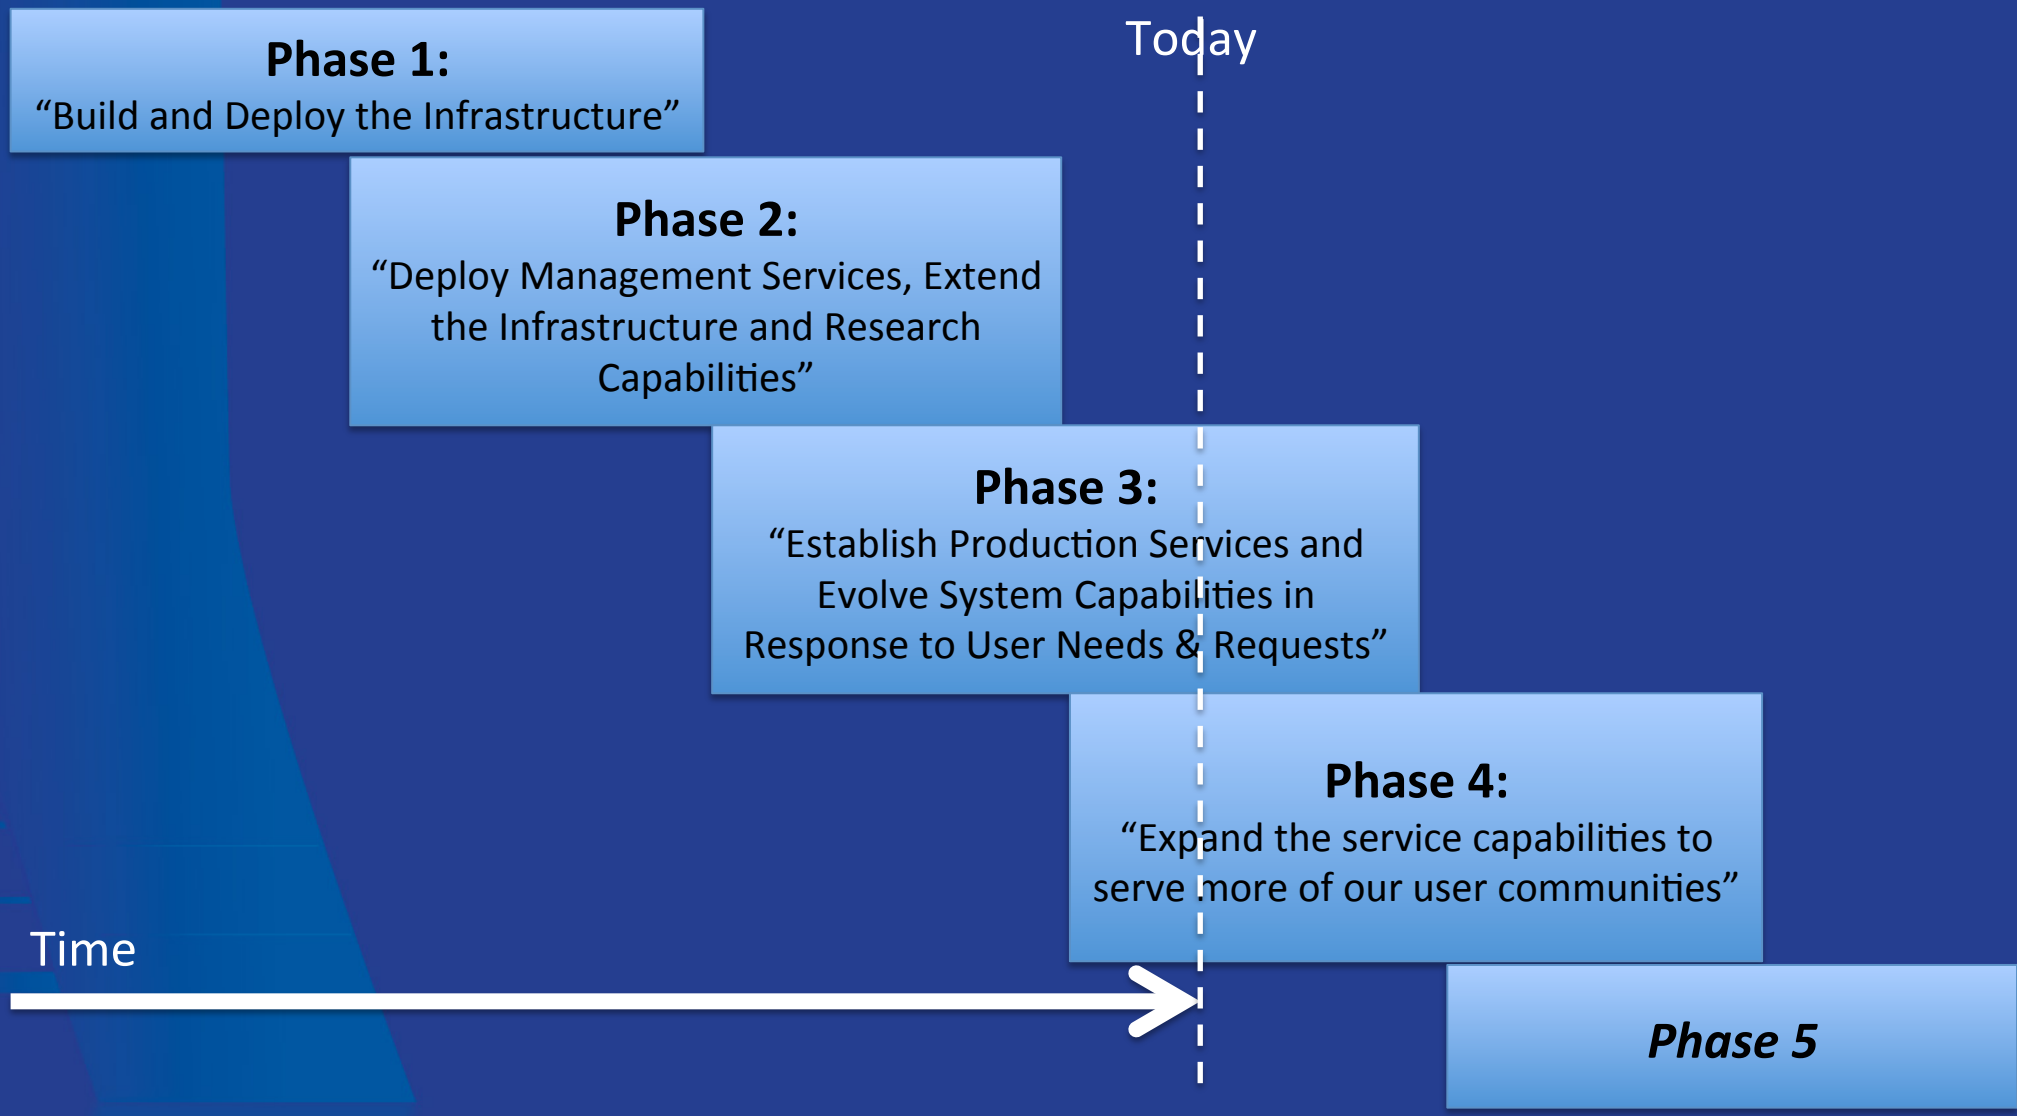

FermiCloud Phase 1: "Build and Deploy the Infrastructure"

- Specify, acquire and deploy the FermiCloud hardware,
- Establish initial FermiCloud requirements and selected the "best" open source cloud computing framework that met these requirements (OpenNebula),
- Deploy capabilities to meet the needs of the stakeholders (JDEM analysis development, Grid Developers and Integration test stands, Storage/ dCache Developers, LQCD testbed).

Completed late 2010

FermiCloud Phase 2: "Deploy Management Services, Extend the Infrastructure and Research Capabilities"

- Implement x509 based authentication (patches contributed back to OpenNebula project and are generally available in OpenNebula V3.2),
- Perform secure contextualization of virtual machines at launch,
- Perform virtualized filesystem I/O measurements,
- Develop (draft) economic model,
- Implement monitoring and accounting,
- Collaborate with KISTI personnel to demonstrate Grid and Cloud Bursting capabilities,
- Perform initial benchmarks of Virtualized MPI,
- Target "small" low-cpu-load servers such as Grid gatekeepers, forwarding nodes, small databases, monitoring, etc.,
- Begin the hardware deployment of a distributed SAN,
- Investigate automated provisioning mechanisms (puppet & cobbler).

Completed late 2012

FermiCloud Phase 3: "Establish Production Services and Evolve System Capabilities in Response to User Needs & Requests"

- Deploy highly available 24x7 production services,
  - Both infrastructure and user services.
- Deploy puppet & cobbler in production,
  - Done.
- Develop and deploy real idle machine detection,
  - Idle VM detection tool written by summer student.
- Research possibilities for a true multi-user filesystem on top of a distributed & replicated SAN,
  - CFS2 on FibreChannel SAN across FCC-3 and GCC-B.
- Live migration becomes important for this phase.
  - Manual migration has been used, Live migration is currently in test, Automatically triggered live migration yet to come.
- Formal ITIL Change Management "Go-Live",
  - Have been operating under "almost" ITIL Change Management for the past several months.

Approaching Completion

FermiCloud Phase 4: "Expand the service capabilities to serve more of our user communities"

- Complete the deployment of the true multi-user filesystem on top of a distributed & replicated SAN,
- Demonstrate interoperability and federation:
  - Accepting VM's as batch jobs,
  - Interoperation with other Fermilab virtualization infrastructures (GPCF, VMware),
  - Interoperation with KISTI cloud, Nimbus, Amazon EC2, other community and commercial clouds,
- Participate in Fermilab 100 Gb/s network testbed.
  - Have just taken delivery of 10 Gbit/second cards
- Perform more "Virtualized MPI" benchmarks and run some real world scientific MPI codes,
  - The priority of this work will depend on finding a scientific stakeholder that is interested in this capability.
- Reevaluate available open source Cloud computing stacks,
  - Including OpenStack,
  - We will also reevaluate the latest versions of Eucalyptus, Nimbus and OpenNebula.

Starting Work ~Now

FermiCloud Phase 5

- Cover the program of work in the (draft) Fermilab-KISTI FermiCloud Cooperative Research and Development Agreement (CRADA),
- We will likely start to move our production FermiGrid services deployment under the FermiCloud umbrella,
- This phase will likely involve work to enable the Fermilab CDF and D0 experiments Run II Data Preservation environment,
- This phase will also incorporate any work or changes that arise out of the Fermilab Scientific Computing Division strategy on Clouds and Virtualization,
- Additional requests and specifications are also being gathered!

Specifications Under Development

Deming Continual Improvement Cycle

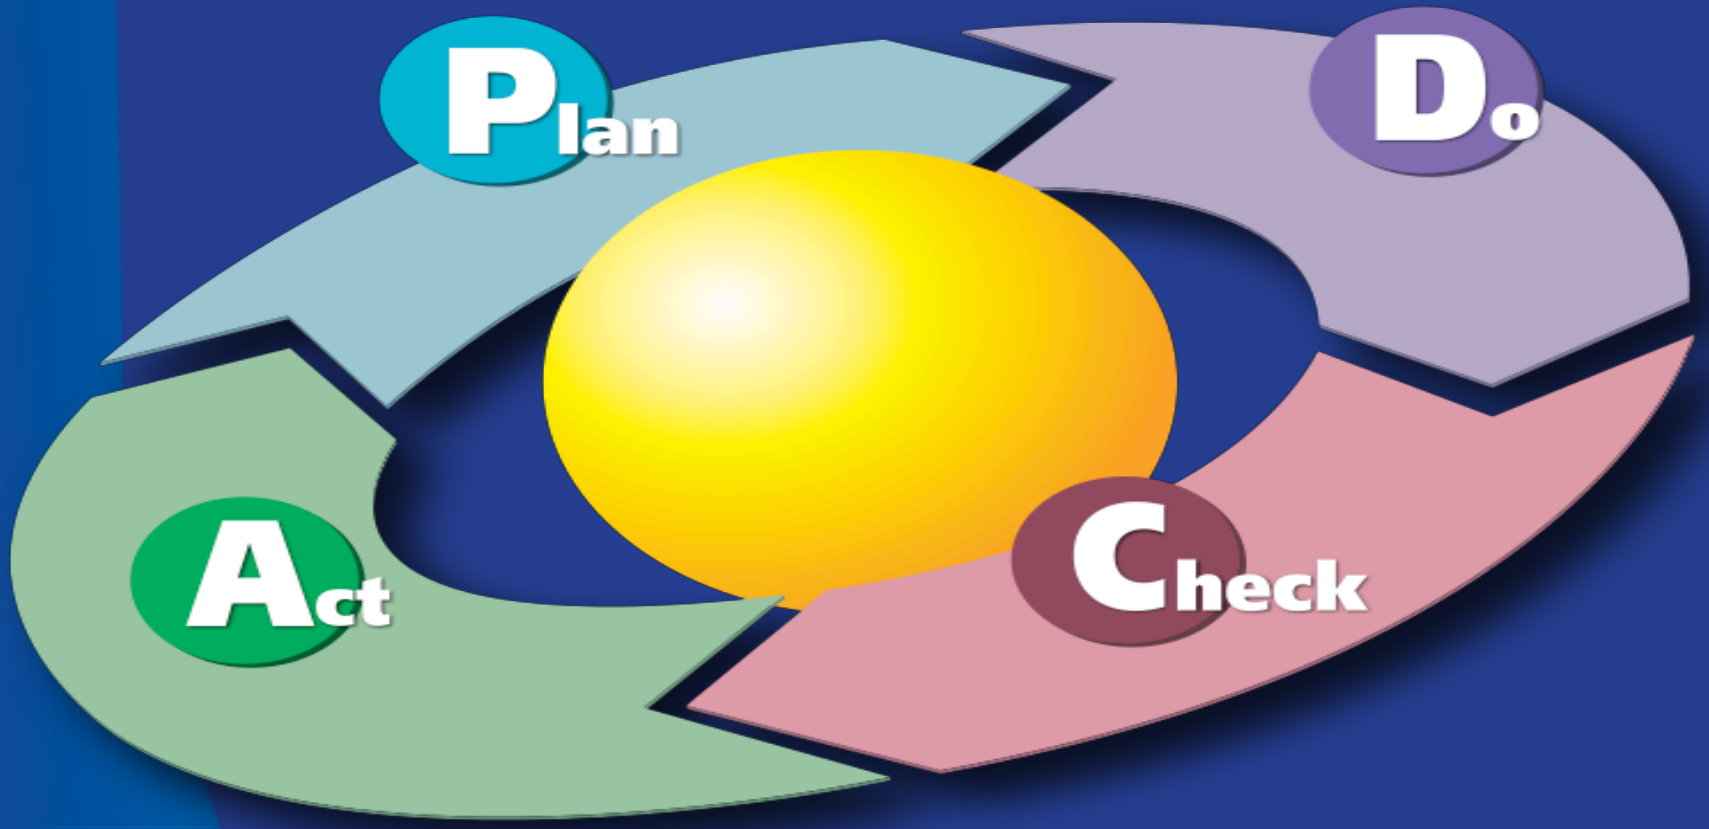

Plan, Do, Check, Act
Source - http://en.wikipedia.org/wiki/PDCA

Lessons from FermiGrid => FermiCloud must be Fault Tolerant

As we have learned from FermiGrid, having a distributed fault tolerant infrastructure is highly desirable for production operations.

Why:

- Users benefit from increased uptime,
- Service operators benefit from resiliency and gain flexibility to schedule routine maintenance.

How:

- Distribute physical systems across multiple buildings and computer rooms,
- Distributed and redundant "head nodes",
- Distributed "cloud services" across locations,
- Distributed network infrastructure,
- Heartbeat across public and private LAN to manage services,
- Distributed shared file system on SAN.

Goal:

- If a building is "lost", then automatically relaunch "24x7" VMs on surviving infrastructure, then relaunch "9x5" VMs if there is sufficient remaining capacity,
- Perform notification (via Service-Now) when exceptions are detected.

Some Recent Facility and/or Network Outages

FCC main breaker 4x (February-November 2010)

FCC-1 network outage 2x (Spring 2011)

GCC-B Load shed events (June-August 2011)

- This accelerated planned move of nodes to FCC-3.

GCC load shed events and maintenance (July 2012).

- FCC-3 cloud was ready just in time to keep server VM's up.

FCC-2 outage (October 2012)

- FermiCloud wasn't affected, our VM's stayed up.

Service Outages on Commercial Clouds

Amazon has had several significant service outages over the past few years:

- Outage in April 2011 of storage services resulted in actual data loss,
- An electrical storm that swept across the East Coast late Friday 29-Jun-20-12 knocked out power at a Virginia data center run by Amazon Web Services.
- An outage of one of Amazon's cloud computing data centers knocked out popular sites like Reddit, Foursquare Pinterest and TMZ on Monday 22-Oct-2012,
- Amazon outage affecting Netflix operations over Christmas 2012 and New Years 2013,
- Outage on Thursday 31-Jan-2013.

Microsoft Azure:

- Leap day bug on 29-Feb-2012.

Google:

- Outage on 26-Oct-2012.

FCC and GCC

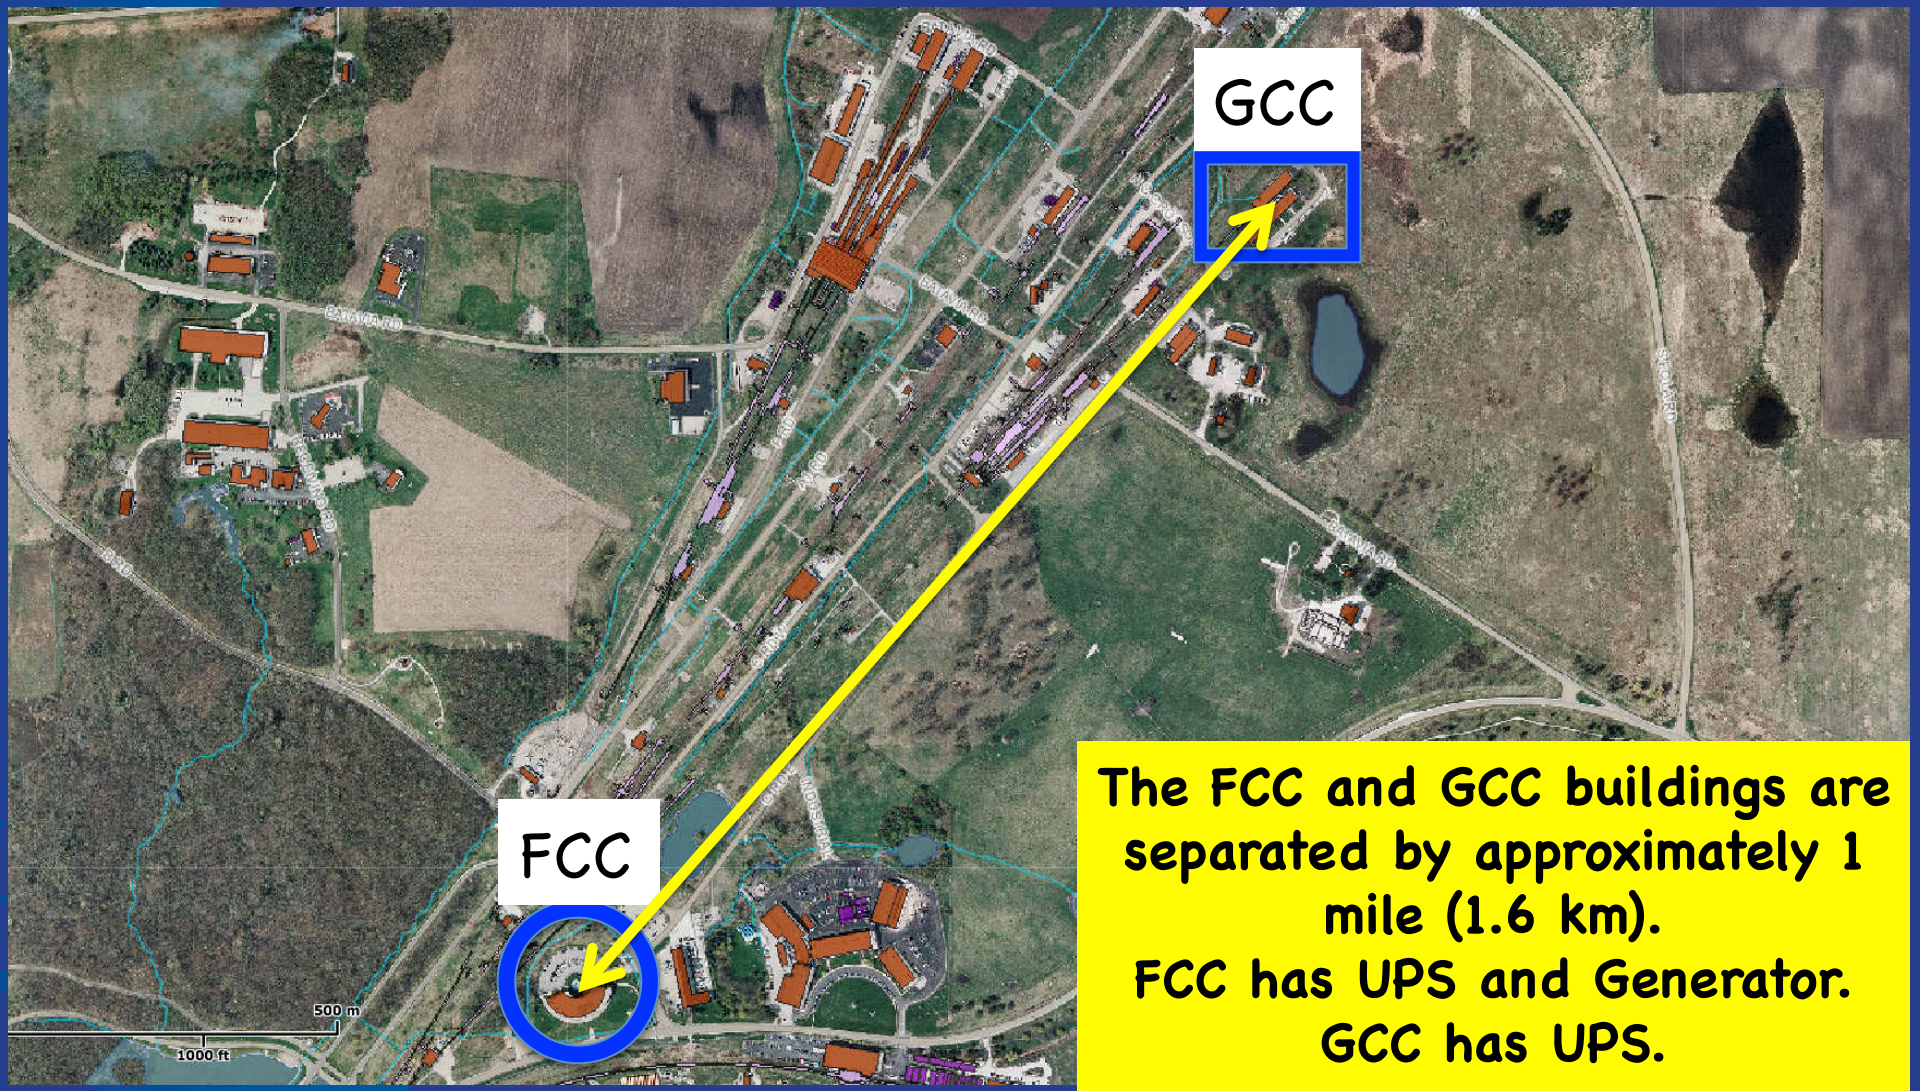

The FCC and GCC buildings are separated by approximately 1 mile (1.6 km).
FCC has UPS and Generator.
GCC has UPS.

Distributed Network Core Provides Redundant Connectivity

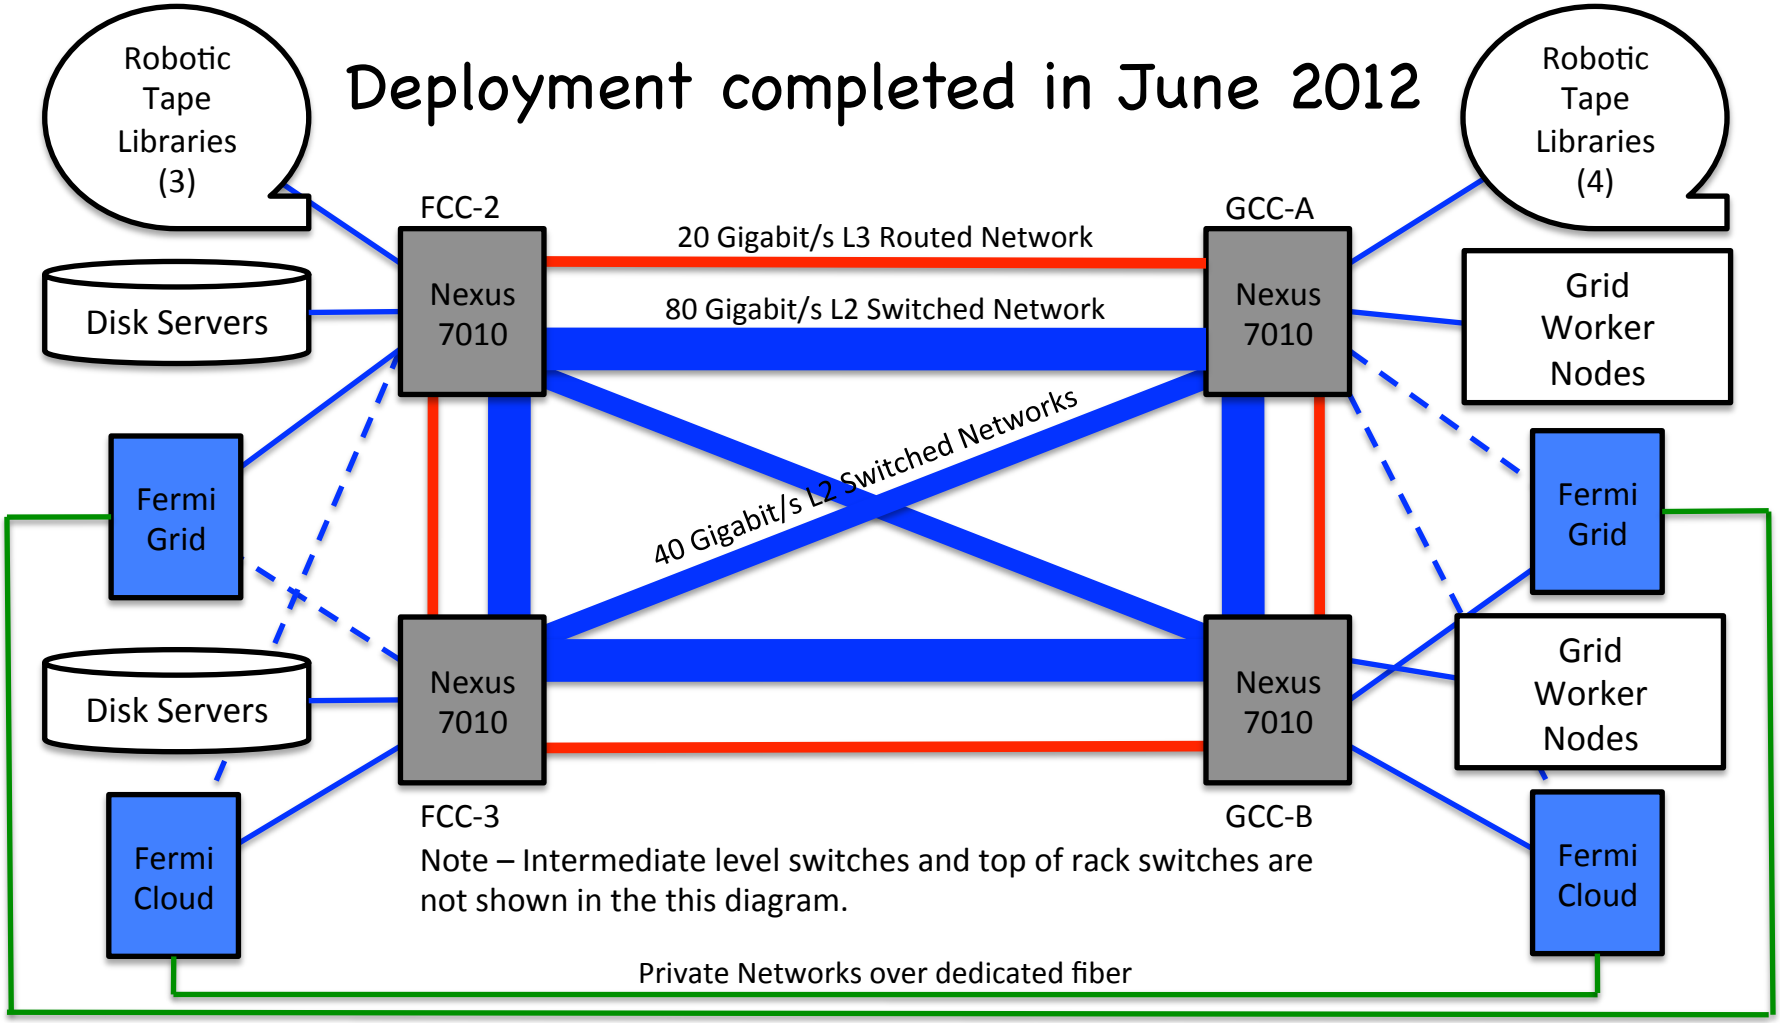

Distributed Shared File System Design

Dual-port FibreChannel HBA in each node,

Two Brocade SAN switches per rack,

Brocades linked rack-to-rack with dark fiber,

60TB Nexsan Satabeast in FCC-3 and GCC-B,

Redhat Clustering + CLVM + GFS2 used for file system,

Each VM image is a file in the GFS2 file system,

Next step—use LVM mirroring to do RAID 1 across buildings. (Documented feature in manual).

FermiCloud – Network & SAN "Today"

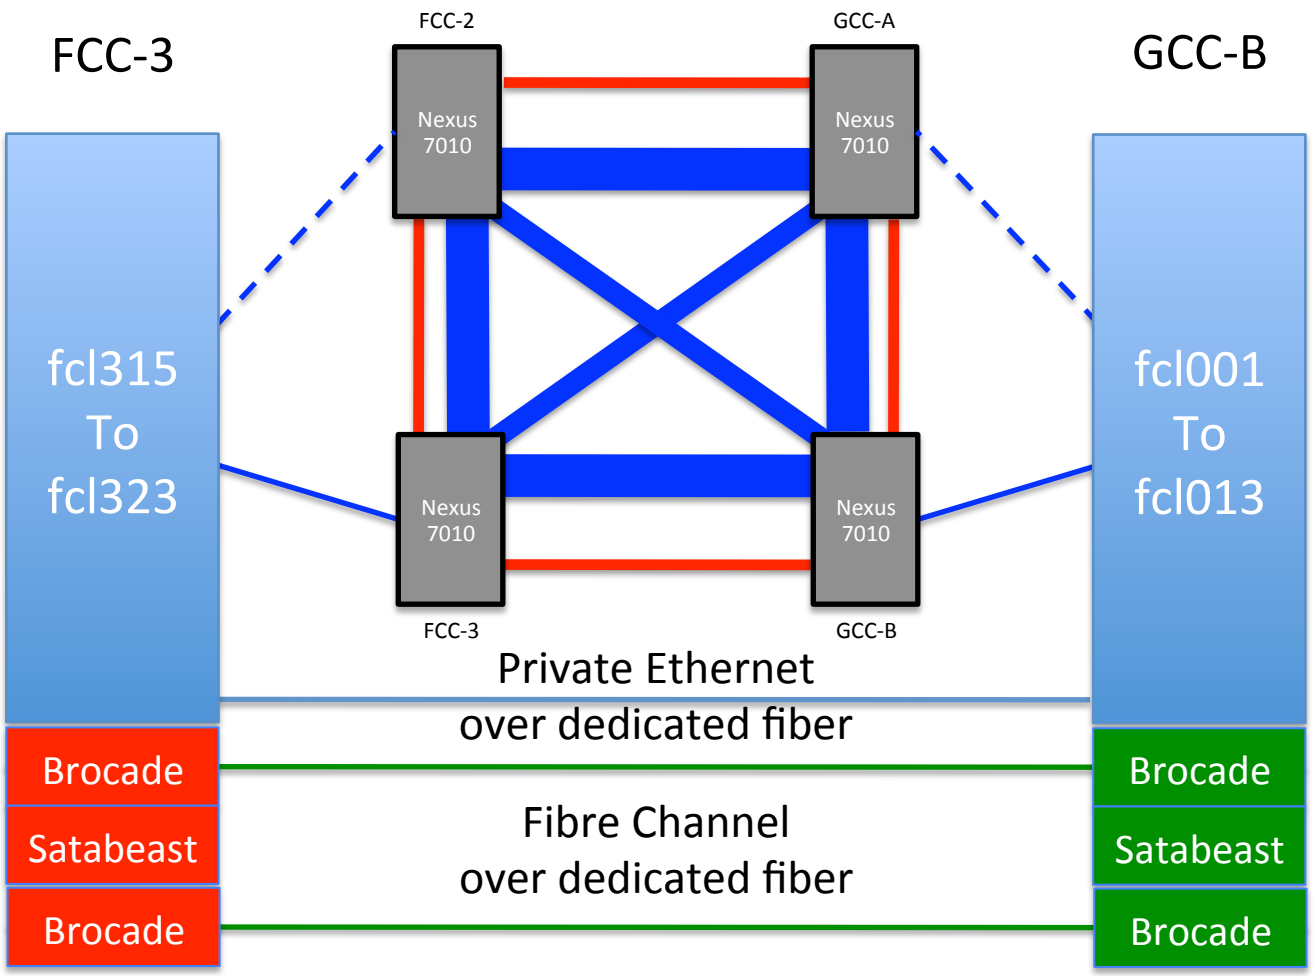

FY2011 / FY2012

FermiCloud – Network & SAN (Possible Future – FY2013/2014)

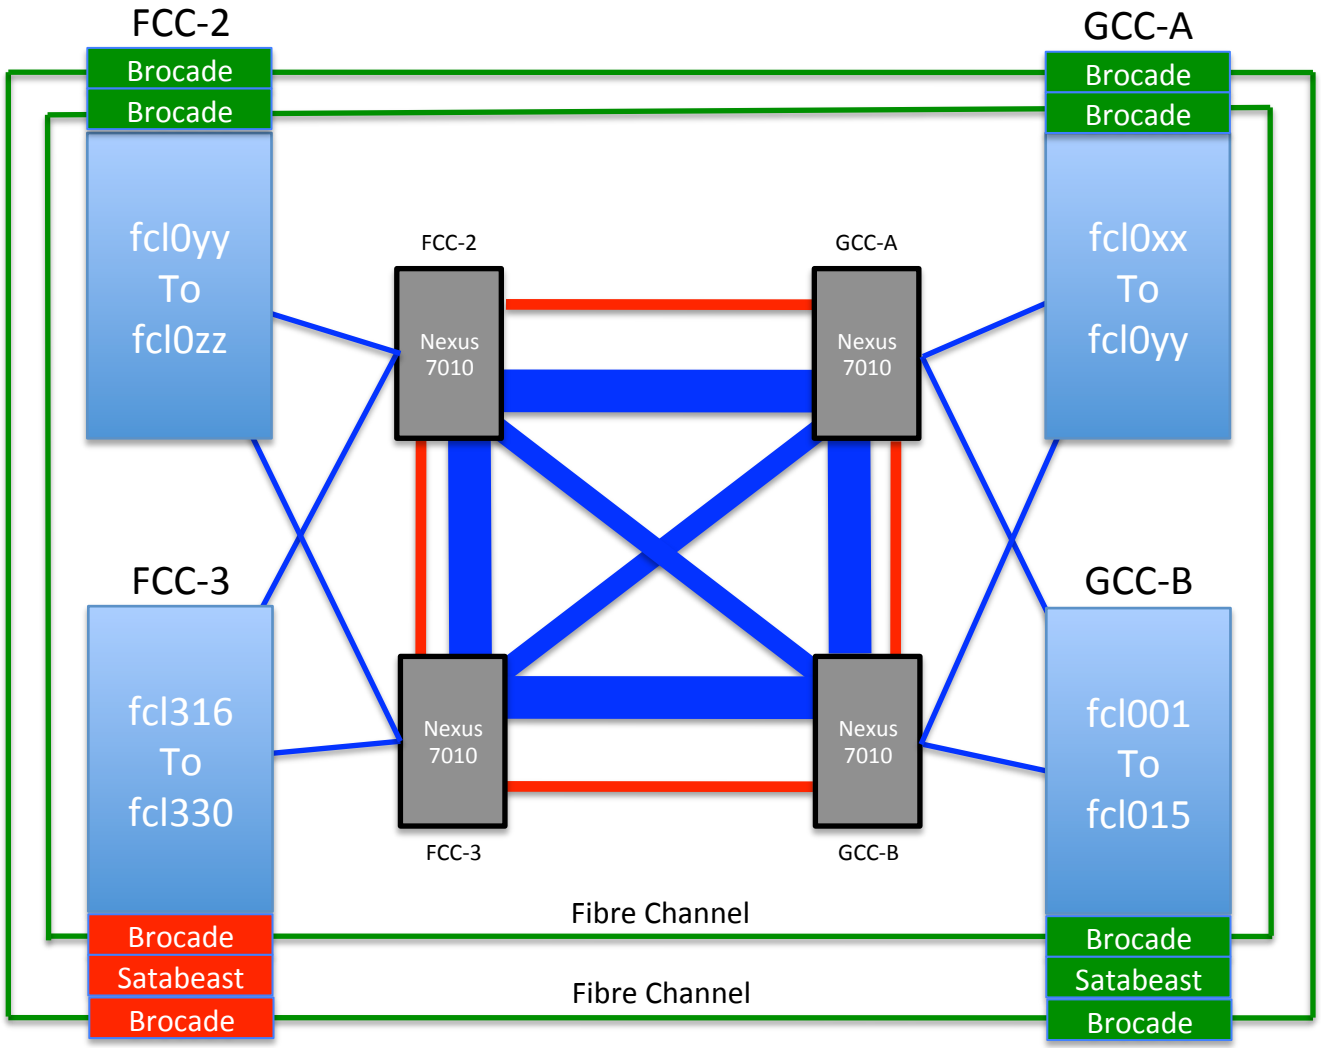

Distributed Shared File System Benefits

Fast Launch—almost immediate as compared to 3-4 minutes with ssh/scp,

Live Migration—Can move virtual machines from one host to another for scheduled maintenance, transparent to users,

Persistent data volumes—can move quickly with machines,

Once mirrored volume in place—can relaunch virtual machines in surviving building in case of building failure/outage,

Head nodes have independent shared file system based on GFS2/DRBD active/active.

FermiCloud-HA Head Node Configuration

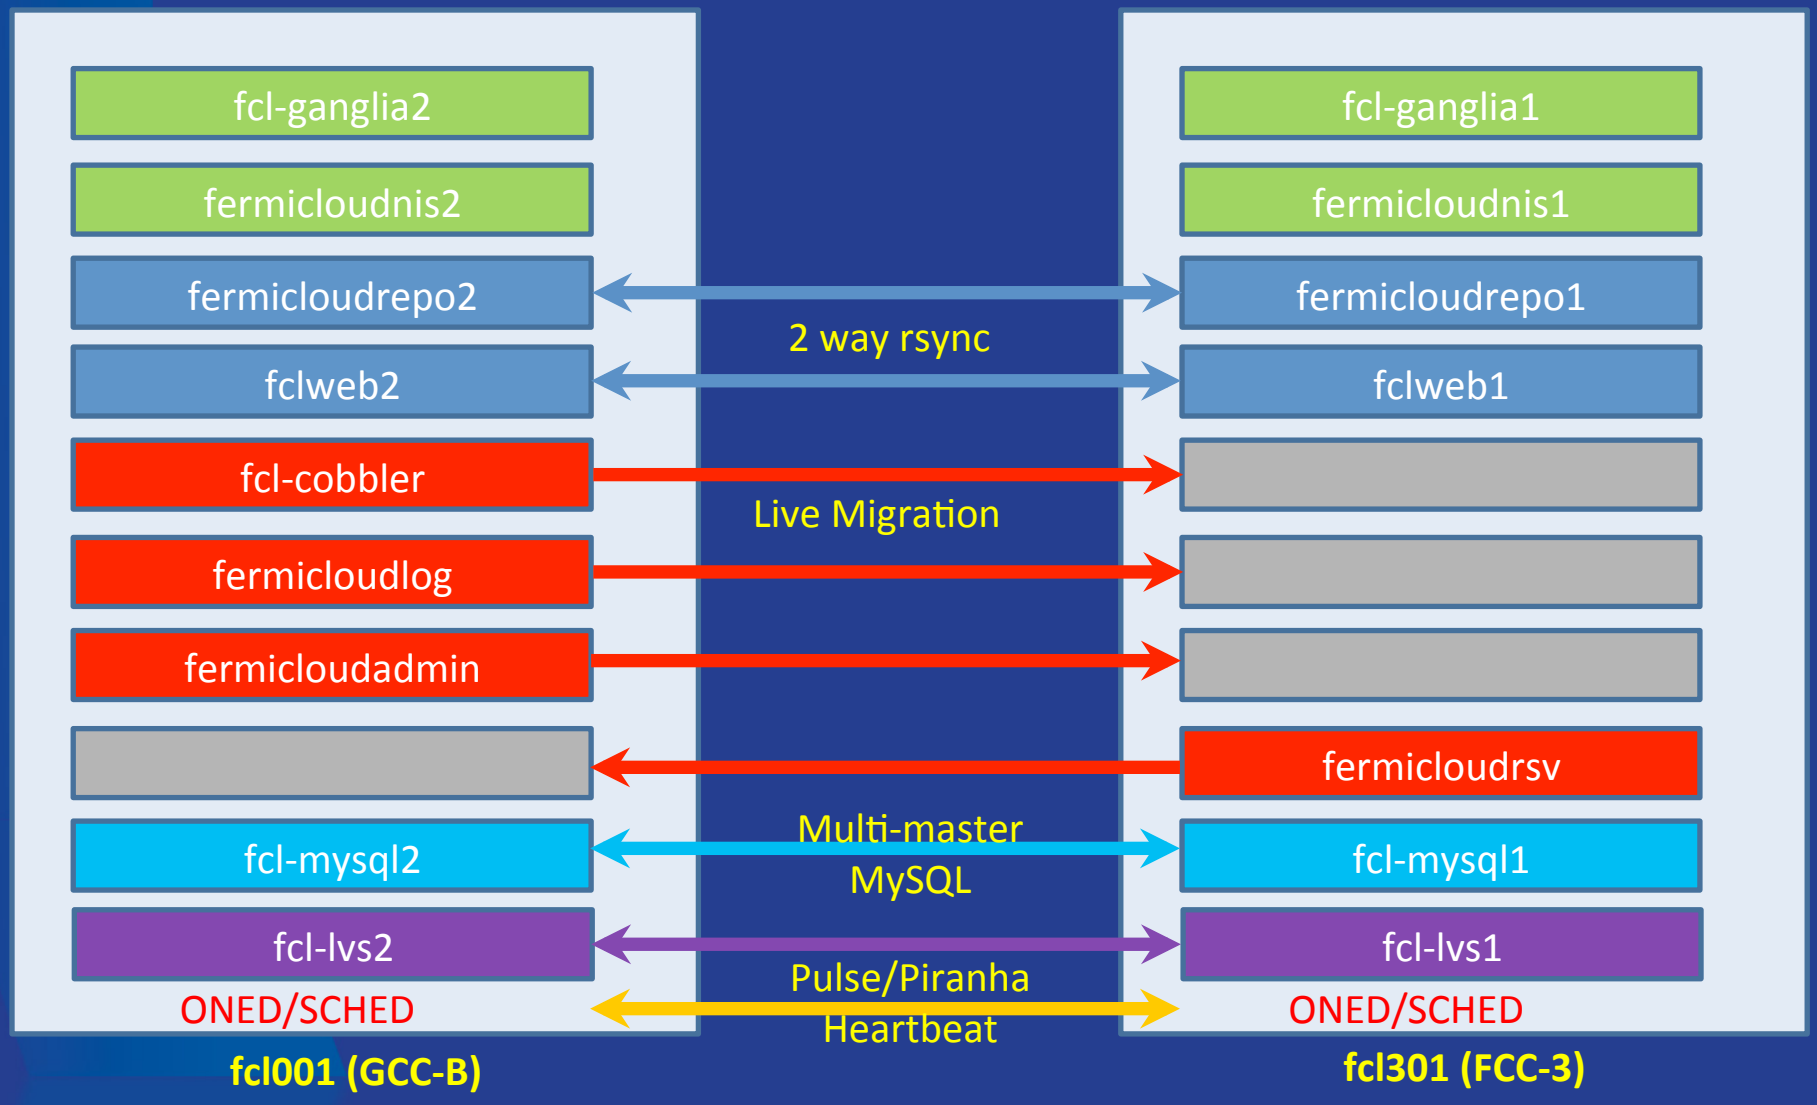

Current FermiCloud Capabilities

Public network access via the high performance Fermilab network,
- This is a distributed, redundant network.

Private 1 Gb/sec network,
- This network is bridged across FCC and GCC on private fiber,

High performance Infiniband network to support HPC based calculations,
- Performance under virtualization is ~equivalent to "bare metal",
- Currently split into two segments,
- Segments could be bridged via Mellanox MetroX.

Access to a high performance FibreChannel based SAN,
- This SAN spans both buildings,

Access to the high performance BlueArc based filesystems,
- The BlueArc is located on FCC-2,

Access to the Fermilab dCache and enStore services,
- These services are split across FCC and GCC,

Access to 100 Gbit Ethernet test bed in LCC (via FermiCloud integration nodes),
- Intel 10 Gbit Ethernet converged network adapter X540-T1.

Typical Use Cases

Public net virtual machine:

- On Fermilab Network open to Internet,
- Can access dCache and Bluearc Mass Storage,
- Shared Home directories between multiple VM's.

Public/Private Cluster:

- One gateway VM on public/private net,
- Cluster of many VM's on private net.

Storage VM:

- VM with large non-persistent storage,
- Use for large MySQL or Postgres databases, Lustre/ Hadoop/Bestman/xRootd servers.

FermiCloud Draft Economic Model

SLA	24x7	9x5	Opportunistic
"Unit" (HT CPU + 2 GB)	$125	$45	$25
Add'l HT core	$125	$125	$125
Add'l memory per GB	$30	$30	$30
Add'l local disk per TB	$40	$40	$40
SAN disk per TB	$475	$475	$475
BlueArc per TB	$430	$430	$430
System Administrator	$750	$750	$750
Specialized Service Support	"Market"	"Market"	"Market"

Note - Costs in the above chart are USD per year

Cloud Computing Environment

FermiCloud Security taskforce recommended to Fermilab Computer Security Board and the Computer Security Team that a new Cloud Computing Environment be established,

- This is currently under consideration.

Normal FermiCloud use is authenticated by Fermi MIT Kerberos credentials, or credentials (x.509) derived from Fermi MIT Kerberos.

Special concerns with Cloud:

- Users have root,
- Usage can be a combination of Grid usage (Open Science Environment) and Interactive usage (General Computing Environment),
- Planning for "secure cloud" to handle expected use cases: Archival systems at old patch levels or legacy OS, Data and code preservation systems, Non-baselined OS (Ubuntu, Centos, SUSE), Non-Kerberos services which can live only on private net,
- Cloud can incorporate non-Fermilab systems (Amazon, Microsoft Azure, etc.)

True Idle VM Detection

In times of resource need, we want the ability to suspend or "shelve" idle VMs in order to free up resources for higher priority usage.

- This is especially important in the event of constrained resources (e.g. during building or network failure).

Shelving of "9x5" and "opportunistic" VMs allows us to use FermiCloud resources for Grid worker node VMs during nights and weekends

- This is part of the draft economic model.

Giovanni Franzini (an Italian co-op student) has written (extensible) code for an "Idle VM Probe" that can be used to detect idle virtual machines based on CPU, disk I/O and network I/O.

- This is the biggest pure coding task left in the FermiCloud project.

Idle VM Information Flow

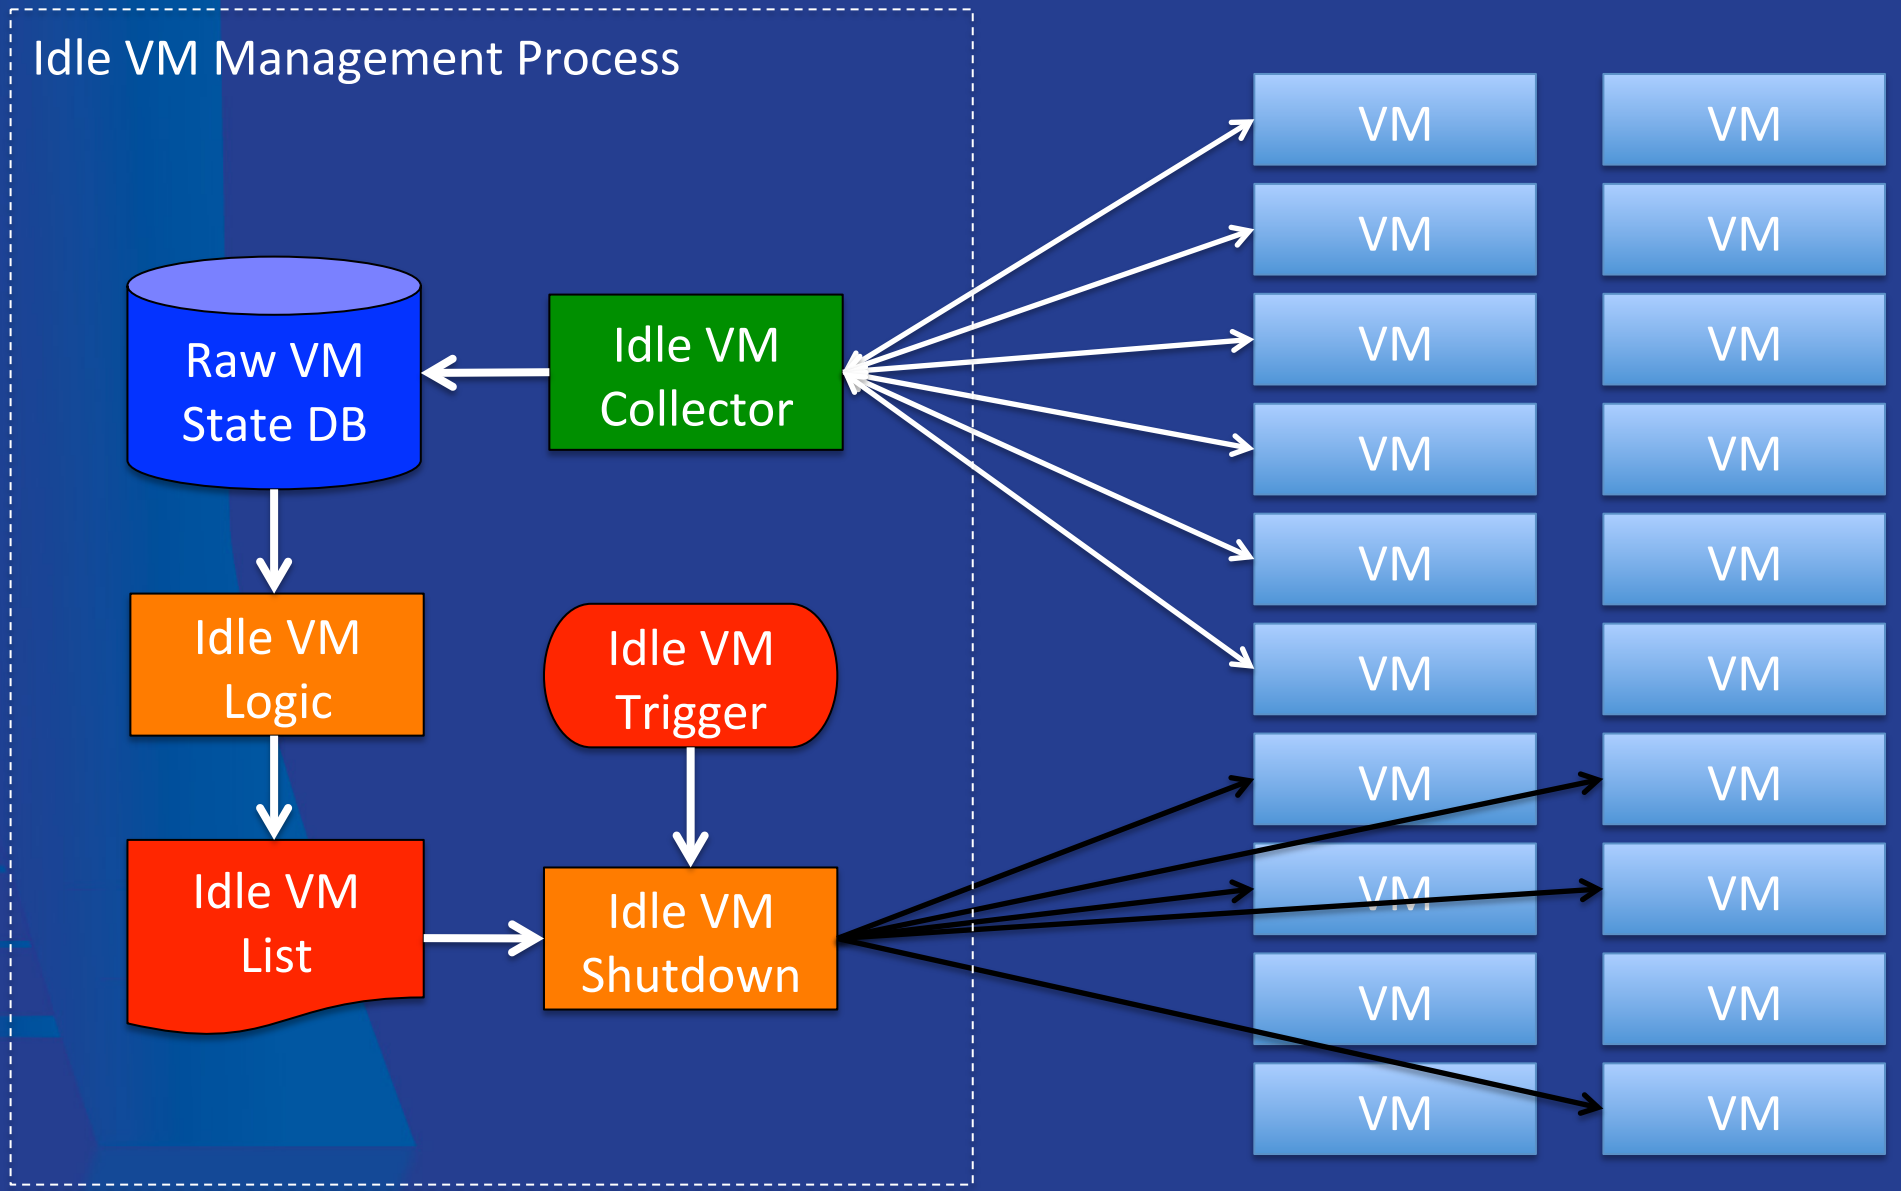

Virtual Infrastructure Automation and Provisioning

Demonstrate Interoperability and Federation:

- Accept Virtual Machines as Batch Jobs via cloud API's (such as native OpenNebula, Amazon EC-2 or OCCI),
- Test interoperability against other public (Amazon, Azure, etc.) and private (KISTI G-cloud, etc.) cloud computing stacks,
- Investigate interoperability with European efforts - EGI Cloud Federation Task Force (OpenNebula, OpenStack, StratusLab) - HelixNebula Federation of commercial cloud providers.

Grid and Cloud Bursting,

- Launch pre-defined user virtual machines based on workload demand,
- Use GlideinWMS to run Monte-Carlo production on public and private clouds.

Seo-Young Noh, KISTI visitor @ FNAL, showed proof-of-principle of "vCluster" in summer 2011:

- Look ahead at Condor batch queue,
- Submit worker node virtual machines of various VO's to FermiCloud or Amazon EC2 based on user demand,
- Machines join grid cluster and run grid jobs from the matching VO.

vCluster at SC2012

FermiCloud Usage

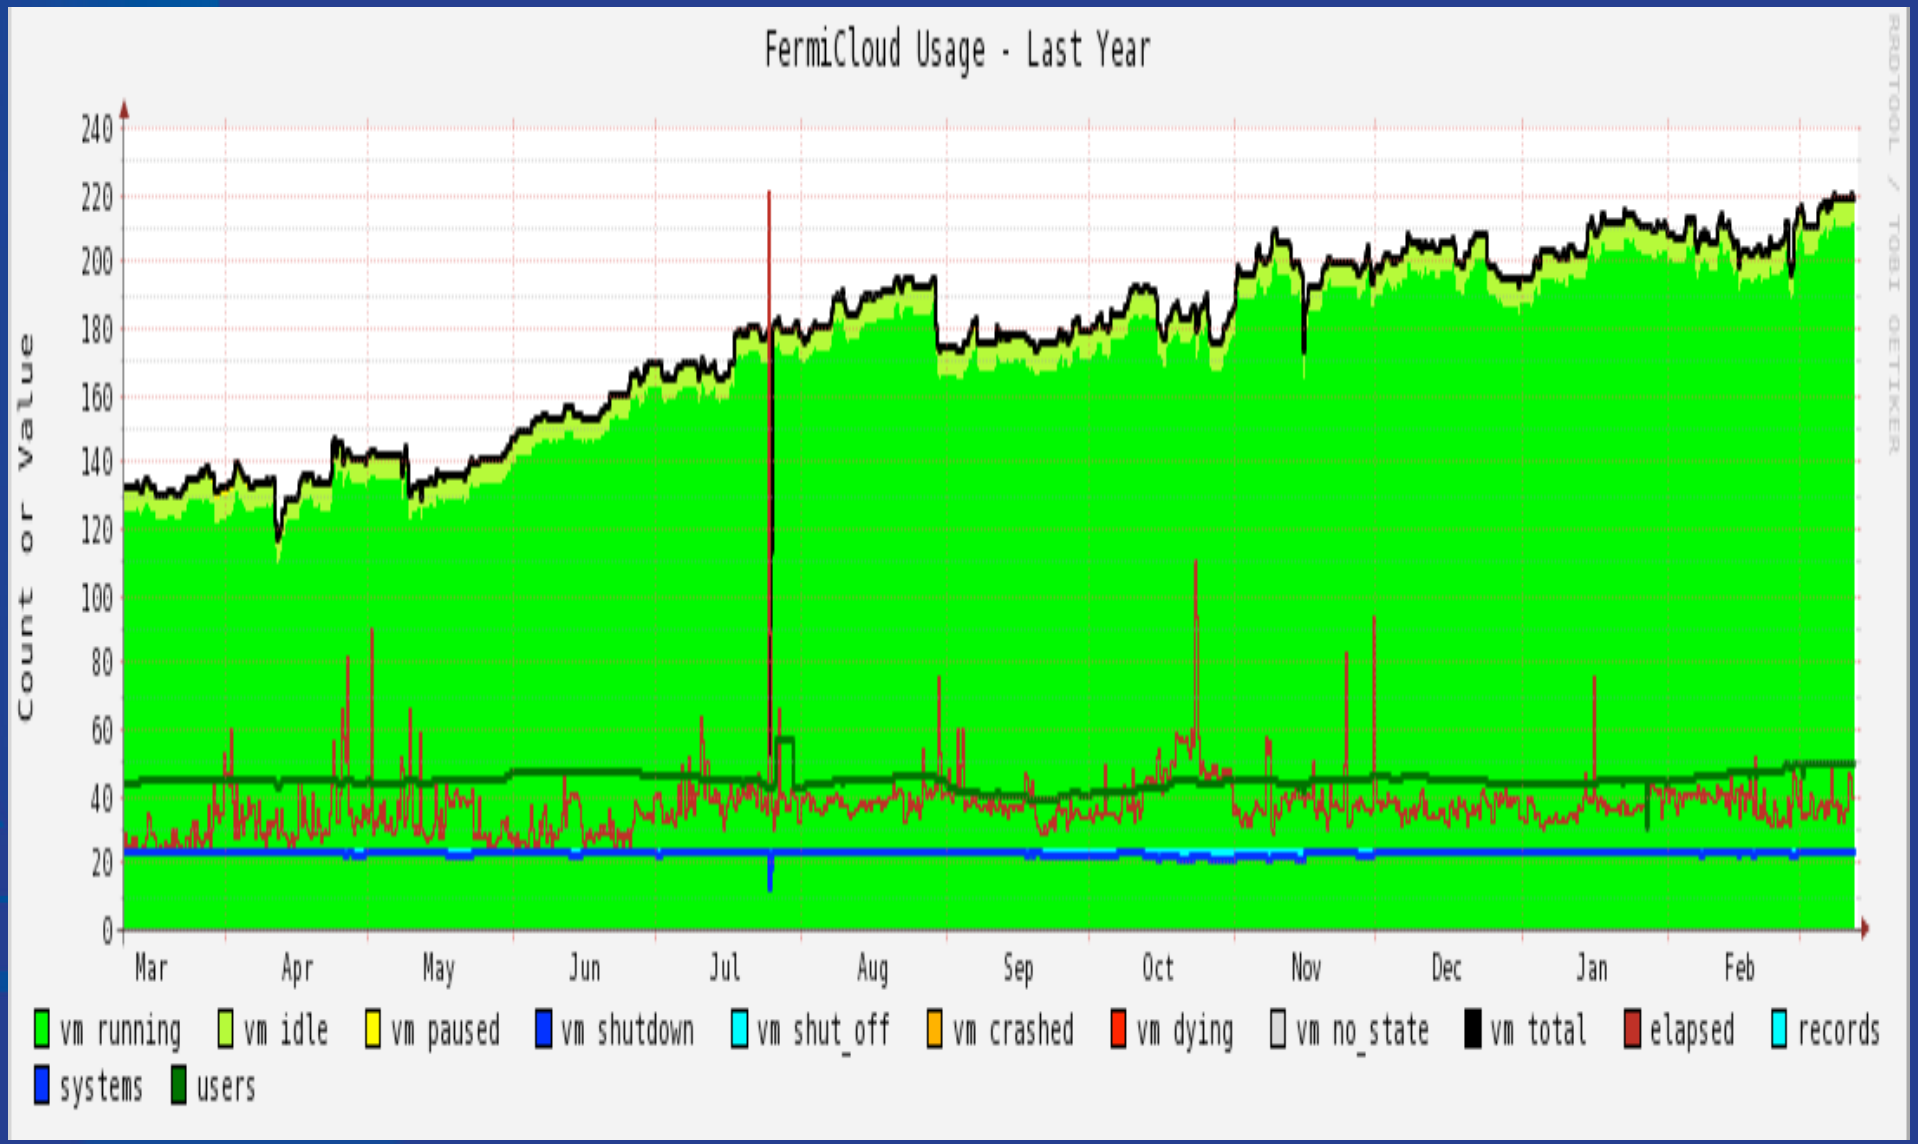

FermiCloud – FTE Months/Year

Fiscal Year	Management	Operations	Project + Development	TOTAL
2010	1.79	4.76	0.02	6.57
2011	2.36	15.69	6.07	24.12
2012	1.56	9.86	3.48	14.90
2013 to Date	0.82	7.03	1.75	9.60
Average	1.87	10.67	3.23	15.77

Note - Fiscal Year starts the 1st of October

FermiCloud FTE Effort Plot

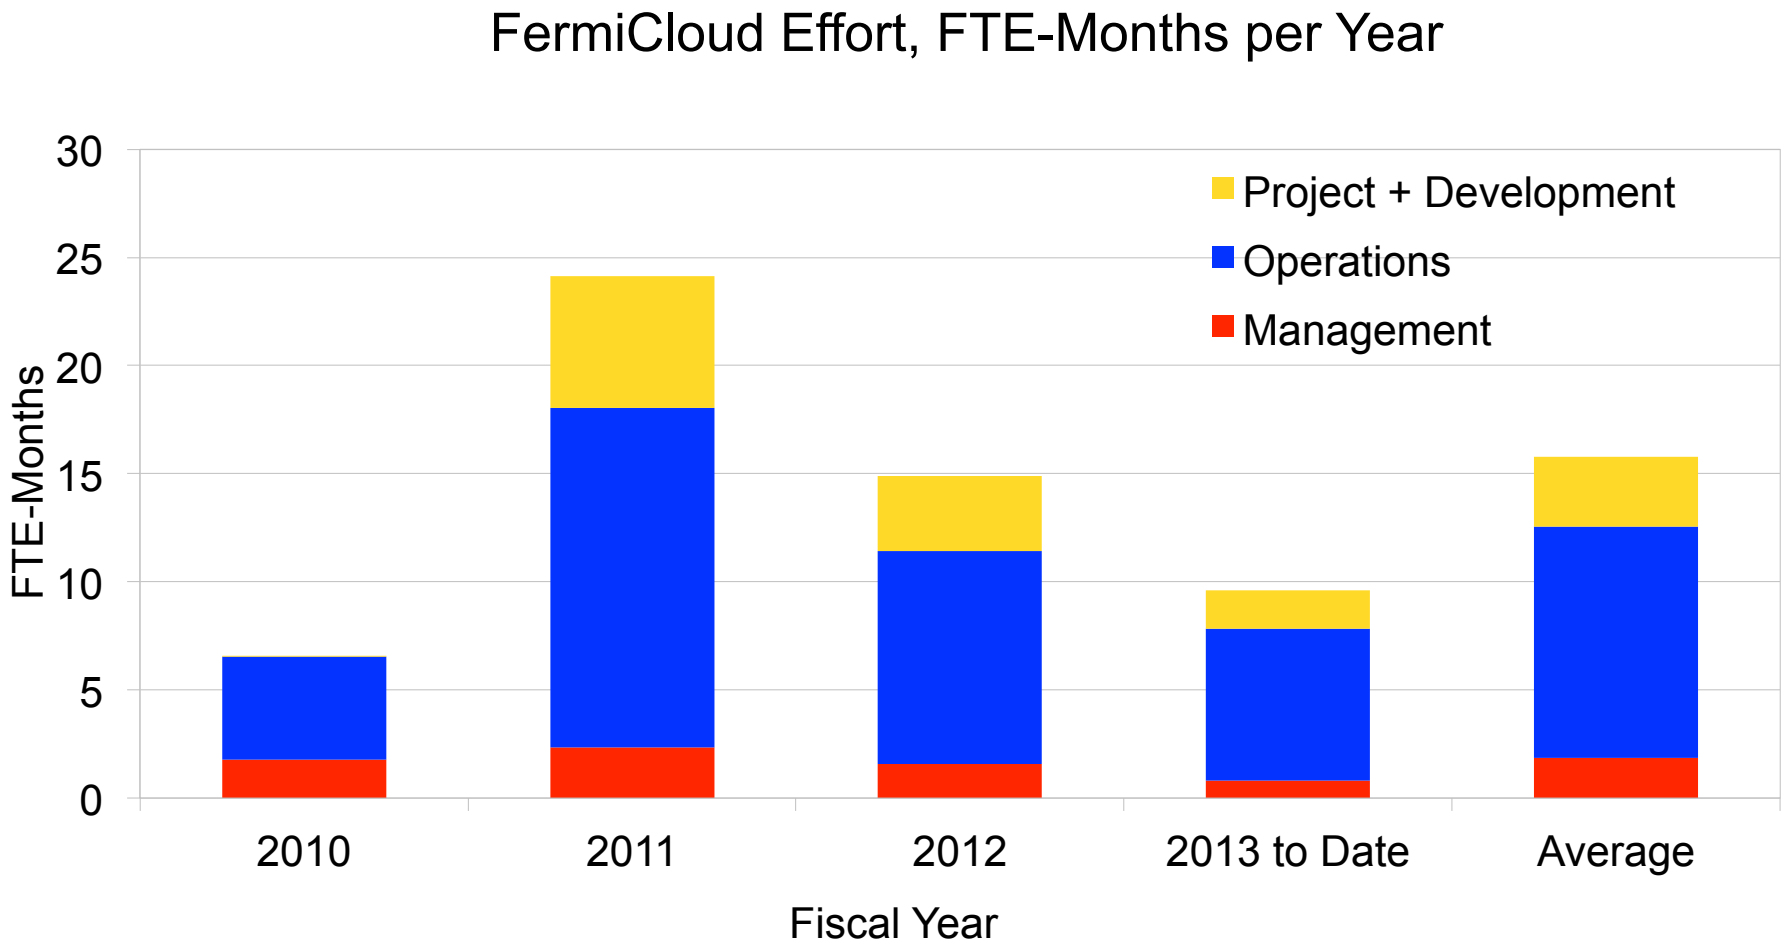

FermiCloud Review

The FermiCloud Project recently underwent a management review that included both internal and external (to Fermilab) reviewers.

The review documents are available at:

- https://indico.fnal.gov/conferenceDisplay.py?confId=6282

Results of the review were very positive:

- Ready for production operations,
- Lots more work to do (see previous slides on project phases),
- Many items for senior line management to work on…

Current Stakeholders

Grid & Cloud Computing Personnel,

Run II - CDF & D0,

Intensity Frontier Experiments,

Cosmic Frontier (LSST),

Korea Institute of Science & Technology Information (KISTI),

Open Science Grid (OSG).

Long Term Vision

[another look at some work in the FermiCloud Project Phases]

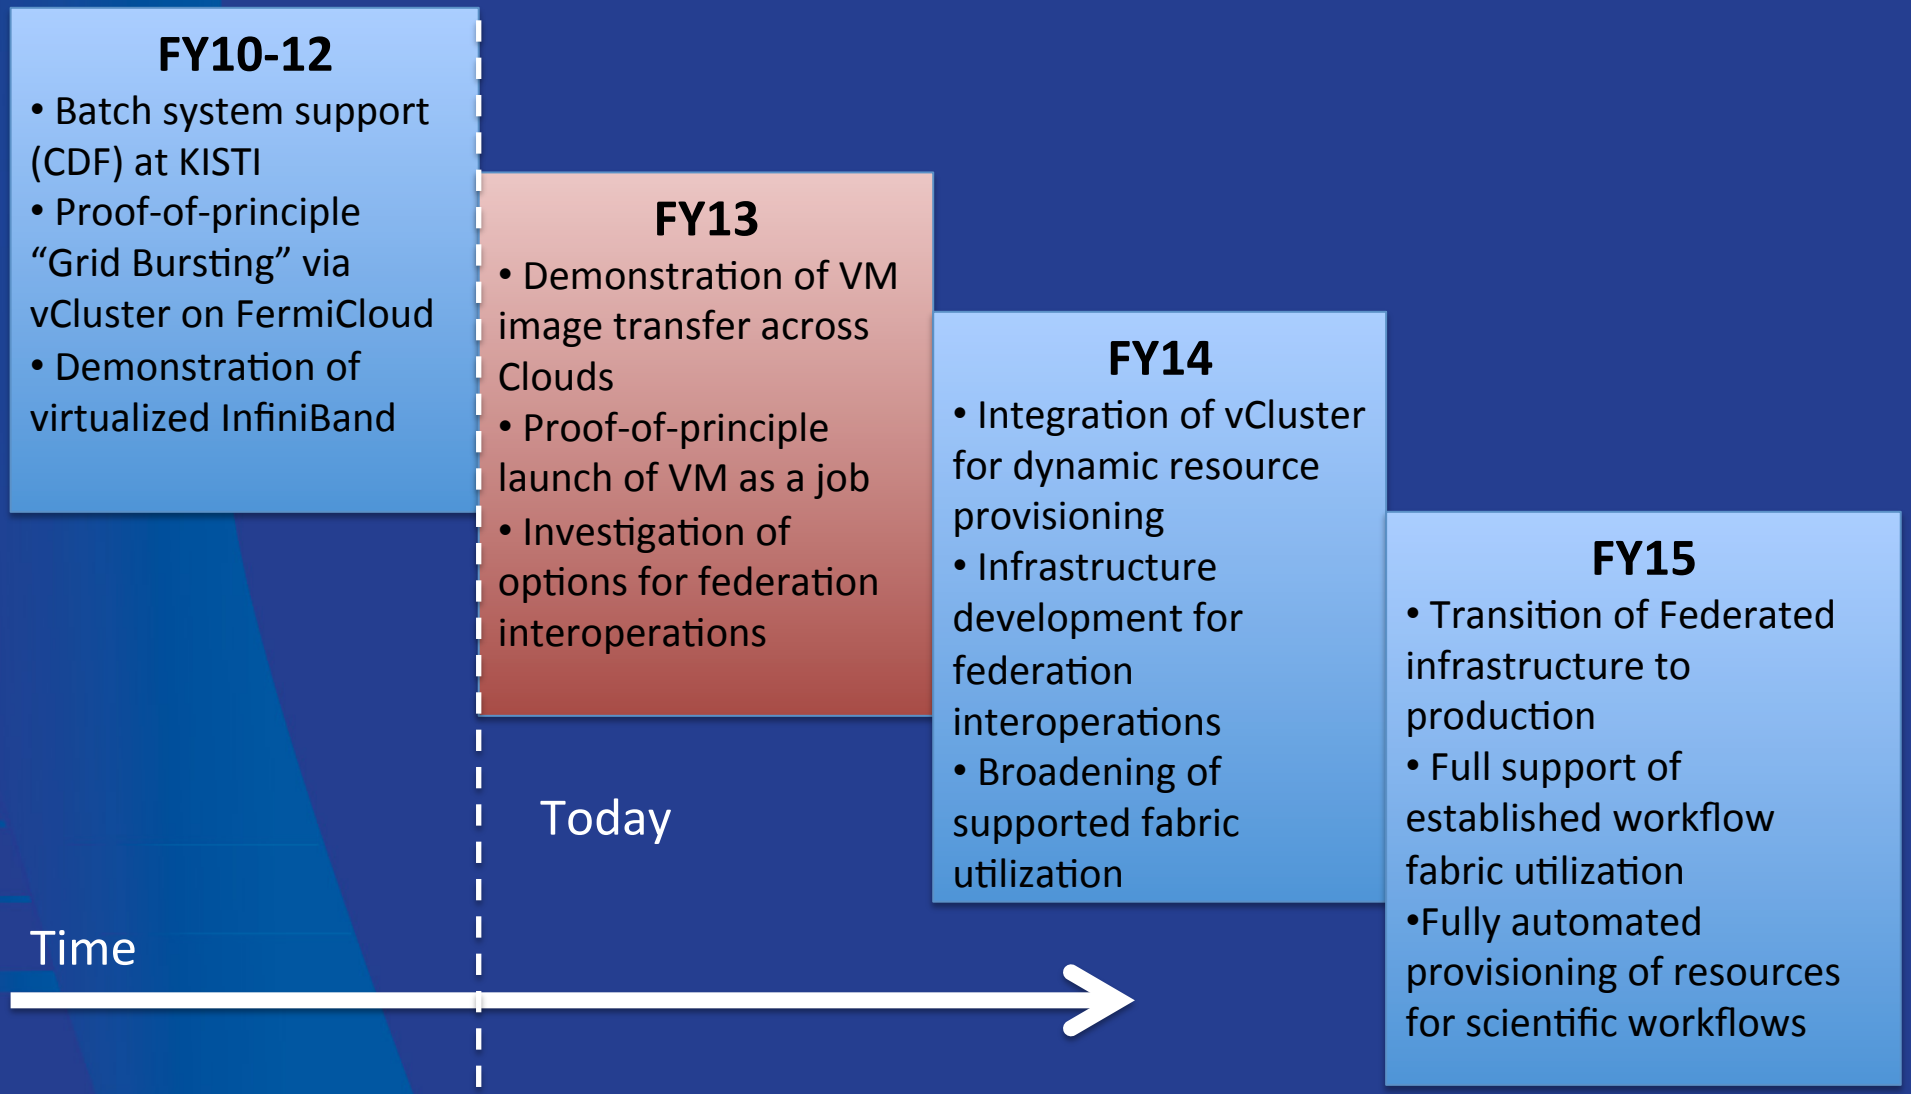

FermiCloud Summary - 1

Science is directly and indirectly benefiting from FermiCloud:

- CDF, D0, Intensity Frontier, Cosmic Frontier, CMS, ATLAS, Open Science Grid,...

FermiCloud operates at the forefront of delivering cloud computing capabilities to support scientific research:

- By starting small, developing a list of requirements, building on existing Grid knowledge and infrastructure to address those requirements, FermiCloud has managed to deliver a production class Infrastructure as a Service cloud computing capability that supports science at Fermilab.
- FermiCloud has provided FermiGrid with an infrastructure that has allowed us to test Grid middleware at production scale prior to deployment.
- The Open Science Grid software team used FermiCloud resources to support their RPM "refactoring" and is currently using it to support their ongoing middleware development/integration.

FermiCloud Summary - 2

The FermiCloud collaboration with KISTI has leveraged the resources and expertise of both institutions to achieve significant benefits.

vCluster has demonstrated proof of principal "Grid Bursting" using FermiCloud and Amazon EC2 resources.

Using SRIOV drivers on FermiCloud virtual machines, MPI performance has been demonstrated to be >96% of the native "bare metal" performance.

FermiCloud Project Summary - 3

FermiCloud personnel are currently working on:

- Finalizing a CRADA with KISTI for future collaboration,
- Completion of the Phase 3 deliverables, including a SAN storage deployment that will offer a true multi-user filesystem on top of a distributed & replicated SAN,
- Getting ready to start work on the Phase 4 deliverables,
- Collaborating on the development of the Phase 5 specifications.

The future is mostly cloudy.

Thank You

Any Questions?