Business Continuity Efforts at Fermilab

Keith Chadwick
Fermilab

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

Outline

- Introduction
- Site Power
- Computing Facilities
- Networking
- Fermilab-Argonne DR Agreement
- DNS & Authentication,
- Data Distribution,
- FermiGrid and FermiCloud,
- Outsourced Services.

Introduction

ITIL defines several Service Management processes that are formally involved in "Business Continuity":

- Availability Management
- Continuity Management
- Capacity Management

The work that I will be discussing in this talk largely predates the adoption of ITIL practices by the Fermilab Computing Sector.

Many people at Fermilab have been involved in the development of these capabilities and plans,

- Please credit them with all the hard work and accomplishments,
- Please blame me for any errors or misunderstandings.

Site Power

- Two substations on the Commonwealth Edison 345 kV electrical grid "backbone":
  - Master substation,
  - Main injector substation.

- Auxiliary 13.6 kV service to Fermilab:
  - Kautz road substation.

Fermilab Computing Facilities

Feynman Computing Center (FCC):

- FCC-2 Computer Room,
- FCC-3 Computer Room(s),
- All rooms have refrigerant based cooling, UPS and Generators.

Grid Computing Center (GCC):

- GCC-A, GCC-B, GCC-C Computer Rooms,
- GCC-TRR Tape Robot Room,
- GCC-Network-A, GCC-Network-B,
- All rooms have refrigerant based cooling, UPS and "quick connect" taps for Generators.

Lattice Computing Center (LCC):

- LCC-107 & LCC-108 Computer Room,
- Refrigerant based cooling,
- No UPS or Generator.

GCC Load Shed Plans

As was mentioned previously, Fermilab has encountered a cooling issue with the GCC-B and GCC-C computer rooms,

- There is work underway to remove the berm that obscures the heat exchangers,

In the event that the berm removal does not completely address the issue(s), we have a three step "load shed" plan agreed with our stakeholders, whole racks have been identified that will participate:

- Phase 1 - 30% load shed racks,
- Phase 2 - 50% load shed racks (strict superset of 30% load shed racks),
- 100% load shed.

If the "load shed" trigger conditions are met:

- During business hours, then the expectation is that the agreed racks will be shut down under program control in 30 minutes,
- During non-business hours, the shut down interval may be longer,
- If the racks are not shut down in the agreed upon time window, then the facilities personnel may elect to throw the corresponding breakers in the electrical distribution system.

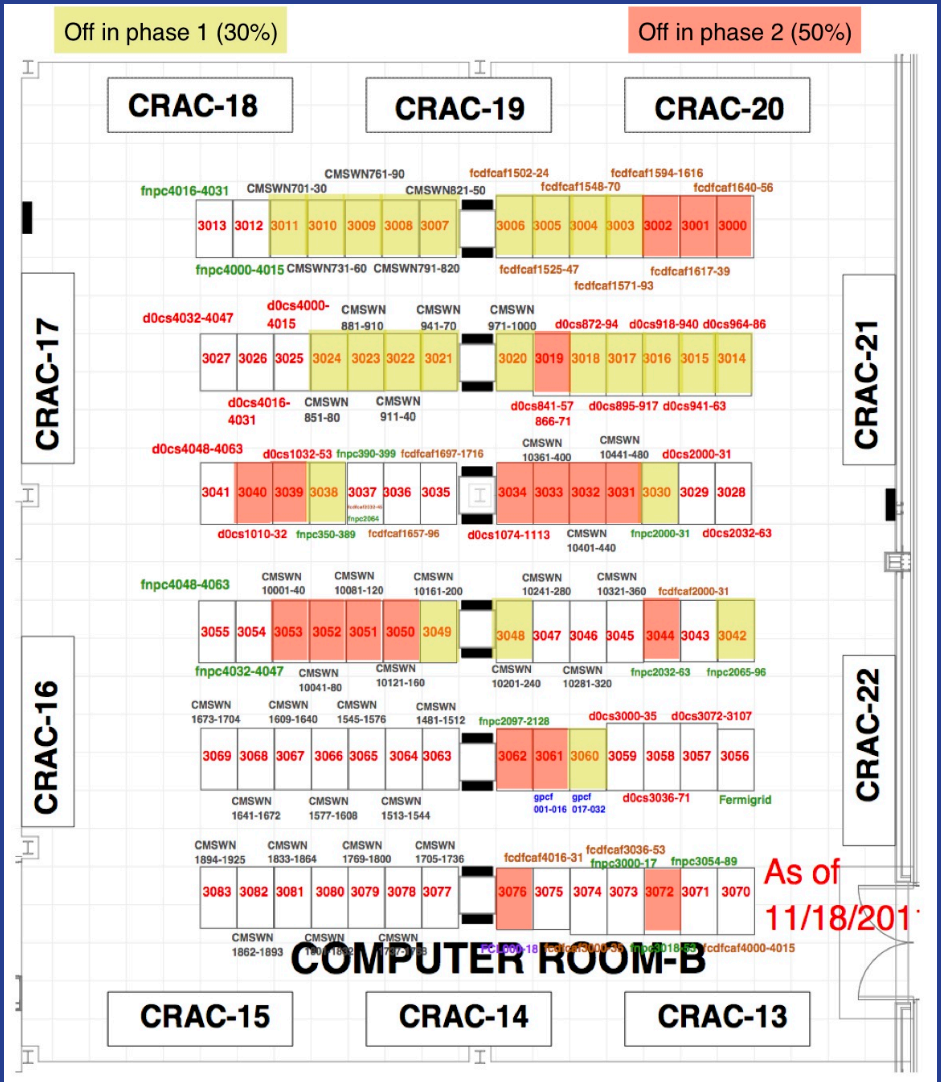

FCC and GCC

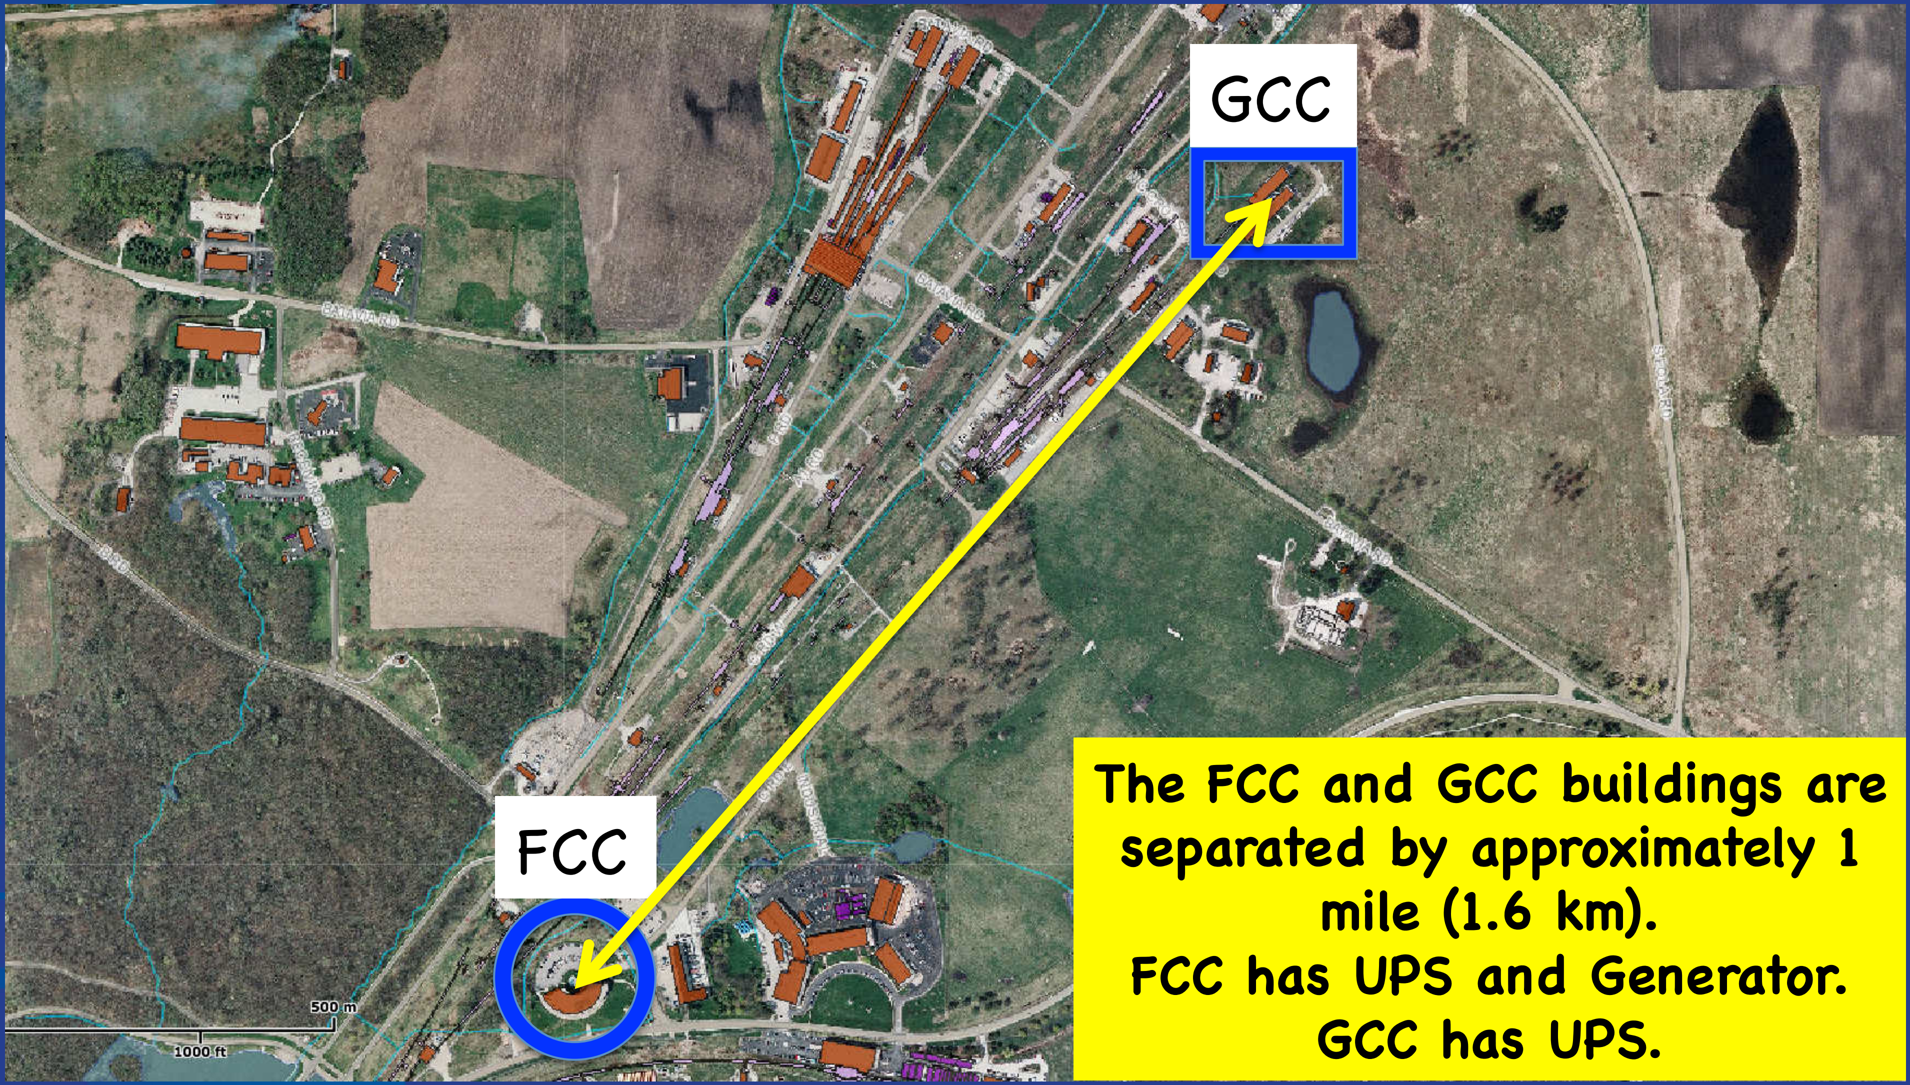

The FCC and GCC buildings are separated by approximately 1 mile (1.6 km).
FCC has UPS and Generator.
GCC has UPS.

Chicago Area MAN

Fermilab has multiple paths to Starlight via the Chicago Area Metropolitan Area Network (MAN).

The Chicago Area MAN lands at two locations in Fermilab:

- FCC-2 Network Area,
- GCC-A Network Room,
- Multiple path buried fiber cross connect between the two MAN locations at Fermilab (see later slide).

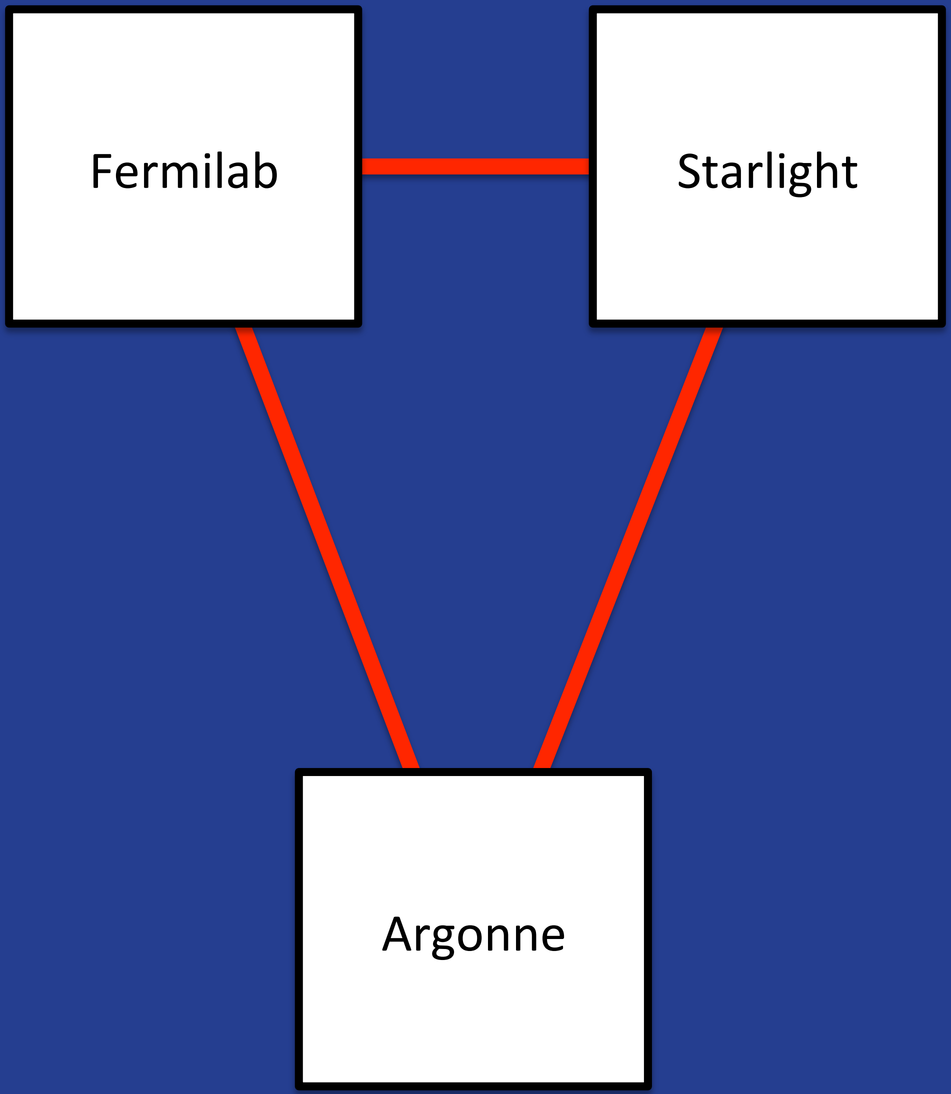

Fermilab – Argonne

~20 miles (33 km) apart = ~24 miles (40 km) driving

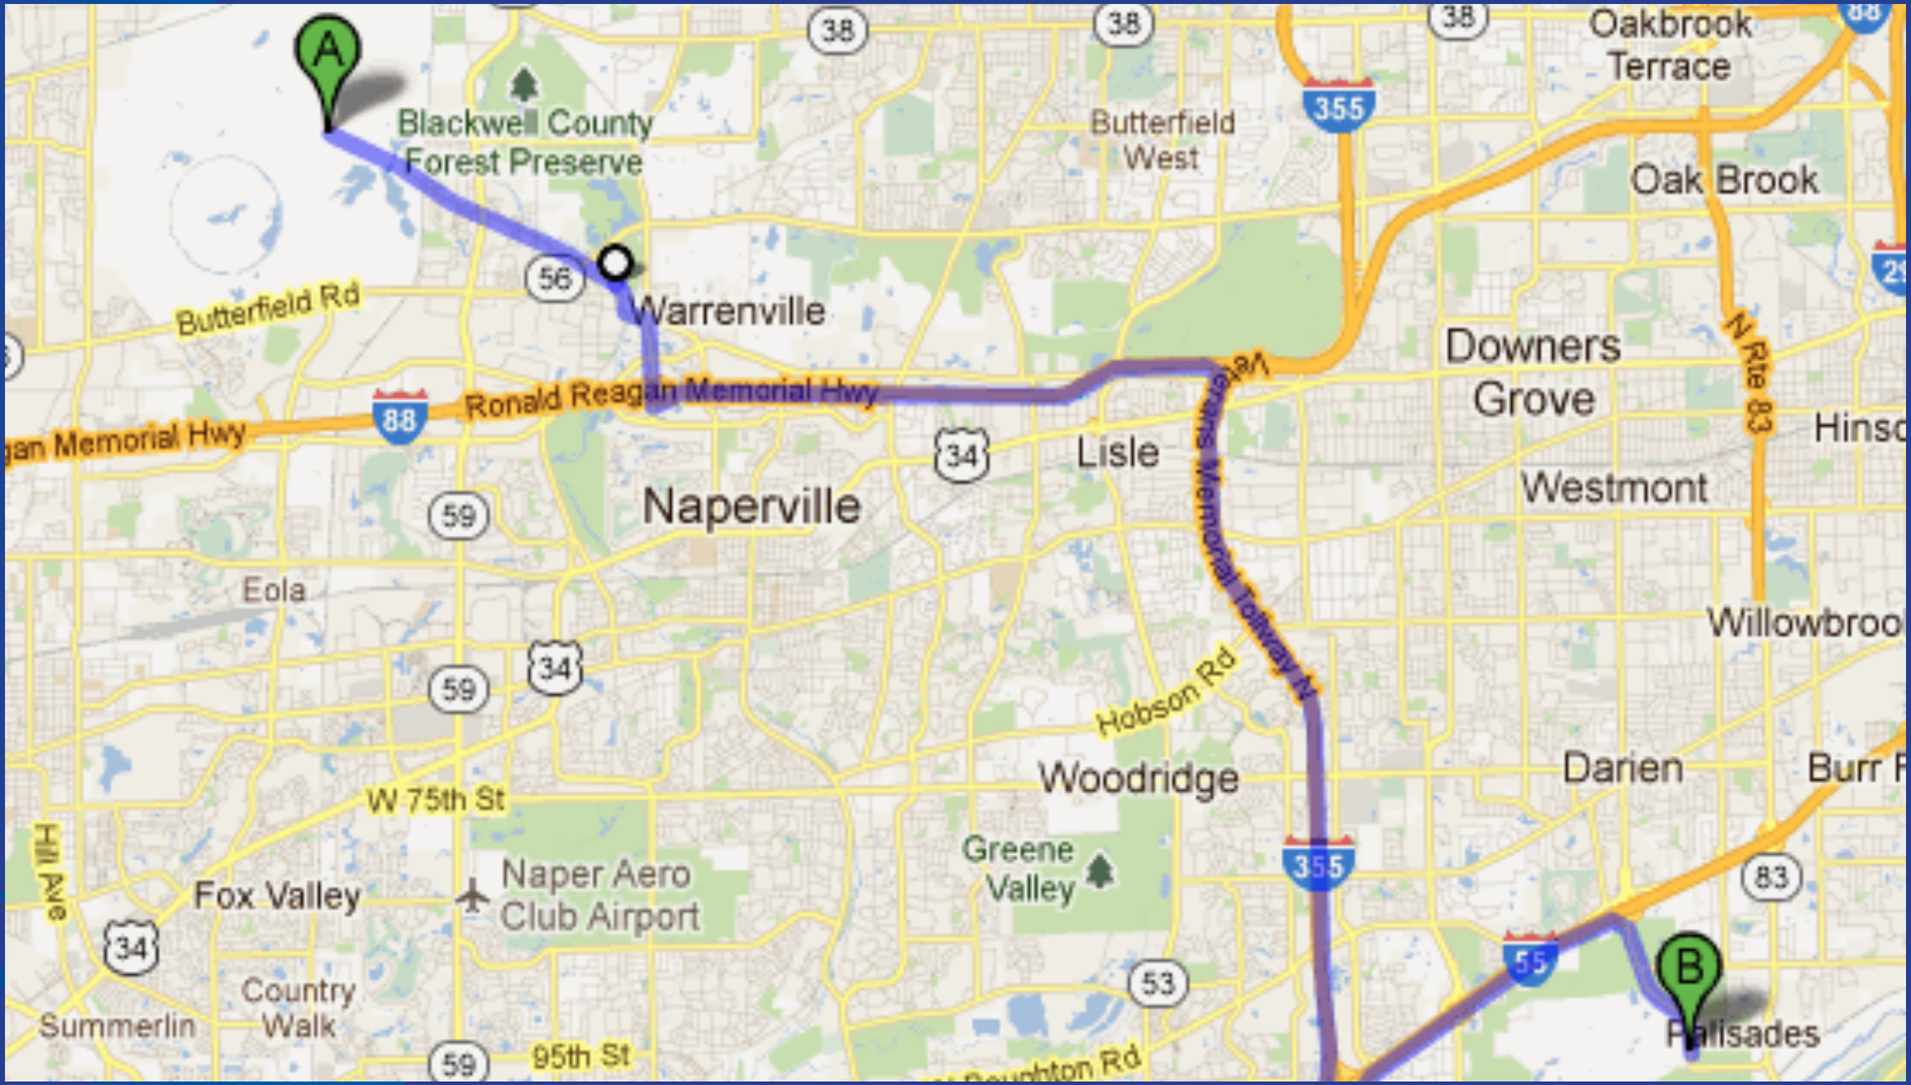

Fermilab – Argonne DR Agreement

- Reciprocal agreement to host Financial DR services between Fermilab and Argonne,
- Agreement has been in place for multiple years,
- Lukewarm "backup" – Requires reinstall from backup tapes shipped offsite every week,
- Fermilab/Argonne personnel visit the corresponding DR site a "handful" of times a year,
- Most work performed via remote administration,
- DR plan is activated ~once a year to verify that it still works.

Buried Fiber Loop Path Diversity (as of Nov-2011)

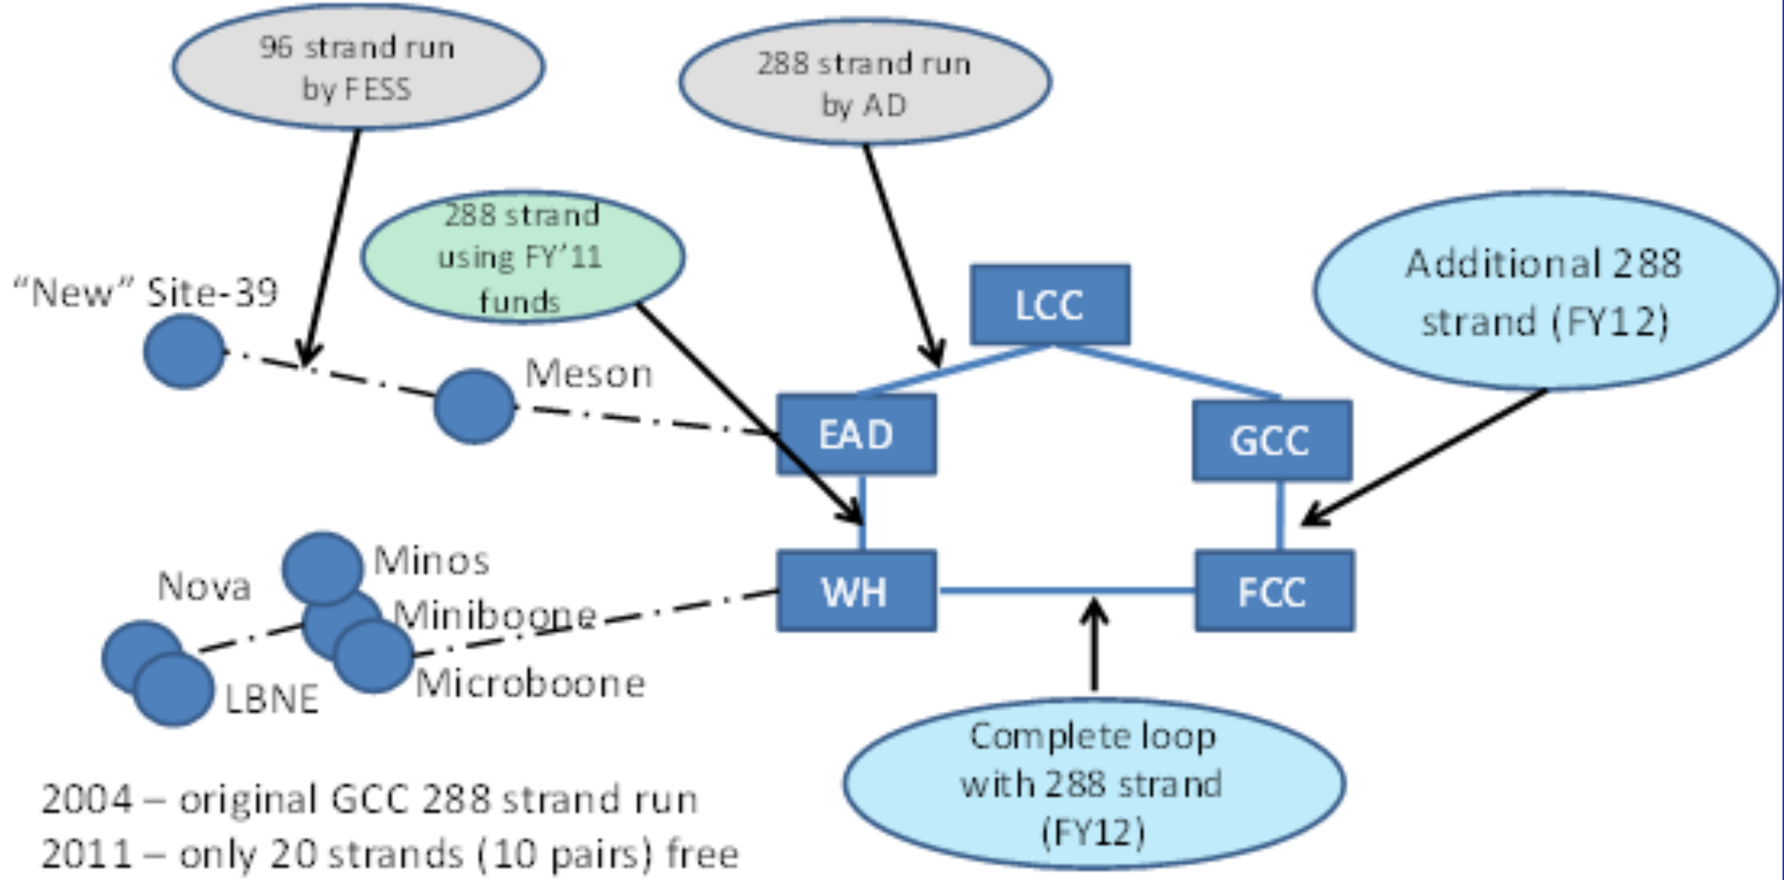

2004 – original GCC 288 strand run

2011 – only 20 strands (10 pairs) free

2012 – Add additional pairs to solve current capacity issues and position ourselves to support the new Frontiers and provide path diversity

2012 – Provides path diversity for computers in FCC and GCC

Distributed Network Core Provides Redundant Connectivity

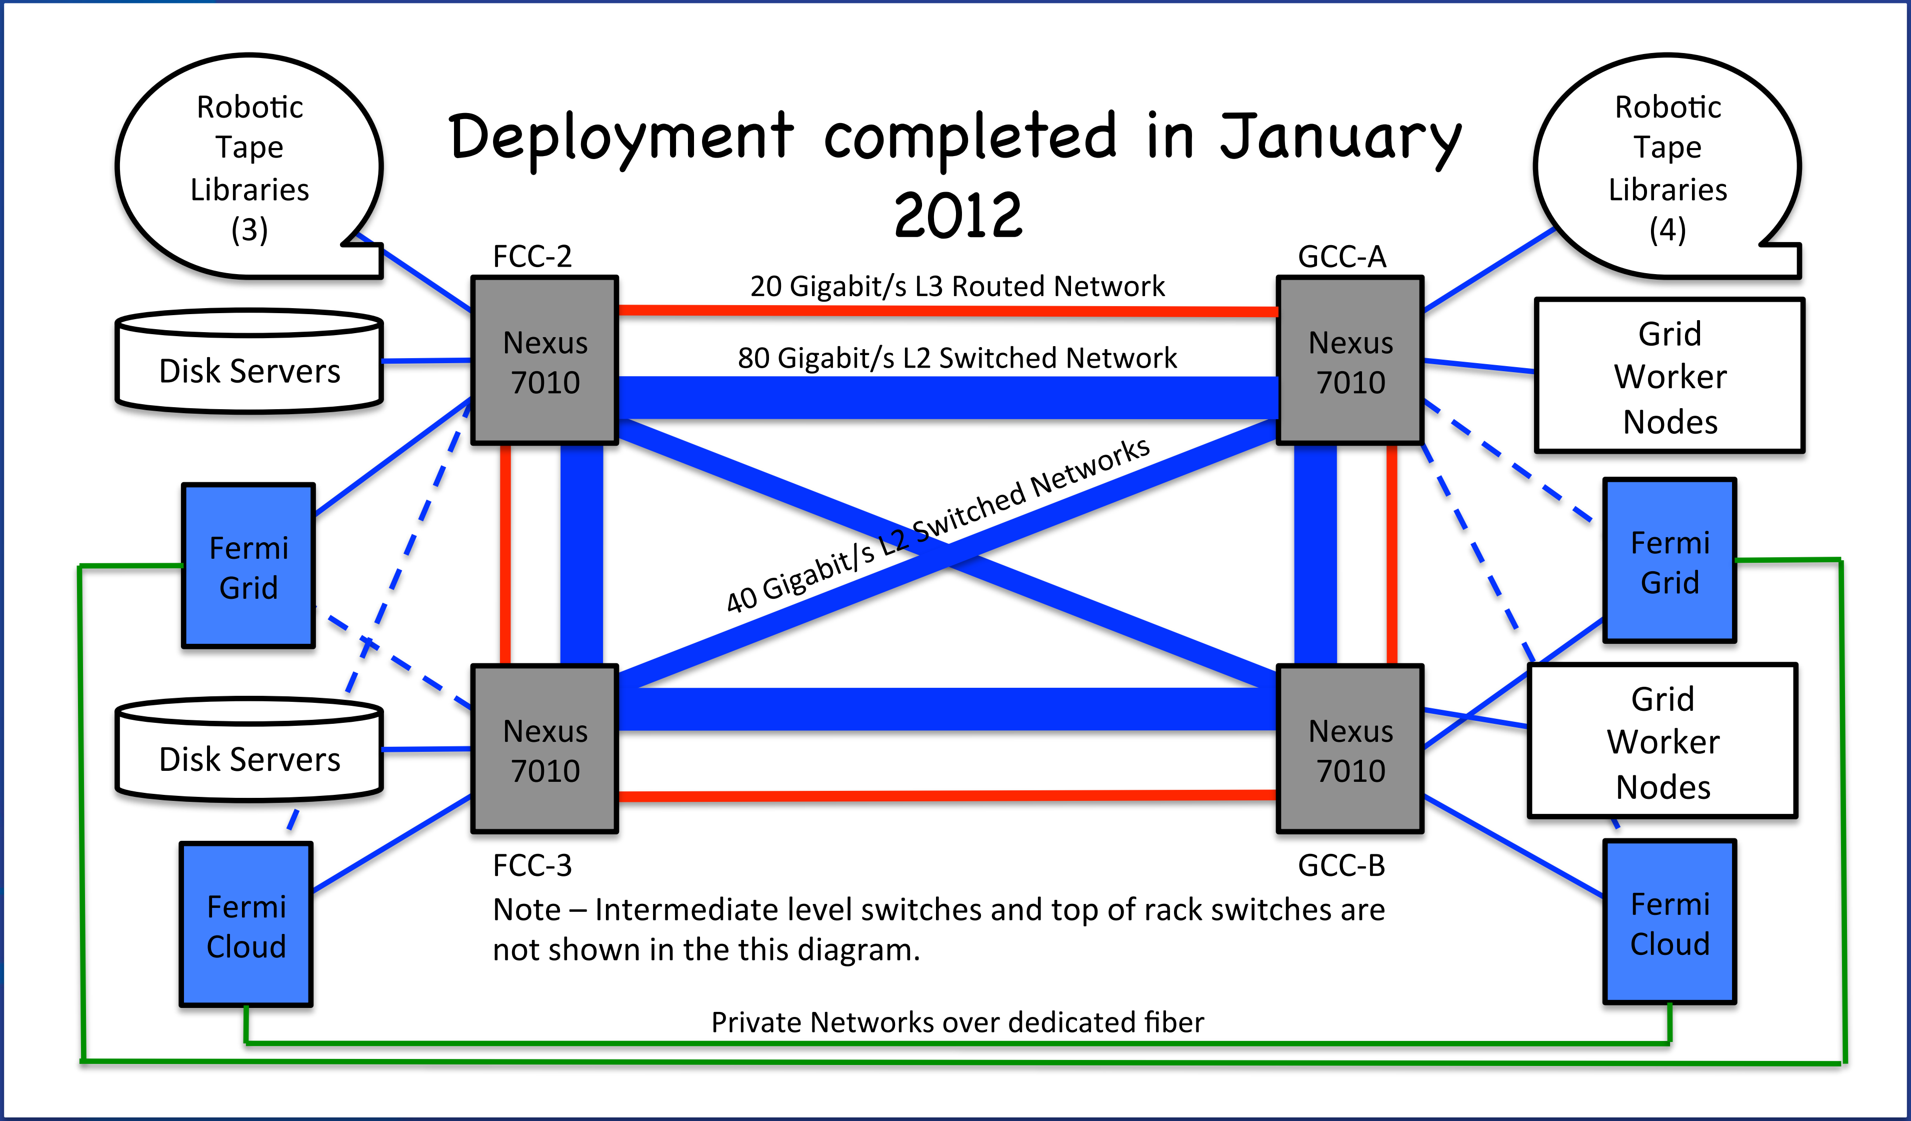
Robotic Tape Libraries (3)
Deployment completed in January 2012
Robotic Tape Libraries (4)
FCC-2
GCC-A
20 Gigabit/s L3 Routed Network
80 Gigabit/s L2 Switched Network
Nexus 7010
Nexus 7010
Disk Servers
Grid Worker Nodes
40 Gigabit/s L2 Switched Networks
Fermi Grid
Fermi Grid
Nexus 7010
Nexus 7010
Disk Servers
Grid Worker Nodes
Fermi Cloud
Fermi Cloud
FCC-3
GCC-B
Private Networks over dedicated fiber

Note – Intermediate level switches and top of rack switches are not shown in the this diagram.

DNS, Kerberos, Windows and KCA

DNS at Fermilab is implemented using Infoblox DNSSEC appliances:

- Master system feeds the primary and secondary name servers,
- The primary and secondary name servers feed a network of distributed "slaves" located at key points around the Fermilab network,
- End systems can either use the primary/secondary name servers, the distributed slave in their area, or use the "anycast" address (131.225.0.254).
- Offsite name servers hosted by DOE ESnet maintained via periodic zone transfer.

MIT Kerberos is implemented with a single master KDC and multiple slave KDCs:

- In the event of a failure of the master KDCs the slave KDCs can take over most functions,
- If the master KDC failure is "terminal", it is possible to "promote" a slave KDC to serve as the master KDC.

The Windows Active Directory domain controllers are likewise distributed around the Fermilab site.

The Fermilab Kerberos Certificate Authority (KCA) is implemented on two (load sharing) systems that are in separate buildings.

Minimal Internet Presence

- In the event of a triggering security incident, Fermilab has a predefined intermediate protection level - Minimal Internet Presence:
  - Only permit essential services that keep the lab "open for business" on the Internet, plus necessary infrastructure to support them,
  - Full disconnection is always an option.
- This plan stops short of "cutting the connection" to the Internet for Fermilab:
  - One web server with static content.
  - Email.
- This plan can be implemented at the discretion of the Fermilab computer security team with the concurrence of the Fermilab Directorate.

Data Distribution

Raw Data Location	Analyzed Data Location
FCC Tape Robot(s)	GCC Tape Robot(s)
GCC Tape Robot(s)	FCC Tape Robot(s)

Virtualization & Cloud Computing

- Virtualization is a significant component of our core computing (business) and scientific computing continuity planning,
- Virtualization allows us to consolidate multiple functions into a single physical server, saving power & cooling,
- Physical to virtual (P2V), virtual to virtual (V2V) and virtual to physical (V2P),
- It's significantly easier and faster (and safer too) to "lift" virtual machines rather than physical machines.

FermiGrid-HA/FermiGrid-HA2

FermiGrid-HA uses three key technologies:

- Linux Virtual Server (LVS),
- Xen Hypervisor,
- MySQL Circular Replication.

FermiGrid-HA2 added:

- Redundant services in both FCC-2 and GCC-B,
- Non-redundant services are split across both locations, and go to reduced capacity in the event of building or network outage.

Deployment has been tested under real world conditions.

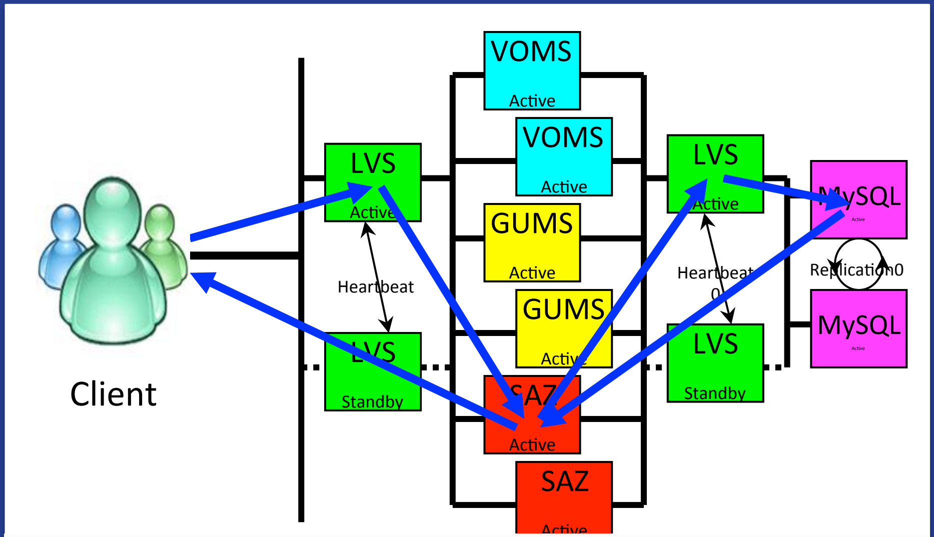

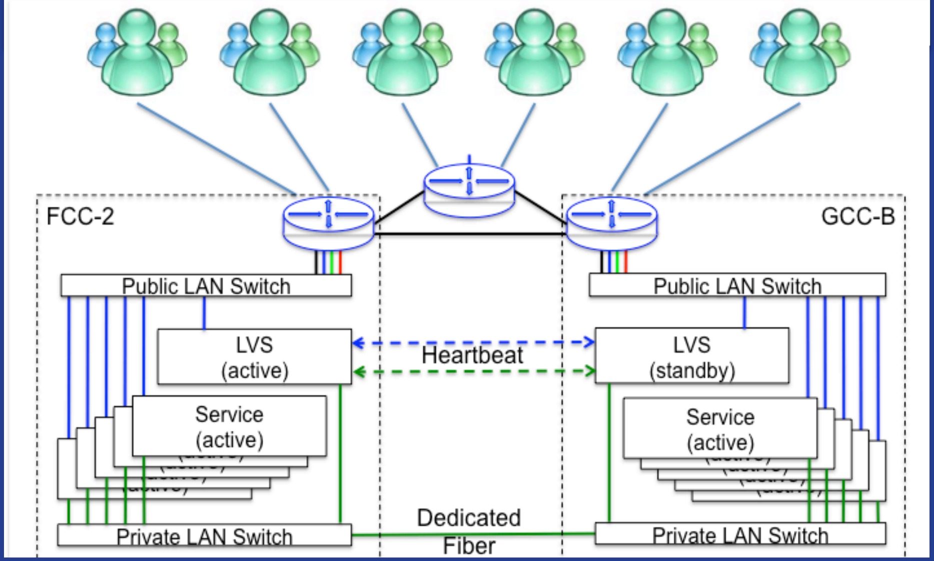

FermiCloud

- Services split across FCC-3 and GCC-B,
- Working towards a distributed & replicated SAN,
- FY2012 SAN hardware has just been delivered and installed in the GCC-B rack,
- We are evaluating GFS.

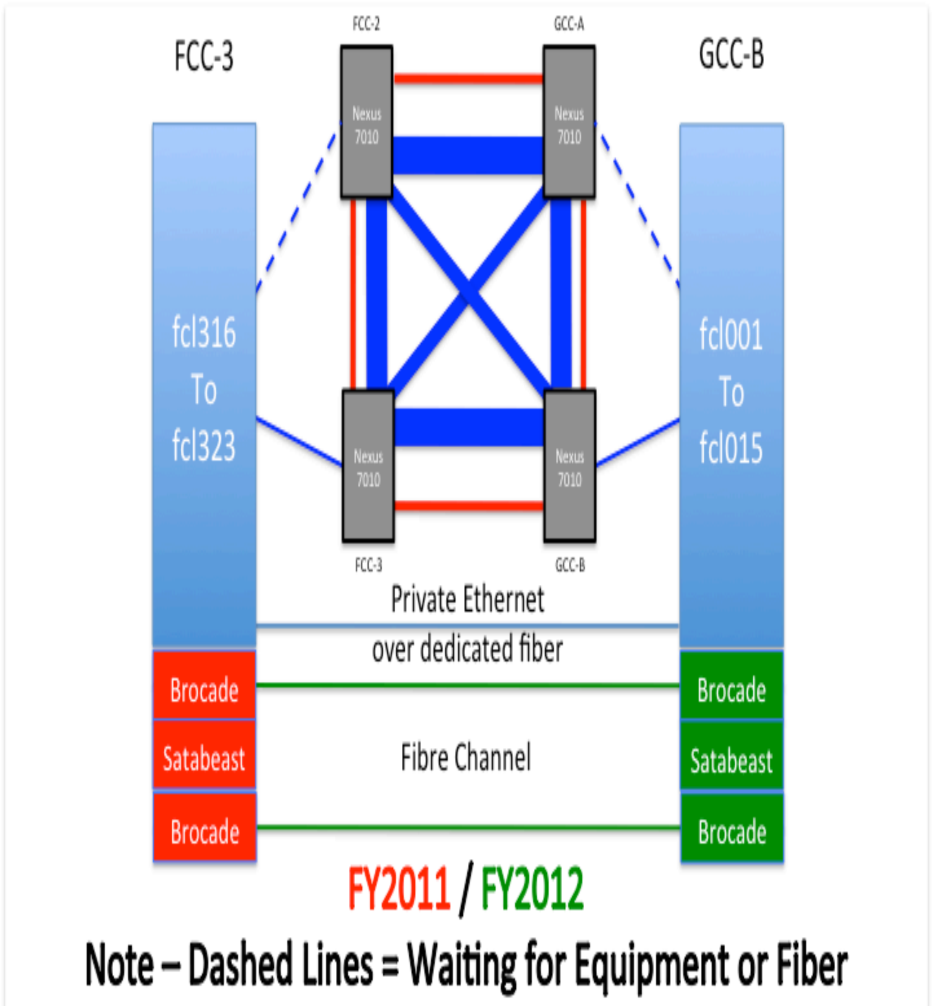

Note – Dashed Lines = Waiting for Equipment or Fiber

Outsourced Services

Fermilab uses several outsourced services, including:

- Kronos – Timecard Reporting,
- Service-Now – ITIL Support.

In each case where Fermilab has contracted for outsourced services, the business continuity plans of the service operators have been carefully evaluated.

Summary

- Utilities (power, cooling, etc.),
- Services distributed over multiple geographic locations,
- Distributed and redundant network core,
- Load sharing services - Active-Active,
- Standby services – Active-Standby,
- Published plan for load shed (if required),
- Virtualization & Cloud Services,
- Careful selection of outsourced services.

Thank You!

- Any Questions?