Networks
Project Status Briefing

August 29, 2006

Phil DeMar
Donna Lamore
Sheila Cisko
Matt Crawford

Networking Project Status Outline

- LAN projects & activities
- WAN projects & activities
- Security Infrastructure activities
- Physical Infrastructure activities
- Video Conferencing activities
- Wide area systems R&D activities

LAN Projects & Upgrade Plans

Donna Lamore

Core Network

- Core network backbone upgrades:
  - 10GE Core Network to CDF (finished)
  - 10GE Core Network to D0 (finished)
  - 2 Gb/s (aggregated) FCC to Wilson Hall
    - 10GE planned for next year
- 2nd 6509 deployed on FCC2:
  - Needed for additional system support
  - FCC2 subnets will be split into 2 groups
  - Routing will be pushed out to FCC2 & FCC1 6509s
    - Will facilitate deployment of redundancy

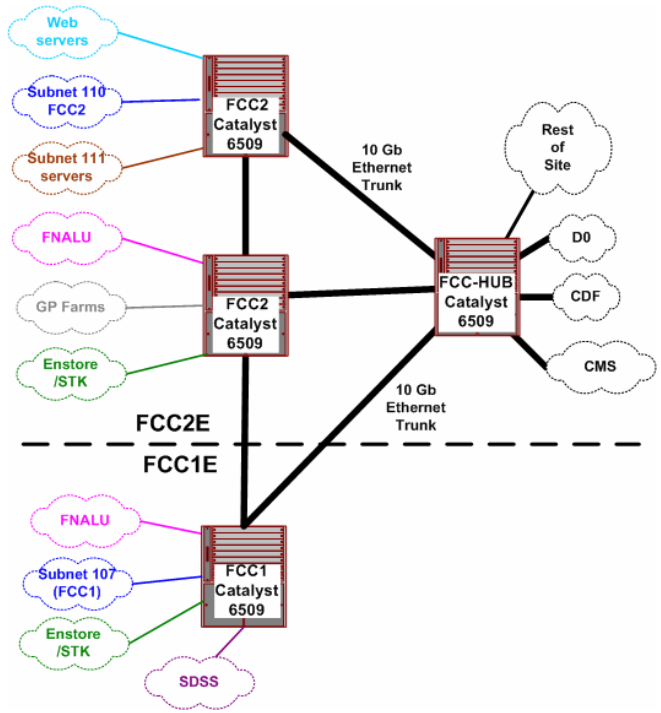

CDF

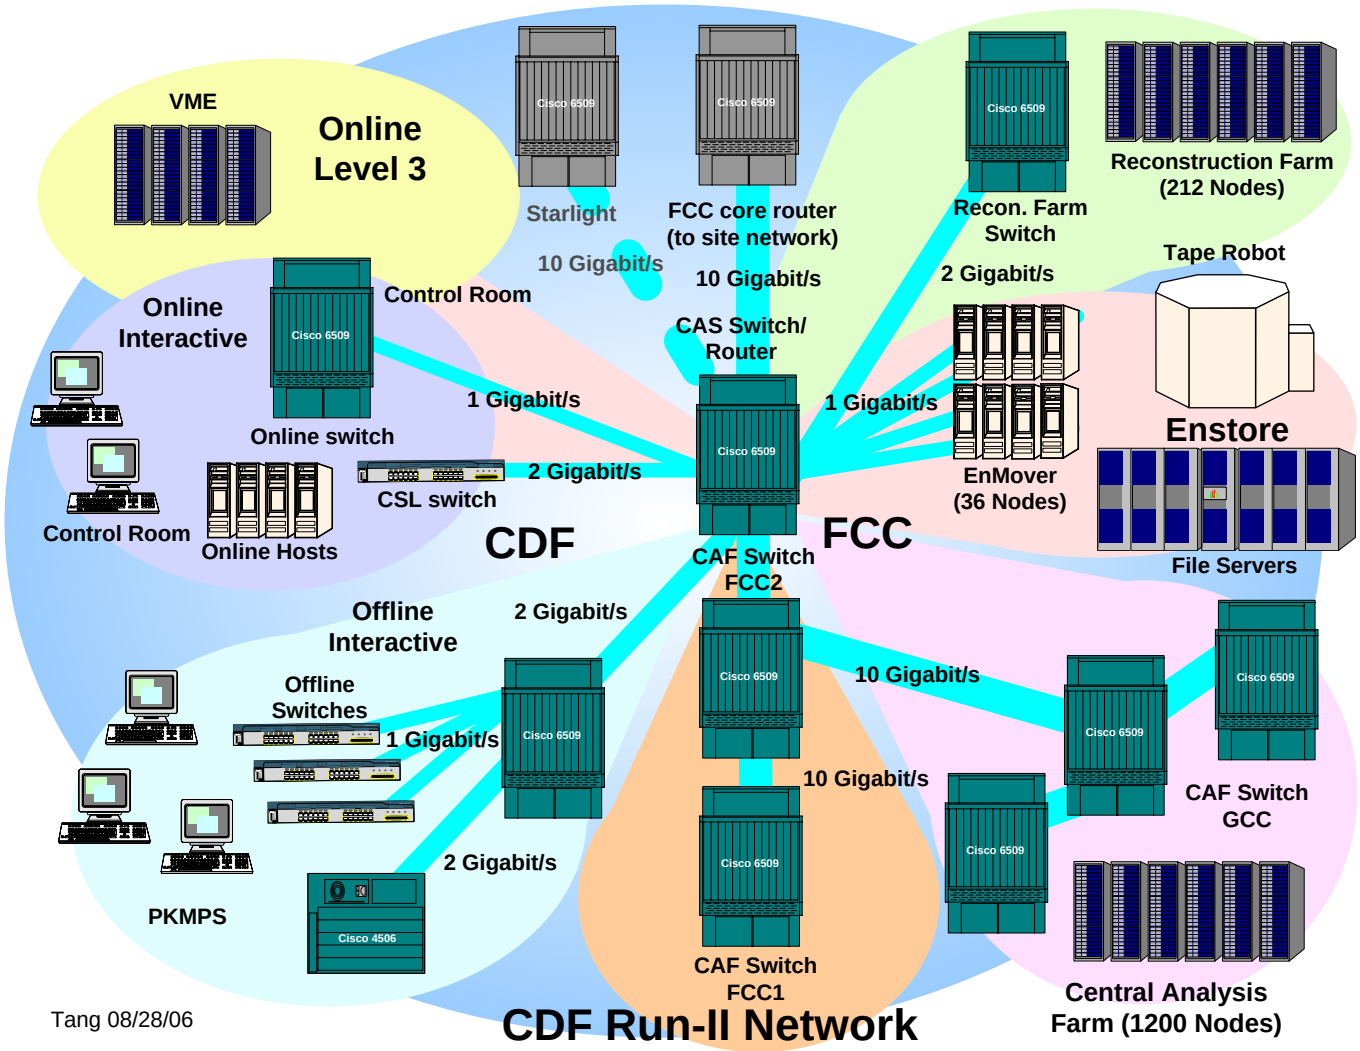
VME
Online Level 3
Cisco 6509
Cisco 6509
Starlight
FCC core router (to site network)
10 Gigabit/s
10 Gigabit/s
Cisco 6509
Recon. Farm Switch
Reconstruction Farm (212 Nodes)
2 Gigabit/s
Tape Robot
Online Interactive
Cisco 6509
Control Room
CAS Switch/ Router
1 Gigabit/s
1 Gigabit/s
Online switch
Cisco 6509
Enstore
2 Gigabit/s
CSL switch
EnMover (36 Nodes)
Control Room
Online Hosts
CDF
FCC
CAF Switch FCC2
File Servers
Offline Interactive
2 Gigabit/s
Offline Switches
Cisco 6509
Cisco 6509
10 Gigabit/s
Cisco 6509
Cisco 6509
1 Gigabit/s
10 Gigabit/s
Cisco 6509
CAF Switch GCC
2 Gigabit/s
Cisco 4506
PKMPS
Cisco 6509
CAF Switch FCC1
Central Analysis Farm (1200 Nodes)
CDF Run-II Network

CDF – Upgrades completed this year

- Grid Computing Center Upgrades:
  - Installed 3rd 6509 switch at GCC
- Core Network link upgraded to 10 Gb
- On-Line Upgrade:
  - Installed Switch & Infrastructure for CSLogger
- CDF Wireless Upgrades:
  - Replaced 17 802.11b APs with 802.11g
    - Completed upgrades to 802.11g
  - Wireless survey completed

CDF – Planned Upgrades

- Upgrade Starlight link to 10 Gb (waiting to be scheduled)
- Network infrastructure for FY ’06 farm upgrades:
  - 10Gb and Copper Gb module upgrade for Recon. Farm switch (waiting for parts, 100K ordered)
  - Copper Gb module upgrade for CAF FCC1 6509 switch
    On hold, determine if systems will be retired
  - New systems 450 total,160-200 in Computer room A
  - New systems waiting for computer room B (250-290)
- Wireless upgrades
  - Install 9 additional APs in portakamps for additional coverage
  - Move antennas in Building 327

CDF – Planned Upgrades

- Upgrades to CDF Off-line network facilities in Portakamps - (waiting for PPD budget)
  - Switch fabric upgrades to Sup720 (720Gb)
  - Expanded 1000B-T support
  - 10 Gb to FCC
- Upgrades to CDF network facilities in Building 327:
  - 4506 switch with 1000B-T support (waiting for budget)
- Moving 6509 from FCC to GCC - (CRB)
  - Module upgrades from 10/100 to 1000

D0

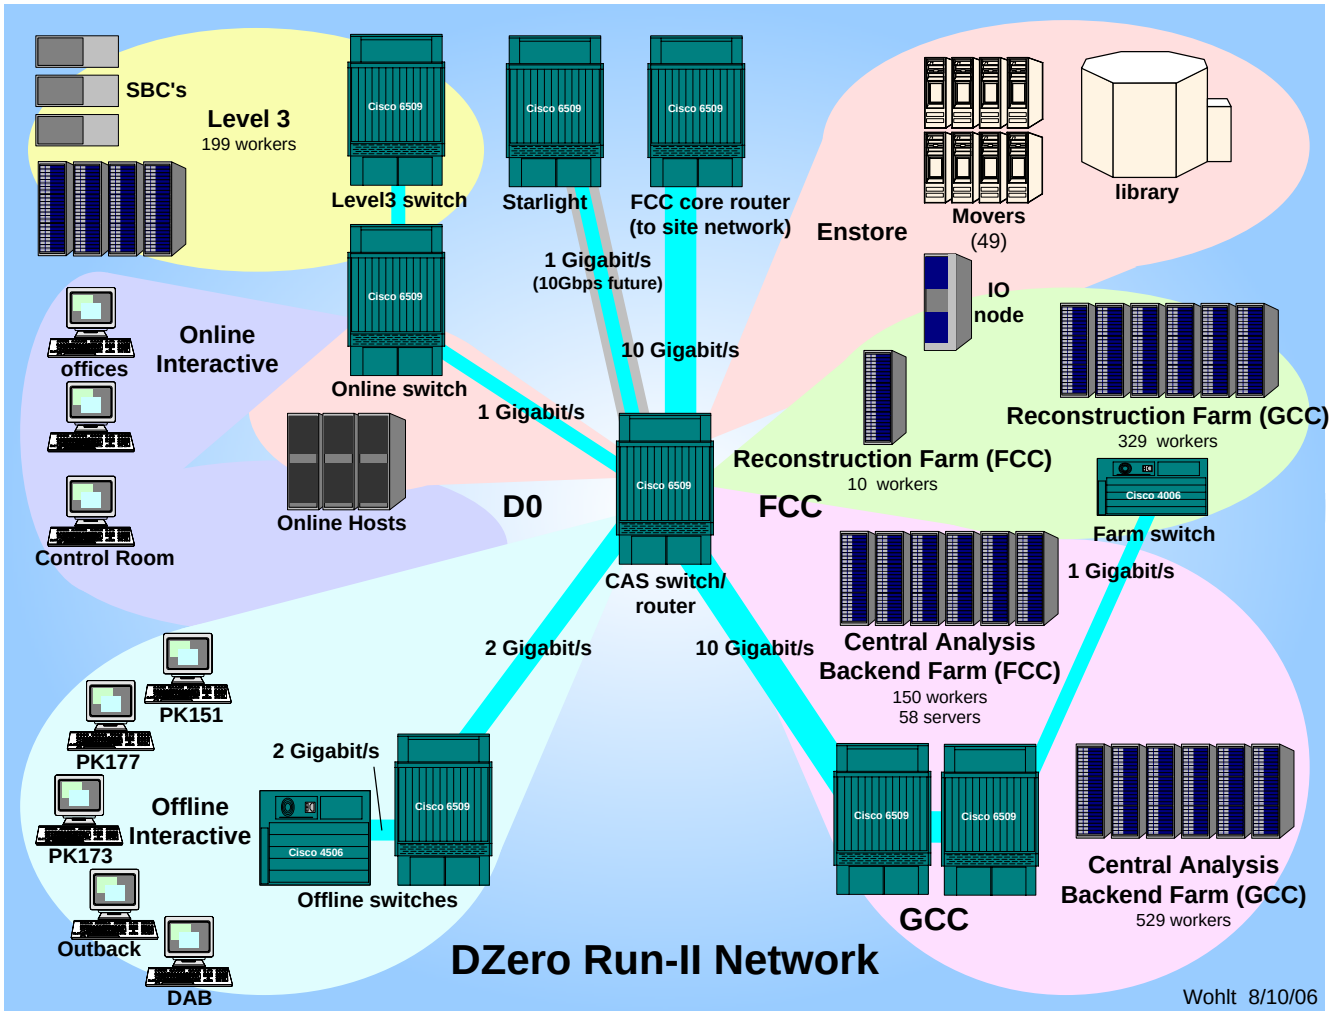
SBC's
Level 3
199 workers
Cisco 6509
Level3 switch
Cisco 6509
Starlight
Cisco 6509
FCC core router
(to site network)
Enstore
Movers
(49)
library
1 Gigabit/s
(10Gbps future)
IO
node
Cisco 6509
Online switch
10 Gigabit/s
offices
Online
Interactive
1 Gigabit/s
Reconstruction Farm (GCC)
329 workers
Reconstruction Farm (FCC)
10 workers
Online Hosts
D0
Cisco 6509
FCC
Cisco 4006
Farm switch
Control Room
1 Gigabit/s
CAS switch/
router
2 Gigabit/s
10 Gigabit/s
Central Analysis
Backend Farm (FCC)
150 workers
58 servers
PK151
PK177
2 Gigabit/s
Cisco 6509
Offline
Interactive
PK173
Cisco 4506
Offline switches
Cisco 6509
Cisco 6509
Central Analysis
Backend Farm (GCC)
529 workers
Outback
GCC
DAB
DZero Run-II Network
Wohlt 8/10/06

D0 – Upgrades complete this year

- Total # of systems supported – 2266
- Installed 3rd 6509 at GCC
- Upgraded all Wireless APs to new model
- Upgraded Network Core connection to 10Gb
- Wireless survey completed

D0 – Planned Upgrades

- Upgrade to 10Gb connection to Starlight
- Add additional 10Gb connection from FCC to GCC
- Install 6509 (currently in storage) in Computer Room A at GCC
  - Number of additional systems to be supported - 320
- Wireless
  - Install 7 additional APs for additional coverage
  - Move DO Outback AP antennas

CMS – Upgrades completed this year

- Installed additional 6509 at GCC – 511 systems supported
- FCC 6509 – systems supported FCC1 - 31, FCC2 - 177
  - Installed 10Gb module
  - Installed 2 additional 48 port 10/100/100 modules
- CMS FCC Core Switch
  - Added additional 10 Gb module
  - Added 3 additional 48 port modules to FCC Core Switch
  - Upgraded Power supplies in FCC Core Switch
- Upgraded connection between FCC to GCC to 40Gb

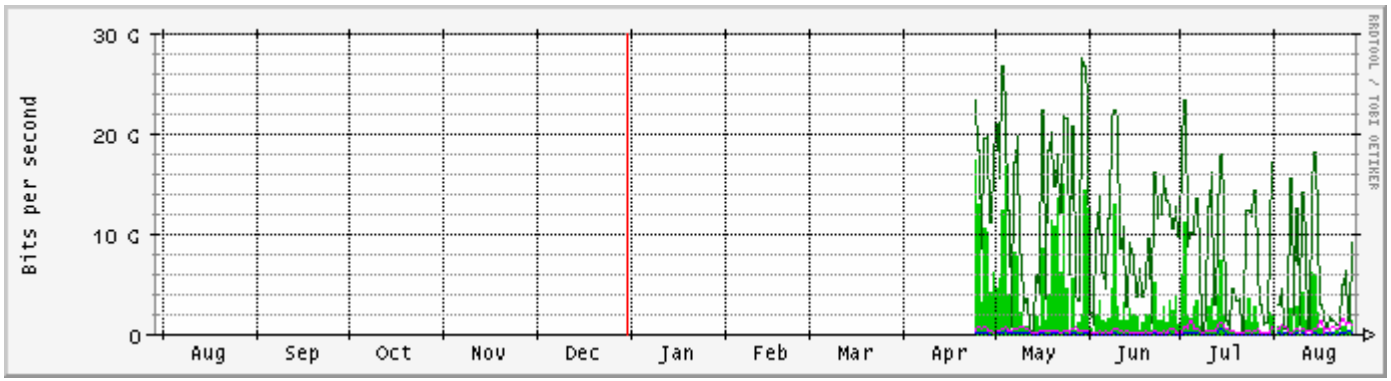

CMS – Upgrades in planning

- All upgrades complete for this year
- Started planning for next year's buys….
  - Issues –
    - Oversubscription of Modules
    - Number of Switches required
    - Oversubscription of Uplinks

CMS Work Group Core based on 2xC6509 for FY07.

08/25/2006

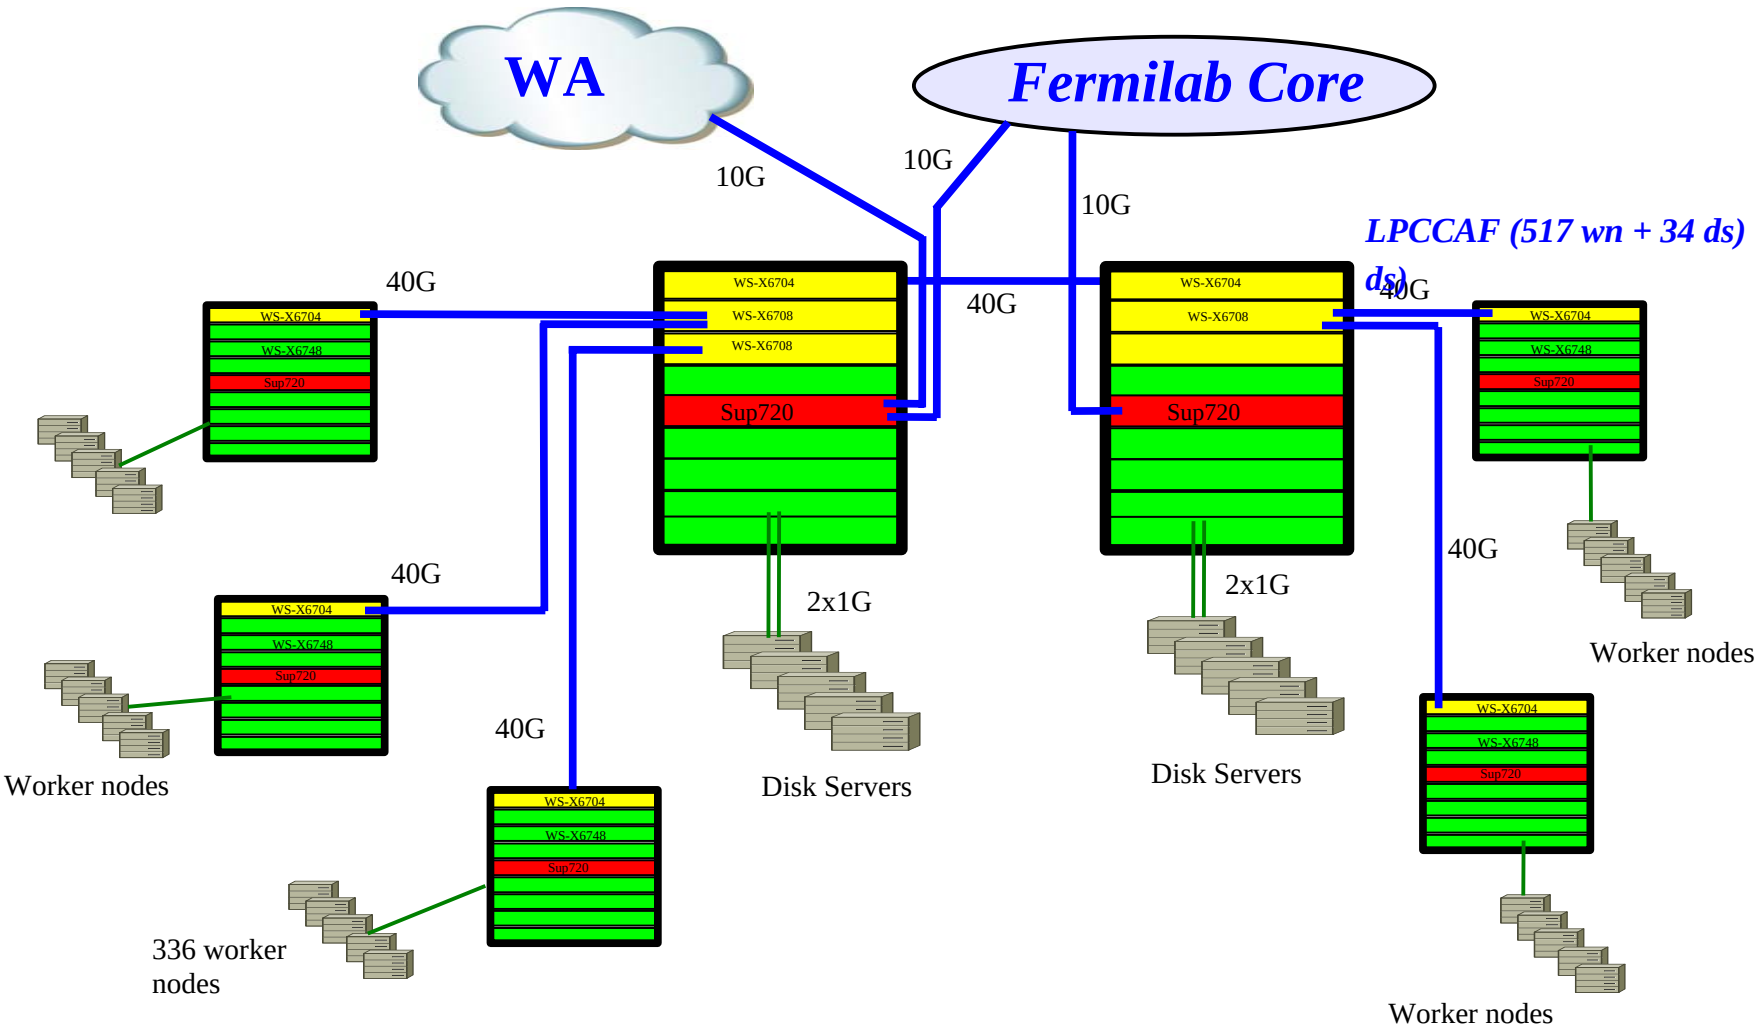

Requirements Total:
worker nodes: 1507
disk servers: 218

Capacity :
worker nodes: 1680 (336x5)
disk servers: 240 at 2x1G

Cost estimates: $468K

Wireless

Upgrade Status –

120 APs upgraded across the site

As of 1/1/06 – 66 new model APs had been installed

Between 1/1 – 5/4 - 55 New APs installed – primarily WH & D0

5/4 to 7/31 – Completed Upgrades – added additional APs – 212 now installed

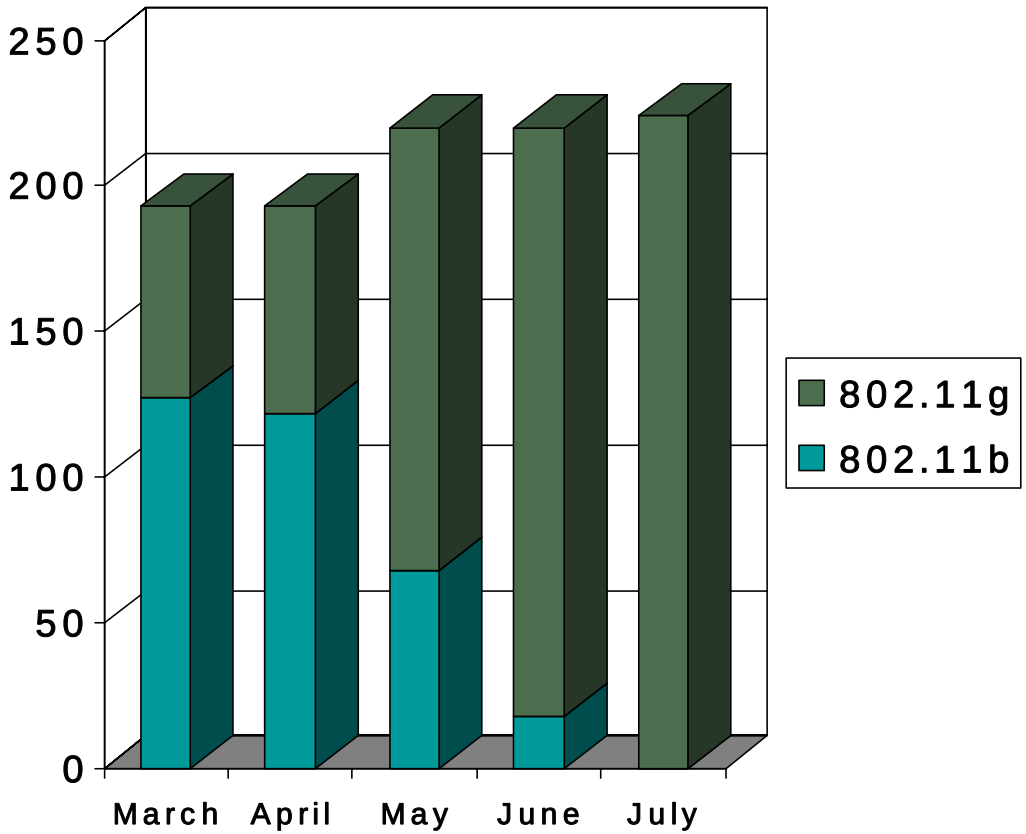

Wireless Deployed (TeraBeam)

- Two base units installed on Wilson Hall 16
  - Currently in service for Site 56, Argonne field sites in progress
  - Plan to Provide service to new areas
    - Site 56, pump houses, Guard shacks
  - Provide for emergency link restoration
- Terabeam to be installed at Lab 8 in the Village
  - Line of sight to 6 houses that have no networking
  - Near line of sight to 4 additional houses that have no networking
  - Near line of sight to User's center – will replace old Westell equipment

Wireless planned (TeraBeam)

- Site 52
  - Line of sight to WH – Very straight forward installation
- Site 29
  - Survey done. No line of sight to WH
  - Recommendation is to keep using copper
  - Looking in to new equipment offerings for small DSL installations
- Gas Shed – Survey to be done
- Guard Shack @ LSC – near line of sight to WH
- Additional Terabeam installs
  - 2 Pine street guard shacks, line of sight OK
  - Wilson street guard shack, will mount on existing utility pole

Wireless planned

- Benefits of WISM
  - Unified visibility and control of all access points
  - Integrated into core network routers
  - Reduced analyst effort spent troubleshooting
- Enhanced rogue detection and location identification – follow MAC address across wireless LAN, monitor signal strength.
  - Only one rogue found, ours, WLSE didn’t recognize it
  - Friendly rogues popping up weekly in surrounding neighborhoods
  - Rogue with SSID “test” reported on lower level WH

Miscellaneous

- GCC Tape Robot Network
  - New Subnet Created For Enstore at GCC:
    - Supported on new tape robot 6509 at GCC
    - 10Gb connection to Core Network – additional10 Gb connections can be easily added
  - Private Robot Controls LAN implemented
    - Includes all nodes needed to control the Robots in FCC & GCC

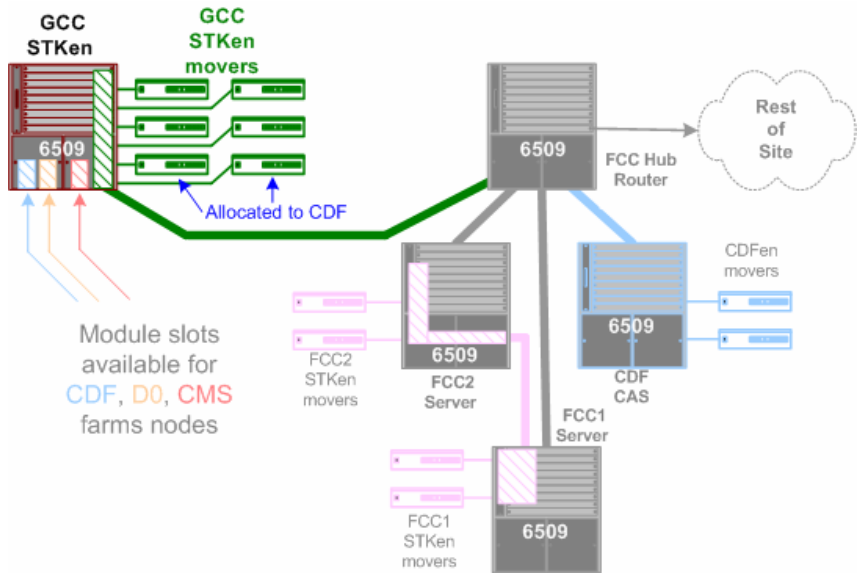

Miscellaneous

- GCC Tape Robot Commissioning Traffic

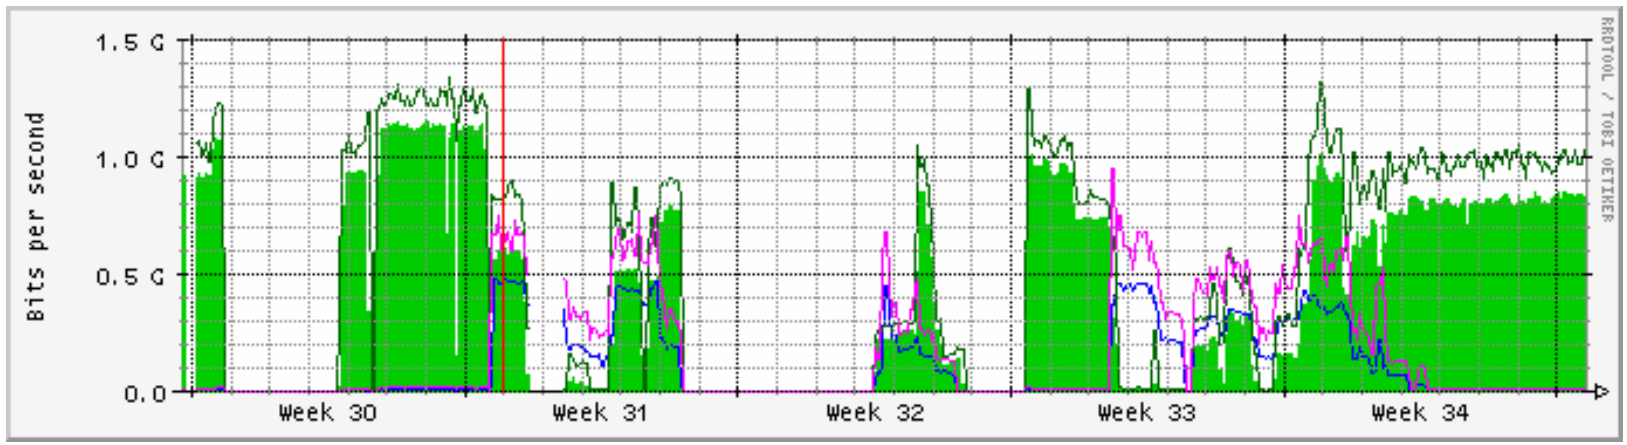

Miscellaneous

- DNS Migration to Linux
  - Testing - New Enterprise server to go on line in parallel this week
  - Primary and Secondary Nameservers to follow shortly
- Dial-in status
  - Down to 4 users in May this year
  - 9 additional users after ISP phase out announcement
    - 6 were help desk support staff
  - Current usage:

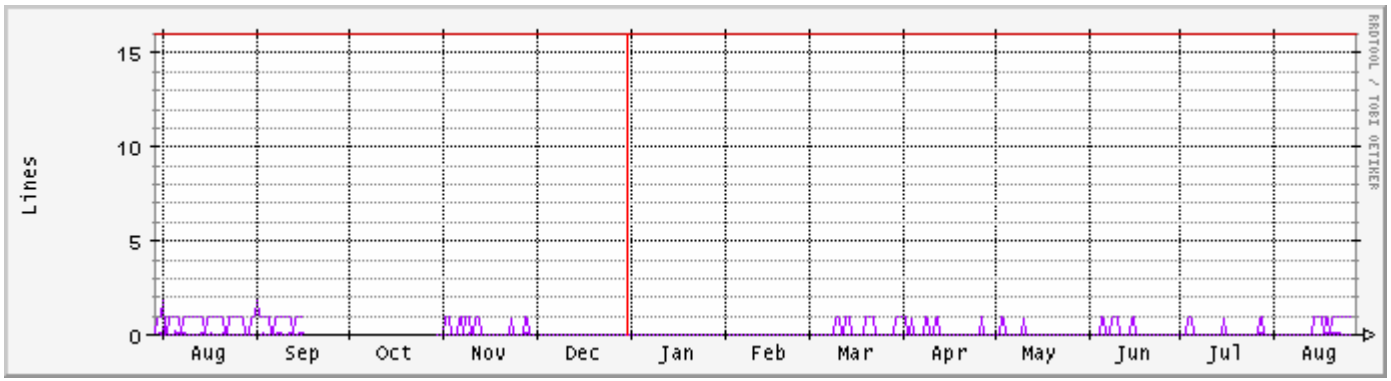

WAN Projects & Upgrade Plans

Phil DeMar

Current WAN status:

- Production IP (ESnet) path:
  - Migrated to one of the two Fermi LightPath 10GE channels on July 11
  - 10 GE WAN paths into ESnet from CMS, CDF, & D0

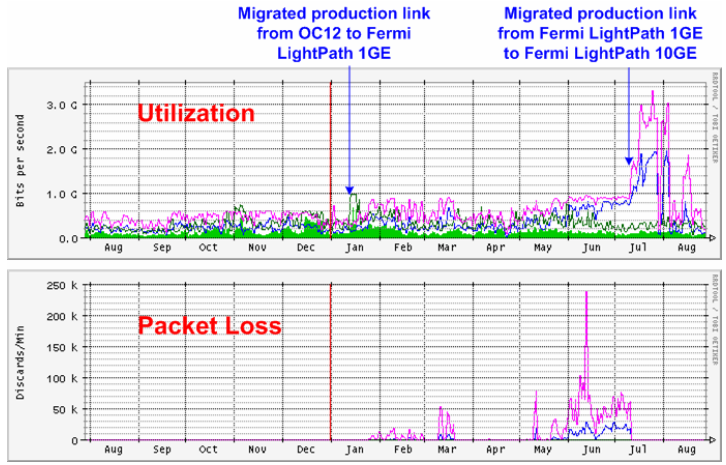

- Fermi LightPath high impact channels:
  - Second 10 GE link dedicated to CMS T0/T1 (CERN) & US-CMS T1/T2 sites
    - Layer 2 path to CERN via US-LHCnet
    - Direct connectivity with U-Wis & Purdue
    - Caltech & UFL connectivity via UltraLight
    - UCSD, UNL, & MIT via ESnet MPLS tunnels
  - Other CMS high impact traffic moved back to default (ESnet) path

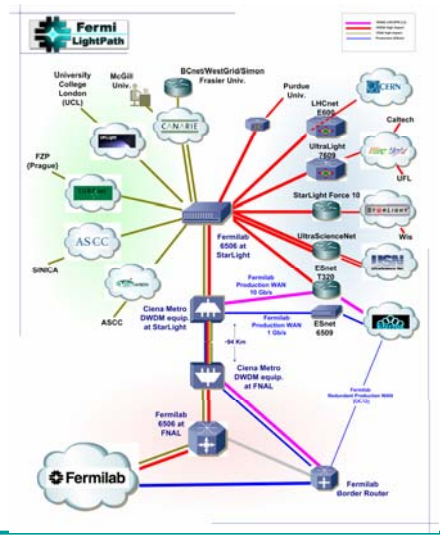

Alternate Path Data Flows

- Alternate path approach:
  - Define high impact traffic flows:
    - Our approach is minimal sized source/destination netblock pairs
  - Implement alternate forwarding via Fermi LightPath channels & local router:
    - Outbound: policy route to alternate path
    - Inbound: dependent on capabilities & policies of remote “end”
      - ESnet MPLS tunnels are useful for doing this
  - Security based on minimal aperture ACLs
  - Redundancy capabilities vary…
    - We try to test failover during maintenance periods

ESnet path

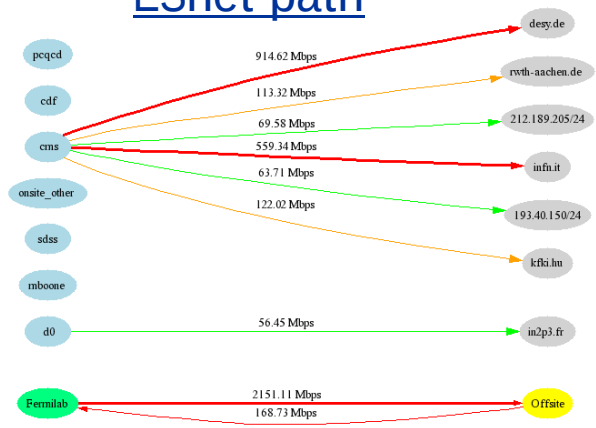

High impact traffic path

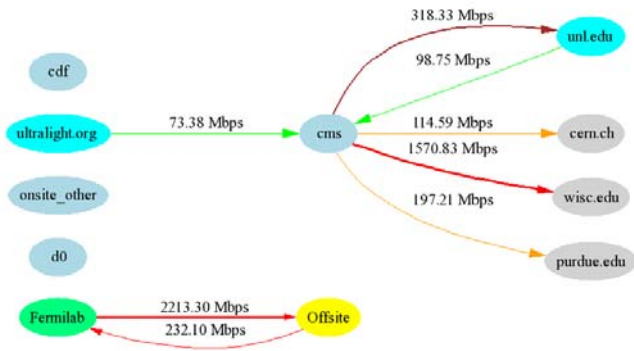

Fermilab

FNAL Outbound Traffic

Fermi LightPath

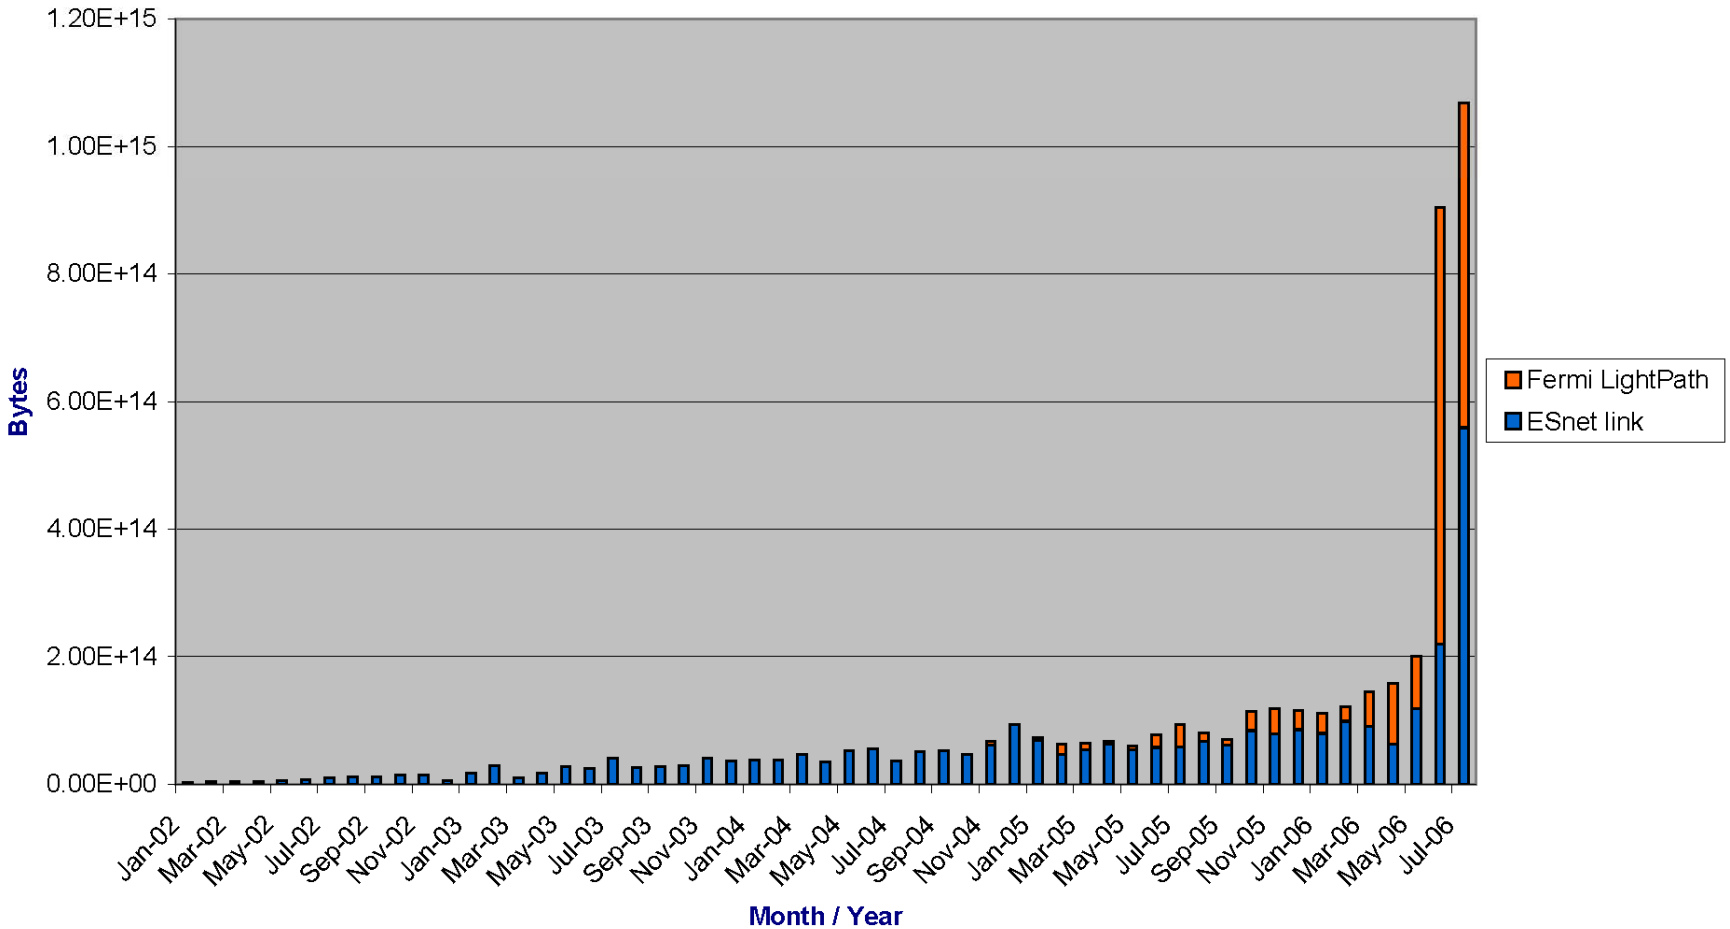

ESnet MAN - Configuration:

- The MAN is a 3-layer architecture:
  - Bottom layer is Fermi LightPath & ANL I-Wire optical network infrastructures
  - Middle layer is Ciena DWDM equipment to provide MAN channels
  - Top layer is ESnet layer 2 & 3 equipment that provides service back to the Labs
- Initial MAN configuration for FNAL – 6x10GE:
  - Redundant 10GE production network links
  - Four 10GE SDN channels:
    - One to be dedicated to CMS T0/T1 (LHCOPN) connectivity (CERN)
    - Second dedicated to US-CMS high impact paths
    - Third for CDF, D0, SDSS, others, high impact paths
    - Fourth, a redundant SDN channel for failover

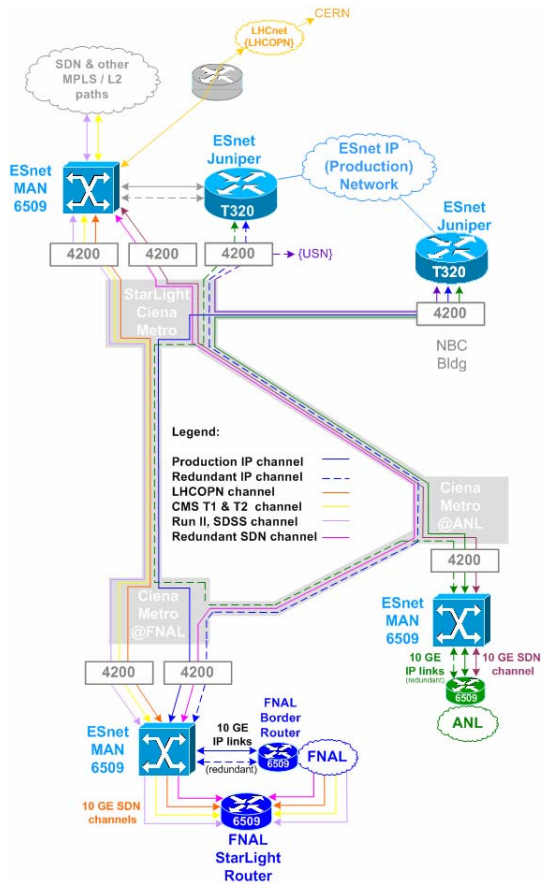

ESnet MAN – Deployment Schedule

{slipping, to no one's surprise…}

- ANL / FNAL fiber segment expected in place by late Sept.
- Ciena DWDM MAN equipment has arrived
  - Preliminary inventory of new equipment completed
  - 4x10GE channels are already in place between FNAL & StarLight, left over from SC2005 product loan
    - Expect to bring these channels up as soon as ESnet equipment is in place
- Kickoff installation with Ciena & ANL held (8/24)
- Two outages of Fermi LightPath will be needed:
  - Rack reorganization in mid-September:
    - Production 10GE link will be down (30-60 min); will force failover to OC12
    - CMS high impact links will be down for up to half-a-day
  - Reconfiguration of Ciena Metro equip. in mid-Oct at StarLight, FNAL, & ANL
    - Full day outage; will need to operate on OC12 for the day
    - Exploring temporary connection to ANL for redundancy for the day

MAN Deployment Status

- MAN install requires addtl rack at StarLight:
  - Existing rack will support Ciena Metro equip:
    - ANL I-Wire modules necessitate 3rd Metro shelf
  - New rack will support Ciena 4200s & FNAL R&D equipment
    - FNAL-managed fiber between the racks
    - Dedicated fiber to ESnet equip. racks

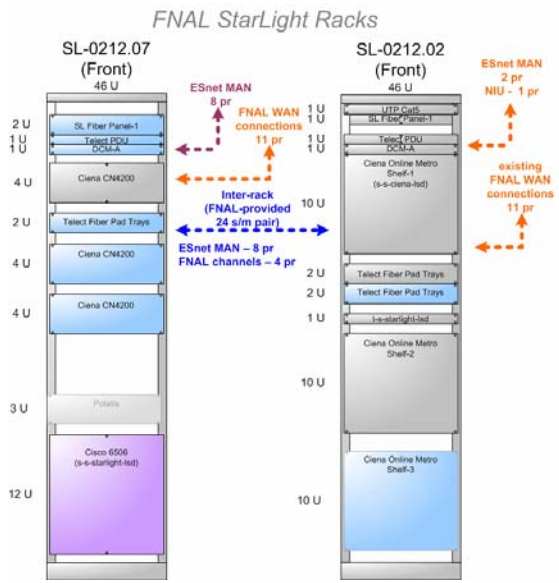

- ESnet's MAN switch at StarLight installed & in production use
- ESnet's MAN switch at FNAL is installed & (almost…) online
  - We will bring up the four 10GE channels used for SC2005 following the rack reorganization

Miscellaneous ESnet Issues

- Post MAN-deployment issues:
  - Detailed procedures for joint FNAL/ANL operational support of MAN infrastructure still need to be worked on
    - Will provide ESnet with per channel 7x24 contacts
  - Ongoing costs & maintenance agreement with ESnet needs to be put in place yet
  - Sparing inventory parts location & access need to be detailed
  - Contingency planning for core MAN equipment failure needs to be done
    - Considering utility & justification of procuring extended reach XENPAKs to bypass damaged Ciena Metro equipment at FNAL, ANL, or StarLight
- Need letter confirming our ESnet Site Coordinators Committee (ESCC) representative

DuPage Tech Park Activities

- Abovenet easement request (orange, right) was submitted to DOE in early August:
  - FESS has received copy of plat
  - Still no final decision on termination pt in DNTP
    - So, no final length on FNAL/DNTP fiber cable
      - Cost est.: $16k for fiber; $12-16k for installation

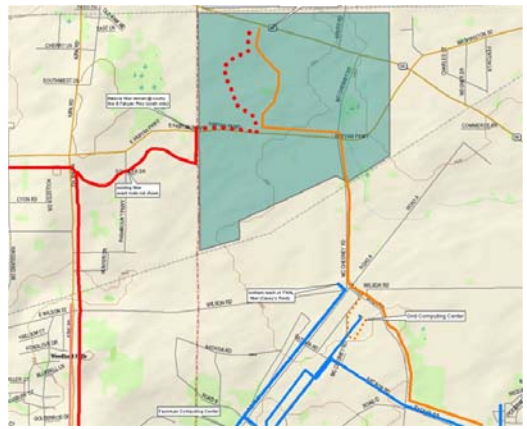

- GCC (red) & LCC (blue) spurs:
  - Plats completed on both spurs
    - GCC spur: 2300’ trench to intercept FCC/GCC duct
    - LCC spur: 1100’ trench in intercept existing duct
    - GCC/LCC fiber cost est.: $16k for fiber & installation
  - $1 req to Center Pt. for GCC spur submitted is in Contracts office

LHC-Support Projects

- LHC@FNAL Operations Center (FCC1):
  - Supported on two Cisco 3750 switches (48 ports)
  - Initial network configuration will be a routed subnet within the FNAL LAN.
  - Preserving the architecture (right) that would allow trusted access to CMS LANs at CERN
  - Deployment begins December 06
- LHCOPN monitoring:
  - Deployed 10GE NIC-equipped performance monitoring system to monitor LHCOPN path to CERN
  - Providing PerfSonar access to Fermi LightPath and internal CMS network counters
  - Strategic direction is to take a leadership position in deploying & developing LHC monitoring infrastructure

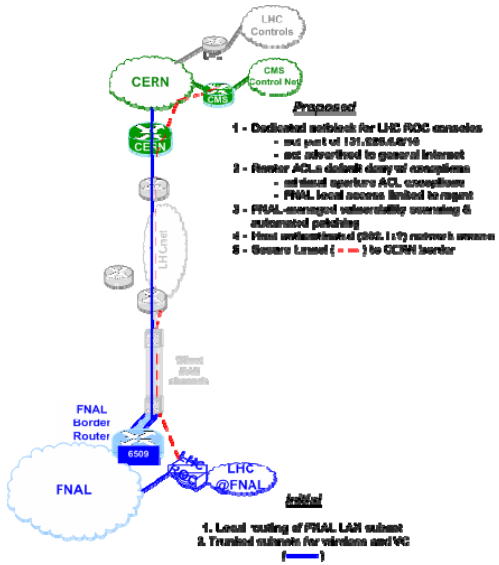

Miscellaneous WAN Activities

- BSS Disaster Recovery Cold Site at ANL:
  - Working with ESnet & ANL on BSS cold site networking
  - Supported by BSS w/ direct link to ESnet:
    - VPN tunnel used to back up data
    - Initially, a 100Mb/s path w/ upgrade to 1 Gb/s when MAN is deployed

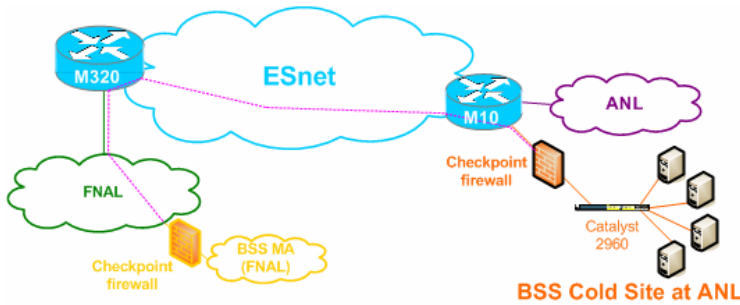

- DMZ LAN upgrade to 10GE:
  - 24-port 10GE Force10 will become new DMZ LAN
    - Satellite Force10 switch for 1GE connections
      - Product delivery delayed until mid-Oct
    - Connections for redundant ESnet IP links
    - Connection to StarLight router for rerouted high impact traffic

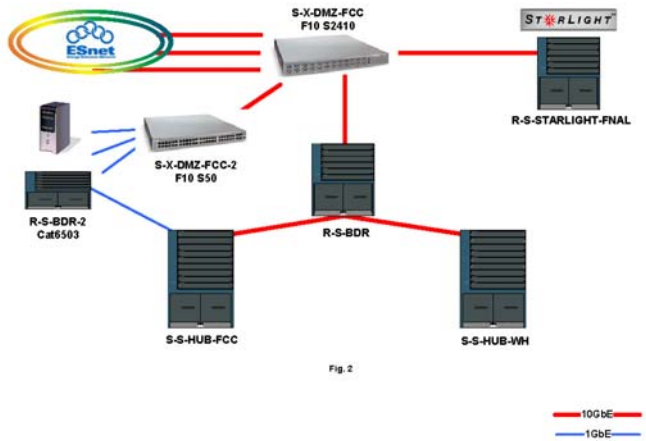

Network Security Infrastructure

Phil DeMar

Multi-level Security Access Zone Project

- Architecture is perimeter defense in depth:
  - Visitor LAN “outside” the border router
  - Green zone is current model of open access w/ exceptions
  - Yellow zone will be default deny on inbound initiated connections
- Alternate path for select high impact traffic
  - Not for general network access…
- Impact tool testing on CD LAN
- Mid sept. target for 2nd FWSM

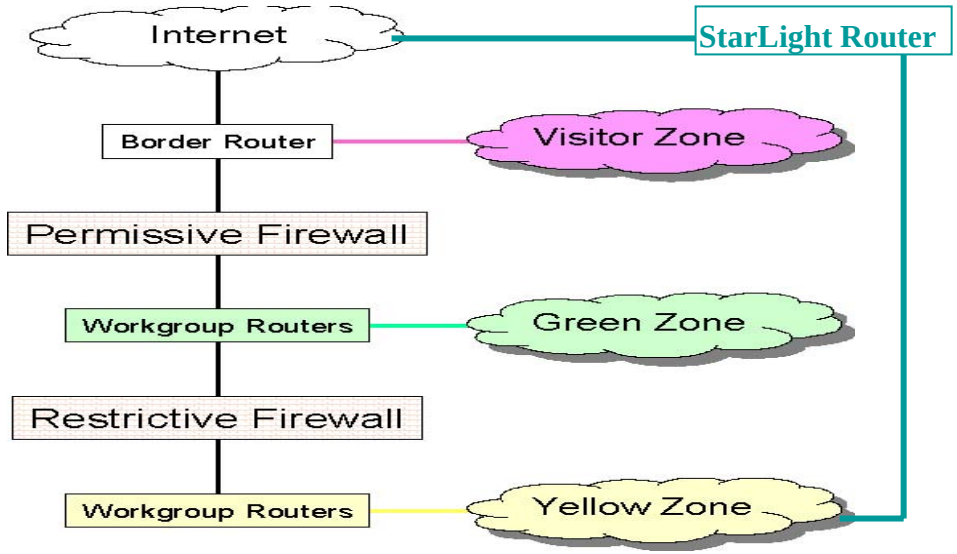

Transition Strategy

- Default is work-group LAN (macro-level) granularity:
  - Difficulty of sub-dividing a work group LAN varies
  - Shared LAN connections (ie., WH non-FTTD) are worst case
- Default zone is green, work group elects to change
  - Cutover involves moving work group uplink to yellow core router
- Server LANs provided on core routers
  - Dual-zone homing allowed for well-managed yellow zone systems
- Transition likely bumpy for first yellow zone work groups
  - Should become smoother as exception issues get ironed out

Status

- Schedule:
  - 2nd FWSM module due in Sept – will be used for Transition Tool
  - Documentation & Web pages in progress
  - Target using Transition Tool – Oct 1
- Transition Tool Status
  - Used for CD lan for last 5 months
  - List of proposed Exceptions created
  - User Web Page created
- Wireless yellow zone investigations
  - Protoype in FCC
  - Ongoing discussions on authentication options
  - Testing Wireless hardware for recommendations
- VPN for Yellow Zone
  - Portion of current VPN pool will be allowed access to Yellow Zone
  - Data Comm testing

Transition Tool - Process

- 1. Move Workgroup behind Transition Tool
  - Firewall NOT turned on
  - Minor downtime while uplink changed
  - Provide Documentation
- 2. Use Transition Tool to identify servers
  - Reports for Data Comm & System Admins to show candidates to move
  - Web Page for users & System Admins to show what WOULD have been blocked
- 3. Move Servers out
- 4. Review – Iterate Steps 2 & 3
- 5. Turn On Firewall
  - Minor downtime while firewall rules turned on

Transition Tool Default Exceptions

DESCRIPTION	SOURCE	SOURCE PORT	DESTINATION	DESTINATION PORT
ICMP	On Site	ICMP port list	any	ICMP port list
CST Scanners	Scanner List	any	any	any
Video Conf.	any	any	any	tcp h323
SSH to router	On Site	any	172.16.1.2	tcp ssh
SNMP	SNMP List	any	any	udp snmp
SMTP	SMTP List	any	any	tcp smtp
Print Server	charming	any	any	any, need to lock down more
AntiVirus	AV Server List	any	any	tcp 2967
Internet Printing Protocol	any	any	any	tcp ipp

Enira Close Blocking

- Enira Blocking Appliance
  - Manual blocks in production for 6 months
  - Writing interface between NIMI/Tissue & Enira for automated blocking
  - Policy Issues
    - AD device blocks
    - Blocks requested for devices no longer alive on the net
    - Multiple blocks for the same device – whack-a-mole
  - Technical Issues
    - Network devices that do not support MAC address blocking – service request for multiple black hole devices to vendor

Miscellaneous Network Computer Security Projects

- Node Verification Tool notifications implemented:
  - Checks for proper MAC & IP address registration for active systems
  - 8 to 10 issues found a day
  - Will implement reminder notices
  - Expected Blocking Parameters – once automated blocking implemented in NIMI –
    - Unregistered IP address – Immediate block
    - Non-registered, Illegal MAC IP pair – block after 3 working day

Miscellaneous Network Computer Security Projects

- Node Verification Tool notifications implemented:
  - Graph of Outstanding Issues:

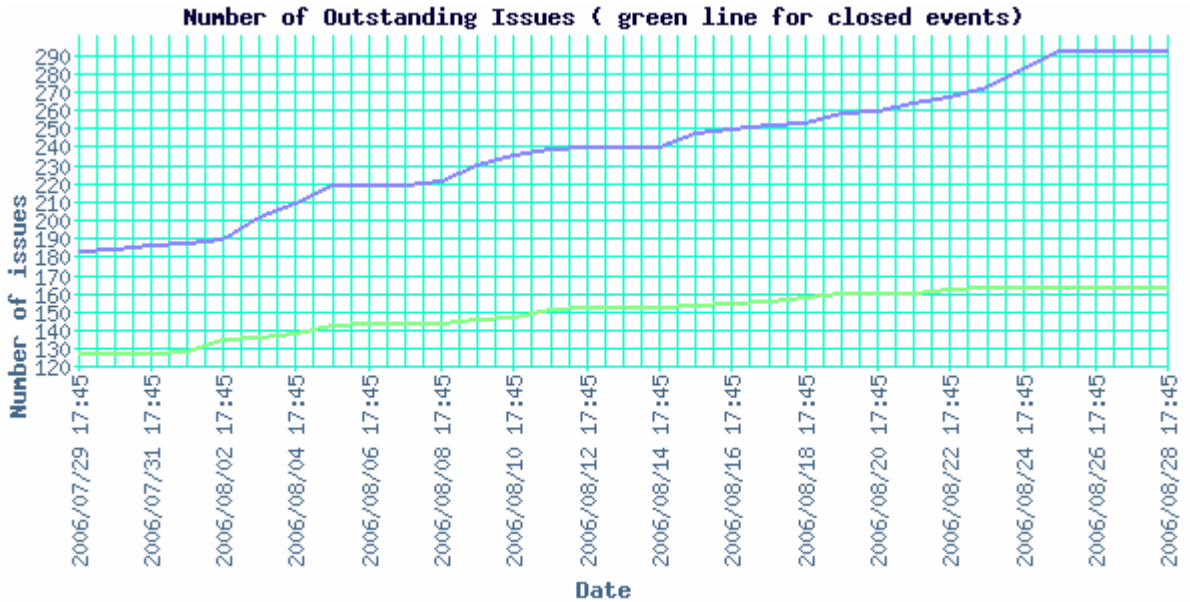

NIMI data collection – network data

- Router ARP tables
- Bridge Forwarding Tables
- VPN access logs
- DHCP address logs
- Network Address Block Assignment
- Autoblocker Events (in Development)
- Subnet Assignment (from MISCOMP – in Development)

ST& E efforts for the Network MA

- Checked compliance with 2nd Qtr 2006 ST&E requirements
  - Border Router ACLS
  - Nessus Scanning of Network Devices
  - Verification of Procedures
- Reviewing ST&E procedures for possible automation using Network Management tools
  - Ciscoworks for configs & compliance
  - Openview for uptime
- Expand Configuration Baselines to include additional NIST recommendations

Physical Infrastructure Projects

Phil DeMar

GCC

- GCC CR-A:
  - New CMS 6509 linked at 4x10GE, UTP fully patched.
  - D0 6509 to be installed in last available network rack space.
  - Unable to service all 77 Farm racks with present switches
    - D0 Catalyst 4500 taking network rack space of 6509; it should be replaced
- GCC Tape Robot:
  - 6509 w/ 10GE to FCC working
    - 2nd 10GE ready
    - UTP wired for 96 connections to Robot room
      - Currently half of the connections are PrivNet.
  - Provided and wired 20+ FiberChannel connections from Tape Robot to servers.

GCC

- GCC CR-B cabling infrastructure planning under way:
  - Propose DCN purchases and installs Farm patch cords:
    - would assist with fully loaded CR-B with racks from day-1 to benefit cooling.
    - would eliminate delays encountered by Integrators ordering non-bulk quantities
    - would streamline installations using Batch-Patch assemblies already deployed in CR-A
    - would realize volume discounts
  - Making contingency for additional 12 Farm rack spaces
  - Fiber duct needs to be put back into Network Rack plans.

LCC & FCC

- LCC/ILC
  - Concentrating network cabling & switches in 107 SE corner:
    - Three racks to be deployed, one each for UTP, fiber, & switches
    - Expect to co-locate all fiber into the facility
      - From Lab B (Cryo); XGAL, and eventually GCC
  - Working on UTP support for ILC offices & work areas
    - Will be supported on general facility network, not AD network
  - Installed switch for AMR move from FCC to room 107.
- FCC
  - Overhead cable tray scheduled to be installed in FCC2E to service network racks area

WH Fiber-to-the-Desktop (FTTD)

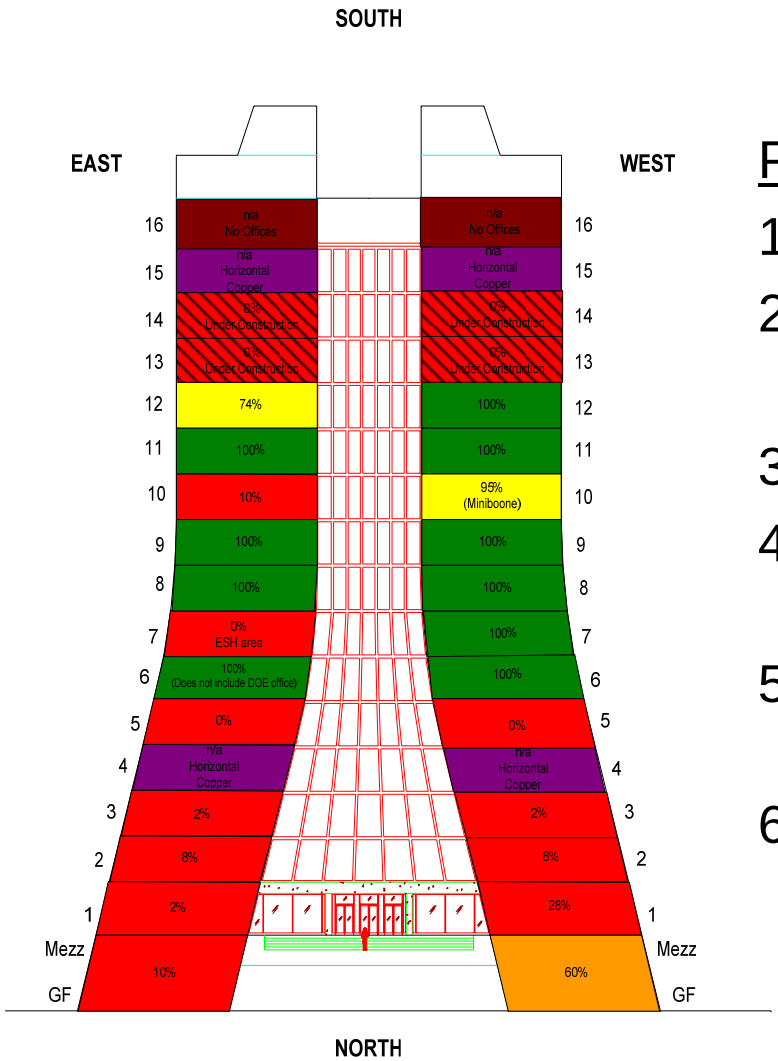

Priorities	Target
1 WH 13/14	Q3
2 Finish partially completed floors	Q4
3 Directorate	Q1 ’07
4 PPD Theory/Astro & Library WH3	Q1-Q2 ‘07
5 ESH WH7e & FESS WH 5e, BSS WH 5w.	Q2 ‘07
6 WHGF/MEZZ,WH1	Q2-Q3 ‘07

Miscellaneous Physical Infrastructure…

- Use of outside contractors (DTI):
  - Replacement effort for DCN Lambda Station work as DCI personnel assume DCN responsibilities
    - Working out well in year 2
      - Getting to know & understand Laboratory locations, procedures, & policies
      - Very effective with large scale projects (WLAN upgrades, GCC cabling, etc)
    - Current model is two DTI contractors for two weeks per month
    - Plan to add one DTI contractor for the other two weeks each month
- Cabling Infrastructure FY07 (& beyond…) Projects:
  - Village fiber: Lab 5 / Lab 8 and Lab 7 / OFS (residence DSL…)
  - FCC1/2 zone cable upgrades (more UTP; s/m fiber)
  - Lab E-B-A restoration
  - LCC to FCC fiber via Neutrino line

Obsolete Equipment Retirement Status

- Catalyst 5000's –
  - Replaced 2 – 5000's – Site 38 & Village this year
  - 2 5500's remaining – will be replaced as part of FTTD in WH
- Catalyst 2924's
  - Ordered 26 replacement switches for all Catalyst 2924's
- Router for TD
  - Ordered integrated switch/router to replace TD router
- Alteons (load balance devices) to be replaced by Cisco AceDirector
- Westell DSL gear – No clear replacement available – will use TeraBeam where feasible

Miscellaneous System Upgrades

- MRTG –
  - Upgraded 2 MRTG boxes
  - 2 more to update
- Name servers –
  - Have systems to upgrade primary and secondary name servers
  - Will update 5 more systems after this
- Monitoring Systems –
  - Moved CiscoWorks to dedicated platform
  - Plan to upgrade HP Openview monitoring system hardware
  - Plan to purchase additional analysis machines

Video Conferencing Support

Sheila Cisko

Video Conference General Operations

- New collaboration@FNAL web pages include
  - room descriptions with peripherals and scheduling information
  - ECS Ad-Hoc & VRVS guides
  - audio conferencing information
  - desktop H323 collaboration guides
  - troubleshooting guides and faq

http://computing.fnal.gov/videoconferencing/

Video Conference General Operations (cont'd)

- Created conference room-centric instructions
  - Provided binders for each room
    - include guides and instructions
    - information also on web
    - will send bi-annual emails to room stake holders for content accuracy

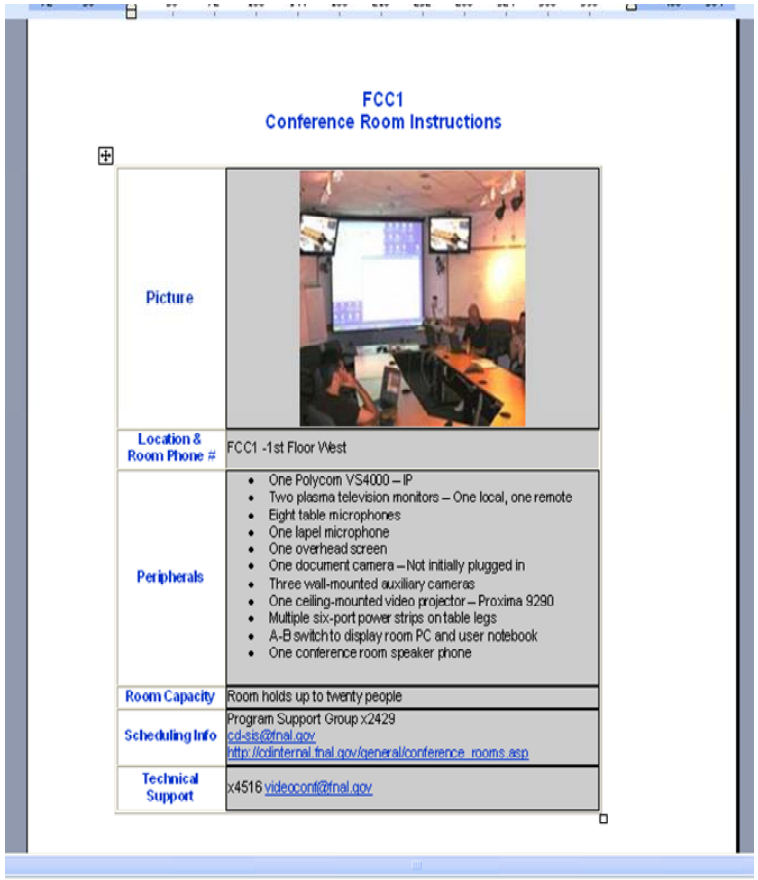

FCC1
Conference Room Instructions

Picture	
Location & Room Phone #	FCC1 -1st Floor West
Peripherals	• One Polycom VS4000 – IP • Two plasma television monitors – One local, one remote • Eight table microphones • One lapel microphone • One overhead screen • One document camera – Not initially plugged in • Three wall-mounted auxiliary cameras • One ceiling-mounted video projector – Proxima 9290 • Multiple six-port power strips on table legs • A-B switch to display room PC and user notebook • One conference room speaker phone
Room Capacity	Room holds up to twenty people
Scheduling Info	Program Support Group x2429 cd-sis@fnal.gov http://cdinternal.fnal.gov/general/conference_rooms.asp
Technical Support	x4516 videoconf@fnal.gov

Video Conference General Operations (cont'd)

- Created codec database
  - contains detailed device information
    - such as location, IP address, installation dates
  - tracks purchasing and maintenance history
  - contains ECS-provided statistics

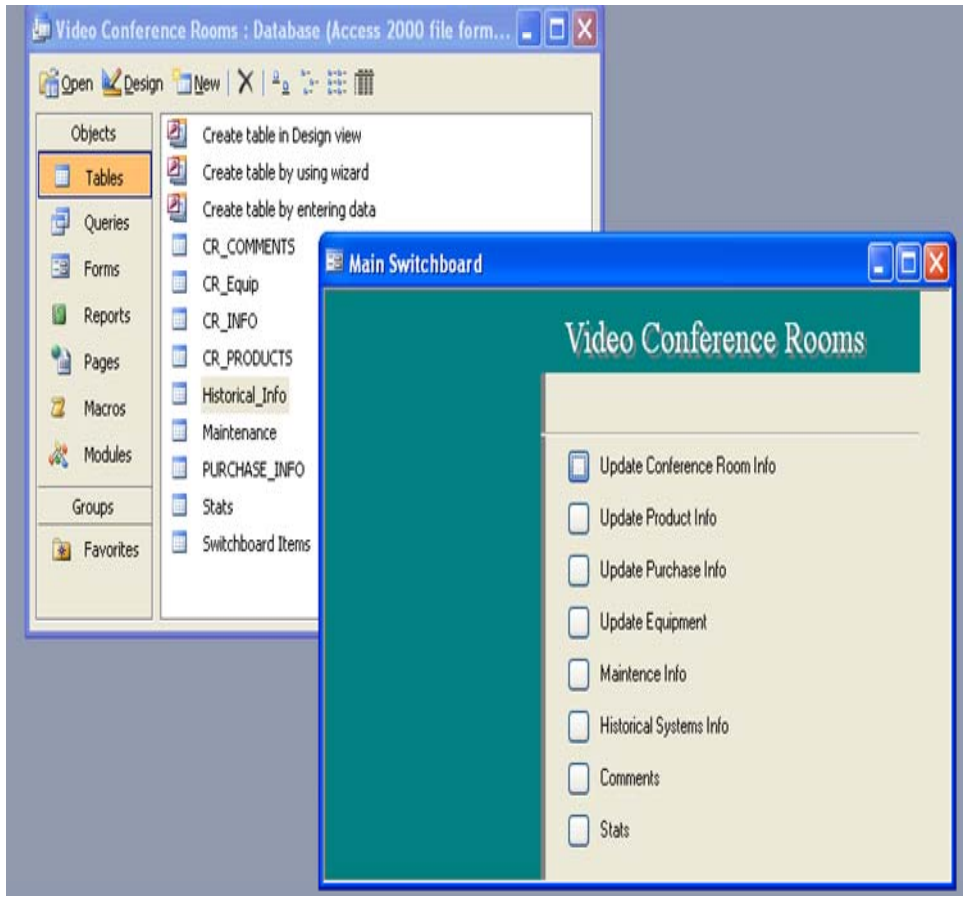

Video Conference General Operations (cont'd)

- ECS audio conference reservations are self-serve
  - Discussing details for Helpdesk to be user back-up if user is not registered
- Security scans causing operational problems with codecs
  - Ongoing investigations with CD/CST
    - Issues:
    - Polycom Global Management System loses contact with device but device may still be pingable and capable of placing calls
    - GMS sends alerts
      - when communication lost between device and gatekeeper
      - for unknown reason, presumably a reboot
    - no control or immediate oversight over non-cd systems
      - users are conditioned to reboot
    - two devices behind ACL always reachable thru GMS and by users

Other VC Activities

- Desktop investigations:
  - Starting MacIntosh video test system set-up
  - Continue urging Polycom to port PVX (desktop H323 software client) to Linux and Mac, thru PUG SIG participation and federal account rep.
  - FNAL-affiliated ECS registered desktop clients (Polycom ViaVideo, Polycom PVX, Gnomemeeting (linux), OhPhoneX (Mac): ~106
  - FNAL-sited ECS registered desktop clients: ~86
    - utilization of desktop clients unknown
- Planning the implementation of H239 (video & data) with new H323 equipment in on-site conference rooms
  - Four non-CD rooms are H239-enabled,
    - FCC2a upgrade and FCC1 upgrade for FY07
  - checking interoperational issues with ECS Ad-Hoc and VRVS
  - write user guides for rooms and web pages
  - advertise users of H239 capabilities
  - inform FNAL room stakeholders of the obsolescence of existing codecs.

Other VC Activities

- Chairing ESnet Remote Conference Work Group
  - Reported RCWG Update to ESCC in July
    - ongoing discussions regarding support for ESnet Ad-Hoc and VRVS
    - http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=1652
  - leading authors of the ECS Workshop 05 Report
    - major recommendations noted in ESCC presentation
- Continued discussions with users about WebEX and domestic & international audio conferencing
  - concerns about proprietary solutions
  - cost
  - FNAL service vs HENP
  - international calls

Videoconferencing Room Appointment

- Pending room projects
  - WH2NW – Black Hole
    - Directorate asked CD to duplicate Snake Pit with features:
      - Table installed, with power & data
      - Enhanced audio, ceiling and table microphones
      - dedicated video projector
      - H239 (video & data) capability
    - majority of a/v equipment and materials on-site
    - installation by subcontractor
  - WH13X – Fish Tank
    - Requested by GDE/ILC
      - video system with Polycom telephone/microphone solution
      - LCD panel for video
      - dedicated video projector
      - H239 (video & data) capability
    - waiting for equipment delivery
    - installation by subcontractor

Videoconferencing Room Appointment (cont'd)

- FCC2a
  - New furniture installed 8/28
    - conference room table with power outlets
    - perimeter tables and hospitality cart
    - Adding additional table microphones
    - LCD panel replaces v/c'g video projector
    - H239 (video & data) capability
  - installation by subcontractor

- WH1E (LHC@FNAL)
  - Planning team – E. Gottschalk, A. Thomas, S. Fry, S. Cisko
    - Video displays ordered and codec on-site
    - projector and audio gear purchase requisitions to be processed by end of FY06
  - FESS to complete electrical & misc. room preparation prior to installation by subcontractor

Video Conferencing Monitoring & Metrics

- Monitoring & metrics project
  - Continuing Polycom management system configuration
    - purchased new license to cover on-site endpoints
    - parsing features such as device traces, system alerts
    - evaluating Tandberg Management Suite solution
  - Creating metrics for FNAL endpoints
    - using PERL scripts to retrieve Polycom Call Detail Reports for
      - utilization time
      - outbound/inbound call information (duration, to/from which site(s))
      - protocols used (H323 transport, G.711 audio)
      - troubleshooting (wrong speed, transport)
    - researching SNMP walks for
      - system uptime
      - # of bytes sent and received
    - investigating optimal management and automation of retrieved data

ESnet Collaboration Services (ECS)

- Discussions began in response to rumor that ECS services would cease at the end of 2007
- Bill Johnston maintains that the rumor's were false:
  - ESnet will refresh current infrastructure hardware now
  - ESnet will budget for full refresh of that hardware in four-year cycles
  - if community wants to significantly expand services then will have to do so by subscription or petition to Program Managers

Wide Area Systems Projects

Matt Crawford

Wide-Area Network R & D

- Research projects are guided by current and anticipated needs of the scientific program.
  - Gross network throughput; troubleshooting; optimization.
  - Investigating emergent kernel behaviors affecting network and computational performance.
  - Tuning system parameters for real workloads.
  - Dynamic allocation of high-performance WAN paths.
  - Pure optical switching and network reconfiguration.
  - Using advanced host capabilities to improve performance in a scalable and deployable way.

Throughput, Troubleshooting, Optimization

- Wide Area Working Group (WAWG) - a forum for investigating & solving WAN performance problems
- Usual symptom is data transfer rates lower than expected
  - Problem causes and solutions vary:
    - host parameter tuning
    - packet loss in network provider's gear
    - application design
    - buffer space in intermediate network devices
    - non-network bottlenecks (disk or memory)
- Meets bi-weekly by video conference:
  - works by email constantly
  - Participation extends beyond FNAL & its external users
- Lead: Demar, Crawford, and many others.

Computational vs. Network-Intensive Tasks

- On a system with compute-bound processes, arriving packets may not be processed by TCP for hundreds of milliseconds due to kernel locks on packet queues.
  - Run-time selectable code fix moving toward inclusion in SL.
- Frequent sleeps & wakes are typical of receiving data from the network. They can overstimulate Linux's interactivity detection – and a *slower* stream more so than a faster one!
  - Code fix in final testing.
- Lead: Wu

Kernel Tuning for Storage Systems

- Problem: Achieving CMS service challenge target rates requires too many concurrent file transfers. Applying the standard formula for high bandwidth $\times$ delay led to system crashes.
  - Analyzing application behavior and studying kernel logic led to many insights apparently not known to the Linux community. Seminar to be given 8/31.
- Lead: Bowden

Dynamic WAN Path Allocation

- Lambda Station project, reported more extensively elsewhere.
- Current state:
  - Lambda Station server prototype in perl fully functional.
  - Re-implementation in Java on Apache Axis/jClarens Web Services platform partially complete.
  - Client calls integrated with dCache/SRM.
  - Integrated dCache being deployed on Caltech CMS T2.
- Lead: Crawford

Pure Optical (“Photonic”) Switching

- The forefront of on-demand network paths for scientific applications is based on switching connections among fibers with micro-electro-mechanical mirrors.
  - Referred to as “photonic switching” largely because the word “optical” was taken by SONET.
- Worldwide interest and activities coordinated through GLIF (or gλif) - Global Lambda Integrated Facility.
- Our interest: reservable clear-channel paths to other sites
  - avoid congestion & effects of limited buffering in routers & switches
- Status: Polatis photonic switch now beginning testing.
  - Will test on-site and deploy at StarLight.
- Lead: Bowden

Advanced Host Capabilities

- Multihomed Linux systems with standard routing tables suffer some reachability failures for incoming connections.
  - Policy-based first-hop selection solves those failures
    - Published to linux-users and DocDB.
- A file transfer sender will commonly send large bursts of large packets. If there's a lower-speed link in the path, network devices in the path will drop packets as buffer space is exhausted. Total sending rate dwindles to a fraction of the link capacity.
  - Solution: use advanced queuing capabilities already in Linux kernel to shape traffic to certain destinations.
- Lead: Bowden, Crawford