
Networks

Project Status Briefing

May 9, 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 Facility Network

□ Increase Backbone 10G Connectivity

- Currently 10Gb
 - from Core switch to CMS and
 - between FCC Core switches
 - To border router
- Add 10Gb capabilities (this year)
 - Core Network to Wilson Hall
 - Core Network to CDF
 - Core Network to D0

□ Add Additional Ports for FCC Computer Room Connections

- Additional 6509 for FCC2
 - Migrate FNALU to new FCC2 switch
-

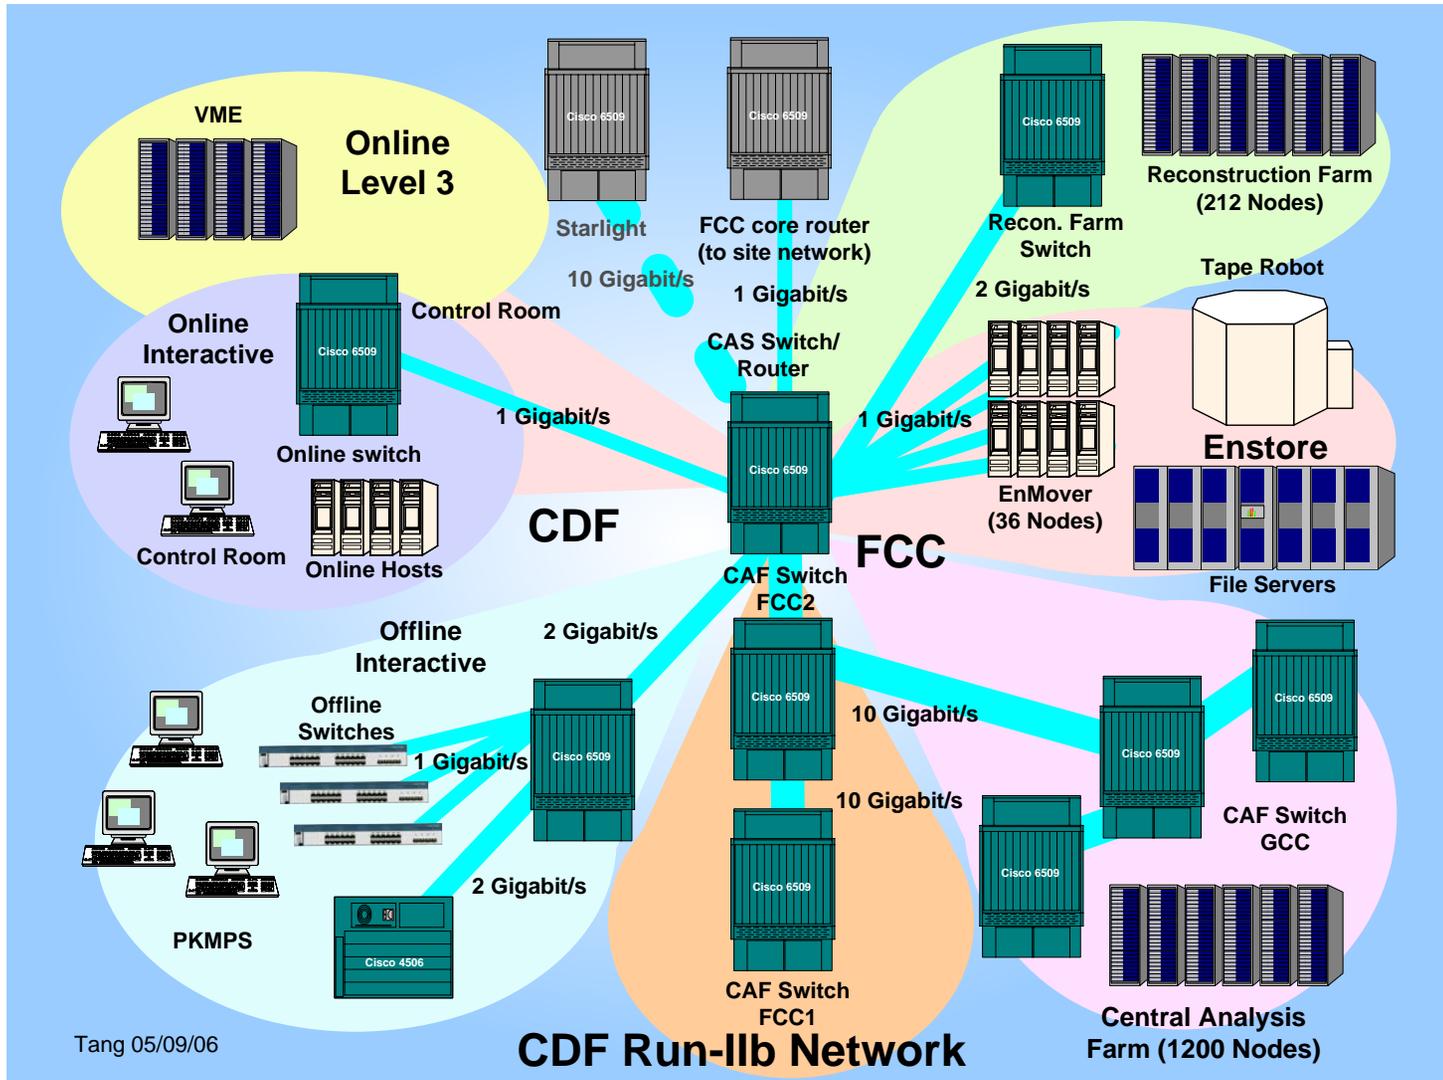
CDF

- Upgrades Completed This Year
 - FCC upgrades:
 - Installed 4000W P/S and 10Gb/s modules in CAS and CAF switches
 - Grid Computing Center:
 - Installed new 6509 switch
 - Installed 4000W P/S and 10Gb/s Modules
 - On-Line Upgrade:
 - Sup720Gb/s, 1000BaseTX ports, 4000W P/S
-

CDF

- ❑ Planned Upgrades
- ❑ Upgrades to CDF Trailers Wireless:
 - Replace 802.11b APs with 802.11g
- ❑ Upgrade Starlight link to 10 gigabits/sec
- ❑ Network infrastructure for FY'06 farms upgrades:
 - 10Gb and Copper Gb/s module upgrade for Recon. Farm switch
 - Copper Gb/s module upgrade for CAF FCC1 6509 switch
- ❑ Upgrades to CDF Off-line network facilities in PK168K:
 - Switch fabric upgrades to Sup720 (720Gb/s)
 - Expanded 1000B-T support
 - 10 Gigabits/sec to FCC
- ❑ Upgrades to CDF network facilities in Building 327:
 - New 4506 switch with 1000B-T support

CDF



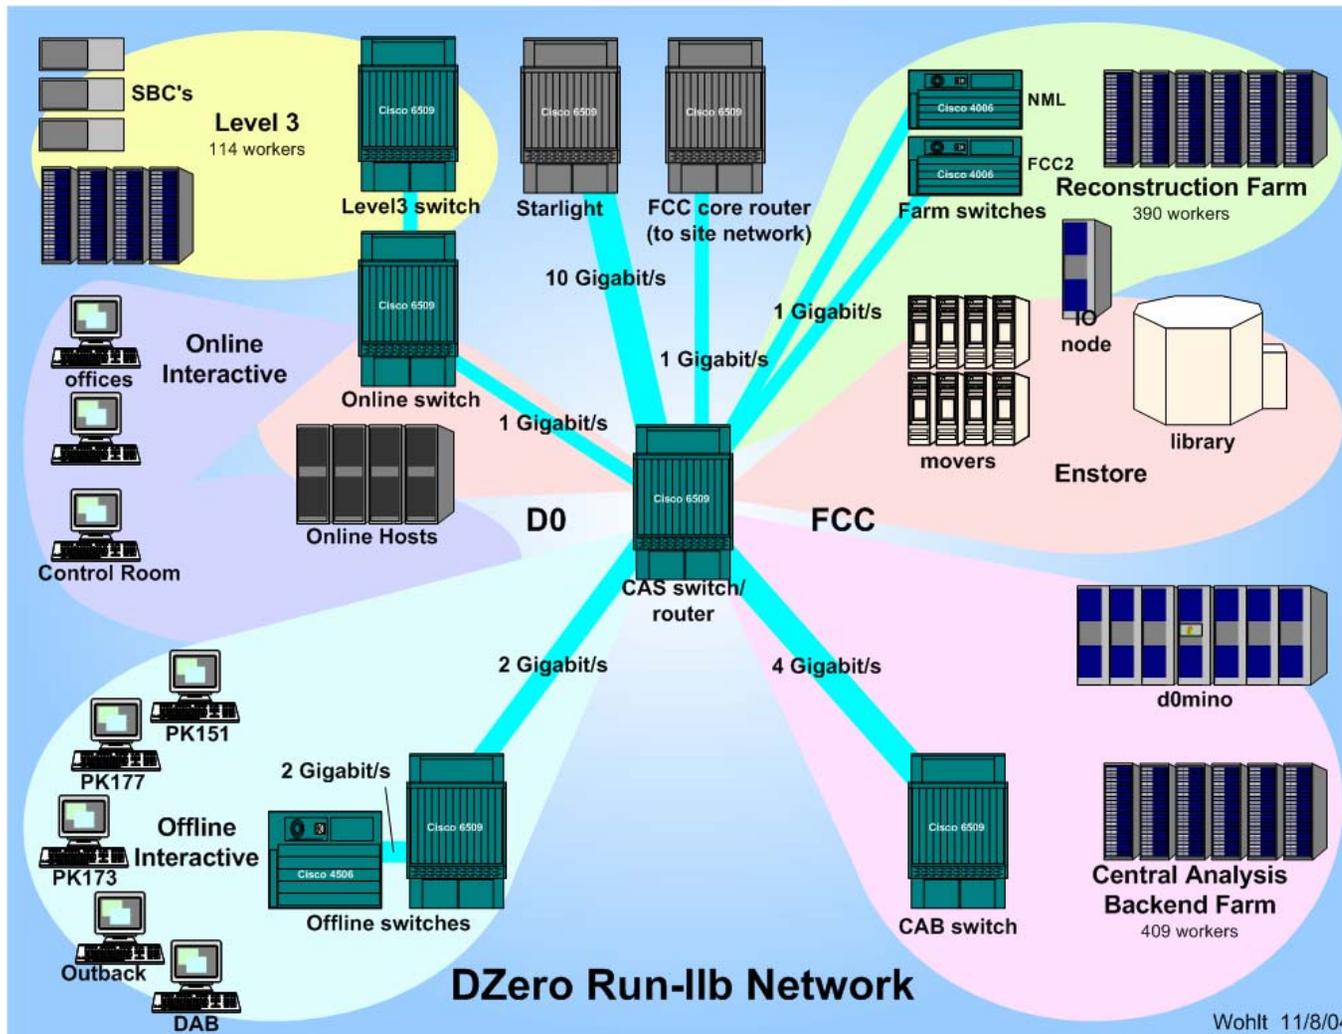
■ Upgrades Completed This Year

- Added 2 - 48 port Gb modules to FCC1 switch
- Added 1 – 48 port Gb module to DAB Offline switch
- Added 1 – 48 port Gb module to DAB online switch – upgraded supervisor module
- Replaced non-managed Hubs in Moving Counting House with managed supported switches
- Upgraded all Wireless APs to new module
- Added 1 Gb Starlight connection

■ Planned Upgrades

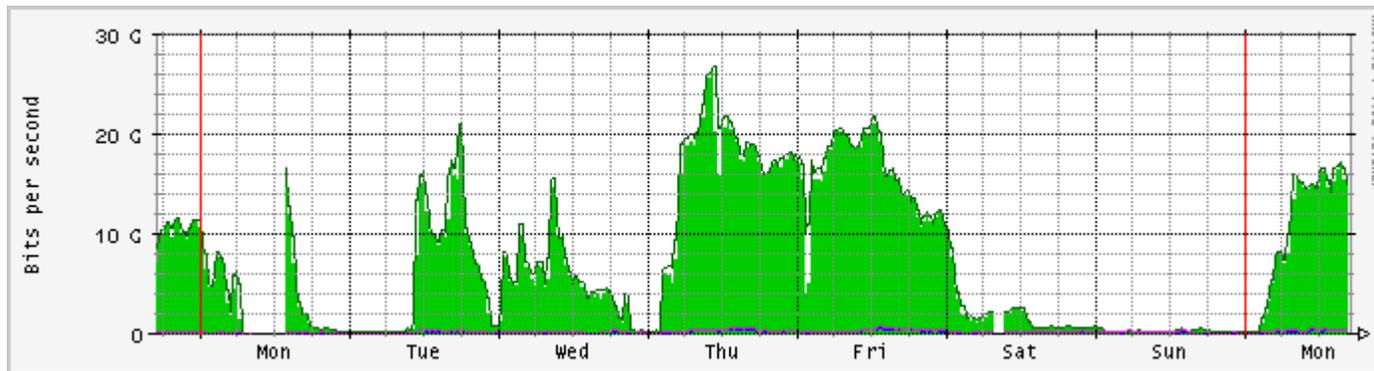
- Upgrade to 10Gb connection to Network Core
- Upgrade to 10Gb connection to Starlight
- Add additional 10Gb connection to from FCC to GCC
- Review Wireless Coverage
- Install 6509 (currently in storage) in Computer Room B at GCC

D0



■ Upgrades Completed This Year

- Installed additional 6509 at GCC
- FCC 6509
 - Installed 10Gb module
 - Installed 2 additional 48 port 10/100/100 modules
- Upgraded connection to Network Core to 10Gb
- Upgraded connection between FCC to GCC to 40Gb

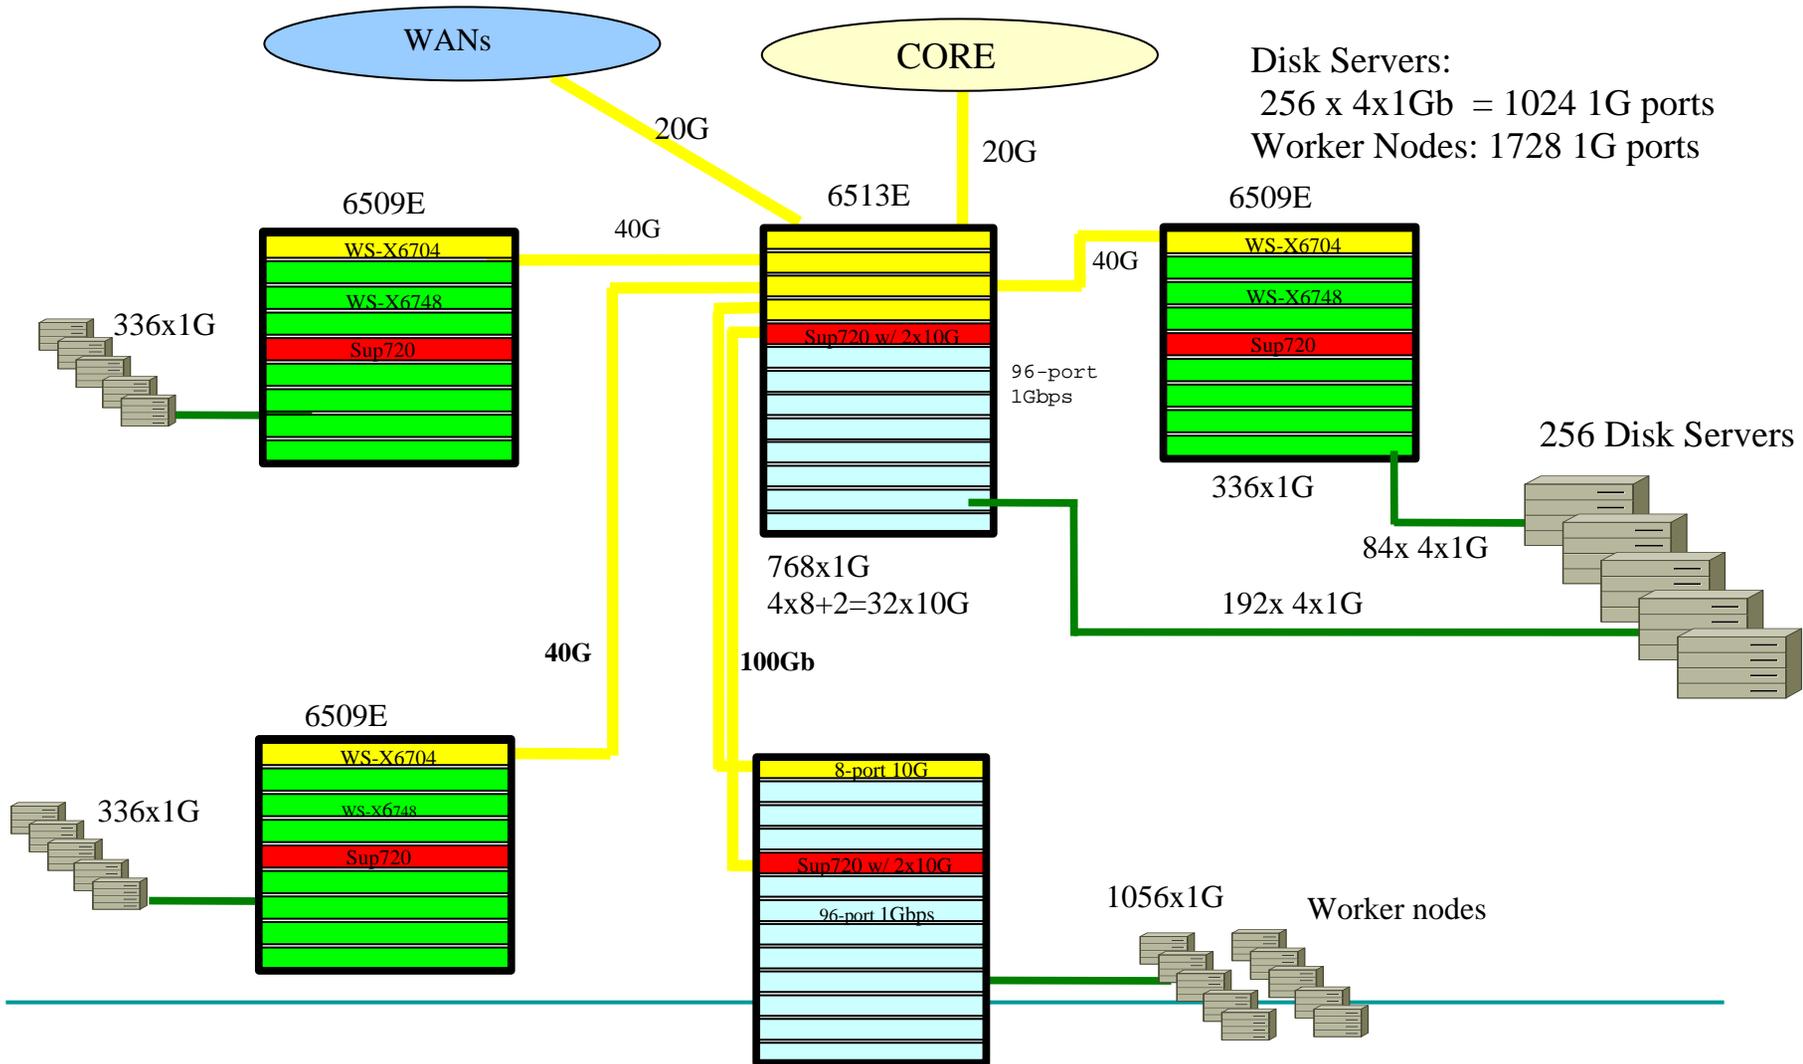


■ Planned Upgrades

□ CMS FCC Core switch

- Add additional 10Gb module
 - Add 3 additional 48 port modules to FCC Core switch
 - Upgrade Power Supplies in FCC Core Switch
-

CMS



Wireless

■ Upgrades –

- Forklift upgrade of all older models – Includes Wilson Hall, CDF, D0, Fixed Target, TD, etc. – about 120 APs to be upgraded
 - New Models -
 - Support 802.11g (54 Mbps radio speed)
 - Support dynamic radio adjustment and network self-healing
 - Support multiple wireless zones (yellow, visitor, green ?) on one AP
 - Support encryption, authentication
 - Add units to increase and/or improve coverage
-

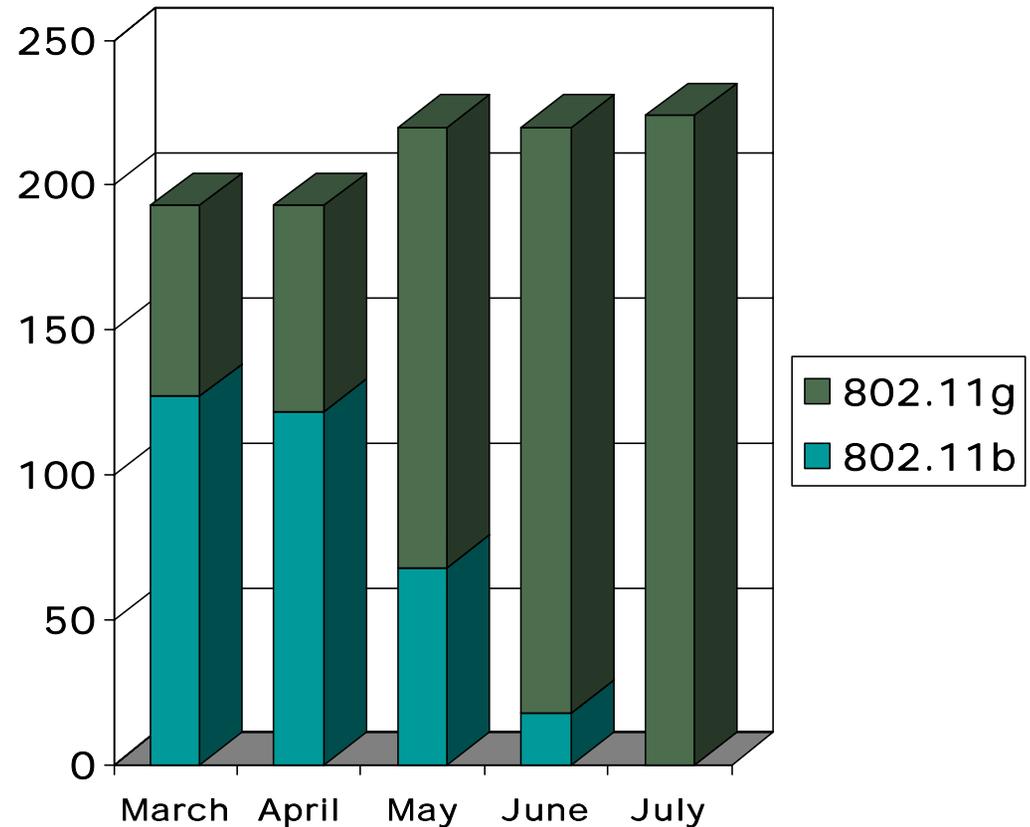
Wireless

Upgrade Status –

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

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

Projected Schedule all APs replaced by end of July



Wireless

TeraBeam

- Point to multi-point wireless bridging solution
 - Intended for infrastructure links
 - No end users (laptops)
 - Two base units atop Wilson Hall
 - Replace current ADSL links
 - Site 52, Site 29, etc.
 - Provide service to new areas
 - Site 56, pump houses, etc.
 - Provide for emergency link restoration
 - One base unit at Lab 8 in the Village
 - Provide service to village houses that currently have no networking
-

Wireless

Futures –

- Implementing encryption and authentication – WPA2
 - Implementing multiple wireless zones per AP
 - Exploring new management software – Cisco LWAP and/or Airwave
 - Enhanced rogue detection and location identification – follow MAC address across wireless lan, monitor signal strength - ~~7~~8 rogues were found by the WLSE in the last year
-

Miscellaneous

■ FAPL Facility

- Test Facility for High Bandwidth Use - installed on FCC1 - Cluster of ~50 1Gb NIC machines, four 10Gb NIC machines
 - Used for SC2005
 - Vanderbilt Machines
 - WAWG testing
 - Networking = 6509 w/ 24 10GE ports (donated by Cisco)

■ DNS Migration to Unix

- In progress, initial machines are here and under test
 - Some concerns about versions of Linux that are supported by vendor

■ MINOS - Soudan Mine Upgrade

- Router upgraded to supported hardware
 - Internal Links Upgraded
 - External link upgraded to DS3
-

Miscellaneous

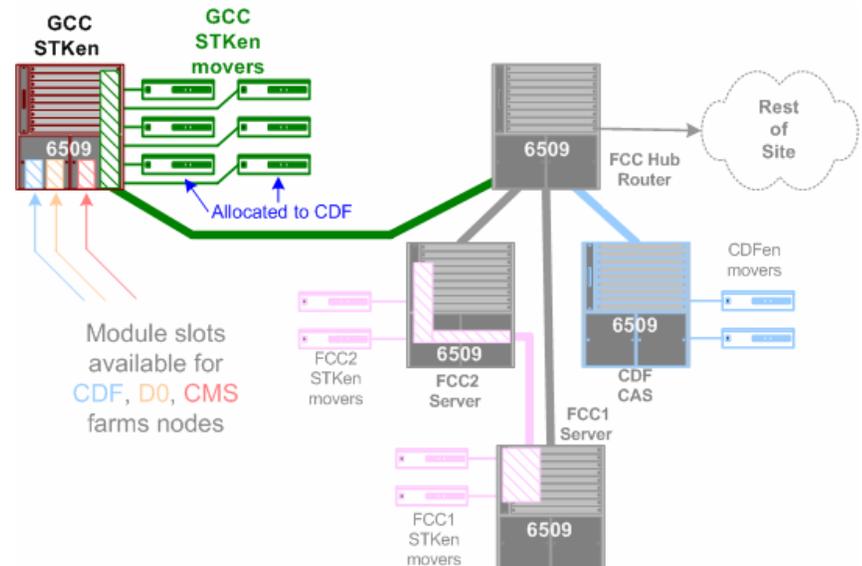
■ GCC Tape Robot Network

□ New Subnet Created For Enstore at GCC:

- Supported on new tape robot 6509 at GCC
- 10Gb connection to Core Network – additional 10 Gb connections can be easily added

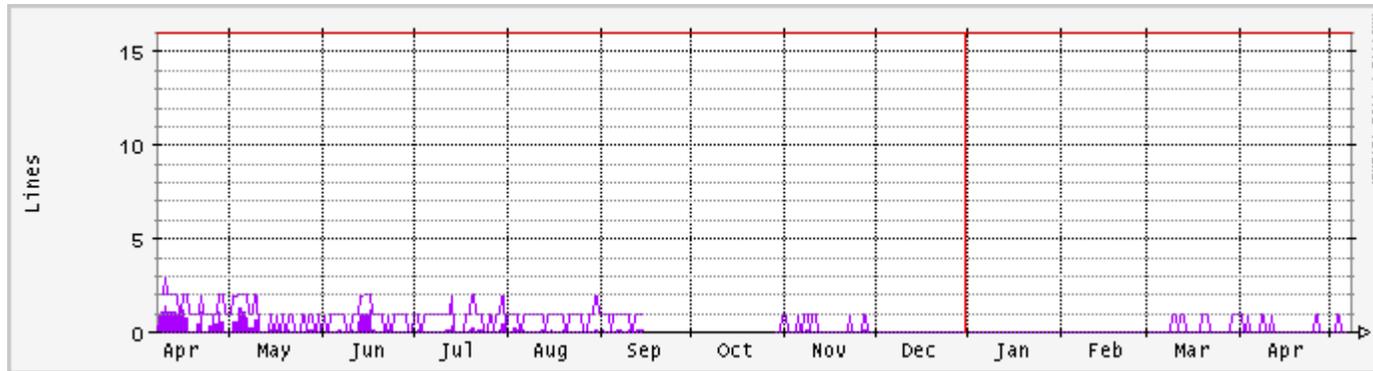
□ Private Robot Controls LAN will be implemented

- Includes all nodes needed to control the Robots in FCC & GCC



Miscellaneous

- Dial-in phase out
 - Down to 4 users in the last year
 - Propose to steer those users to commercial service – July 1 ?

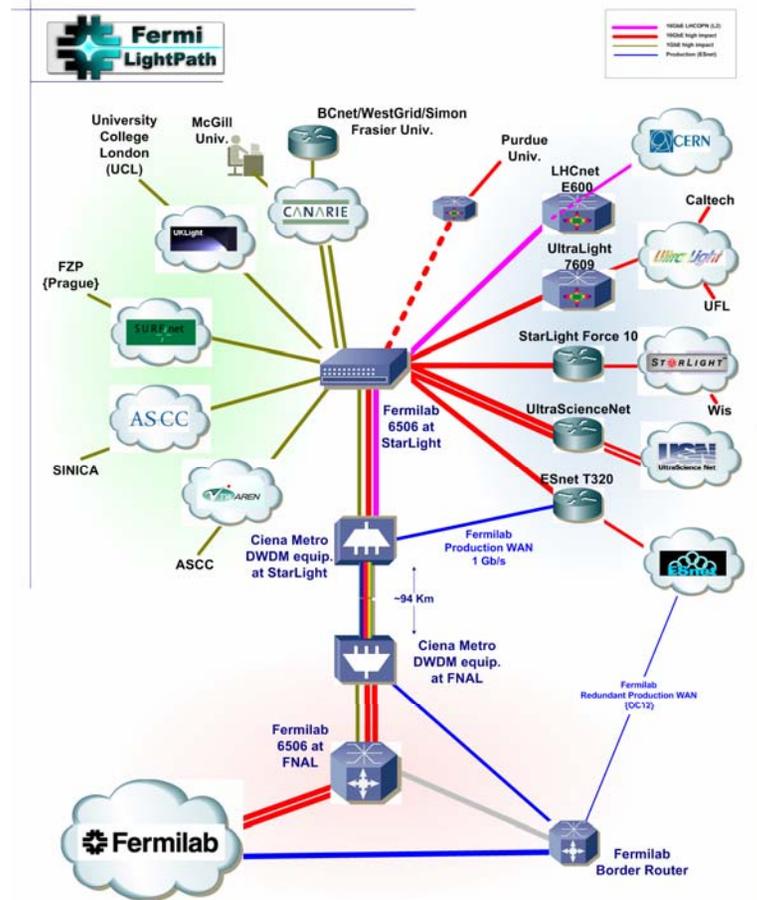


WAN Projects & Upgrade Plans

Phil DeMar

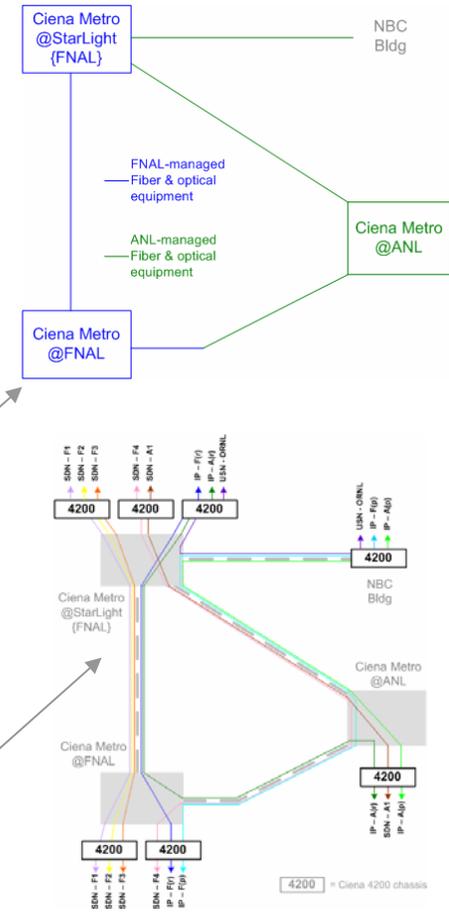
Fermi LightPath Status

- Current channel configuration:
 - 10GE link dedicated to LHCOPN
 - 10GE link for designated CMS Tier-2 high impact traffic
 - 1 GE for production (ESnet) traffic
 - OC12 now serves a redundant role
 - 1 GE for CDF & D0 overflow traffic
- Switch connections at StarLight:
 - 10 GE: ESnet, LHCnet (CERN), USN (2), UltraLight, StarLight, Purdue
 - 1GE: CAnet (2), UKlight, Surfnet, ASnet, Twaren
- Also supporting NIU (2x1GE) connectivity to StarLight



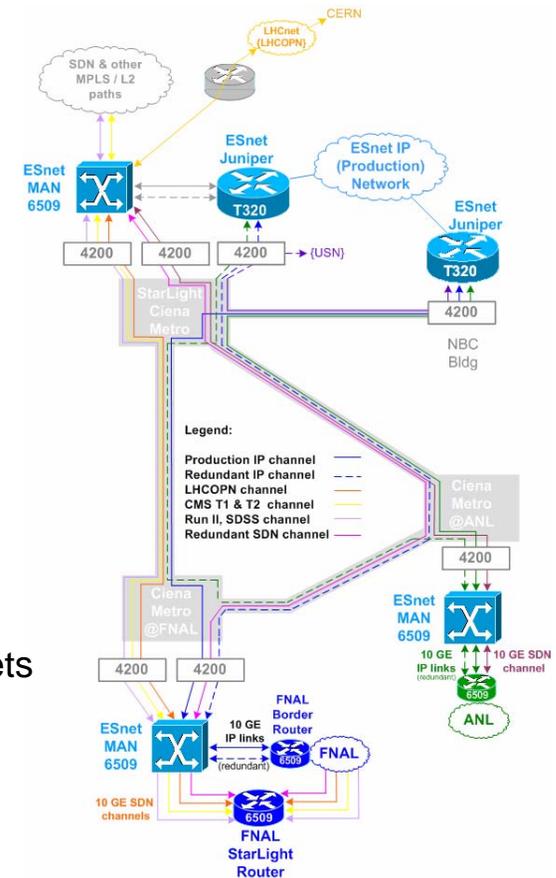
ESnet MAN Deployment

- Off-Site connectivity migrating to very high bandwidth MAN, based on optical network (DWDM) technology:
 - Provides 10 GE channel(s) for general off-site network connectivity
 - Supports separate data (SDN) channels for high impact data movement
- Three-layer architecture:
 - Bottom layer is existing Fermi LightPath & ANL I-Wire infrastructure
 - Augmented with fiber between FNAL & ANL
 - Managed independently by FNAL & ANL
 - Middle layer is new DWDM equipment (Ciena 4200 chassis) that provides MAN channels
 - Managed jointly by FNAL & ANL



ESnet MAN Deployment (cont)

- Three-layer architecture (cont):
 - Top layer is ESnet layer 2 & 3 equipment that provides service back to the Labs
 - Managed by ESnet
- Initial MAN configuration for FNAL:
 - Redundant 10GE production network links
 - To different ESnet Chicago area PoPs
 - With potential to load distribute for ~20Gb/s
 - Four (maybe only three) 10GE SDN channels:
 - Utilized for designated high impact traffic
 - Policy routed, based on source & destination subnets
 - Redundant SDN channel for failover
- Existing Fermi LightPath 10GE & 1 GE channels are not part of the MAN:
 - Will be available for wide-area R & D projects



ESnet MAN Deployment (cont)

■ Deployment schedule:

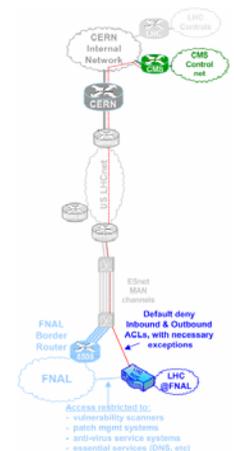
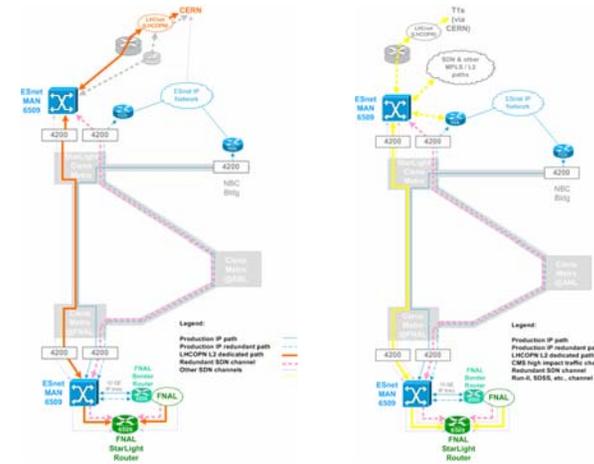
- ANL / FNAL fiber connection expected in place by July
- DWDM equipment under procurement, also expected in by July
- Some of the new DWDM equipment (between FNAL & StarLight) already in place, left over from SC2005 product loan
 - Expect to work with ESnet to utilize this equipment as soon as PO for it is placed

■ Post-deployment issues:

- Procedures for joint operational support of MAN infrastructure by FNAL & ANL need to be worked out
 - Model for ongoing costs & maintenance not in place yet
 - Contingency plan needed for core MAN equipment failure
-

LHC-Related Support Efforts

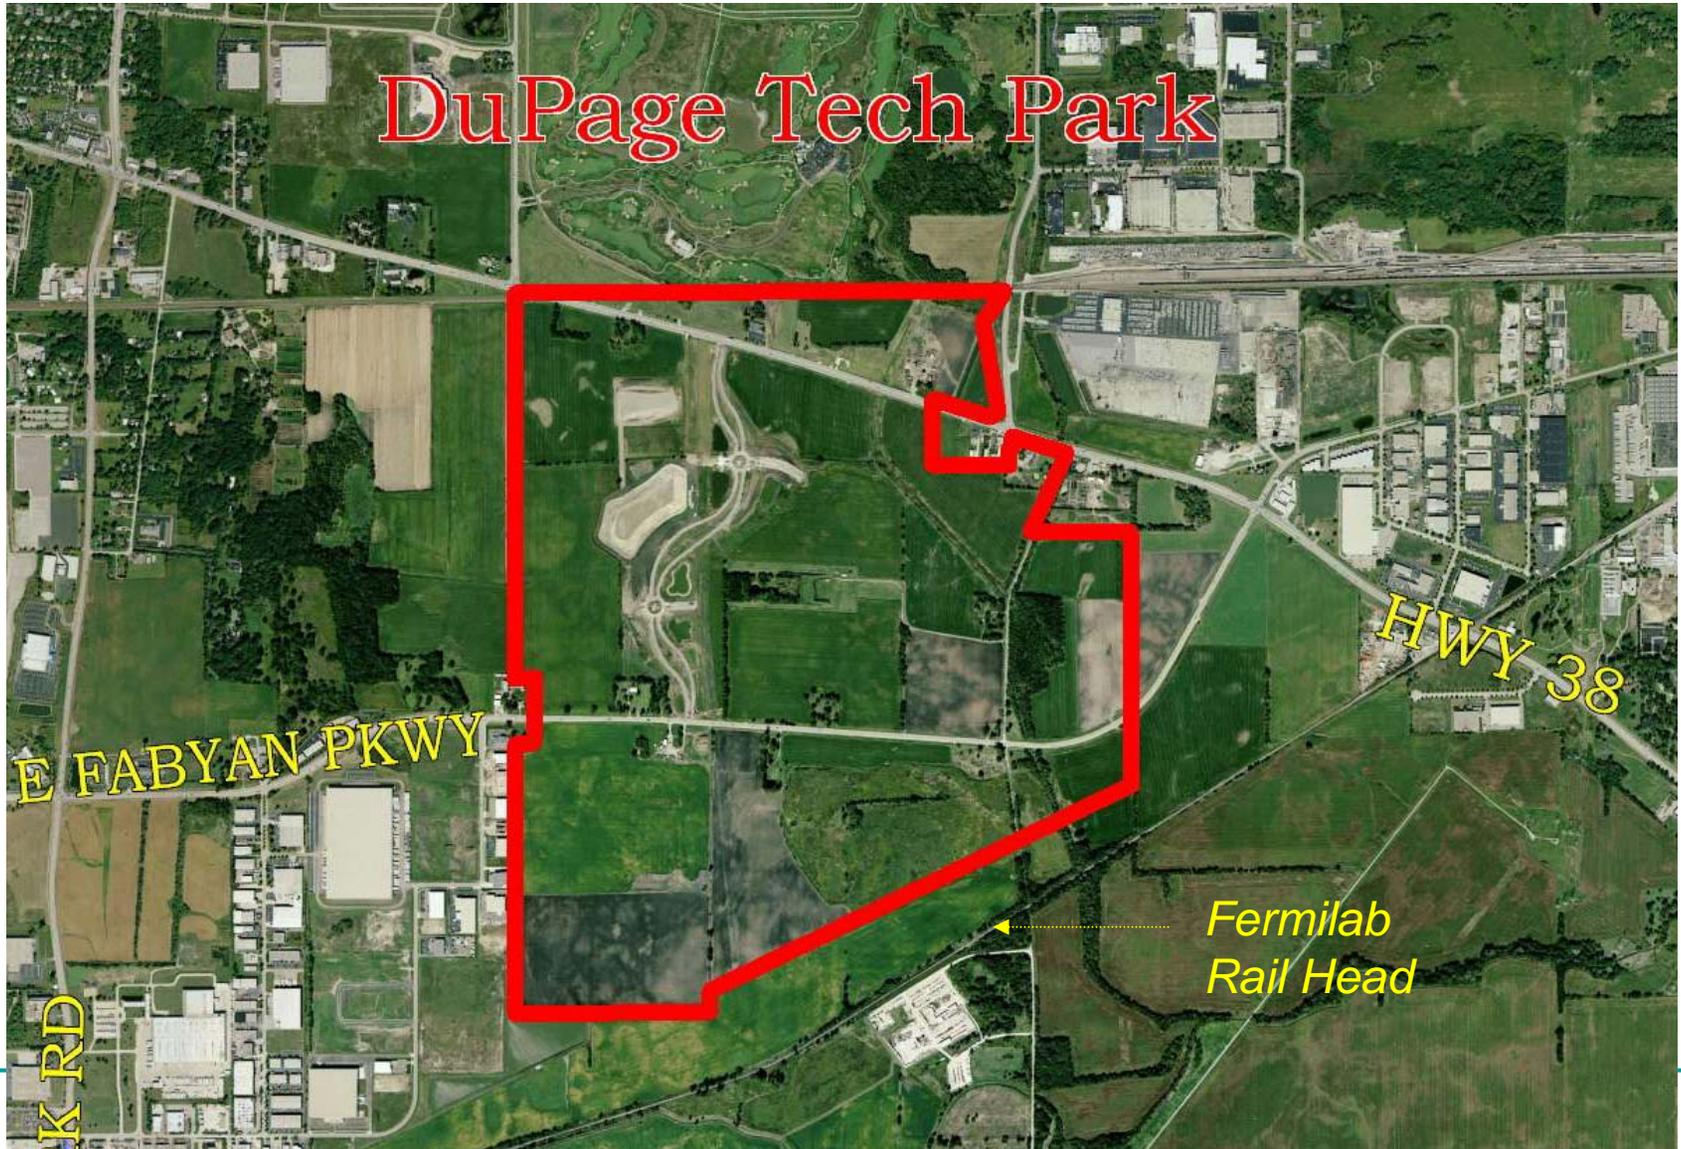
- LHCOPN – Logical, dedicated network supporting LHC T0/T1 data movement
 - Will deploy 10GE layer 2 path (orange; right) between FNAL & CERN for LHCOPN
- CMS T1/T2 & T1/T1 support:
 - Dedicated 10GE channel (yellow; far right) for designated CMS high impact traffic
- LHCOPN Monitoring Efforts:
 - Deploying a 10GE NIC monitoring system supporting SNMP & flow data access to our LHCOPN infrastructure
- LHC@FNAL Remote Operations Center
 - Engineering “private” network for LHC@FNAL consoles
 - Objective: trust relationship w/ CERN on console access to CMS control network resources



DuPage National Technology Park (DNTP)

- Technology-based business park adjacent to FNAL's north boundary
 - County land owned by the DuPage Airport Authority, +34 M\$ in state money +Hastert's support = DNTP
 - High bandwidth connectivity to R&E network community is a major selling point for the Park
 - Center Point Properties is now the Park's "master developer"
 - CenterPoint has contracted with AboveNet to provide optical fibers from the Park to Starlight and to a commercial PoP in Chicago
 - Abovenet fiber terminating on FNAL property is easiest path to the Park
 - CenterPoint would provide duct path to tie FNAL (GCC) into the Park
 - FNAL would provide & install fiber between GCC and the Park
 - Fermilab has been asked to consider managing the Parks' R&E PoP
-

DNTP Location



The Path

- **Yellow** path is the requested easement
- **Red** path is inner duct build-out to GCC



Implications & Status of DNTP Project

- Status:
 - Abovenet expects to submit final easement request to DOE by 6/1
 - The DNTP fiber light-up date is still targeted for 9/1
 - Implications & Potential Opportunities for FNAL:
 - DNTP fiber offers an alternate, non-MAN path to StarLight
 - Opportunities to explore off-site co-location computing facility space
 - Potential for high bandwidth collaboration w/ future Park tenants
 - Cooperative possibilities to lead a regional metro network consortium
 - Using non-FNAL facilities...
 - 80% of a duct path from GCC to LCC (New Muon)
 - Issues for FNAL:
 - Procurement & installation of fiber to Park PoP
 - Design, implementation, & support effort for Park R&D WAN facilities
-

Miscellaneous Operational WAN Efforts

- DMZ LAN upgrade to 10GE:
 - 10GE-based ESnet MAN makes existing 1GE DMZ LAN obsolete
 - Plan to implement “low cost” 10GE DMZ with non-modular Force10 switch:
 - Planned deployment in line with ESnet MAN (but not necessary...)
 - Open question of what to do with 1GE DMZ-connected systems
 - Probable satellite 1GE switch w/ 10GE uplink

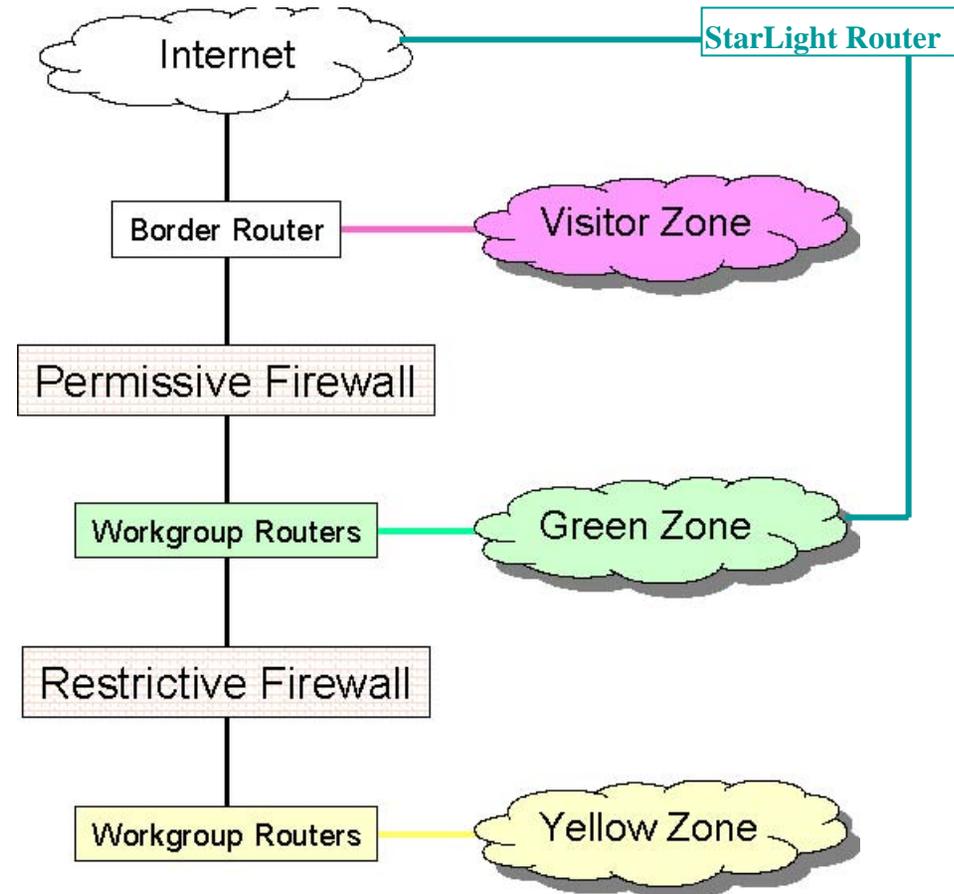
 - Cold Site Support:
 - Supporting a remote cold site for ANL’s disaster recovery planning
 - Planning to support a reciprocal cold site at ANL for BSS’s disaster recovery plans.
 - Functionally a protected, private network w/ secure tunnel back to FNAL
-

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...



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:
 - Infrastructure & migration procedures to be in place by 6/1/06
 - “Some” work group LANs transitioned to protected zone by 12/1/06
 - Basic network infrastructure in place:
 - Firewall & yellow zone router (FCC)
 - But no work groups are behind it
 - Needs to be expanded to include additional firewall & router (WH)
 - Migration procedures based on having transition tool to reveal consequences of migration to yellow zone
 - Transition tool to facilitate migration being worked on
 - Wireless yellow zone investigations & implementation proceeding independently
-

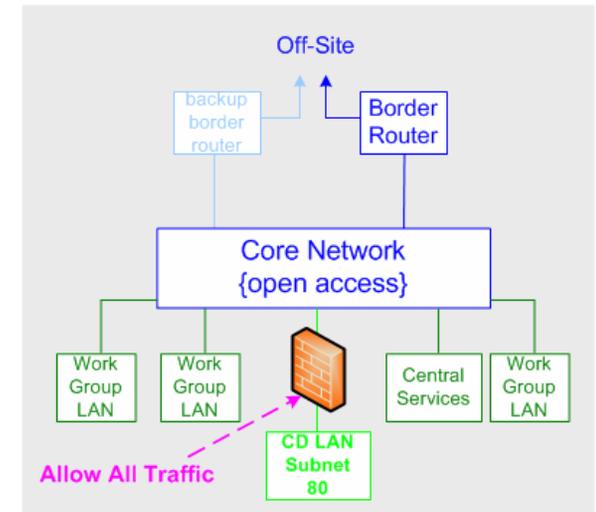
Transition Tool Development

■ Concept:

- ❑ Run work group traffic through firewall configured to allow everything
- ❑ Process firewall logs to determine impact of converting to default deny inbound configuration
- ❑ Identify & mitigate obvious problems
- ❑ Provide iterative process to identify other inbound connection patterns

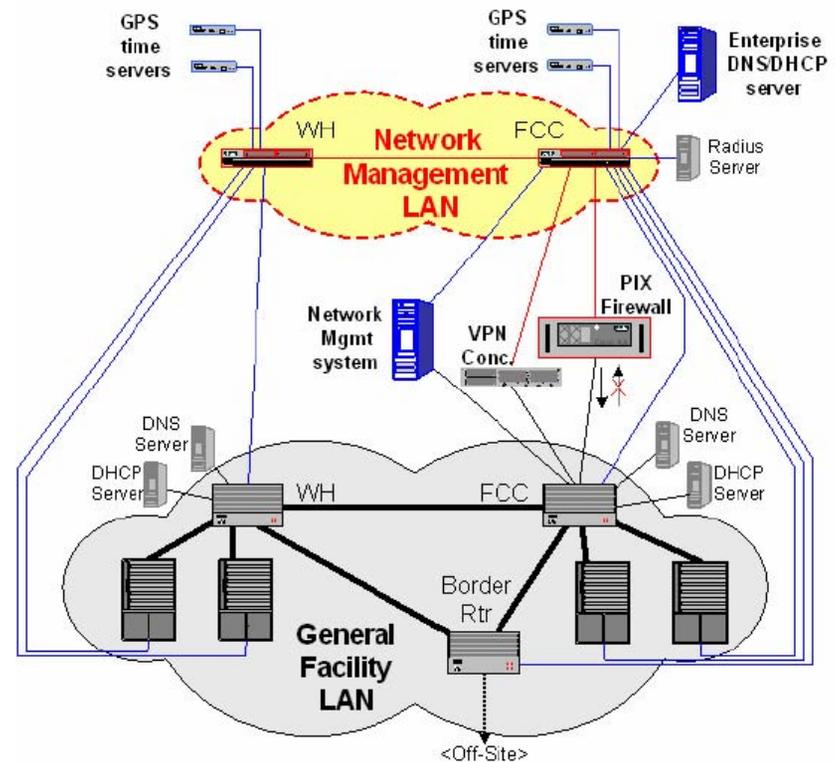
■ Status

- ❑ Testing the analysis tool on CD LAN now
 - Just testing the tool, not taking steps to migrate CD to yellow zone
- ❑ Development is proceeding, but its still a reasonably clumsy tool



Network Management LAN Project

- Network management LAN protects core network devices & servers in FCC & WH
 - But other network devices remain directly attached to general facility LAN
 - Will be expending effort to migrate all network devices to protected network mgmt vLAN
 - Servers for central network services will remain on the general facility LAN



Enira Close Blocking

■ Enira Blocking Appliance

- Installed on Network Management LAN
 - All Production Network Devices accessible
 - Except CDF/D0 Major Application Devices
 - “Close” blocks – MAC address blocks at local port
 - Layer 2 Blocks only
 - Testing on Production Connected Devices
 - Illegal IP address users
 - CD NIMI requested Blocks
 - “Test” Users
 - Working with Vendor on Issues
 - Multiple MAC addresses sharing the same switch port have been a continual problem
-

Miscellaneous Network Computer Security Projects

- Computer Security Protection Plan Certification & Auditing
 - Network System Major Application risk, security, contingency plans finalized
 - Significant effort needed on implementing & ongoing support of ST&E
 - Node Verification Tool in final stages of development:
 - Checks for proper MAC & IP address registration for active systems
 - In initial deployment, will send out notification on discrepancies
 - Expect to employ automated blocking of unregistered or improperly registered systems at some point:
 - How rapidly to block with a given circumstance still under consideration
 - Auto-Blocker related activities:
 - Application recognition capability has been built into blocking algorithm
 - Should alleviate (not eliminate) blocking of legitimate applications (BitTorrent)
 - Skype work-around to auto-blocker appears to work w/ Windows
 - Still investigating MAC/OS & Linux
-

Physical Infrastructure Projects

Phil DeMar

GCC & LCC

■ GCC CR-A:

- Original installation:
 - 72 racks - ~2k/rack (\$50/system)
 - Utilized jack-cord technology developed in-house (Orlando...)
 - Cat 6 cable; capable of 10Gb/s at 55m
- Completed wiring for five additional racks (77 total)
- New CMS 6509 ready for new farms buy, partially patched
 - Space for one or two more 6500s, to finish out racks in room

■ GCC Tape Robot:

- Minimal wiring required
- Overhead cable trays in place to accommodate UTP & fiber needs
 - Fiber capability to stream tape drive data back to movers in FCC
 - 144 pair of s/m fiber in place between GCC & FCC

GCC & LCC (cont)

■ GCC CR-B:

- Planning for cabling infrastructure under way
 - 84 racks - estimating at \$2.7k/rack (\$67/system)
 - Will utilized jack-cord technology again
 - Cat 6A (augmented...) cable; capable of 10Gb/s at 100m
 - But the cost of Cat 6 cable has gone up too...
 - Anticipating needing to support an additional 12 racks (total=96)
 - Modest cost (~\$5k) fiber infrastructure planned between CR-B & CR-A

■ LCC/ILC

- No defined requirements yet for computer room floor cabling
- Working on UTP support for ILC offices & work areas
 - Will be supported on general facility network, not AD network
- Working on options for additional fiber out to New Muon
 - Preference would be extension of FCC/GCC fiber via DNTP ducting

WH Fiber-to-the-Desktop (FTTD)

- FTTD is replacement cabling model for shared WH media
 - Needed to facilitate more granular yellow/green zone designations
 - Currently there are 550 FTTD office drops in WH
 - FTTD conversion is ~55% completed building-wide
 - 87 pillars remain w/ 10 or 10/100 mb/s hubs
 - This represents approximately 450 office drops
 - Excludes UTP infrastructure on WH 4/5W (BSS) & WH 15 (LSS)
 - Equipment in hand or under procurement for two additional WH floor FTTD upgrades
 - Proposing WH13 and WH14 for the next series of upgrades
 - PPD, AD and ILC occupy these floors.
-

Status of FTTD by Floor

<u>FLOOR</u>	<u>EST Drops</u>	<u>% Complete</u>	<u>Occupant(s)/Remarks</u>
<u>WHGF/Mezz</u>	65	35%	Machine shop, Mail Cage, Xerox Room, FESS/T&M, D1
<u>WH 1</u>	10	15%	Public Information, Users office
<u>WH 2</u>	60	0%	Directorate
<u>WH 3</u>	80	0%	Astro, Theory, Library
<u>WH 4</u>	0	n/a	Horizontal Copper (BSS)
<u>WH 5</u>	20	½=0%; ½=n/a	FESS, Internal Audit (subnet 19) East
<u>WH 6</u>	0	100%	DOE area not included
<u>WH 7</u>	30	60%	ESH on 7E.
<u>WH 8</u>	0	100%	Floor Complete
<u>WH 9</u>	0	100%	Floor Complete
<u>WH 10</u>	30	50%	PPD, East side. West (Miniboone, PPD) Complete
<u>WH 11</u>	0	100%	Floor Complete
<u>WH 12</u>	6	85%	North Xover, PPD
<u>WH 13</u>	75	0%	PPD, AD (subnet 18) Many 10 Mb hubs.
<u>WH 14</u>	70	0%	PPD, AD (subnet 18) Many 10 Mb Hubs.
<u>WH 15</u>	0	n/a	Horizontal Copper (LSS)

Proposal to Complete FTTD in Wilson Hall (Excludes 2 Floors in Pipeline)

Total Pillars w/Hubs->	58		
Total FTTD Drops->	301		
Fiber to Desktop Model			
<u>Est. Cost - Office End-></u>	301	@	\$105,350
<u>New Switch Blades-></u>	5	@	\$47,500
<u>Fibrmax Cards-></u>	51	@	\$11,220
<u>Fiber Fanouts-></u>	51	@	\$6,120
<u>Cisco 4506 Chassis-></u>	1	@	\$3,600
<u>Cisco Supervisor Mod-></u>	1	@	\$9,500
<u>Est. Total Cost FTTD-></u>			\$183,290

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 Wish List Projects:
 - Village fiber: Lab 5 / Lab 8 and Lab 7 / OFS (residence DSL...)
 - FCC1/2 zone cable upgrades (more UTP; s/m fiber)
 - Last segment of inter-building coax (PS6-PAB) replaced
-

Video Conferencing Support

Sheila Cisko

Video Conference General Operations

- Video resource scheduling no longer required, as of 2/1/04
 - Audio conferences continue to require scheduling
 - Most users obtain own accounts to self schedule
 - Should this be moved to Helpdesk?
 - Security scans causing operational problems with codecs
 - Users are learning to reboot
 - Investigating security issues with Polycom, CD/CST, et al
 - Improving documentation for users is a high priority
 - VC@FNAL web page under development.
 - Provide room-centric instructions (on web also)
 - Write self-help documentation
 - Continued need for back-up support
 - emphasis on technical rather than operational
-

Videoconferencing Room Appointment

- Currently 29 video conferencing rooms
 - Increasing at rate of ~4 rooms/year
 - Also doing audio enhancements - this year: 2 for CDF, 2 for D0 enhanced and generic installations, several mobile units.
 - Pending room projects
 - WH2NW (DIR), WH13X (GDE/ILC), FCC2a, WH1E (LHC@FNAL)
 - Efforts to outsource enhanced installations, starting FY06, have been modestly successful.
 - Installation for audio upgrades were done by PPD/CDF or D0 under CD/DCI guidance
 - Pending projects likely to require more of subcontractor's installation services
-

Miscellaneous Other Activities

- Desktop investigations:
 - Expanding to include Mac & Linux desktop demo capabilities
 - Discussions with Polycom to port PVX to Linux and Mac
 - Chairing ESnet Remote Conference Work Group
 - Involved in technology investigations, troubleshooting, guiding ECS, interacting with VRVS, vendors, et al
 - WH1E Conference Room for LHC@FNAL Ops Center
 - CCF consulting with LHC@FNAL committee and FESS
 - Room layout and equipment placement has been suggested
 - Gathering cost estimates for equipment and installation
 - assuming DIR provides funds, not in CD/VC budget
 - FESS considers WH1E separate project from LHC@FNAL
-

Video Conferencing Monitoring & Metrics

■ Monitoring & metrics

- Installed Polycom Global Management System
 - lacks usable accounting information
 - not stable with security scans.
 - Evaluating Tandberg's management system as an alternative
 - Generally, comprehensive monitoring tools are not available for endpoints or MCU service provider infrastructure
 - Intention is to create metrics for local endpoints
 - Currently relying upon ECS and VRVS for utilization statistics
-

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.
 - Addressing the antagonistic interaction between computational and network-intensive tasks – Resilient dCache on compute elements.
 - Dynamic configuration & 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 a cast of thousands.

Computational vs. Network-Intensive Tasks

- Linux kernel was made preemptible in 2.5 - 2.6.
 - Locking was added to protect network socket state when TCP receive processing is preempted
 - Division lines for processing of incoming packets are poorly set, so that on a system with compute-bound processes, arriving packets may not be processed by TCP for hundreds of milliseconds
 - Papers in CHEP06 and in DocDB
 - SciDAC funding applied for.
 - Lead: Wu
-

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.
 - Need to enlist more testing and deployment sites!
 - 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: photonic switch on order, due this month
 - Will test on-site and deploy at StarLight.
 - Lead: Bowden
-

Advanced Host Capabilities

- Previous solved problem –
 - Multihomed Linux systems with standard routing tables suffered from some reachability failures for incoming connections.
 - Policy-based first-hop selection solves those failures
 - Published to linux-users and DocDB.
 - Current problem –
 - A host on the sending side of a file transfer will commonly send large bursts of large packets. If there's a lower-speed link in the path, or if several hosts are sending concurrently, network devices in the path will drop packets as buffer space is exhausted. Total sending rate dwindles to a fraction of the link capacity.
 - Approach: use advanced queuing capabilities already in Linux kernel to shape traffic to certain destinations.
 - Lead: Bowden, Crawford
-