Campus Research Network Overview

Chris Griffin
Chief Network Architect
University of Florida & Florida LambdaRail

5/6/2013

Agenda

- Research Networking at UF
  - A brief history
  - CRNv2
- Florida LambdaRail
  - What is it?
  - How it provides for UFL's Research Needs
- Internet2 Innovation Platform
  - Design
  - Requirements
  - Use cases

A Brief History @UF

- Historically, all research was done over the production UF network.
  - Late 90s: ATM, then 100Mbps Ethernet backbone in a ring topology.
  - Early 2000's: Gigabit Ethernet backbone in a star topology.
  - Today: Moving toward ubiquitous dual 10G backbone links over the next 6 months.
- Security was not as big of an issue.
  - Public IP use was high.
  - No central firewall.
  - Some network access control in use.
  - IDS was able to handle campus traffic.

A Brief History @UF (cont)

- External R&E focused connectivity was limited, but stayed ahead of demand.
  - OC-3 to Internet2 (vBNS backbone.)
  - OC-12 to Internet2 Abilene Backbone (some of this was carved off for Internet.)
  - Much was still likely going over somewhat congested Internet paths.

Need for a change

- Aggregation of research computing resources.
  - HPC Center
  - CMS Tier 2
- Increased focus on moving data around campus and externally
  - LHC (CMS)
  - Ultralight
- New regional and national network resources emerging
  - NLR
  - FLR

Need for a change (cont)

- Increase in expected data rates exceed what typical security devices could handle.
  - Do we slow the research or allow unchecked traffic deep into the campus production network?
- NSF MRI grant provided for the first "Campus Research Network" (CRN).
  - CASTOR: A High-Performance Communication and Storage Backbone for Data-Intensive Scientific and Engineering Computing
- Planning began in 2003. Grant submitted 2004. First light was in 2005.

Campus Research Network v1

- CRNv1 provided for
  - 2*10Gbps research oriented network backbone to CISE, Physics, and Larsen Hall.
  - 1*10Gbps connectivity to another location in Larsen, as well as Benton.
  - 1*10Gbps dedicated research connection to FLR via a 10Gbps wave to Jacksonville, FL.
  - 1*10Gbps connectivity to the campus production network.
- CRN was connected "off to the side" of the production network, and outside of the security boundary.
  - This is now called a "Science DMZ"

CRNv1 Topology - 2006

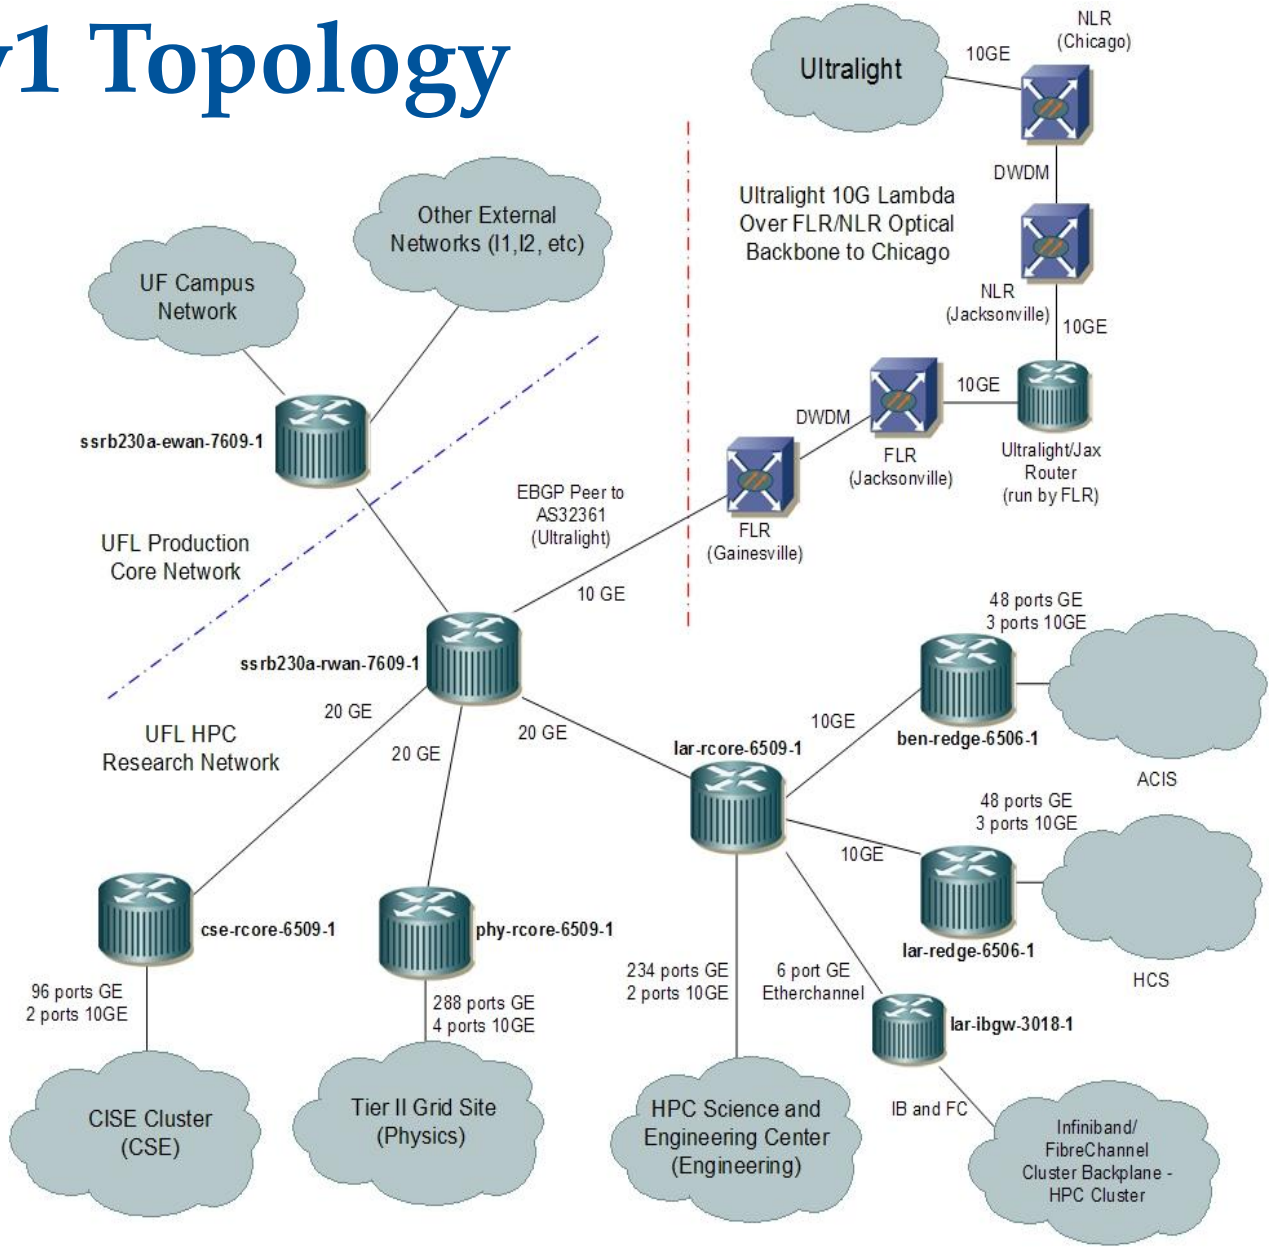

CRNv1 Topology – Late 2012

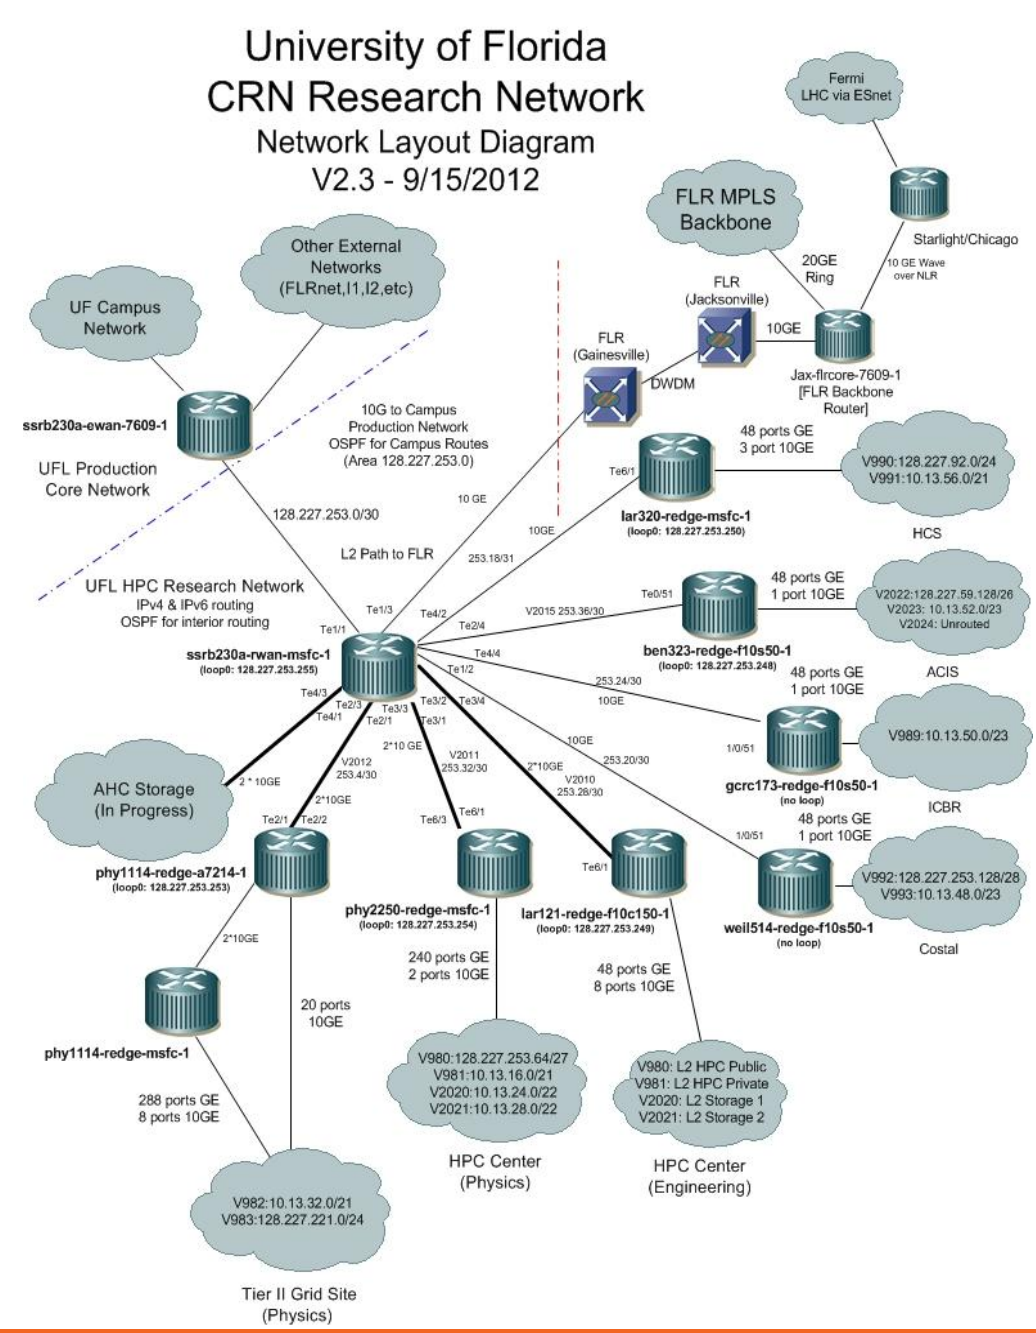

Time for an Upgrade

- Over time, new sites have been added, sites have moved around, and there have been very minor equipment upgrades.
- Most of the equipment originally purchased in 2005 is still in use (7+ years!)
- Some of the equipment is no longer under support contract due to EOL schedule.
- Nearly all will become EOL this year.
- Starting to see some contention on certain 10G links.

Time for an Upgrade (cont)

- 10G Research Wave to FLR becoming a limitation.
- We use "Etherchannel" to bond links together. We have no control over what traffic goes into what link, so congestion is possible with < 20Gbps of traffic.
- New UFDC/EC site and HPC cluster coming online.
- New network paradigms such as software defined networking emerging.
- Did I mention the current gear is **old**?

CRNv2 Goals

- 200Gbps backbone to the primary research compute and data storage locations on campus.
- Multiple 40Gbps or 10Gbps backbone to smaller research sites on campus.
- 100Gbps connection to FLR and other research networks.
- 10G and 40G edge ports.
- Flexibility to support advanced network paradigms such as SDN/Openflow.
- Preserve the "Science DMZ" style topology we enjoyed with CRNv1.

CRNv2 Design

- CRN Sites:
  - Tier 1 (200Gbps): Larsen 121, UFDC/EC.
  - Tier 2 (40/80Gbps): Physics 2250, AHC, Benton, GCRC.
  - Tier 3 (20/10Gbps) Weil, Physics 1114, Larsen 320.
  - Wan (100G) to FLR Jacksonville.
  - Main Campus (10 to 20Gbps).
- Native Capabilities
  - Line rate IPv4 and IPv6 routing.
  - Multi-Protocol Label Switching and VLAN support to provide transparent Layer 2 transport for off campus circuits and the Openflow testbed.
  - VRF support to enable multiple virtual networks.
  - Redundant route modules, power supplies, fan trays, etc.
  - Many many other features you would expect of high end Enterprise class equipment.

CRNv2 Design (cont)

- Openflow Testbed:
  - Native: SSRB, NEB, CISE, NS Lab
  - Tunneled: Lar121, UFDC/EC
  - Eventually all Tier 1 and Tier 2 sites will natively support Openflow via code updates. Full support for both SDN and "classical" network operations on the same hardware.
  - Flowvisor and test controller housed within CNS.
  - We are also deploying an Openflow testbed on the FLR backbone.
    - One Openflow switch in each FLR core pop, 10Gbps connected.
    - 10G Statewide SDN ring.
  - The combination of Openflow testbeds along with the connection to the Innovation Platform permits UF to participate in both regiona and national SDN efforts.

CRNv2 Design (cont)

- New Hardware:
  - Brocade MLXe16s at Tier 1 sites
  - Brocade ICX6650s at Tier 1 and Tier 2 sites
  - Existing Hardware at Tier 3 sites (Cisco/Force10/Arista)
  - Pica8/Pronto Switches at Openflow testbed sites (3295/3920)
  - Dell R620 servers for perfsonar.

CRNv2 Topology

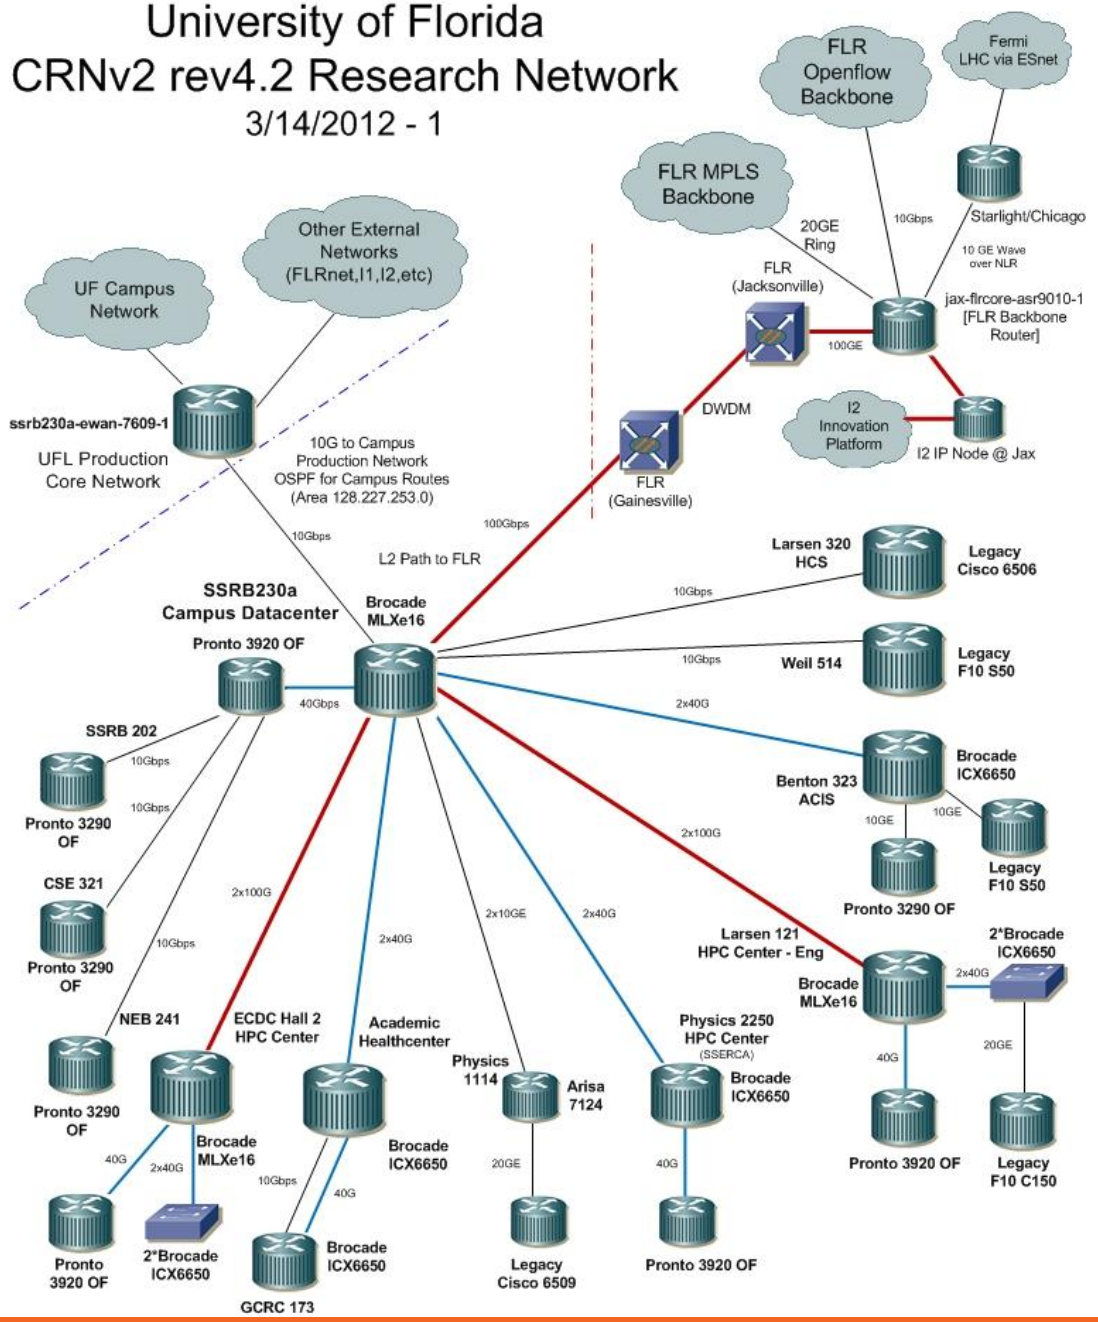

CRN Upgrade

- Three distinct projects/grants made the CRNv2 upgrade possible
  - 100G Research Wave/Innovation Platform Connection. Elias Eldayrie, CIO, University of Florida.
  - NSF CC-NIE Network Infrastructure: Dr. Erik Deumens PI
  - NSF MRI "Acquisition of Gatorcloud": Dr. Andy Li PI

CRNv2 Upgrade Status

- 100G circuit over FLR to Jacksonville, Innovation Platform 100G connection in Jacksonville and installation of central CRNv2 node at SSRB was completed and operational on January 30th.
- UFDC/EC and Lar121 Tier 1 sites have been installed and are in production.
- Tier 2 site installation will be commencing in the next few weeks pending some fiber installation.
- Openflow switches will start deployment in the next few weeks.

Florida LambdaRail

- First Light 12/17/2004. Fully Production 4/2005.
- Member owned, facilities based network.
- 12 equity owners including UF.
- Connects all Florida SUS institutions, most major private schools, as well as many Florida College System institutions.
- UFL is also the site of the FLR Network Operations Center.
  - Engineering.
  - 24x7 support.

Florida LambdaRail (cont)

- Optical Network
  - "FLRwave"
  - Dense Wave Division Multiplexing (DWDM) optical network.
  - Up to 32 channels (100Ghz spacing.)
  - 1540 miles of dark fiber around the state.
  - Provides 1G, 2.5G,10G, 40G, 100G waves between optical add/drop sites.
  - Based on Cisco 15454 optical nodes.
  - 28 locations around the state.
  - 38 active optical shelves.

FLR Optical Topology

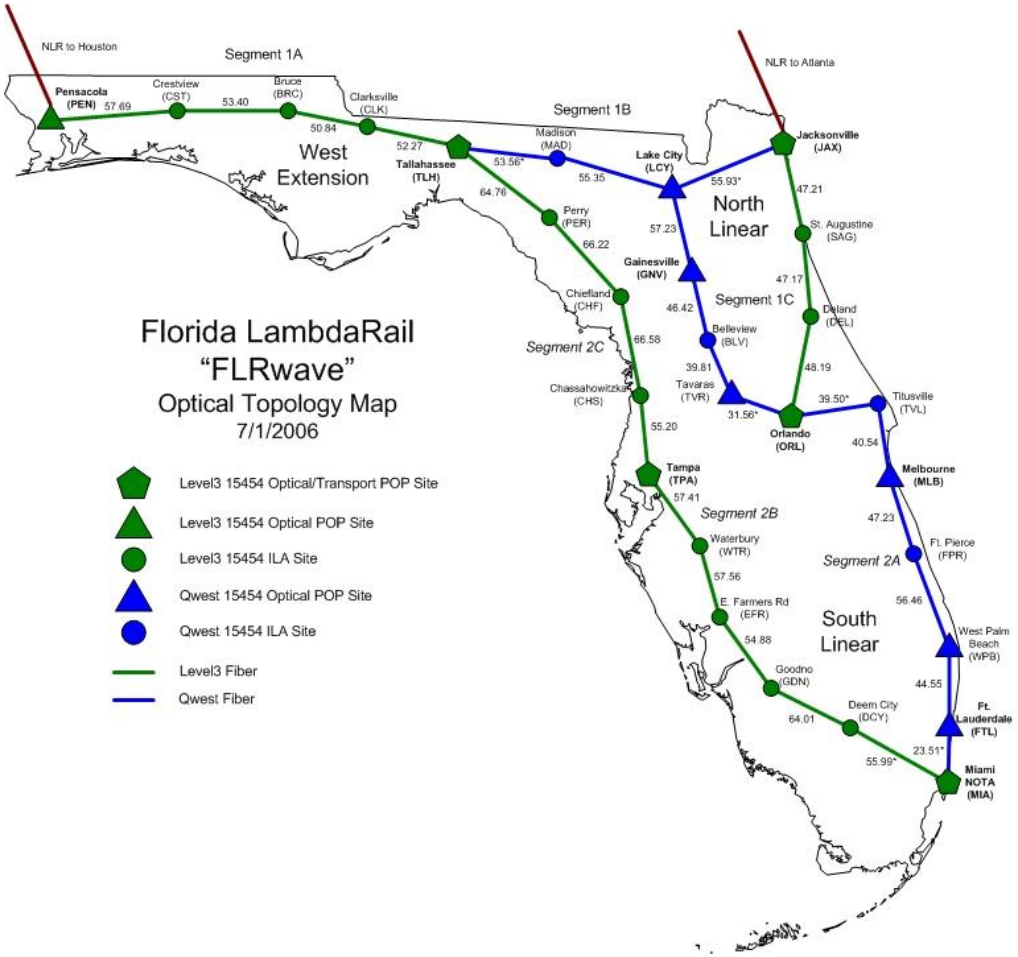

Florida LambdaRail (cont)

- Transport Network
  - "FLRnet"
  - 5 Core nodes (Jax, Orl, Mia, Tpa, Tlh). Several 2nd tier nodes as well.
  - MPLS based transport network. Provides Layer 2 and Layer 3 services.
  - Point to Point Ethernet Circuits (Pseudowires)
  - Commodity Internet Services.
  - Commodity Peering.
    - Google
    - Microsoft
    - Netflix
    - etc

Florida LambdaRail (cont)

- R&E Networking
  - Internet2
    - I2 Innovation Platform
  - National LambdaRail
  - R&E Peering @ Starlight
- Layer 3 VPN service ("Virtual Backbone")
  - UF Statewide Backbone VRF
  - Research Backbones (XSEDE, etc)
- Based on Cisco ASR 9010 routers
  - 10G, 40G, 100G capability.
  - 4M IPv4 routes.

About FLRnet

- 20G Backbone (dual 10G waves)
- 150G of Research Peering Capacity
  - 100G Internet2 Innovation Platform/Advanced L2 Service
  - 10G Primary, 10G Secondary to NLR L3
  - 10G to Internet2 direct to Atlanta.
  - 10G to Starlight (used for Esnet peering and other research).
  - Peering with Cwave at 10G.
- >52G+ of Commodity Peering Capacity.
  - 10G at the NOTA.
  - 10G at TIE.
  - 20G to Google.
  - 1G with Transitrail, 1G+ with CPS.
  - 10G to Akamai AANP.
- 55G of Internet Capacity.

FLR Optical Node

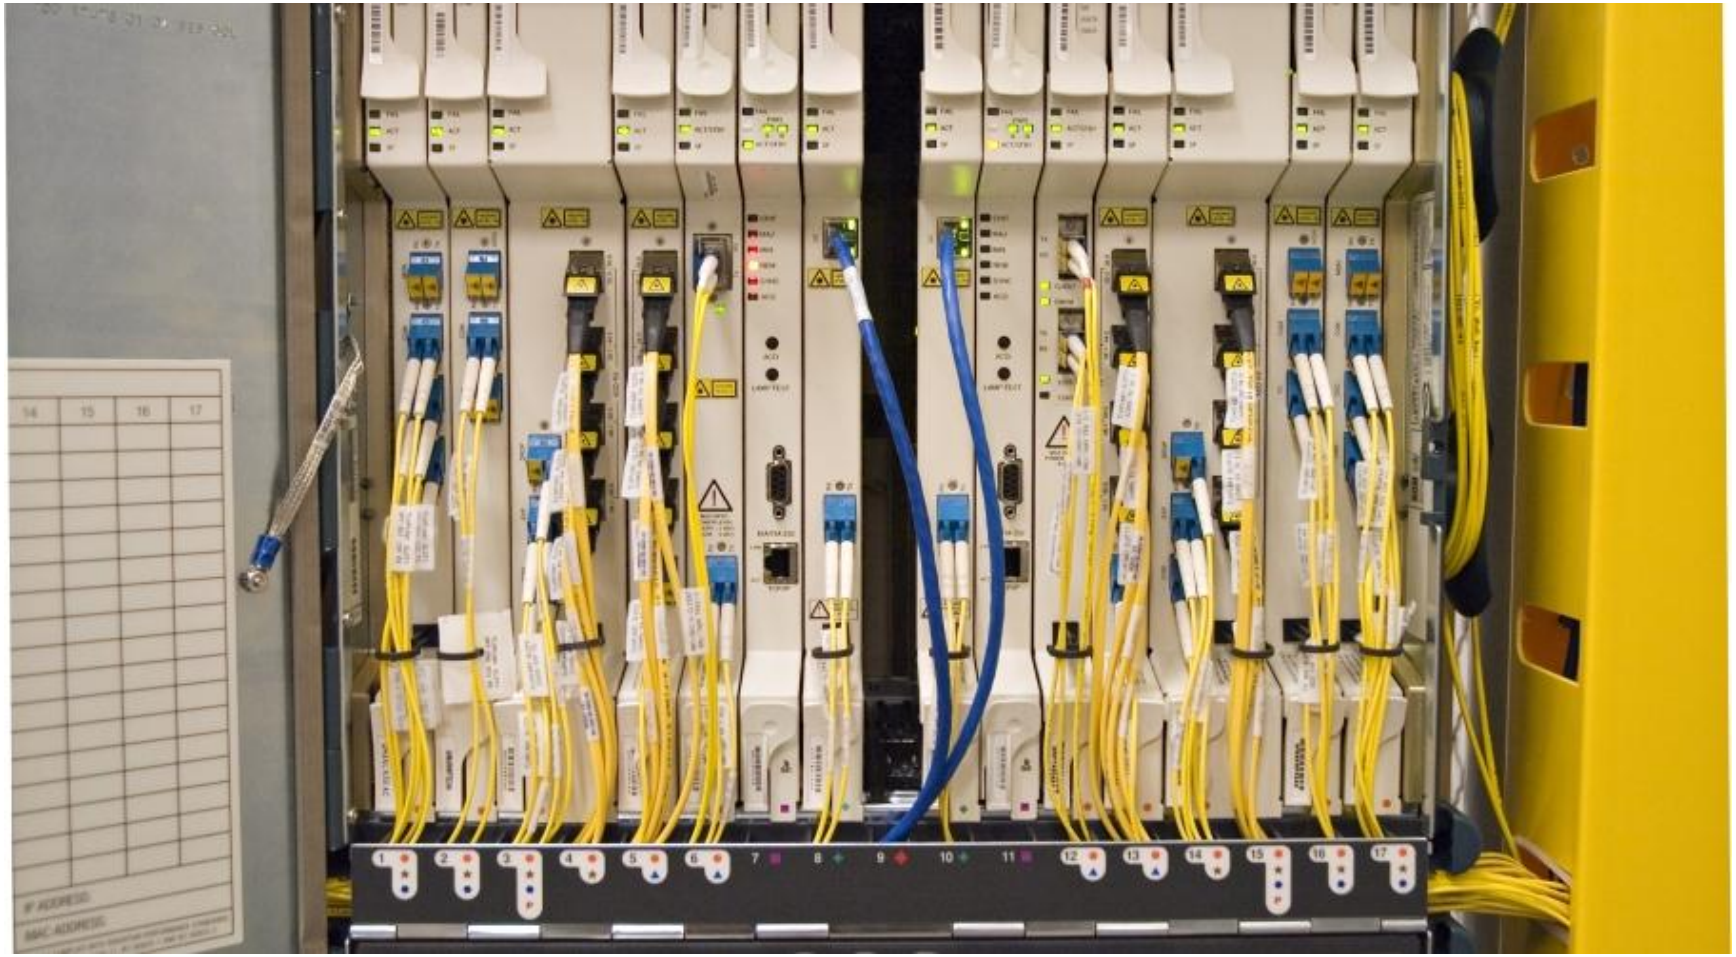

Typical FLRnet Node

Internet2 Innovation Platform

- What is it?
  - Next generation Internet2 Network
    - 100G connected.
    - Lower cost nodes (thus more of them, less backhaul)
    - SDN (Openflow) capable
    - Based on Brocade MLXe and Juniper MX Platforms. Most nodes are currently Brocade.
    - Will provide the Advanced Layer 2 Service (AL2S) on day 1.

Internet2 Innovation Platform (Cont)

- Operates over Internet2's upgrade optical network provided by the BTOP award.
- Current production I2 network will continue to operate on the existing I2 backbone. Will eventually migrate to IP based on publicized plans of I2.

◦ Architecture

- Internet2 is encouraging a standard R&E networking environment on campuses, at least through the initial phase of the project, through a set of requirements in order to connect.

Internet2 Innovation Platform (Cont)

- Science DMZ
  - Don't run your research network traffic "at scale" through your production network, and thus through your production firewall, ids, etc.
  - Create a purpose built network and connect it outside of the typical campus security parameter.
  - Within the Science DMZ, you must run PerfSonar.
  - Sounds a lot like the UF CRN!
- Layer 2 connectivity at 100G.
  - The intent is for "unlimited bandwidth" or as close as you can come to it. Prevent congestion from being a limiting factor

Internet2 Innovation Platform (Cont)

- Layer 2 paths provide for more deterministic behavior by the network. Typical "hop by hop" routing is replaced by pre-configured paths through the network.
- CRNv2 will provide for this by Jan 2013.
- FLR will be connected to the IP at 100G, thus giving UF a 100G path from the CRN to the IP node in Jax at layer 2.

- SDN Network Capabilities
  - The IP itself will be SDN based. Goal is to provide for ubiquitous end to end SDN capabilities.
  - CRNv2 will have an Openflow testbed on day 1.
  - All Tier 1 and Tier 2 CRNv2 sites will have hardware which is capable of Openflow once we decide to turn it on.
  - CRN is also currently running Dynes, which is another form of SDN.

Internet2 Innovation Platform (Cont)

- Philosophy
  - Unlimited Bandwidth.
  - Ubiquitous Deployment/Environment.
  - End to End.
    - Capabilities
    - Monitoring
  - Disruptive.
    - Researchers free to do new and interesting things.
  - "Friction Free."
    - Clear of as many network impediments as possible.
  - At Scale.

Internet2 Innovation Platform (Cont)

- Innovation Platform is being introduced as a limited pilot.
- University of Florida is 1 of around 14 sites which are part of the pilot program.
- Connected January 2013.
- University of Florida is the first site nationwide to complete all three requirements for the Innovation Platform Pilot.
  - 100G connection
  - Science DMZ
  - SDN capability (Dynes project)

Innovation Platform Pilot

What will UF do with the IP

- Initial use cases center around AL2S.
  - Static high speed L2 circuits between UF and other IP connectors.
    - Research
    - Production Traffic (future)
    - Perfsonar used to monitor health of the network, end to end.
  - General dynamic capability integrated into the CRN.
  - Openflow test/research.
  - Currently planning a project with FNAL to test end to end network/systems/storage performance at scale.
- Whats next?
  - Really up to the research community.
  - IP + SDN capable CRNv2 are foundational infrastructures for disruptive and innovative applications. Please commence disrupting!

Thanks!