



September 27th 2011, LHCOPN

Eric Boyd, Internet2

LHCONE in North America

Motivation: Internet2's Seven focus areas

Advanced network and network services leadership

Services at scale:
Services “above the network”

U.S. UCAN – Community Anchor Network Program

Industry partnership development and engagement

Global reach and engagement

Research community development and engagement

National/Regional Partnership



Internet2 BTOP-funded Upgrade



- Upgraded 100 Gbps IP/MPLS/ION Network with 10 Juniper T1600's
- Upgraded peering service network with 6 Juniper MX960's
- Deployment of a new Layer 2 service on NDDI/OS3E network
- Enhanced research programs and support

Agenda

- **LHCONE**
- LHCONE-NA Status
 - Internet2 Network
 - New York
 - Chicago
 - Canada
 - Open Issues
- DYNES, ION, and LHCONE
- MAN LAN and WIX
- NDDI & OS³E
- SC Planning

LHCONE Status

- LHCONE is a response to the changing dynamic of data movement in the LHC environment.
- It is composed of multiple parts:
 - North America, Transatlantic Links, Europe
 - Others?
- It is expected to be composed of multiple services
 - Multipoint service
 - Point-to-point service
 - Monitoring service

LHCONE Multipoint Service

- Initially created as a shared Layer 2 domain.
- Uses 2 VLANs (2000 and 3000) on separate transatlantic routes in order to avoid loops.
- Enables up to 25G on the Trans-Atlantic routes for LHC traffic.
- Use of dual paths provides redundancy.

LHCONE Point-to-Point Service

- Planned point-to-point service
- Suggestion: Build on efforts of DYNES and DICE-Dynamic service
- DICE-Dynamic service being rolled out by ESnet, GÉANT, Internet2, and USLHCnet
 - Remaining issues being worked out
 - Planned commencement of service: October, 2011
 - Built on OSCARS (ESnet, Internet2, USLHCnet) and AUTOBAHN (GÉANT), using IDC protocol

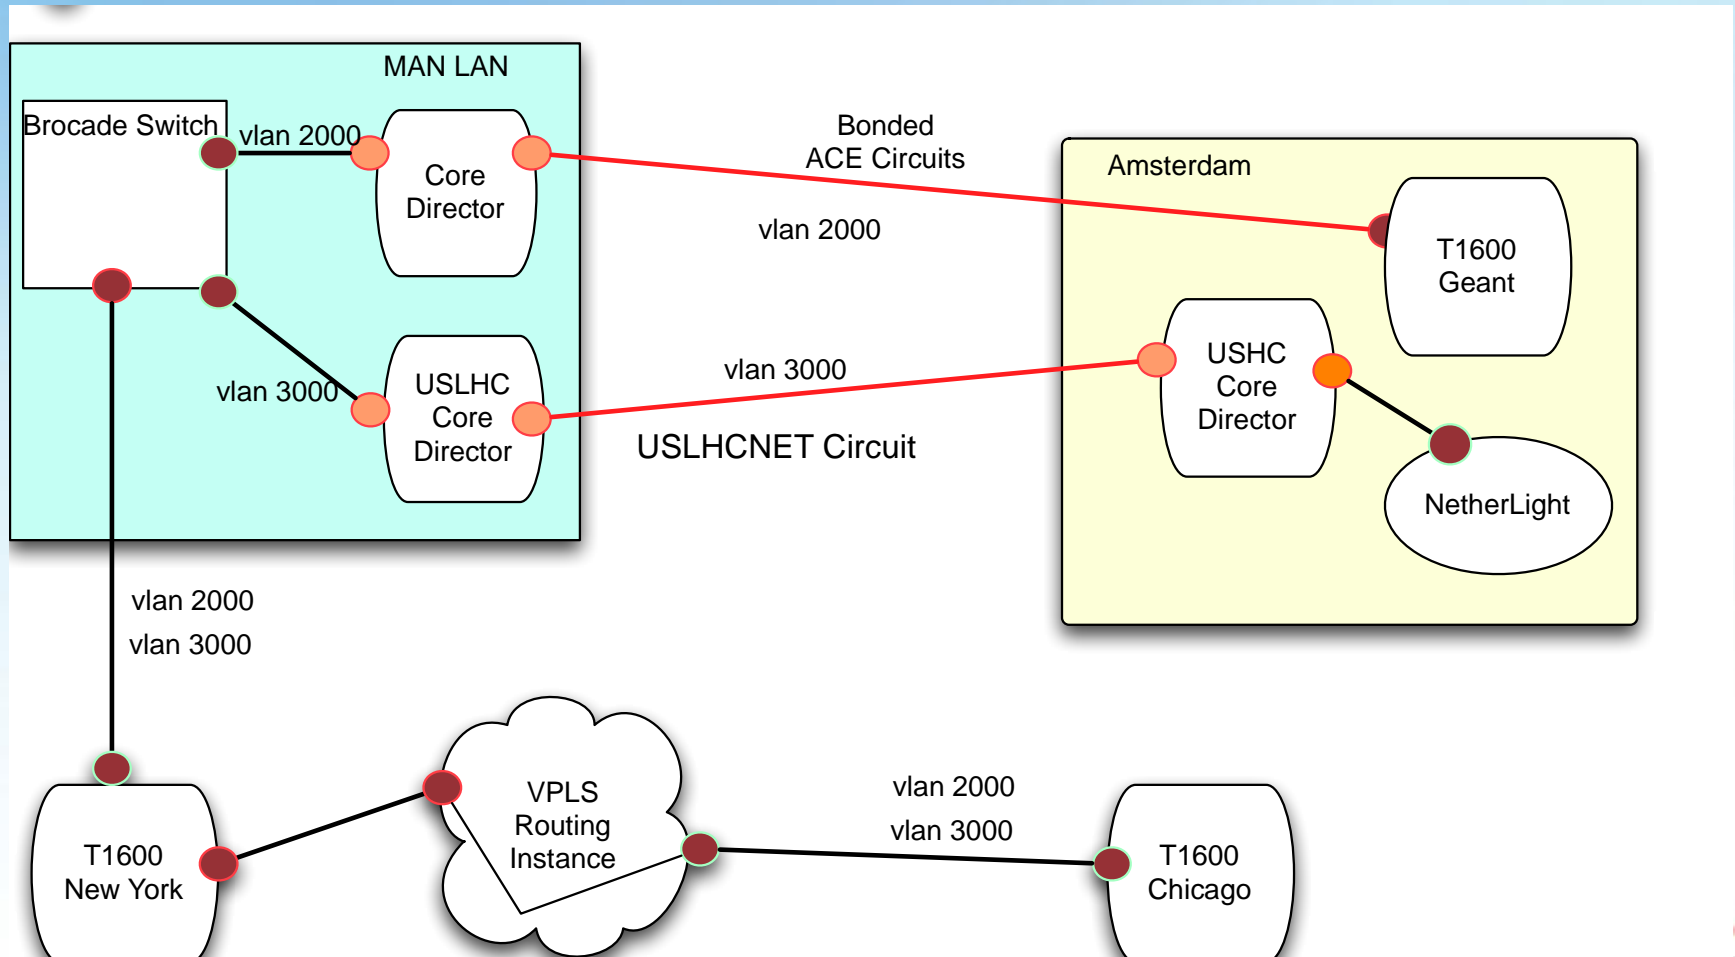
LHCONE Monitoring Service

- Planned monitoring service
- Suggestion: Build on efforts of DYNES and DICE-Diagnostic service
- DICE-Diagnostic service, being rolled out by ESnet, GÉANT, and Internet2
 - Remaining issues being worked out
 - Planned commencement of service: October, 2011
 - Built on perfSONAR

Agenda

- LHCONE
- **LHCONE-NA Status**
 - Internet2 Network
 - New York
 - Chicago
 - Canada
 - Open Issues
- DYNES, ION, and LHCONE
- MAN LAN and WIX
- NDDI & OS³E
- SC Planning

LHCONE-NA / Multipoint service



LHCONE-NA MAN LAN

- LHCONE-TA transatlantic links are fully operational.
- Traffic has been moving through MAN LAN.
 - Internet2, CANARIE, and transatlantic paths all participate in VLANs 2000 and 3000.

LHCONE-NA Internet2 Backbone

- Configuration on the backbone is in place between Chicago and New York.
 - Currently implemented as a Layer 2 overlay on the IP network
- Extended to other core sites as we gain more experience.
 - In part, dictated by participation requests.

LHCONE-NA Status – Internet2 NEWY

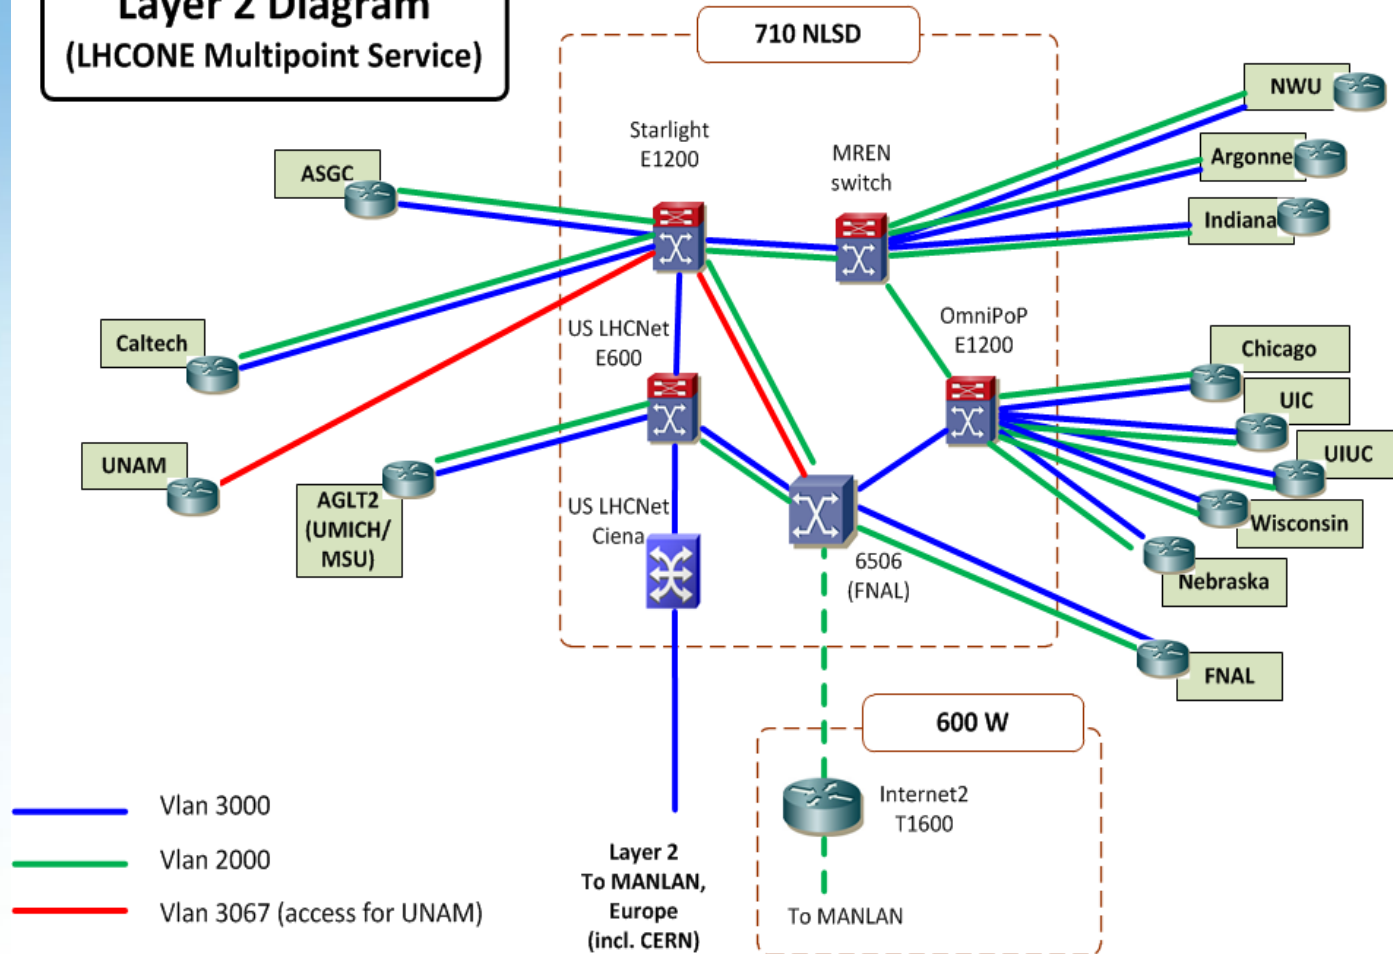
- VLANS 2000 and 3000 for the multipoint service are configured.
 - Transatlantic routes, Internet2, and CANARIE all are participating in the shared VLAN service.
- New switch will be installed at MAN LAN in October.
 - Will enable new connection by BNL
- Peering with Univ of Toronto through the CANARIE link to MAN LAN is complete
- End sites that have direct connections to MAN LAN are:
 - MIT
 - BNL
 - BU/Harvard

LHCONE-NA Status – Internet2 CHIC

- VLANS for multipoint service configured on 9/23.
 - Correctly configured shortly thereafter to prevent routing loop 😊
 - Testing on the link can start any time.
- Status of FNAL Cisco.
 - Resource constraints on the Chicago router have prevented this from happening.
 - Port availability is the issue.
- End Sites
 - See diagram from this summer

LHCONE-NA Status – Internet2 CHIC

**Summer 2011
Layer 2 Diagram
(LHCONE Multipoint Service)**



LHCONE-NA Status - Canada

- The link to Toronto was configured early the week of September 19th.
 - The current configured lightpath is provisioned over existing circuits with 2G capacity sharing the same 10GE to the MANLAN switch
 - CANARIE would be able to upgrade the lightpath to a higher capacity or provide a dedicated wavelength if capacity requirement increases.
- Canarie/UofT have established BGP sessions to the route servers.
- Ian Gable will be providing more updates on this.

LHCONE-NA Status – Open Issues

- Multipoint
 - Late getting the VLANs up, but we're getting there 😊
 - Extending to other routers.
 - Adding new interfaces and thus extending the service to other backbone nodes is straightforward at this point.
 - Extensions will be driven by requests for the service.

Agenda

- LHCONE
- LHCONE-NA Status
 - Internet2 Network
 - New York
 - Chicago
 - Canada
 - Open Issues
- **DYNES, ION, and LHCONE**
- MAN LAN and WIX
- NDDI & OS³E
- SC Planning

ION

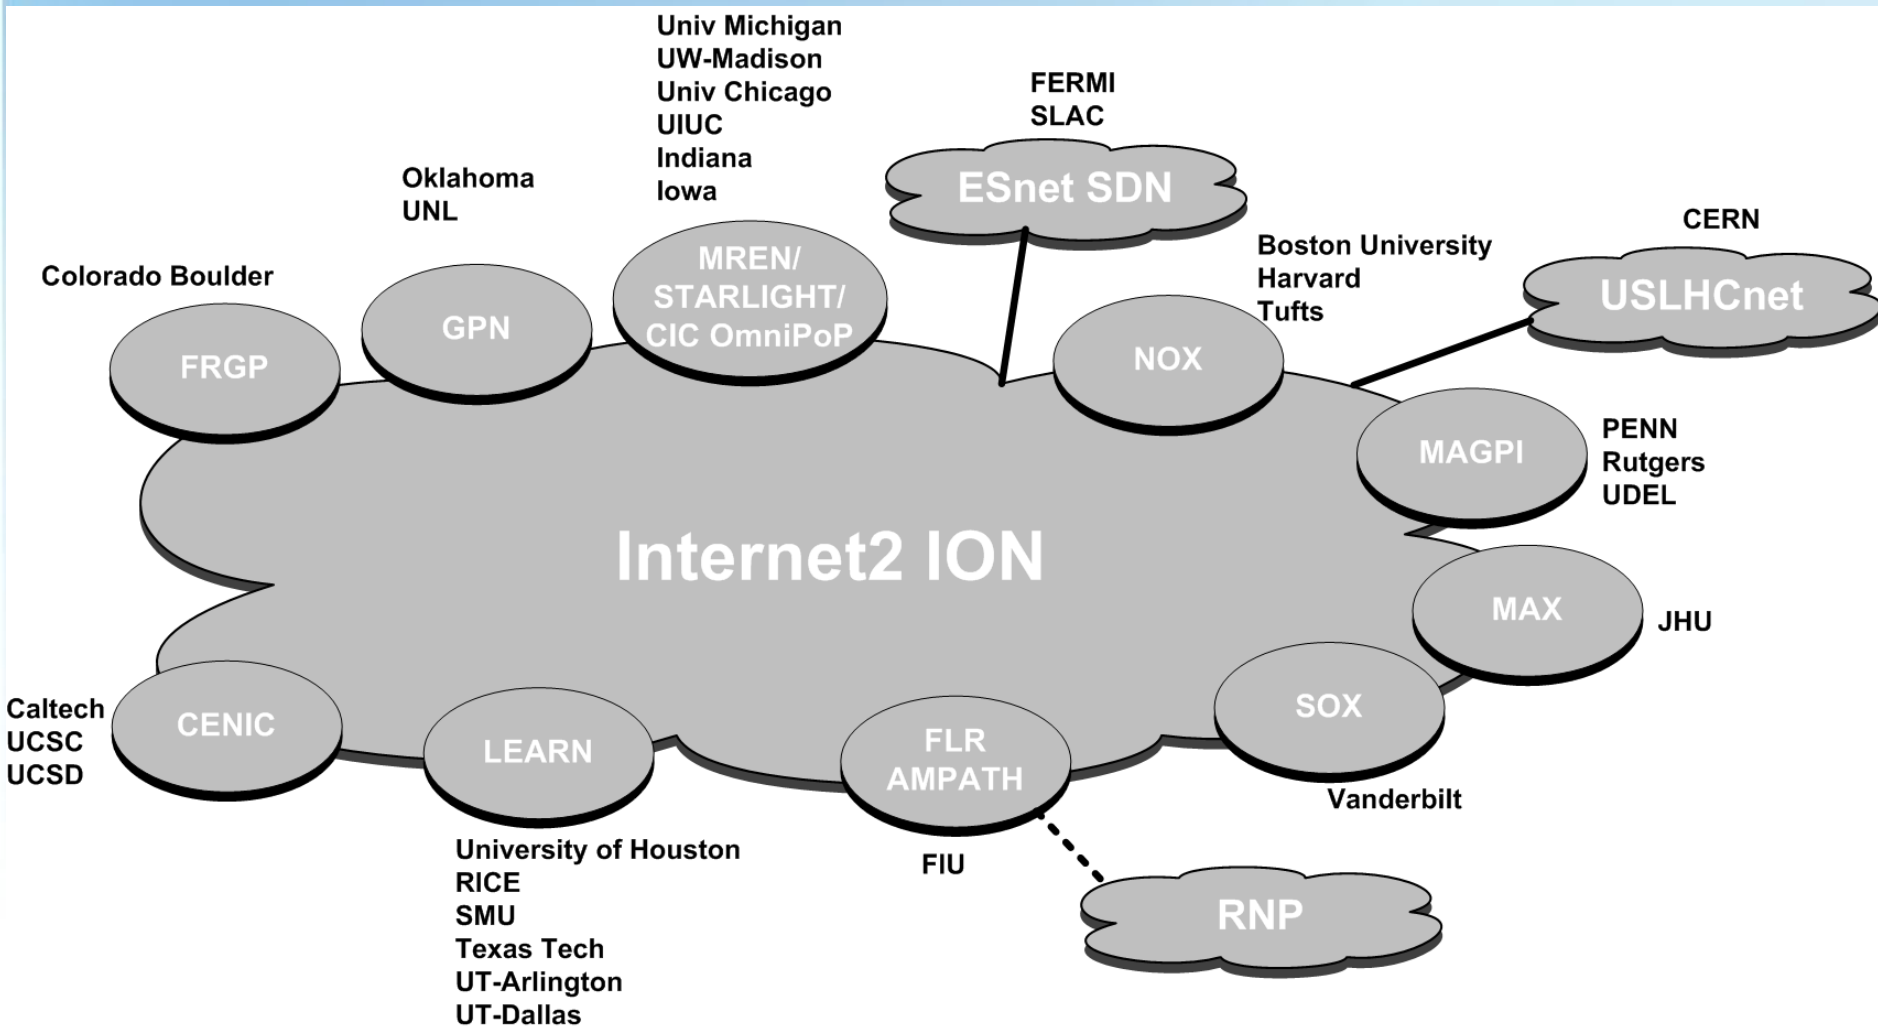
- ION
 - Shared VLAN service across the Internet2 backbone
 - Implemented as combination dedicated / scavenger service atop the Layer 3 infrastructure
 - Implements IDC protocol
 - Implemented with OSCARS and perfSONAR-PS
- What's new (October, 2011)
 - Modifying ION to be a persistent VLAN service
 - Integrated with ESnet SDN, GÉANT AUTObahn, and USLHCnet as part of DICE-Dynamic Service
- What's planned (late 2011)
 - As DYNES service rolls out, ION is the backbone linking the various regional deployments
 - Peer ION with OS3E in a few locations

DYNES

- DYNES (NSF #0958998) = Enable dynamic circuit services end to end
 - Deploy equipment at the regional and campus levels
 - Based on OSCARS to control circuits, FDT to move data, perfSONAR to monitor performance
 - Funding – May 2010 – May 2013
 - Emphasis to enable this service for scientific use
- Current Status
 - Through with our first deployment group, into testing of the software and hardware
 - Configuring second group, shipments have started
 - Third group in planning stages

DYNES Projected Topology (Fall 2011)

- Based on applications accepted
- Showing peerings to other Dynamic Circuit Networks (DCN)



DYNES Deployment Status

- Group A – Fully deployed, still undergoing testing (Caltech, Vanderbilt, UMich, MAX, Rutgers, UDel, JHU, SOX, AMPATH)
- Group B – Configuring now, deployment expected fall (TTU, UTA, UTD, SMU, UH, Rice, LEARN, MAGPI, UPenn, MREN, UChicago, UWisc, UIUC, FIU)
- Group C – Late Fall/Winter configuration expected, deployment and testing into next year (UCSD, UCSC, UNL, OU, Ulowa, NOX, BU, Harvard, Tufts, FRGP, UColorado)

DYNES/ION and LHCONE

- Simple to integrate DYNES/ION and LHCONE point-to-point service
- Possible to integrate DYNES/ION and LHCONE multipoint service?
 - DYNES / LHCONE Architecture team discussing ways to integrate DYNES functionality with LHCONE
 - It is expected that point to point connections through DYNES would work ...
 - Possible to position DYNES as an ‘onramp’ or ‘gateway’ to the multipoint service?
 - Glue a dynamic connection from a campus (through a regional) into the VLAN 2000/3000
 - Requires some adjustments to the DYNES end-site addressing and routing configurations to integrate into LHCONE multipoint layer2 environment
 - Would allow smaller T3 sites in the US instant access as soon as they get their DYNES gear.

Agenda

- LHCONE
- LHCONE-NA Status
 - Internet2 Network
 - New York
 - Chicago
 - Canada
 - Open Issues
- DYNES and LHCONE
- **MAN LAN and WIX**
- NDDI & OS³E
- SC Planning

MAN LAN

- New York Exchange Point
- Ciena Core Director and Cisco 6513
- Current Connections on the Core Director:
 - 11 OC-192's
 - 9 1 Gig
- Current Connection on the 6513
 - 16 10G Ethernets
 - 7 1G Ethernet

MAN LAN Roadmap

- Switch upgrade:
 - Brocade MLXe-16 was purchased with:
 - 24 10G ports
 - 24 1 G ports
 - 2 100G ports
 - Internet2 and ESnet will be connected at 100G.
- The Brocade will allow landing transatlantic circuits of greater than 10G.
- An IDC for Dynamic circuits will be installed.
 - Comply with GLIF GOLE definition

MAN LAN Services

- MAN LAN is an Open Exchange Point.
- 1 Gbps, 10 Gbps, and 100 Gbps interfaces on the Brocade switch.
 - 40 Gbps could be available by 2012.
- Map dedicated VLANs through for Layer2 connectivity beyond the ethernet switch.
- With the Brocade the possibility of higher layer services should there be a need.
 - This would include OpenFlow being enabled on the Brocade.
- Dynamic services via an IDC.
- perfSONAR-ps instrumentation.

WIX

- WIX = Washinton DC International Exchange Point
- Joint project being developed by MAX and Internet2 and transferred for MAX to manage once in operation.
- WIX is a state-of-the-art international peering exchange facility, located at the Level 3 POP in McLean VA, designed to serve research and education networks.
- WIX is architected to meet the diverse needs of different networks.
- Initially, WIX facility will hold 4 racks, expandable to 12 racks as needed.
 - Bulk cables between the existing MAX and Internet2 suites will also be in place.
- WIX is implemented with a Ciena Core Director and a Brocade MLXe-16.

WIX Roadmap

- Grow the connections to existing Exchange Points.
- Expand the facility with “above the net” capabilities located in the suite.
 - Allows for easy access both domestically and internationally
- Grow the number of transatlantic links to insure adequate connectivity as well as diversity.

WIX Services

- Dedicated VLANs between participants for traffic exchange at Layer 2.
- WDC-IX will be an Open Exchange Point.
- Access to Dynamic Circuit Networks such as Internet2 ION.
- With the Brocade, there exists the possibility of higher layer services, should there be a need.
 - Possibility of OpenFlow being enabled on the Brocade
- 1 Gbps, 10 Gbps, and 100 Gbps interfaces are available on the Brocade switch.
- 40 Gbps could be available by 2012.
- perfSONAR instrumentation

Agenda

- LHCONE
- LHCONE-NA Status
 - Internet2 Network
 - New York
 - Chicago
 - Canada
 - Open Issues
- DYNES and LHCONE
- MAN LAN and WIX
- **NDDI & OS³E**
- SC Planning

Network Development and Deployment Initiative (NDDI)

Partnership that includes Internet2, Indiana University, & the Clean Slate Program at Stanford as contributing partners. Many global collaborators interested in interconnection and extension



Builds on NSF's support for GENI and Internet2's BTOP-funded backbone upgrade



Seeks to create a software defined advanced-services-capable network substrate to support network and domain research [note, this is a work in progress]



Components of the NDDI Substrate

- 30+ high-speed Ethernet switches deployed across the upgraded Internet2 network and interconnected via 10G waves
- A common control plane being developed by IU, Stanford, and Internet2
- Production-level operational support
- Ability to support service layers & research slices



64 x 10G SFP+ 1.28 Tbps non-blocking
4 x 40G QSFP+ 1 RU



The NDDI Control Plane

- The control plane is key to placing the forwarding behavior of the NDDI substrate under the control of the community and allowing SDN innovations
- Eventual goal to fully virtualize control plane to enable substrate slices for community control, research and service development
- Will adopt open standards (e.g., OpenFlow)
- Available as open source (Modified Berkley Apache 2.0 License)

Open Science, Scholarship and Services Exchange (OS3E)

- An example of a community defined network service built on top of the NDDI substrate.
- The OS3E will connect users at Internet2 POP's with each other, existing exchange points and other collaborators via a flexible, open layer 2 network.
- A nationwide distributed layer 2 “exchange”
 - Persistent layer 2 vlans with inter-domain support
- Production services designed to support the needs of domain science (e.g., LHCONE, etc.)
- Will support open interdomain standards
- Available as open source (Modified Berkley Apache 2.0 License)

Timeline

- | | |
|----------------|---|
| April, 2011 | Early Program Announcement |
| May-September | Hardware, Controller selection
Substrate development |
| October, 2011 | First Deployment and Domestic Demo
Link Policy & funding discussion
Next site group selection
iNDDI engagement |
| November, 2011 | Expanded Deployment
Inter-domain capabilities |
| January, 2011 | Large scale domestic deployment |

Support for Network Research

- NDDI substrate control plane key to supporting network research
 - At-scale, high performance, researcher-defined network forwarding behavior
 - virtual control plane provides the researcher with the network “LEGOs” to build a custom topology employing a researcher-defined forwarding plane
- NDDI substrate will have the capacity and reach to enable large testbeds

Making NDDI global...

- Substrate will support IDC (i.e., it will be inter-domain capable)
 - Expect interconnection with other OpenFlow testbeds as a first step (likely statically)
- While the initial investors are US-based, NDDI seeks global collaborators on the substrate infrastructure as well as control plane features
- Currently collecting contact information for those interested in being a part of NDDI
 - please send e-mail to nddi-intl@internet2.edu

“Open”


- Although it may be disruptive to existing business models, we are committed to extending a policy-free approach
- Each individual node should function like an “exchange point” in terms of policy, cost, capabilities
- Fully distributed exchange will operate as close to exchange point as possible given constraints: i.e. transport has additional associated costs
 - Inter-node transport scalability and funding needs discussion
 - Initially, an open, best effort service
 - Potential to add a dedicated priority queuing feature
- Internet2 would like to position this service on the forefront of pushing “open” approaches in distributed networks.



Courtesy of CERN

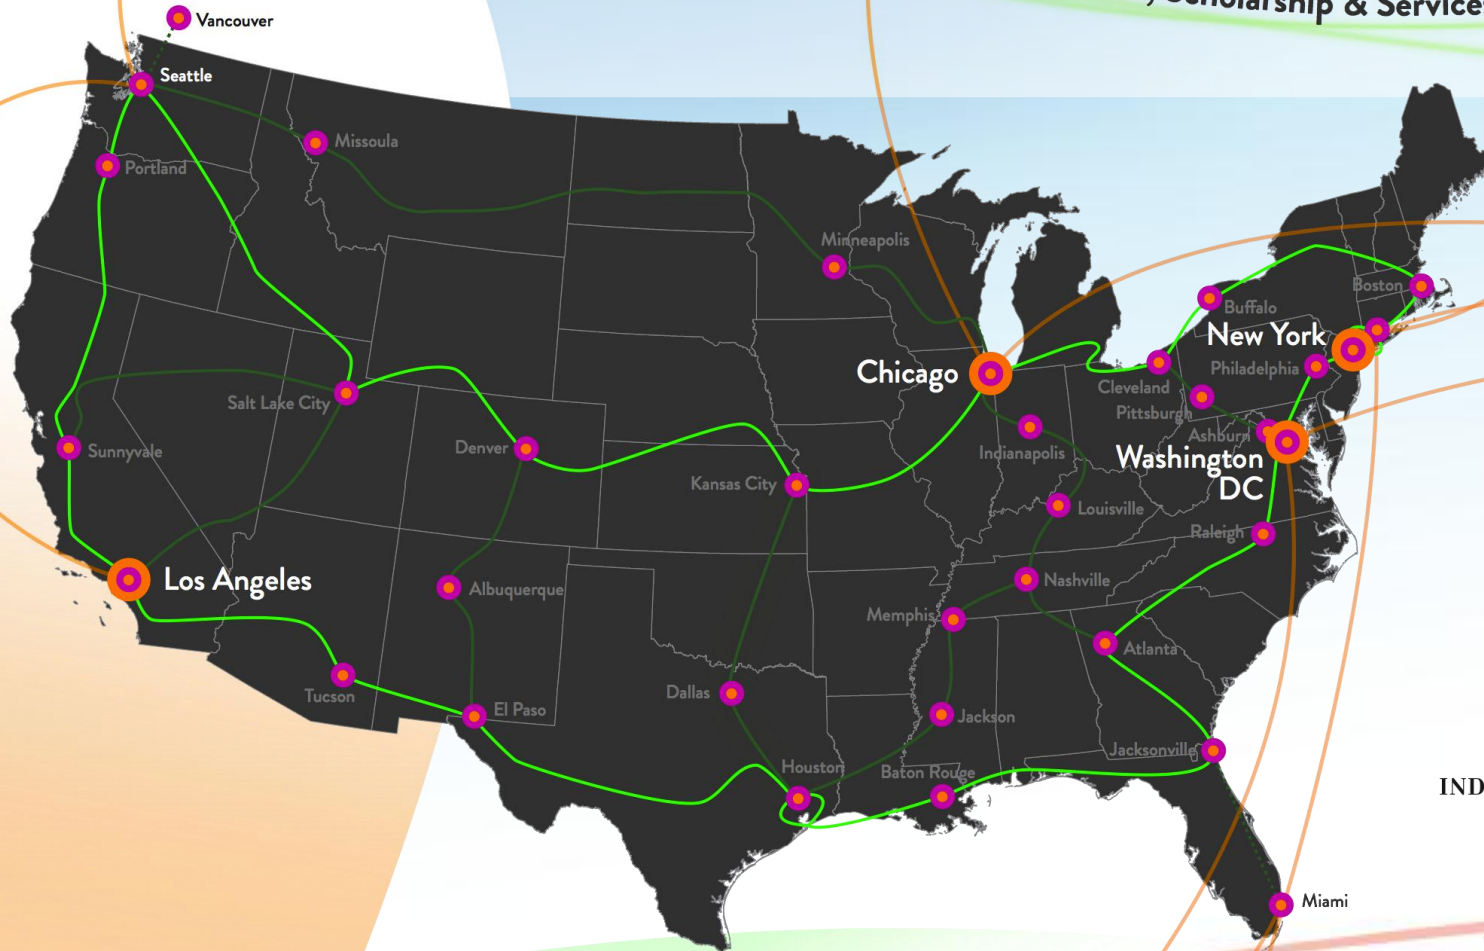


Courtesy of LBNL



S³E

The Open Science, Scholarship & Services Exchange






Courtesy of CERN

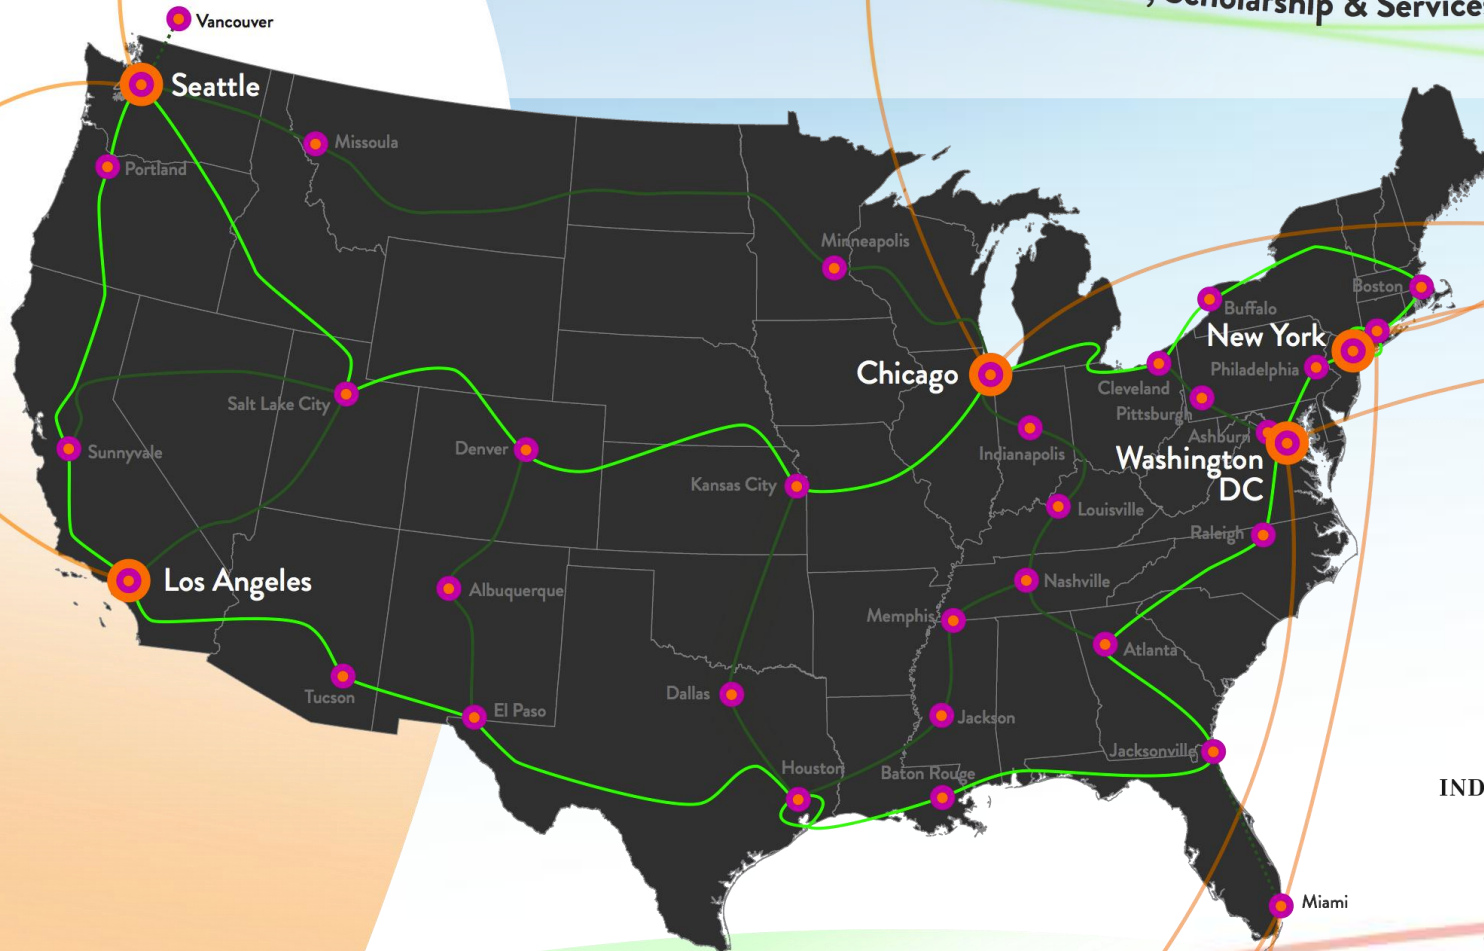


Courtesy of LBNL



O S³E

The Open Science, Scholarship & Services Exchange






Courtesy of CERN

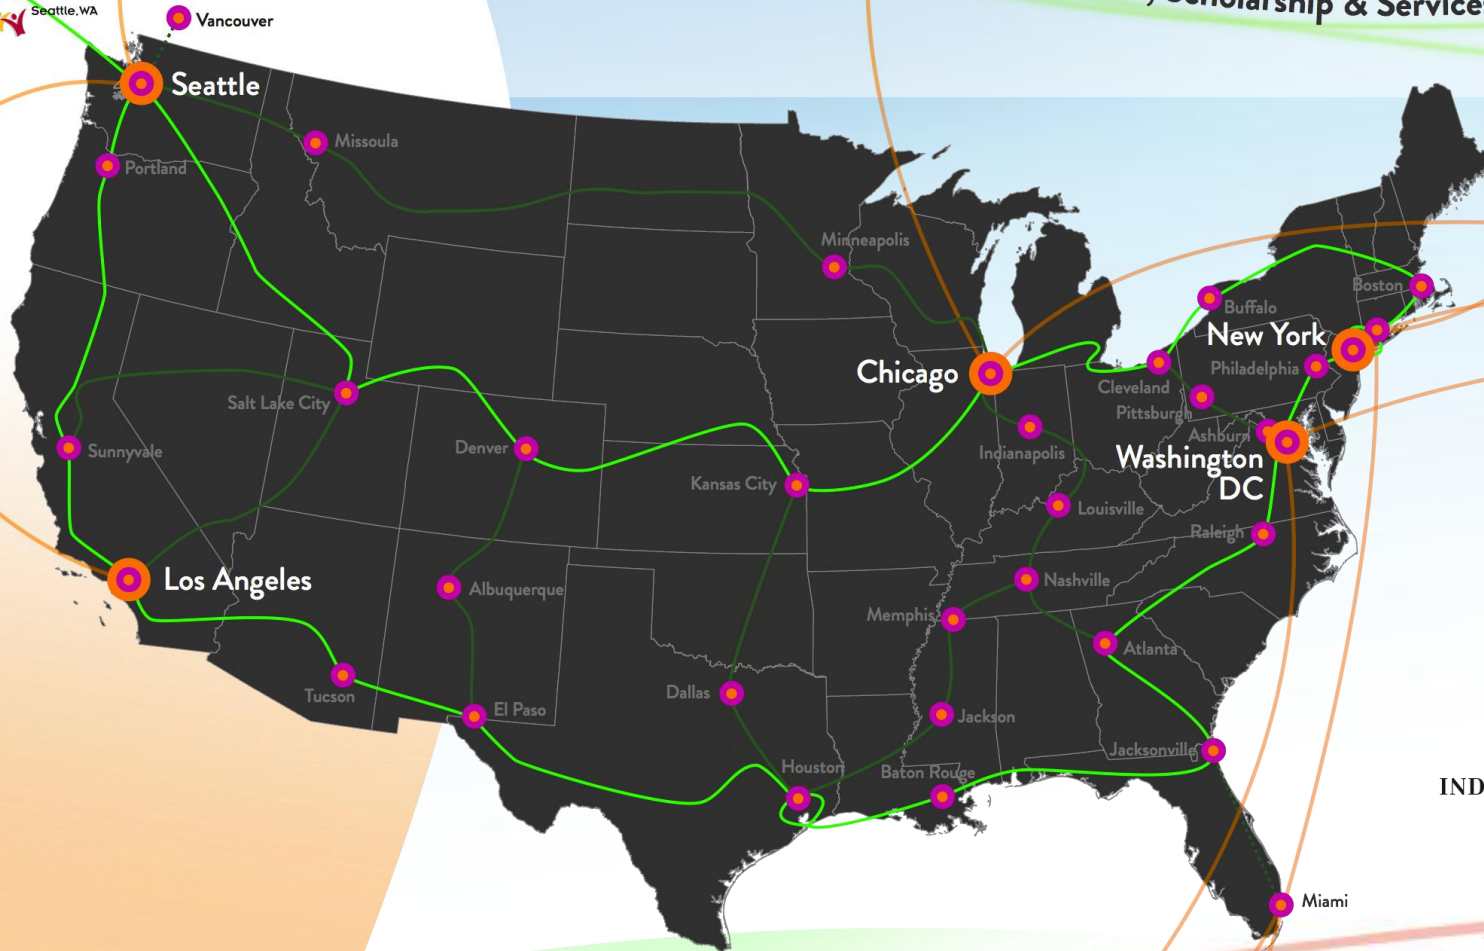


Courtesy of LBNL



O S³ E


The Open Science, Scholarship & Services Exchange



NDDI / OS3E Service Description

- This service is being developed in response to the request from the community as expressed in the report from the NTAC and subsequent approval by the AOAC.
- Fundamentally it is a best effort service with long term reservations.
 - It is at Layer 2
 - Different price points for hairpin service and inter-node service
 - It has a completely open access policy
 - Underlying wave infrastructure will be augmented as needed using the same general approach as used in the IP network.

9/27/2011



INTERNET
2

NDDI / OS3E Implementation Status

- Deployment
 - NEC G8264 switch selected for initial deployment
 - Chicago node installed
 - 4 nodes by Internet2 FMM
 - 5th node (Seattle) by SC
- Software
 - NOX OpenFlow controller selected for initial implementation
 - Software functional to demo Layer 2 VLAN service (OS3E) over OpenFlow substrate (NDDI) by FMM
 - Software functional to peer with ION (and other IDCs) by SC11
 - Software to peer with SRS OpenFlow demos at SC11
 - Open source software package to be made available in 2012

OS3E Fees

- There will likely be graduated fees:
 - A fee for connectors only wishing to peer with other connectors on the same switch.
 - A greater fee for connectors wishing to utilize the network interconnecting these exchange facilities.
- It is hard at this point to suggest exact fees, they could be adapted depending on the adoption levels.
- This discussion is more about gathering information from the community.

Agenda

- LHCONE
- LHCONE-NA Status
 - Internet2 Network
 - New York
 - Chicago
 - Canada
 - Open Issues
- DYNES and LHCONE
- MAN LAN and WIX
- **NDDI & OS³E**
- SC Planning

SC Planning

- Planning demos for the Internet2 booth, in collaboration with GÉANT
 - Showing how LHCONE is fostering scientific activity
 - Trying to avoid showing ‘just network monitoring’, want to show LHC data movement tools, visualizations
 - Speakers from the community (ATLAS, CMS, and networking experts) will be asked to come and give a short presentation in the booth
 - Demonstrate peering between OS3E and ION
- DYNES demo with Caltech, Vanderbilt, UMich, and Internet2
 - Show use of infrastructure, even into South America (via AMPATH)
 - Connects SPRACE and HEPGrid. Could be a good first example of an LHCONE-SA since these sites are connected via RNP OSCARS deployment.



LHCONE-NA Update

September 27th 2011, LHCOPN

Eric Boyd, Internet2

For more information, visit <http://www.internet2.edu/>