



May 24th 2012 – CHEP 2012

Jason Zurawski, Senior Network Engineer - Internet2

The DYNES Instrument: A Description and Overview

Agenda

- DYNES Motivation and Overview
- DYNES Hardware and Software
- Current Status
- Conclusion

DYNES Motivation

- Data movement to support science:
 - Increasing in size (100s of TBs in the LHC World, approaching PB sizes)
 - Becoming more frequent (multiple times per day)
 - Reaching more consumers (VO sizes stand to increase, more VOs)
 - Time sensitivity (data may grow “stale” if not processed immediately)
 - Almost always “multi-domain”



[1]

INTERNET
2

DYNES Motivation

- Traditional networking:
 - R&E or Commodity “TCP/IP” connectivity is subject to congestion by other users
 - TCP is sensitive to network use as well as physical flaws
 - Primary choice for application developers (reliability)
 - Supporting large sporadic flows is challenging for engineers
 - Need to worry about your network, as well as the networks of others (e.g. the end-to-end problem)
 - Can we ‘see’ how a network (or networks) are performing?
 - Can we dynamically change behavior and patterns?
- The End Game?
 - Many disciplines require a stable data transfer mechanism
 - Campuses/regionals have a duty to their customers to manage network traffic and deliver required bandwidth



The View From The Ivory Tower

- Recent calls to action
 - [CC-NIE](#) (NSF)
 - [“Big Data”](#) (NSF/NIH)
- End goal will be to make the campus and regional infrastructure ready for next generation of Networking:
 - Architecture Changes (e.g. “Science DMZ”)
 - Capacity Upgrades to 100G (and beyond)
 - New Technologies (Software Defined Networking, “Layer 2” solutions)
 - Monitoring
- What is the current state?

Reality Check: Expectations

- The following shows how long it (should*) take to transfer 1 Terabyte of data across various speed networks**:
- 10 Mbps network :
 - 300 hrs (12.5 days)
- 100 Mbps network :
 - 30 hrs
- 1 Gbps network :
 - 3 hrs
- 10 Gbps network :
 - 20 minutes



417

Expectation Failed

***Can your network do this?**

****Assumes running at 100% efficiency, no performance problems**

DYNES Motivation

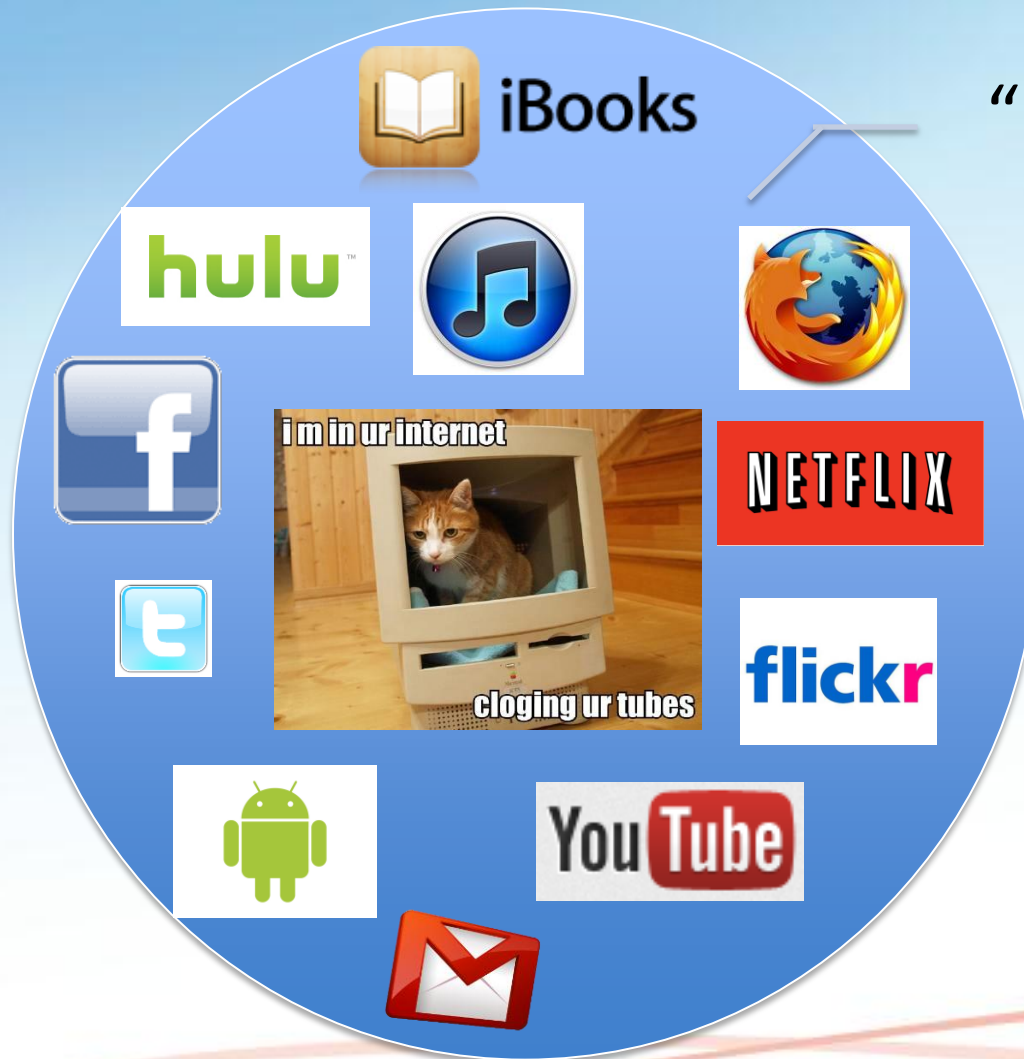
- Solution
 - Reserved bandwidth (over the entire end to end path) to move scientific data
 - Invoke this “on demand” instead of relying on permanent capacity (cost, complexity)
 - Exists in harmony with traditional IP networking
 - Ability to use ‘alternate’ protocols besides TCP in your private Layer 2 circuit
 - Connect to facilities that scientists need to access
 - Integration with data movement applications
 - Invoke the connectivity when the need it, based on network conditions
- Proposed Deployment:
 - Software and hardware support spanning domain boundaries
 - Campus
 - Regional
 - Backbone
 - Integration with existing technologies and deployments



DYNES Motivation

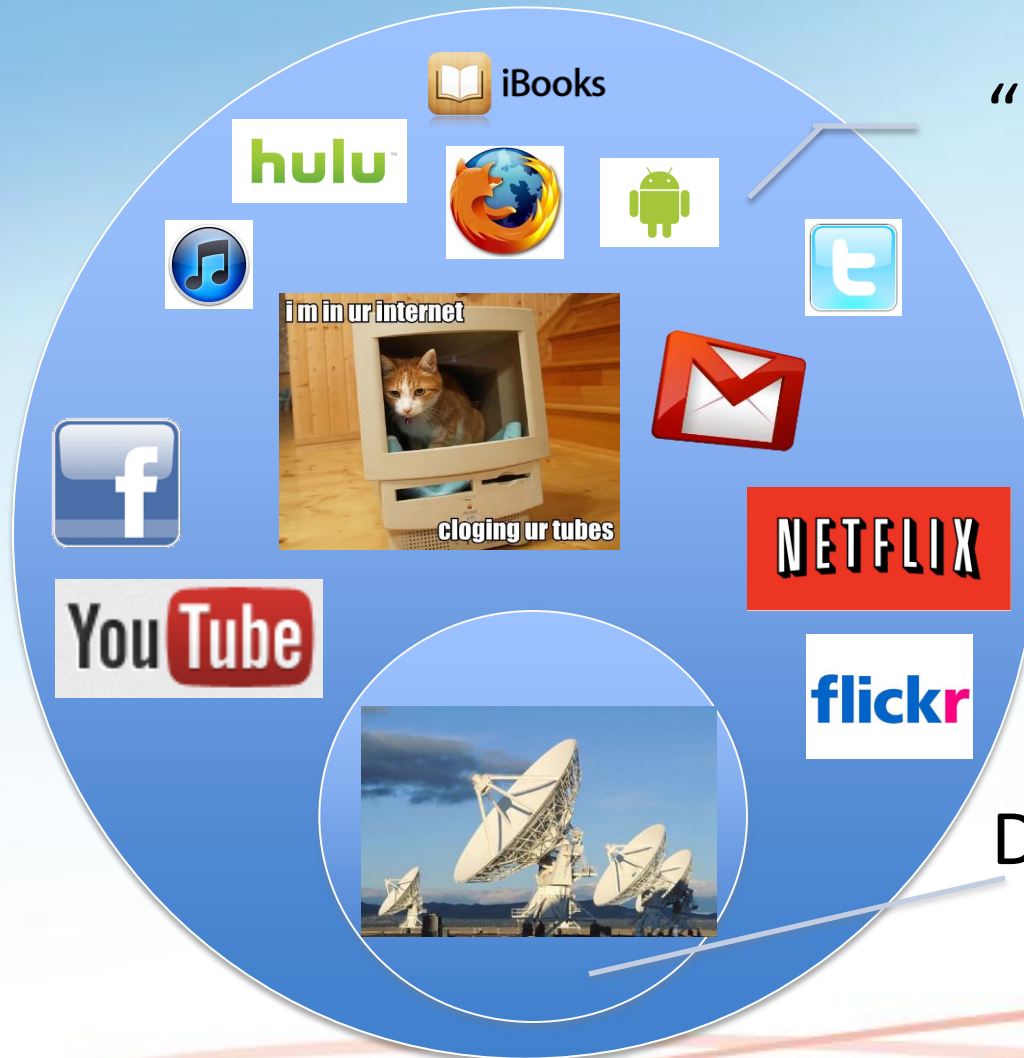
- “What if there was a way to avoid campus/regional/national congestion with a direct point to point link”
 - Option 1: Buy point to point links. Infeasible
 - Option 2: Have the ability to create overlay networks across the R&E infrastructure
- Software Defined Networking
 - Dynamically adding logic into the network fabric
 - End to end “VLANs”
 - Add/subtract bandwidth to a given use as needed
 - Protection from Congestion

Clogging Ur Tubes



"Internets"

Its Tubes All the Way Down



“Internets”

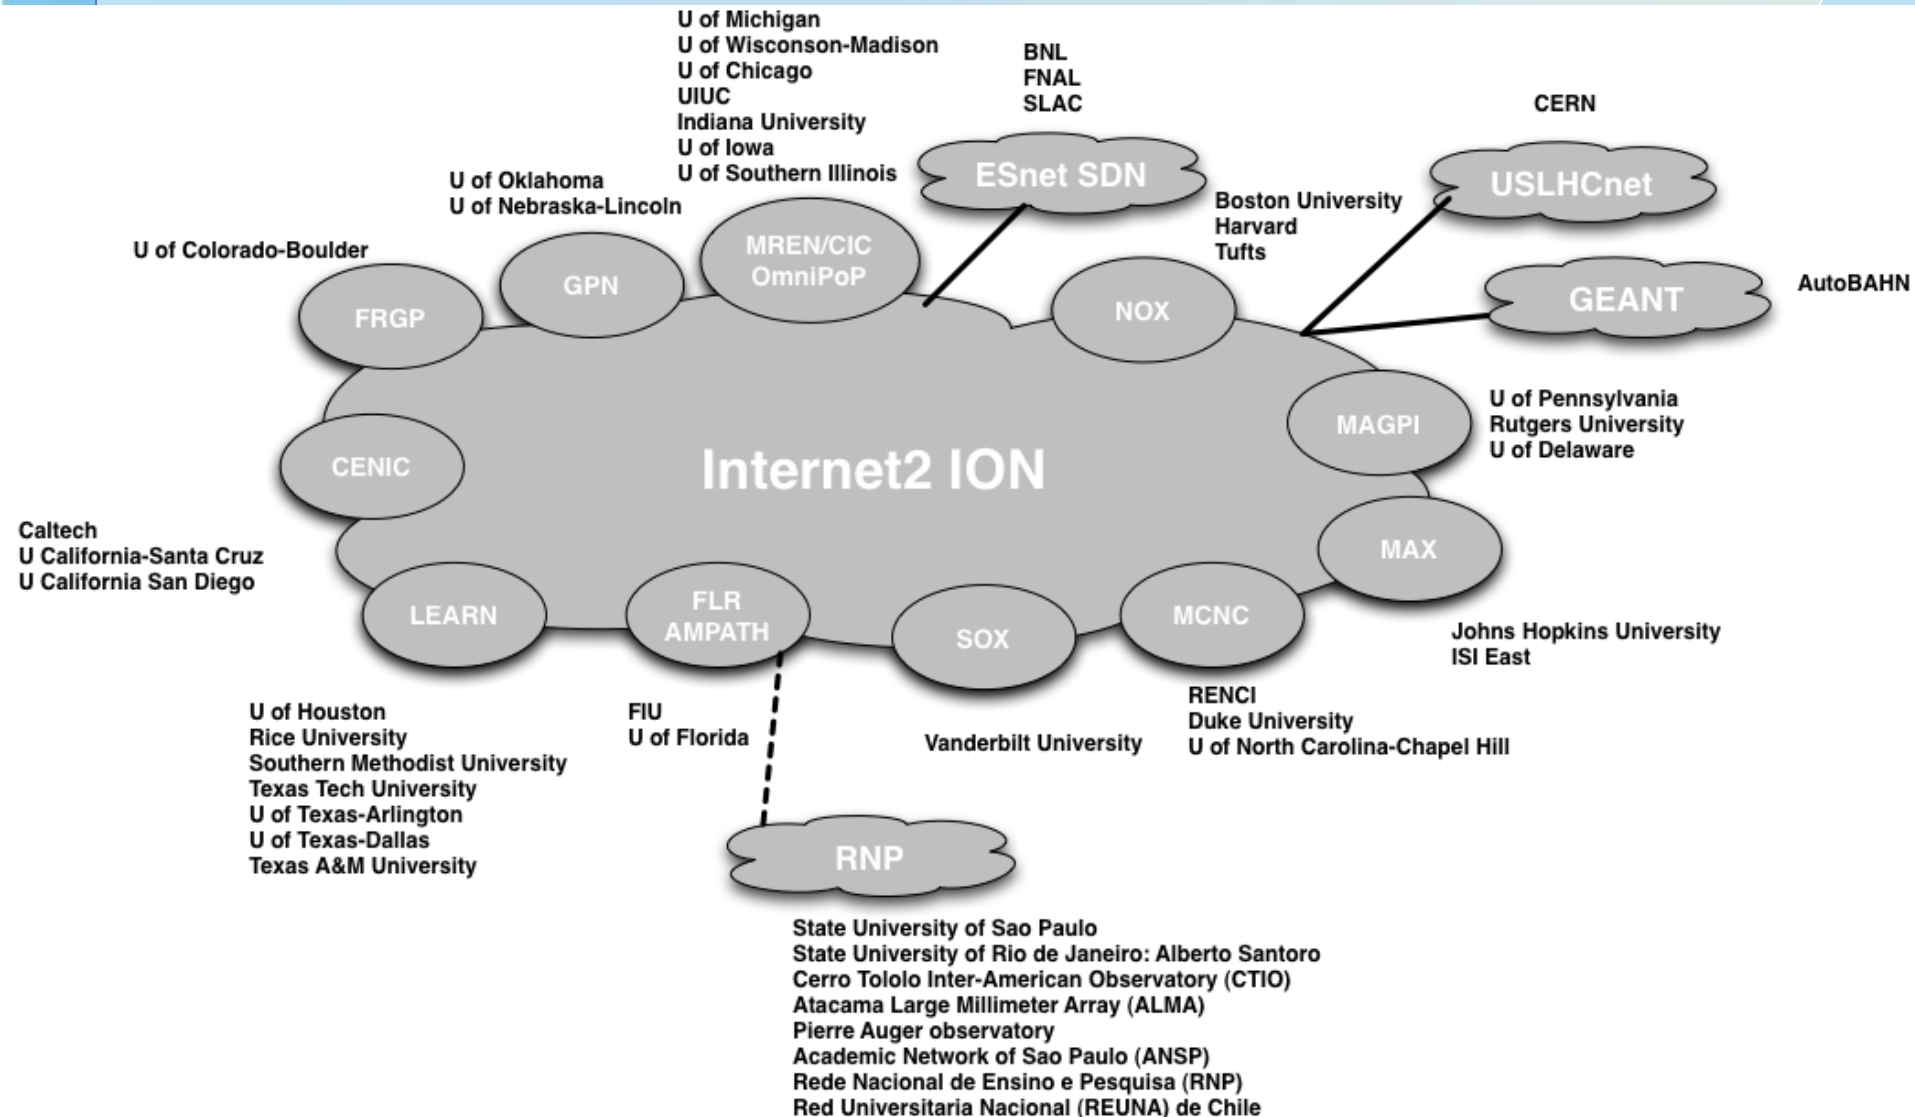
DYNES Circuit

INTERNET
2

DYNES Summary

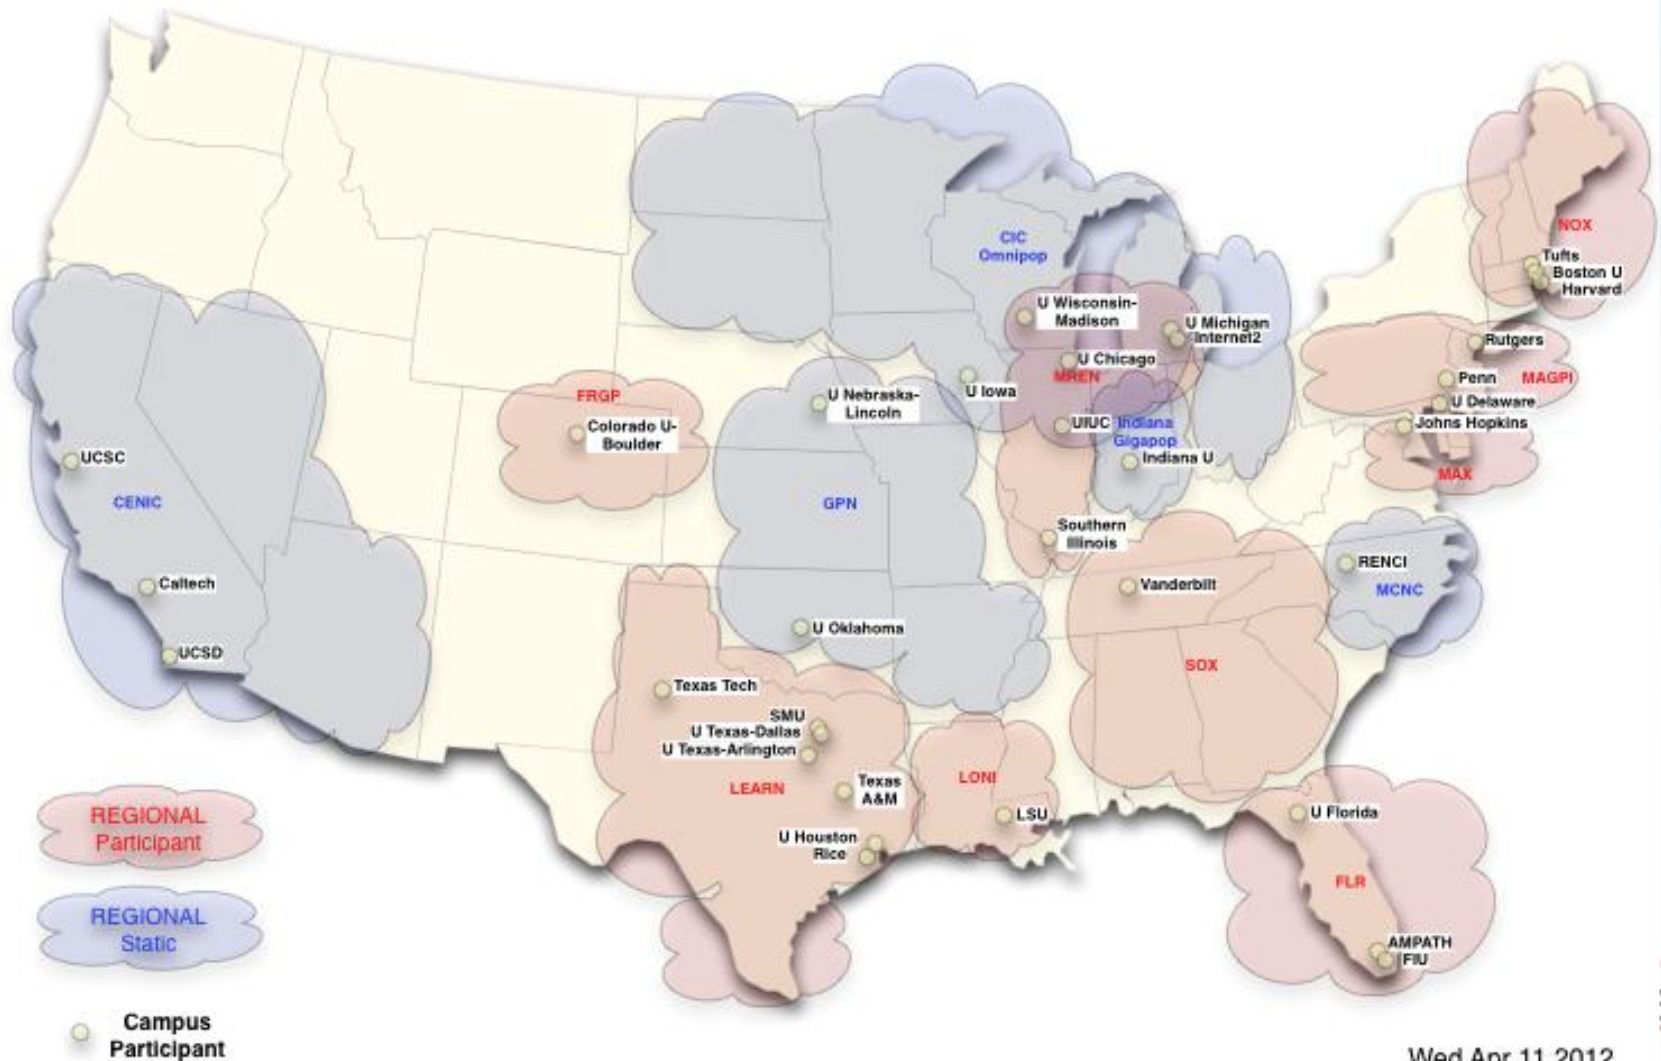
- What is it?:
 - NSF Funded Project (#0958998)
 - A nationwide cyber-instrument spanning ~40 US universities and ~14 Internet2 connectors
 - Extends Internet2's ION service into regional networks and campuses, based on OSCARS implementation of IDC protocol (developed in partnership with ESnet)
 - High-performance file store at sites
- Who is it?:
 - A collaborative team including **Internet2**, **Caltech**, **University of Michigan**, and **Vanderbilt University**
 - Community of regional networks and campuses
 - LHC, astrophysics community, OSG, WLCG, other virtual organizations

DYNES Logical Topology (Spring 2012)



DYNES Deployment Footprint (Spring 2012)

DYNES Participants

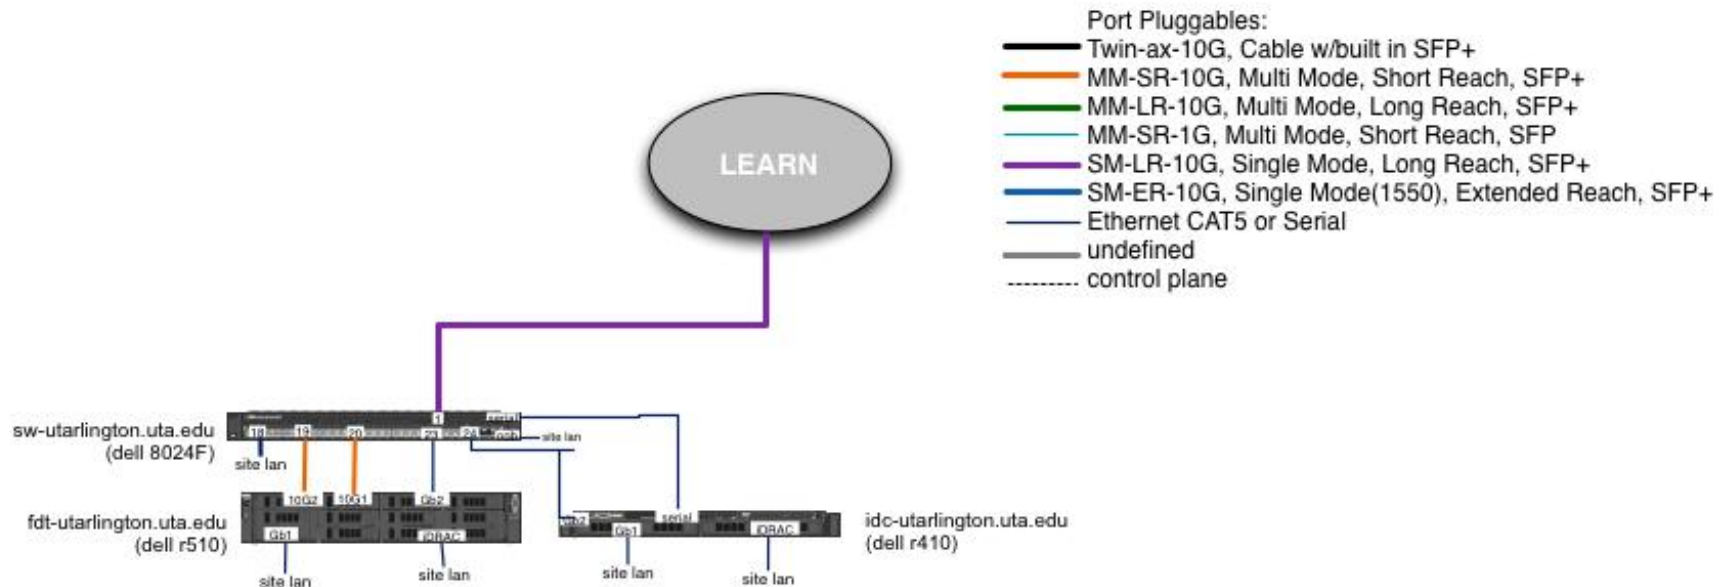


Agenda


- DYNES Motivation and Overview
- **DYNES Hardware and Software**
- Current Status
- Conclusion

Deployment Strategy

- “Black Box”
 - Designed to plug in seamlessly
 - WAN/LAN fiber into the switch/data mover
 - Other local RJ45s
 - Remote management (iDRAC devices)
 - Peering w/ IDCs (national, local)
 - All functionality is built/configured
 - We help to implement Science DMZ ideas on installation, fix performance issues if we see them



DYNES Network Placement – Science DMZ



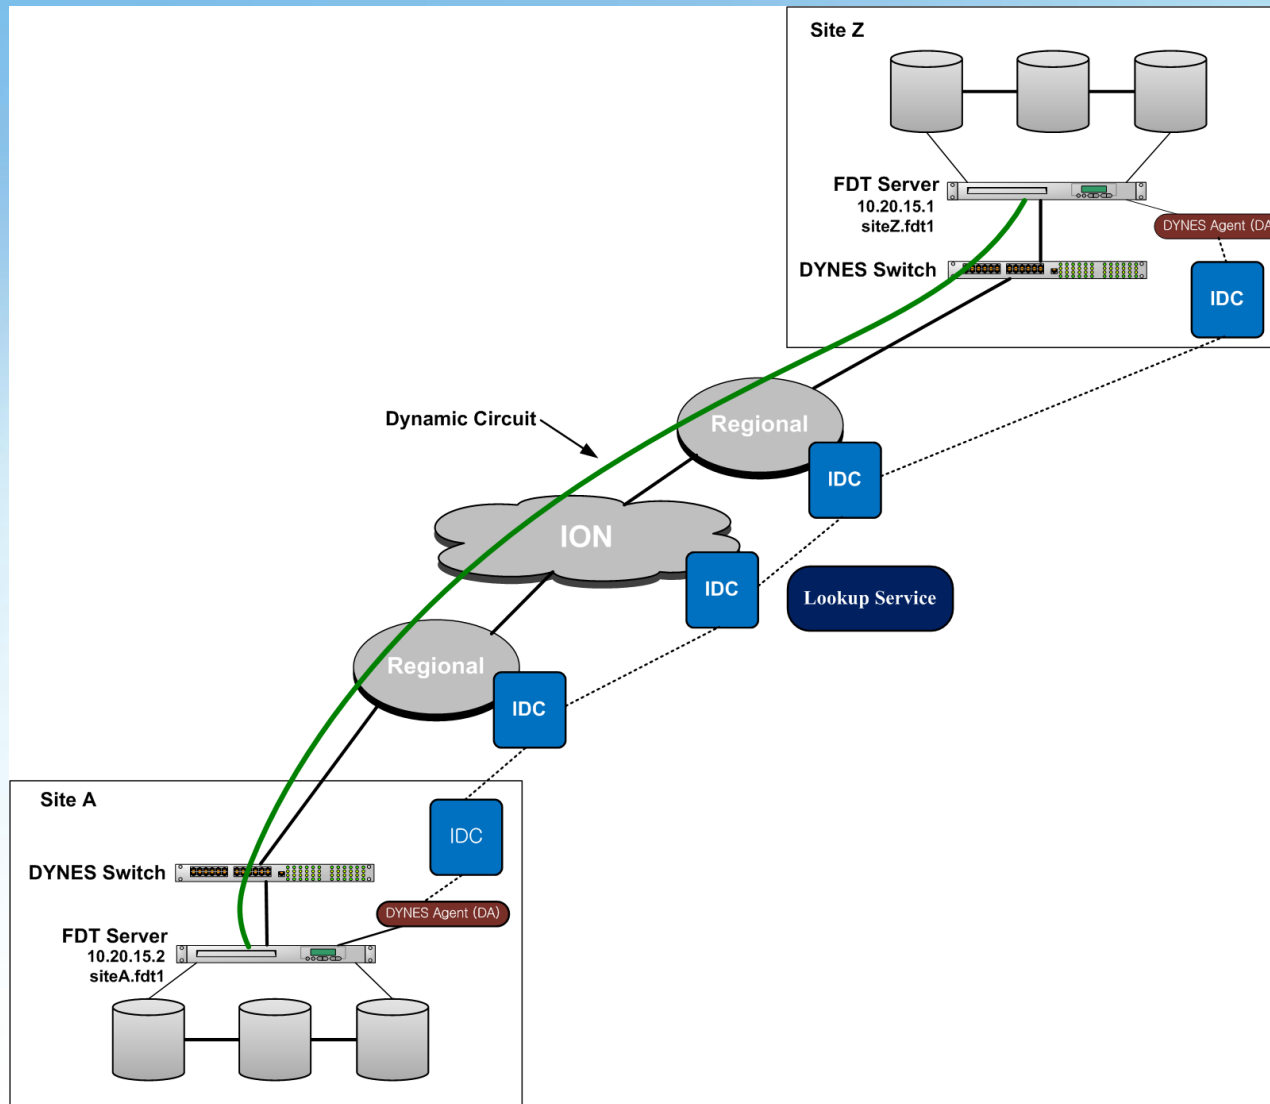
Standard Equipment Overview

- IDC Server
 - Inter-domain/Domain controller. Speaks with OSCARS instances in other domains to arrange circuit management
 - Contains passive measurement tools (e.g. Circuit Monitoring)
- FDT Server
 - Primary data movement server
 - Available active measurement tools (OWAMP, BWCTL)
- Switch
 - Connects FDT and other resources, controlled by IDC server.

Standard Equipment Overview

- <http://www.internet2.edu/ion/hardware.html>
- IDC
 - Dell R410 1U Server
 - Dual 2.4 GHz Xeon (64 Bit), 16G RAM, 500G HD
 - <http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/R410-Spec-Sheet.pdf>
- FDT
 - Dell R510 2U Server
 - Dual 2.4 GHz Xeon (64 Bit), 24G RAM, 300G Main, 12TB through RAID
 - <http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/R510-Spec-Sheet.pdf>
- Switch
 - Dell 8024F, 6248, or S4810
 - 10G vs 1G Sites; copper ports and SFP+; Optics on a site by site basis
 - http://www.dell.com/downloads/global/products/pwcnt/en/PC_6200Series_proof1.pdf
 - <http://www.dell.com/downloads/global/products/pwcnt/en/switch-powerconnect-8024f-spec.pdf>
 - <http://www.dell.com/us/enterprise/p/force10-s4810/pd>

DYNES Data Flow Overview



DYNES Software

- Dynamic Circuit Control
 - OSCARS
 - ION Service
- Monitoring
 - perfSONAR Circuit Monitoring
- Data Movement
 - FDT

DYNES Software – ION/OSCARS

- OSCARS v0.5.4
 - Released March 14 2011
 - Features
 - VLAN translation to allow integration into existing network deployments
 - Robust handling of circuit creation and failures
 - Numerous Bugfixes
 - Additional Documentation/Installation Guidance
 - Security enhancements
- OSCARS v0.6
 - Anticipated Spring 2012
 - Features:
 - Major re-write of the underlying codebase by ESnet
 - Modular, web-services based design
 - Integration with perfSONAR monitoring framework
- DYNES is deploying OSCARS v0.5 and already transitioning to OSCARS v0.6

DYNES Software – Monitoring Circuits

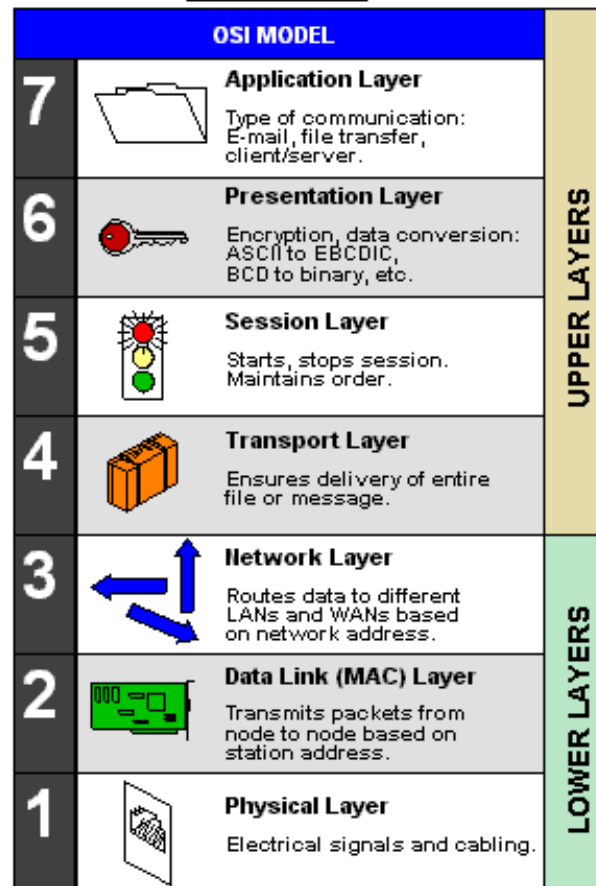
- Goal: to enable users to get measurements in their circuits while allowing domains to provide as much or as little information to the user as the domain wants
- Develop a solution in collaboration with other groups and organizations including DANTE, ESnet, the Network Markup Language Working Group and the Network Measurement Control Working Group
 - Broad agreement ensures that users can monitor their circuits, no matter what domains they traverse
- Multi-faceted approach
 - Enable domains to export monitoring data about circuits
 - Enable users to discover the domains that make up their circuit, and the monitoring data those domains contain about the circuit
- Leverage the standard perfSONAR infrastructure when available

DYNES Software – FDT


- The DYNES Agent (DA) will provide the functionality to request the circuit instantiation, initiate and manage the data transfer, and terminate the dynamically provisioned resources. Specifically the DA will do the following:
 - Accept user request in the form of a DYNES Transfer URLs indicating the data location and ID
 - Locates the remote side DYNES EndPoint Name embedded in the Transfer URL
 - Submits a dynamic circuit request to its home InterDomain Controller (IDC) utilizing its local DYNES EndPoint Name as source and DYNES EndPoint Name from Transfer URL as the destination
 - Wait for confirmation that dynamic circuit has been established
 - Starts and manages Data Transfer using the appropriate DYNES Project IP addresses
 - Initiate release of dynamic circuit upon completion

DYNES Software - OSI Stack

From Computer Desktop Encyclopedia
© 2004 The Computer Language Co., Inc.



Layer Jumping e.g. OSCARS,
FDT controlling the behavior
dynamically



INTERNET
2

Agenda

- DYNES Motivation and Overview
- DYNES Hardware and Software
- **Current Status**
- Conclusion

DYNES Project Schedule

- ~~Phase 1: Site Selection and Planning (Sep-Dec 2010)~~
 - ~~33 Were Accepted in 2 categories~~
 - 8 Regional Networks
 - 25 Site Networks
 - Additional sites under consideration for supplemental deployments
- ~~Phase 2: Initial Development and Deployment (Jan 1-Jun 30, 2011)~~
 - Initial Site Deployment Complete: February 28, 2011
 - Caltech, Vanderbilt, University of Michigan, MAX, USLHCnet
 - Initial Site Systems Testing and Evaluation complete: April 29, 2011
 - Longer term testing (Through July)
 - Evaluating move to CentOS 6
 - New functionality in core software:
 - OSCARS 6
 - perfSONAR 3.2.1
 - FDT Updates

DYNES Project Schedule

- **Phase 3: Scale Up to Full-scale System Development (14 months) (July 1, 2011-August 31, 2012)**
 - Phase 3-Group A Deployment (9 Sites): March 1-Fall, 2011
 - Phase 3-Group B Deployment (13 Sites): July 31-Late Fall, 2011
 - Phase 3-Group C Deployment (11 Sites): Fall 2011- Spring, 2012
 - Full-scale System Development, Testing, and Evaluation (Spring 2012-August 31, 2012)
- **Phase 4: Full-Scale Integration At-Scale; Transition to Routine O&M (12 months) (September 1, 2012-August 31, 2013)**
 - DYNES will be operated, tested, integrated and optimized at scale, transitioning to routine operations and maintenance as soon as this phase is completed

Phase 3 Group A Members

- AMPATH
 - Mid-Atlantic Crossroads (MAX)
 - The Johns Hopkins University (JHU)
 - Mid-Atlantic Gigapop in Philadelphia for Internet2 (MAGPI)*
 - Rutgers (via NJEdge)
 - University of Delaware
 - Southern Crossroads (SOX)
 - Vanderbilt University
 - CENIC*
 - California Institute of Technology (Caltech)
 - MREN*
 - University of Michigan (via MERIT and CIC OmniPoP)
- Note: USLHCNet will also be connected to DYNES Instrument via a peering relationship with DYNES

* temp configuration of static VLANs until future group

Phase 3 Group B Members

- Mid-Atlantic Gigapop in Philadelphia for Internet2 (MAGPI)
 - University of Pennsylvania
- Metropolitan Research and Education Network (MREN)
 - Indiana University (via I-Light and CIC OmniPoP)
 - University of Wisconsin Madison (via BOREAS and CIC OmniPoP)
 - University of Illinois at Urbana-Champaign (via CIC OmniPoP)
 - The University of Chicago (via CIC OmniPoP)
- Lonestar Education And Research Network (LEARN)
 - Southern Methodist University (SMU)
 - Texas Tech University
 - University of Houston
 - Rice University
 - The University of Texas at Dallas
 - The University of Texas at Arlington
- Florida International University (Connected through FLR)



Phase 3 Group C Members

- Front Range GigaPop (FRGP)
 - University of Colorado Boulder
- Northern Crossroads (NoX)
 - Boston University
 - Harvard University
 - Tufts University
- CENIC**
 - University of California, San Diego
 - University of California, Santa Cruz
- CIC OmniPop***
 - The University of Iowa (via BOREAS)
- Great Plains Network (GPN)***
 - The University of Oklahoma (via OneNet)
 - The University of Nebraska-Lincoln
- FLR
 - University of Florida
- MCNC***
 - RENCI
- Louisiana Optical Network Initiative (LONI) ***
 - Louisiana State University

Phase 3 Group D (Planned) Members

- Texas A&M University
- National Capital Aggregation Facility (NatCap)
 - University of Virginia
- Southern Illinois University
- MD Anderson Cancer Center
- SUNY Stony Brook
- Utah Education Network (UEN)
 - Utah State University
 - University of Utah
- Colorado State University
- Princeton University
- Temple University
- Three Rivers Optical Exchange (3RoX)
 - Drexel University
- Michigan State University
- Oregon GigaPoP
 - University of Oregon
 - Oregon State University
- Mississippi Optical Network (MissION)
 - Mississippi Optical Network

Agenda

- DYNES Motivation and Overview
- DYNES Hardware and Software
- Current Status
- **Conclusion**

DYNES Additional Activities

- We anticipate that we will be able to add a some more sites. Additional applications are being collected from those that are interested.



- Send email to dynes-questions@internet2.edu if interested

DYNES References

- Documentation:
 - <http://www.internet2.edu/dynes>
- DYNES
 - <http://www.internet2.edu/dynes>
- OSCARS
 - <http://www.es.net/oscars>
- DRAGON
 - <http://dragon.east.isi.edu>
- DCN Software Suite (DCNSS)
 - <http://wiki.internet2.edu/confluence/display/DCNSS/>
- FDT
 - <http://monalisa.cern.ch/FDT/>
- perfSONAR-PS
 - <http://psps.perfsonar.net>



The DYNES Instrument: A Description and Overview

May 24th 2012 – CHEP 2012

Jason Zurawski, Senior Network Engineer - Internet2

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