



# ESnet

ENERGY SCIENCES NETWORK

# Science DMZ Implementation at the Advanced Photon Source

Jason Zurawski – [zurawski@es.net](mailto:zurawski@es.net)

Science Engagement Engineer, ESnet

Lawrence Berkeley National Laboratory

GlobusWorld 2016

April 22<sup>nd</sup> 2016



U.S. DEPARTMENT OF  
**ENERGY**

Office of Science



# Outline

- Preliminaries
- The Science
- The Project
- The Results

# The Science DMZ in 1 Slide



# Science DMZ now recognized as a global best practice.



NSF is investing ~\$130M to promote adoption by US universities (among other CI goals). Fourth funding round underway.

>120 universities in the US have deployed or are considering this DOE architecture.

In addition: USDA, NIH – with NASA, NOAA investigating.

Australian, Canadian universities following suit.




# Outline

- Preliminaries
- **The Science**
- The Project
- The Results



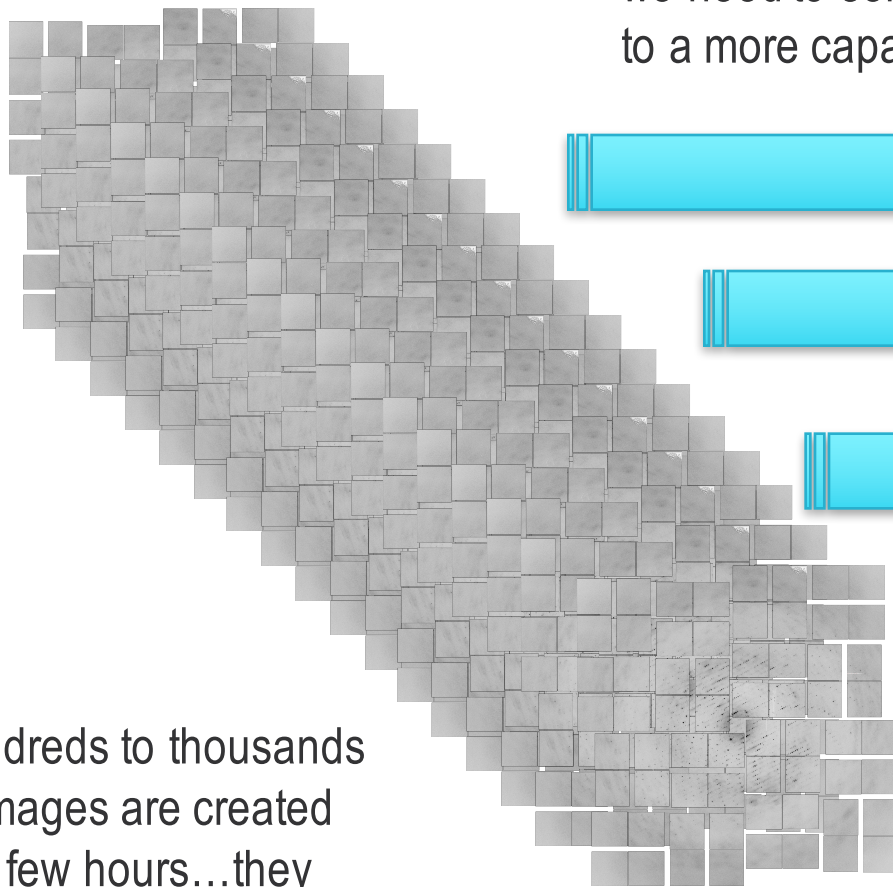
# GM/CA

- General Medical Sciences and Cancer Institutes Structural Biology Facility @ APS
  - Focused on the study of crystallographic structure determination of biological macromolecules by X-ray diffraction
- Typical User:
  - Visits and spends some allocated time with samples and the beamline machinery
  - Mails samples, controls device remotely
  - In either case – data has to go ‘somewhere’



# E Pluribus Unum

Processing on this order of magnitude can't be done locally – we need to send (over a network) to a more capable facility

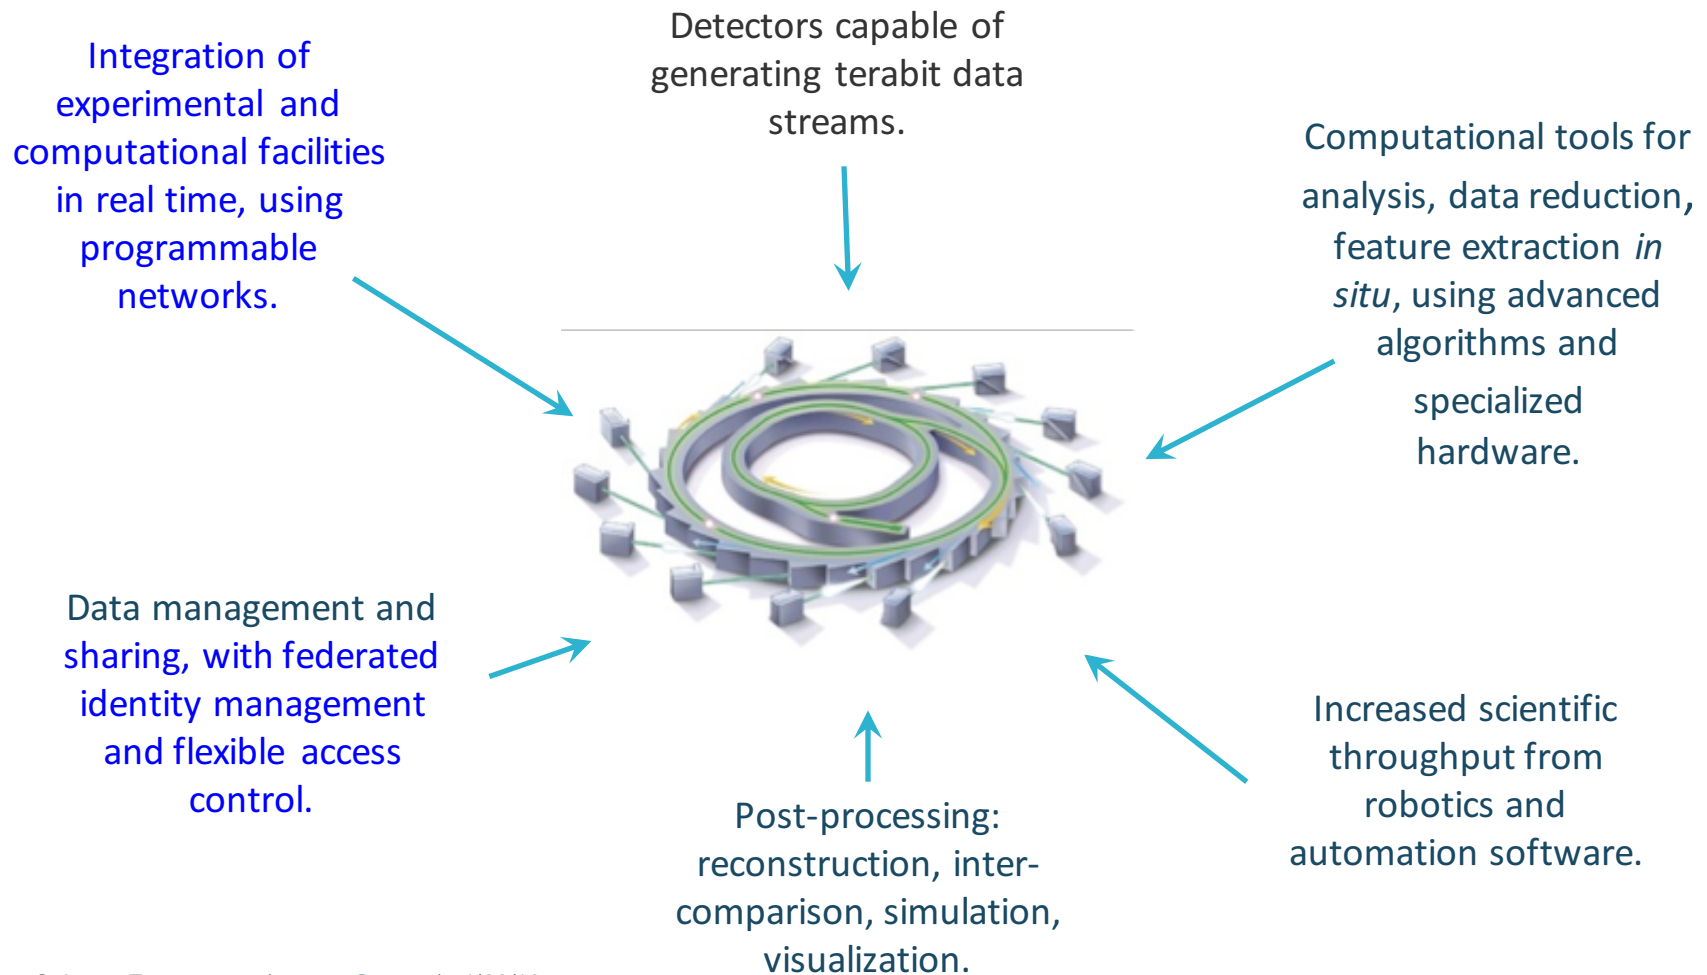


Hundreds to thousands of images are created in a few hours...they can range in size from MB to TB

After processing on a supercomputer, models are created.

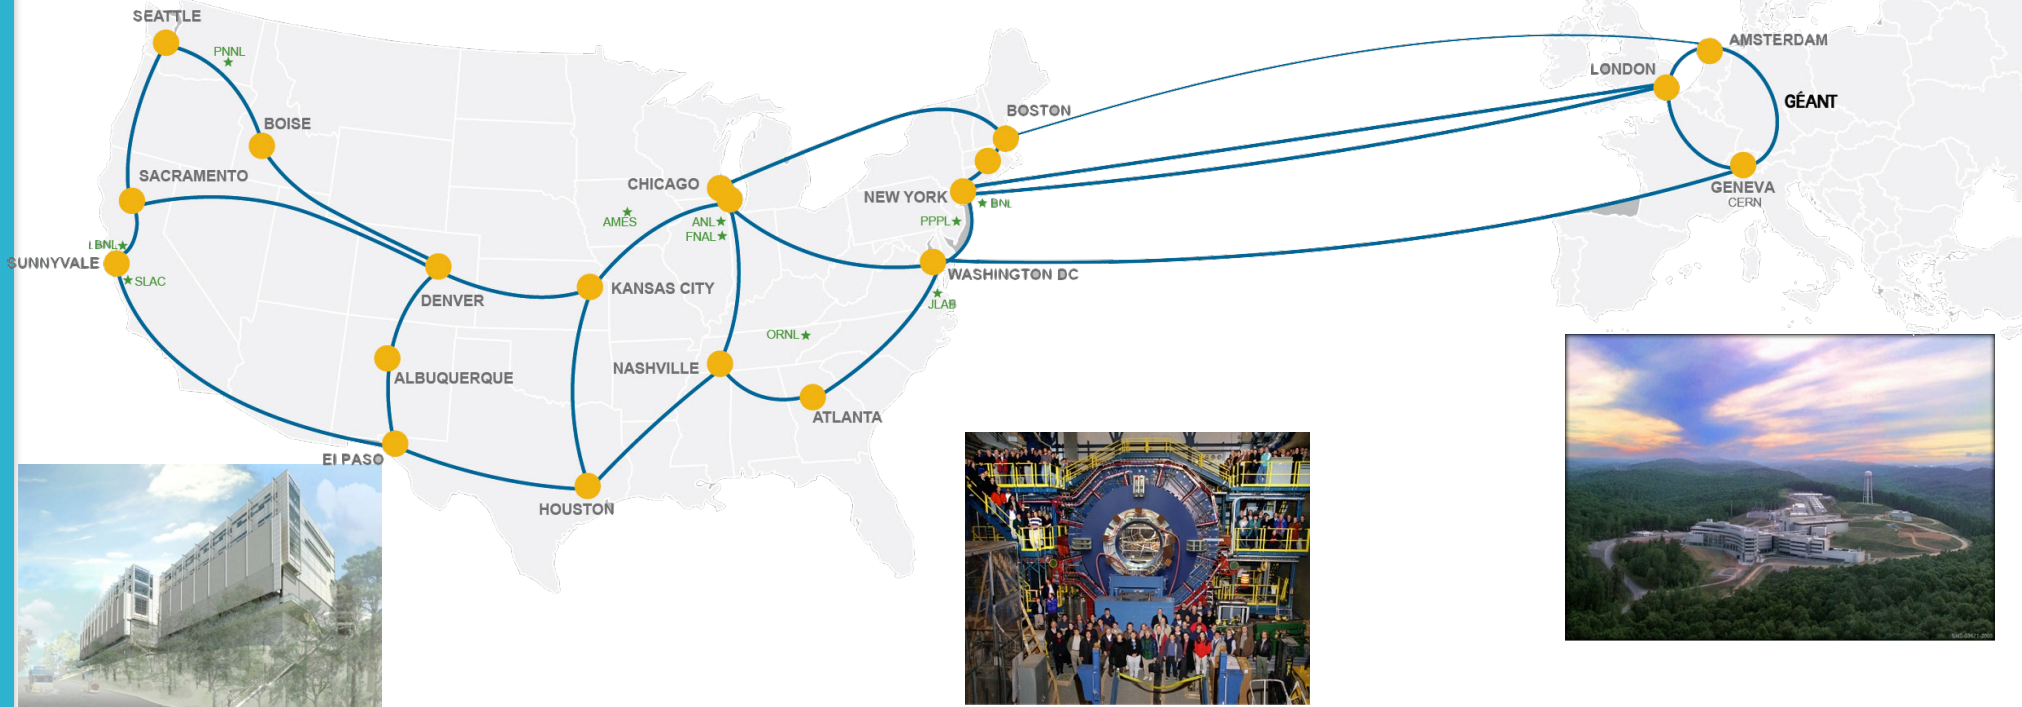


# Superfacility instances multiply.

By 2025, DOE light sources will be transformed by high-resolution detectors, advanced mathematical analysis techniques, robotics, software automation, programmable networks.





# Network as Infrastructure *Instrument*



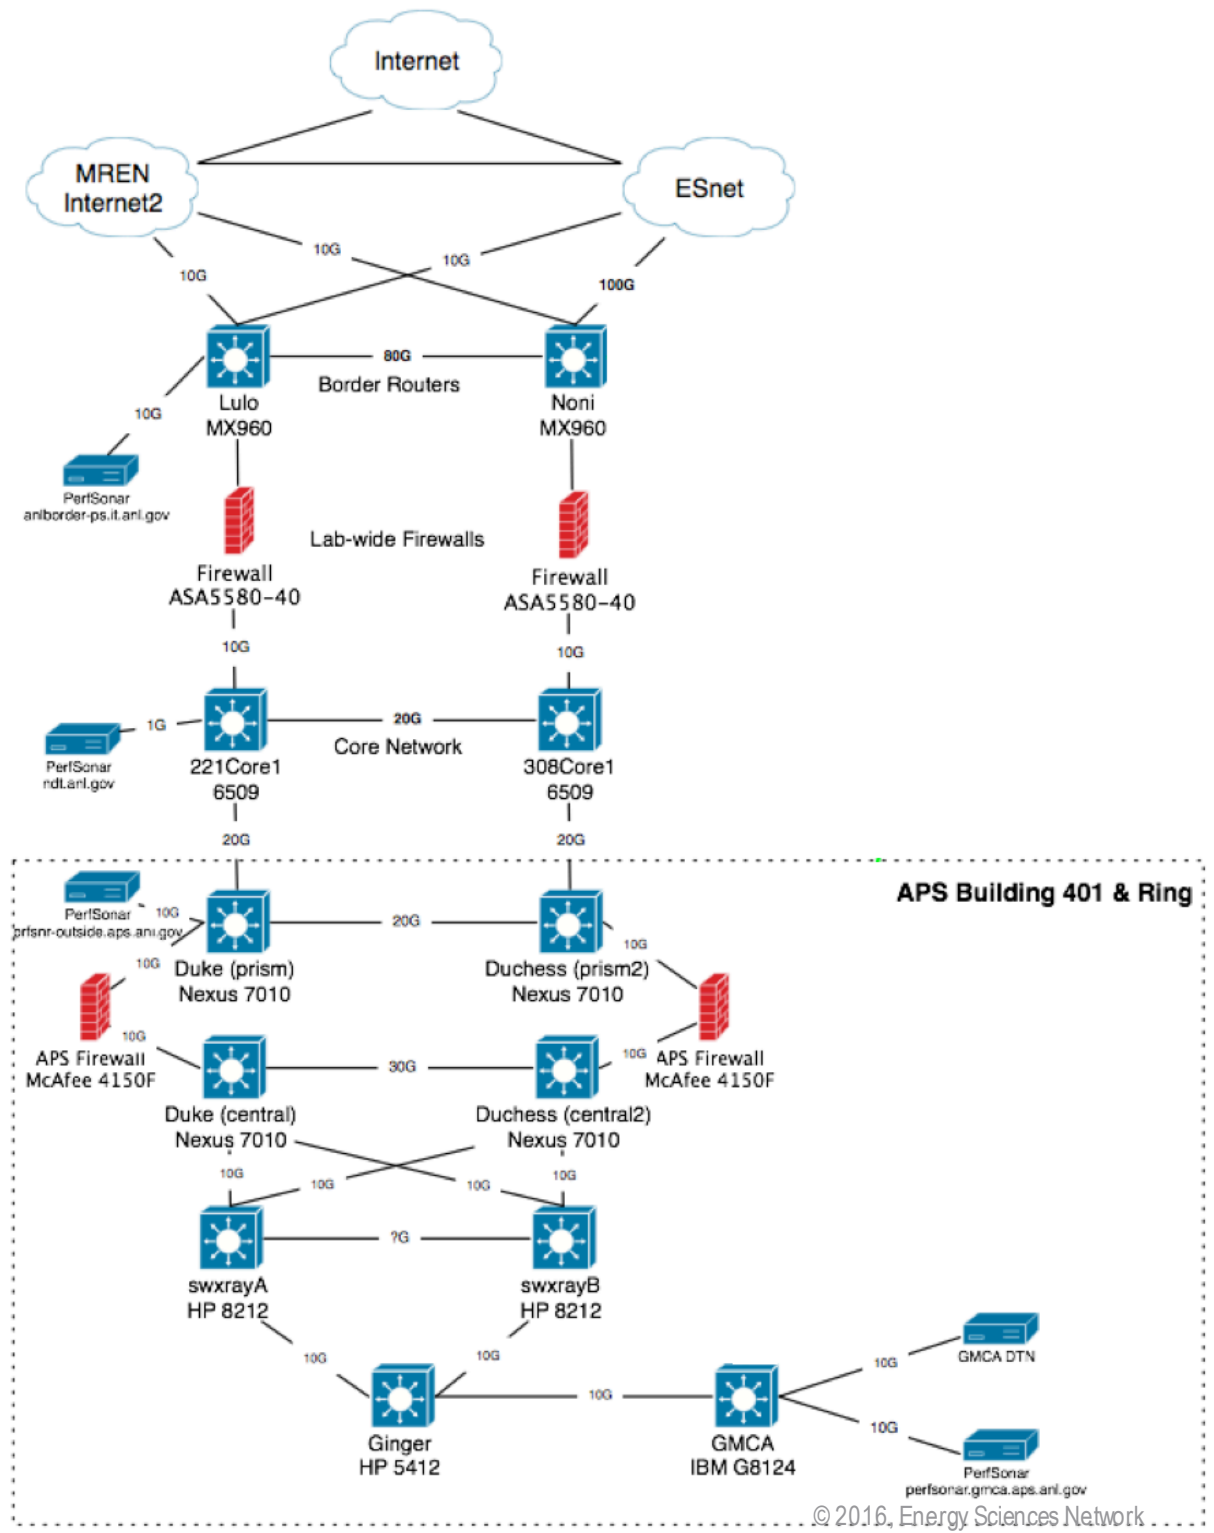
**ESnet Vision:** Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.



# Outline

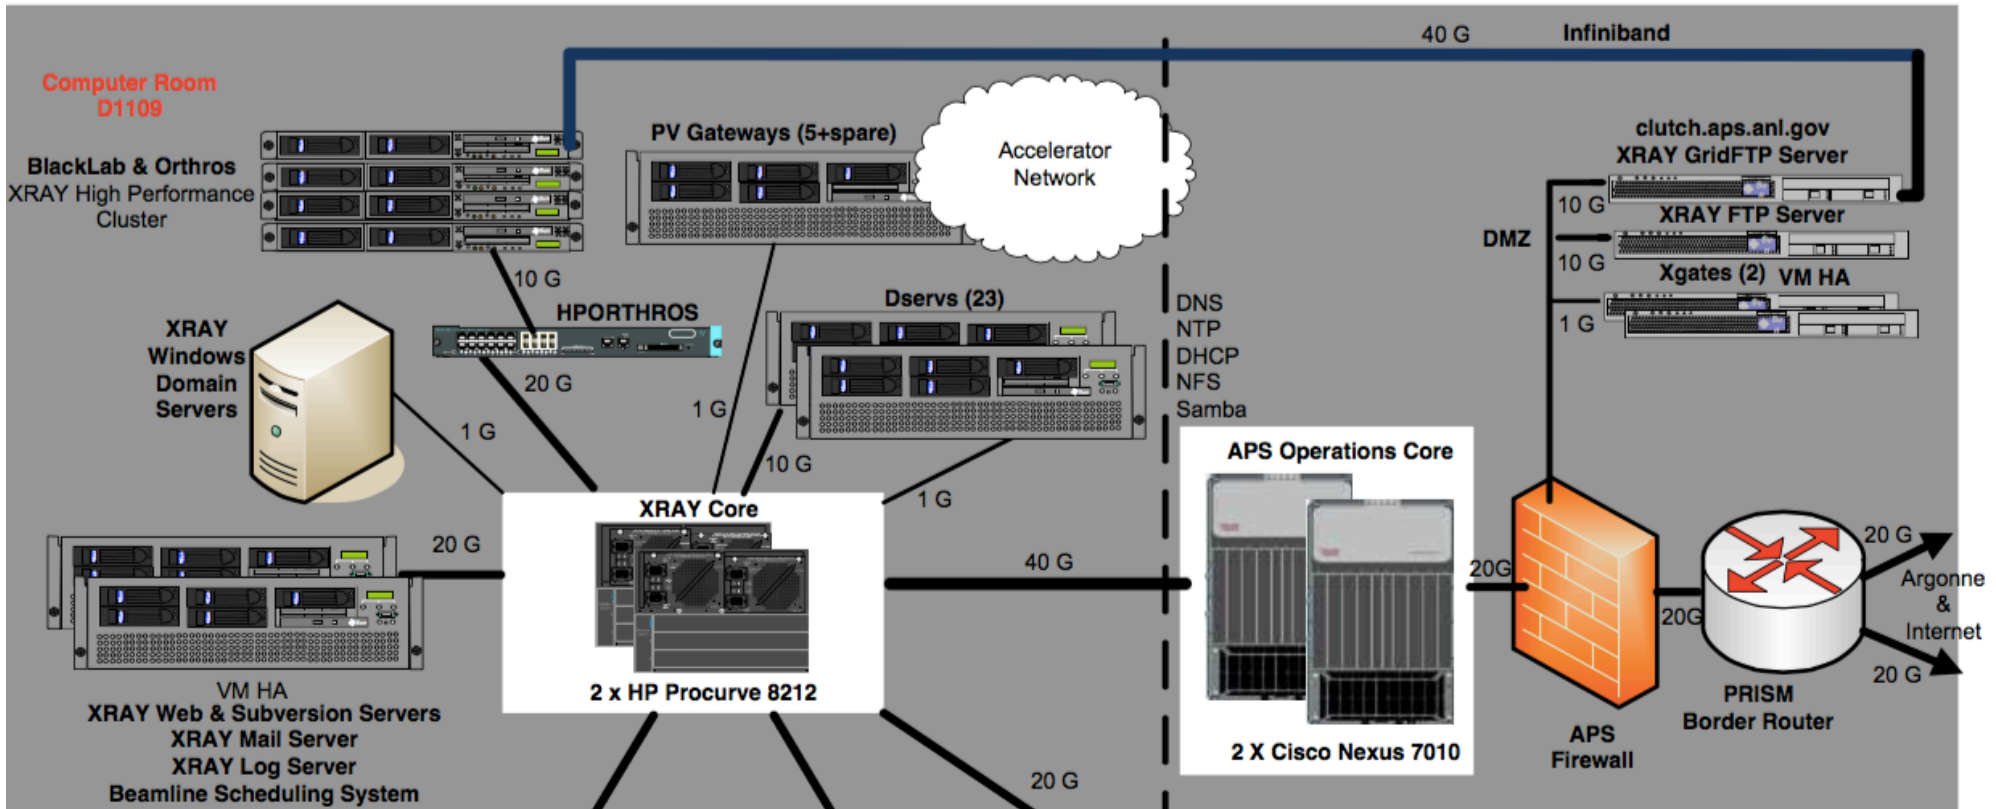
- Preliminaries
- The Science
- **The Project**
- The Results

# Current Design



# Current Design

## XRAY Beamline Logical Network



# ANL Border to ESnet New York (perfSONAR Test)

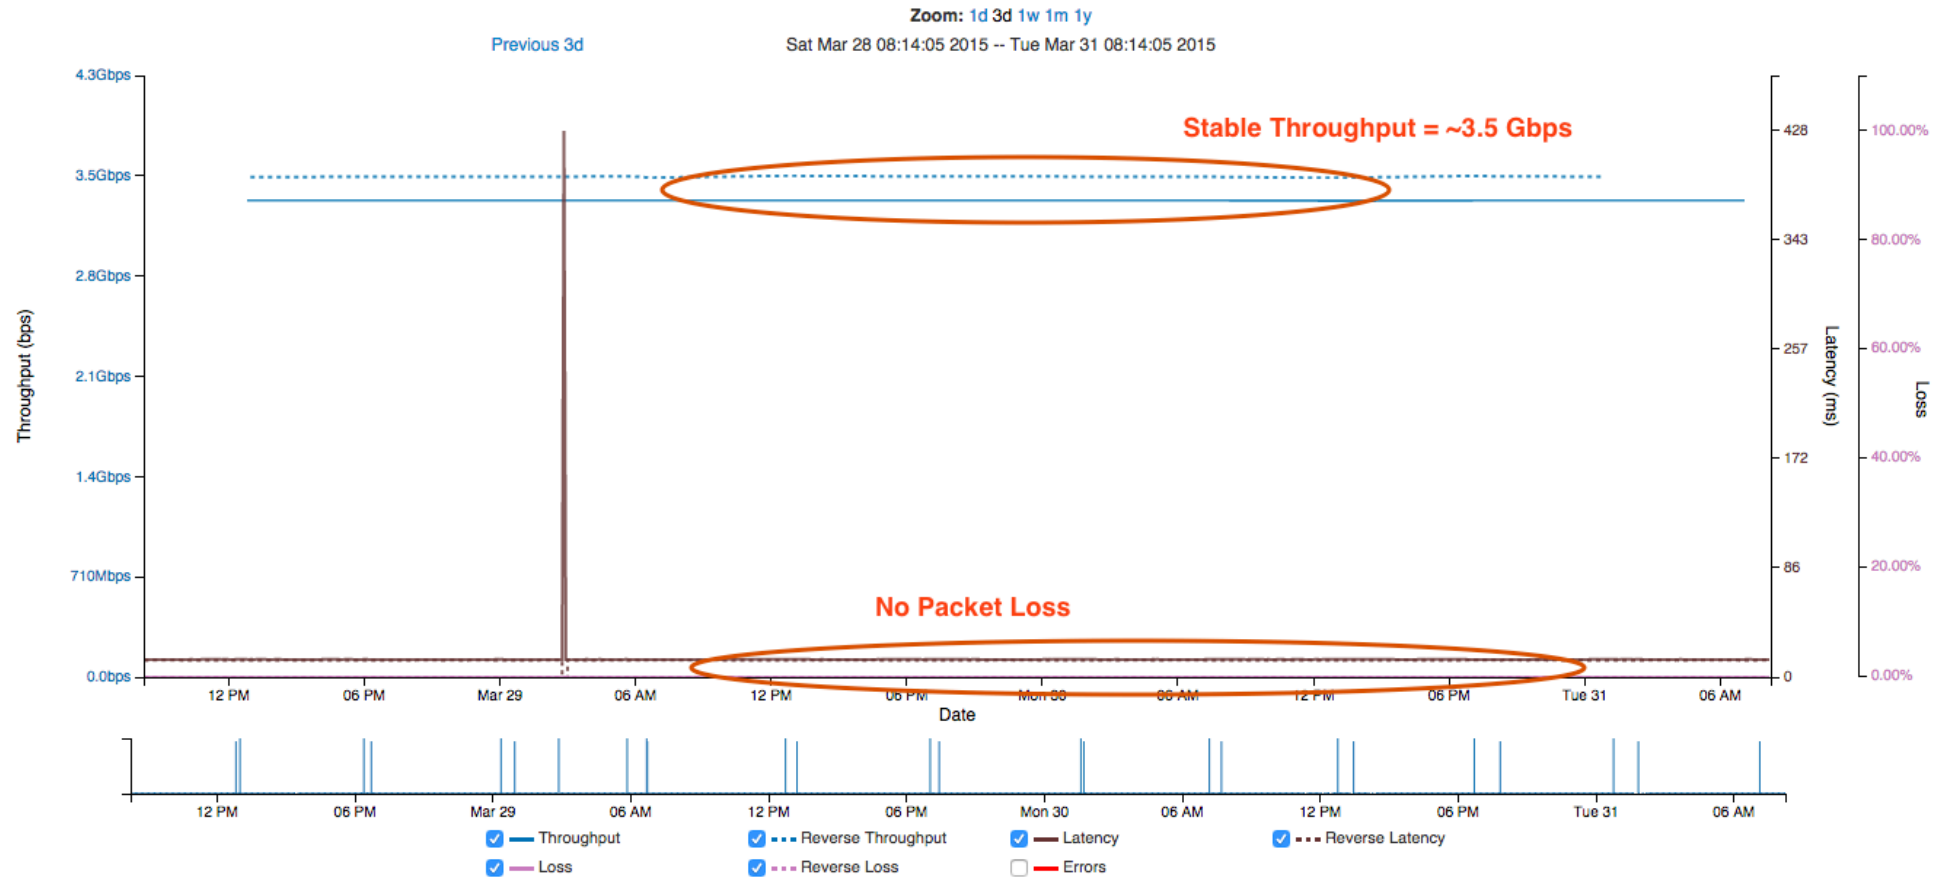
**Source**  
**anlborder-ps.it.anl.gov** - 130.202.222.58  
 Capacity: 10G MTU: 9000

**Destination**  
**newy-pt1.es.net** - 198.124.238.54  
 Capacity: 10G MTU: 9000

[Link to this chart](#)

**anlborder-ps.it.anl.gov** - 130.202.222.58 -- **newy-owamp.es.net** - 198.124.252.158 [traceroute]  
 Capacity: 10G MTU: 9000 Capacity: 1.0G MTU: 1500

Negative latency values found in the reverse direction. Typically, this occurs when one or both hosts' clocks are out of sync, or the hosts are very close together.



# GM/CA to ESnet New York (perfSONAR Test)

**Source**  
 perfsonar.gmca.aps.anl.gov - 164.54.103.141 -- Capacity: 10G MTU: 1500  
**Destination**  
 newy-pt1.es.net - 198.124.238.54 -- Capacity: 10G MTU: 9000  
[Link to this chart](#)

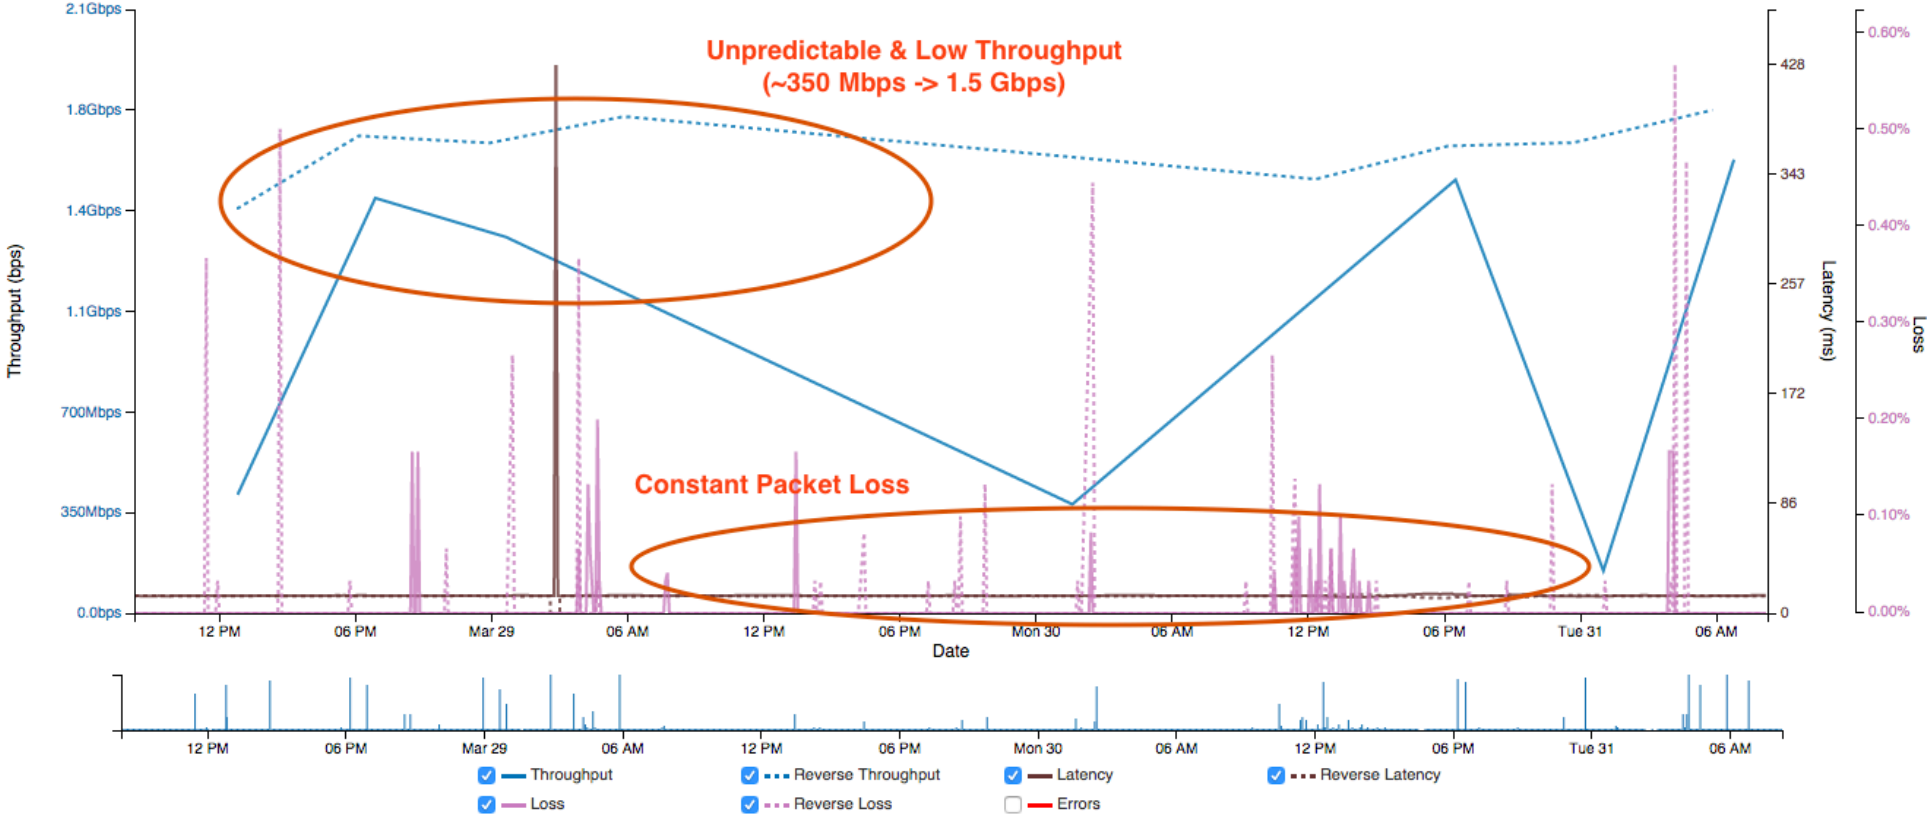
perfsonar.gmca.aps.anl.gov - 164.54.103.141 -- Capacity: 10G MTU: 1500  
 newy-owamp.es.net - 198.124.252.158 -- Capacity: 1.0G MTU: 1500

Negative latency values found in the reverse direction. Typically, this occurs when one or both hosts' clocks are out of sync, or the hosts are very close together.

Zoom: 1d 3d 1w 1m 1y

Sat Mar 28 08:14:42 2015 -- Tue Mar 31 08:14:42 2015

Previous 3d



# GM/CA to ANL Border (perfSONAR Test)

Source: **personar.gmca.aps.anl.gov** - 164.54.103.141 -- Destination: **anlborder-ps.it.anl.gov** - 130.202.222.58  
Capacity: Unknown MTU: Unknown

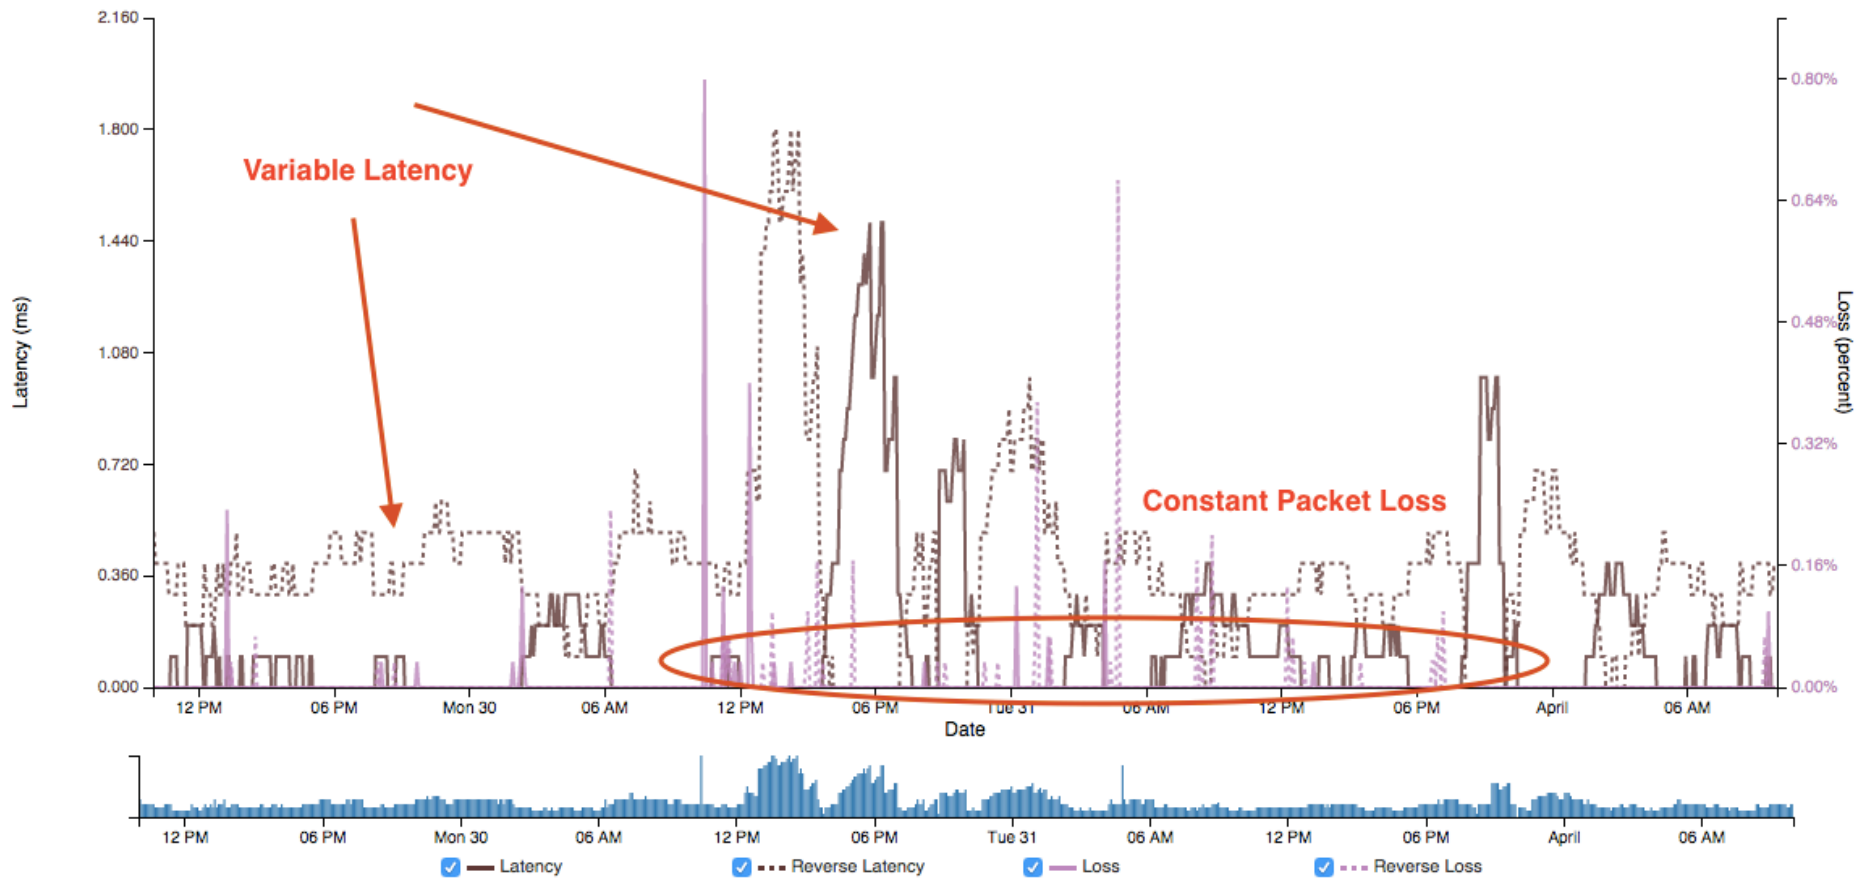
[Link to this chart](#)

Negative latency values found in both directions. Typically, this occurs when one or both hosts' clocks are out of sync, or the hosts are very close together.

Zoom: 1d 3d 1w 1m 1y

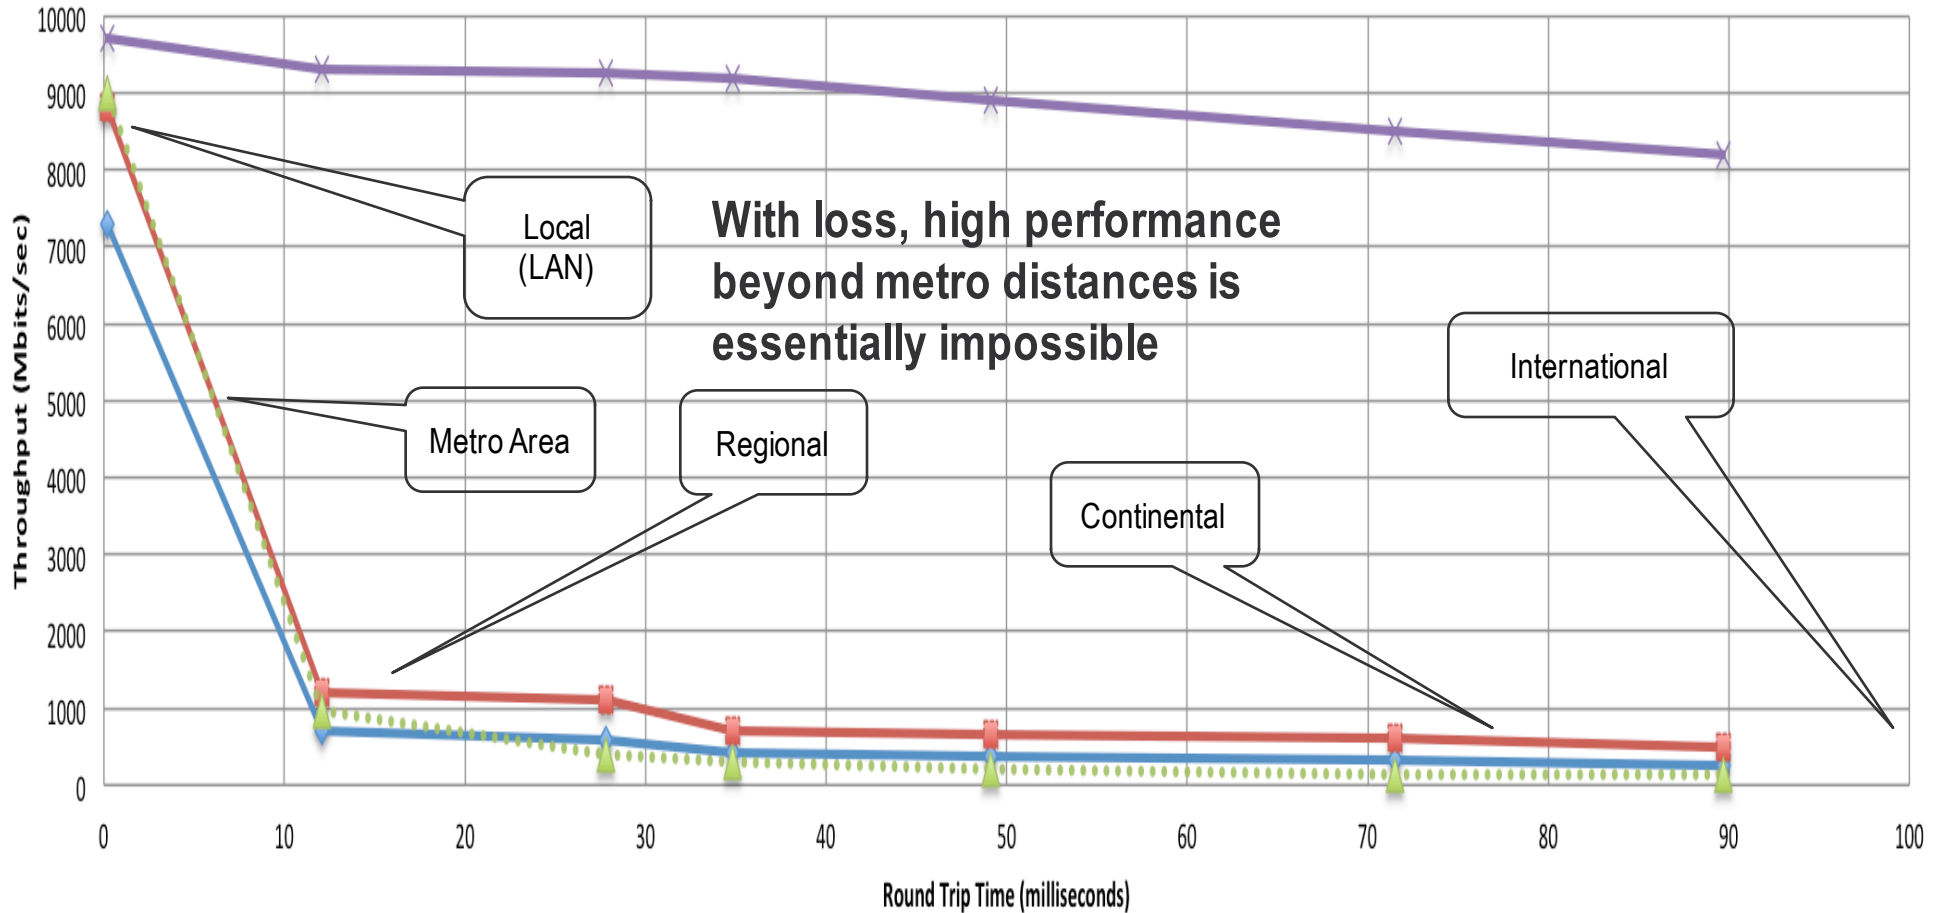
Previous 3d

Sun Mar 29 09:57:59 2015 -- Wed Apr 1 09:57:59 2015



# A small amount of packet loss makes a huge difference in TCP performance

Throughput vs. Increasing Latency with .0046% Packet Loss



Measured (TCP Reno)

Measured (HTCP)

Theoretical (TCP Reno)

Measured (no loss)

$$\text{Throughput} \leq \frac{\text{MSS}}{\text{RTT} \sqrt{P_{\text{loss}}}}$$

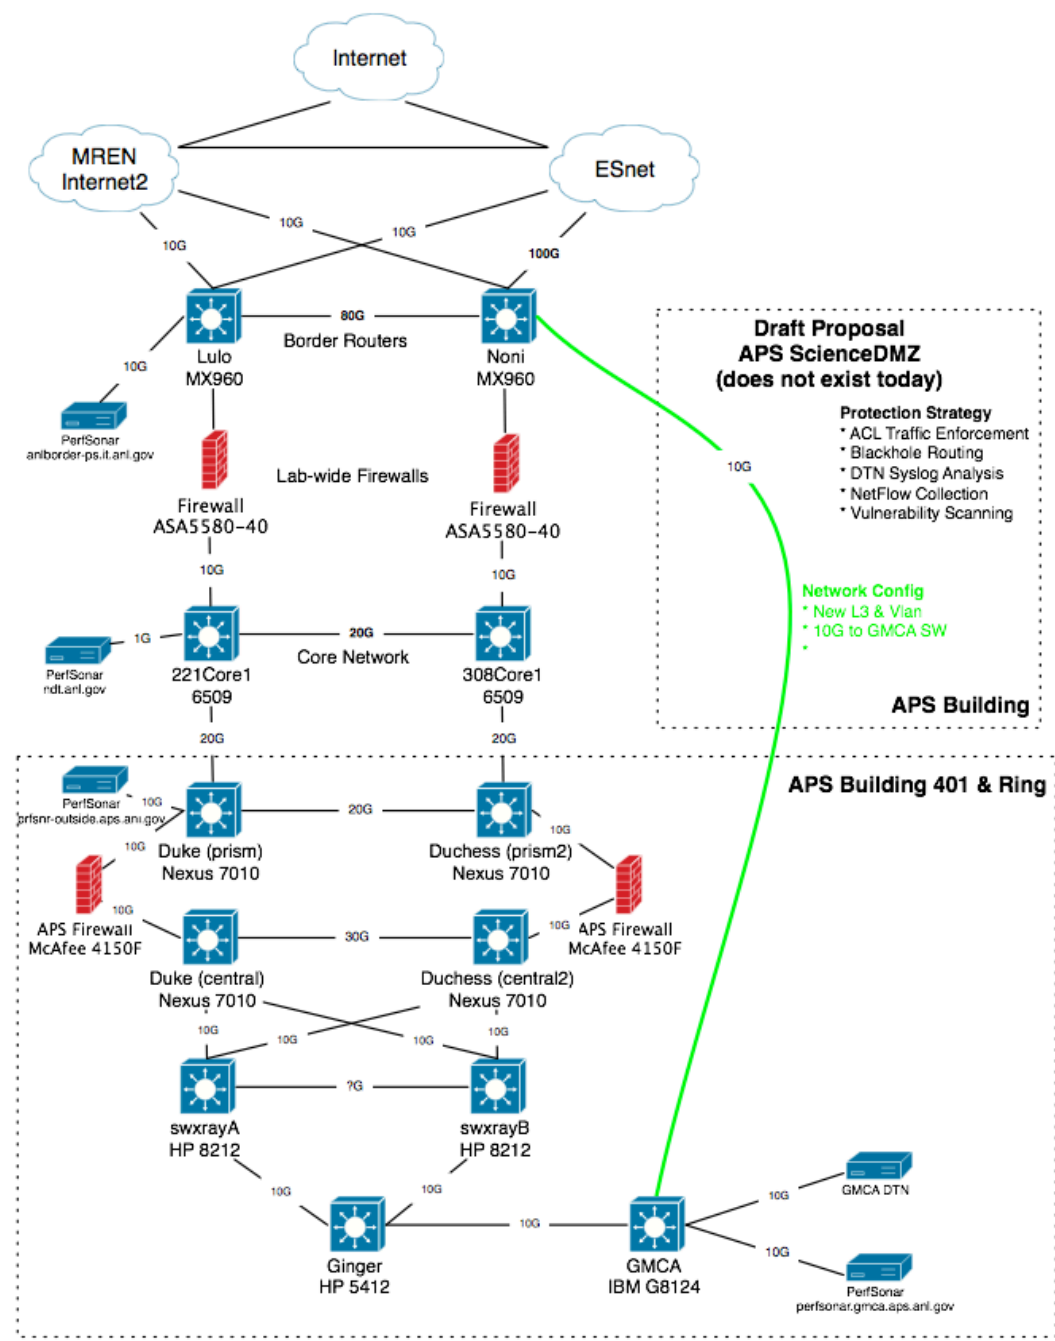


# Outline

- Preliminaries
- The Science
- The Project
- **The Results**

# Pilot Network

- Plan:
  - Use campus and building fiber resources for 2<sup>nd</sup> path
  - Applies to only a limited set of resources (perfSONAR, DTN)
- Benefits
  - Bypass congested local infrastructure
  - Apply targeted (vs. blanket) security policy
- Cautions:
  - Prevent just ‘anything’ from using fast path – policy to control this
  - Still need to figure out cause of local issues (e.g. this isn’t a pave over)



# perfSONAR Results

Source	Destination
192.5.180.130 - 192.5.180.130 Capacity: Unknown MTU: Unknown	-- elpa-pt1.es.net - 198.129.254.82 [traceroute] Capacity: 10G MTU: 9000
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	-- elpa-pt1.es.net - 198.129.254.82 Capacity: 10G MTU: 9000

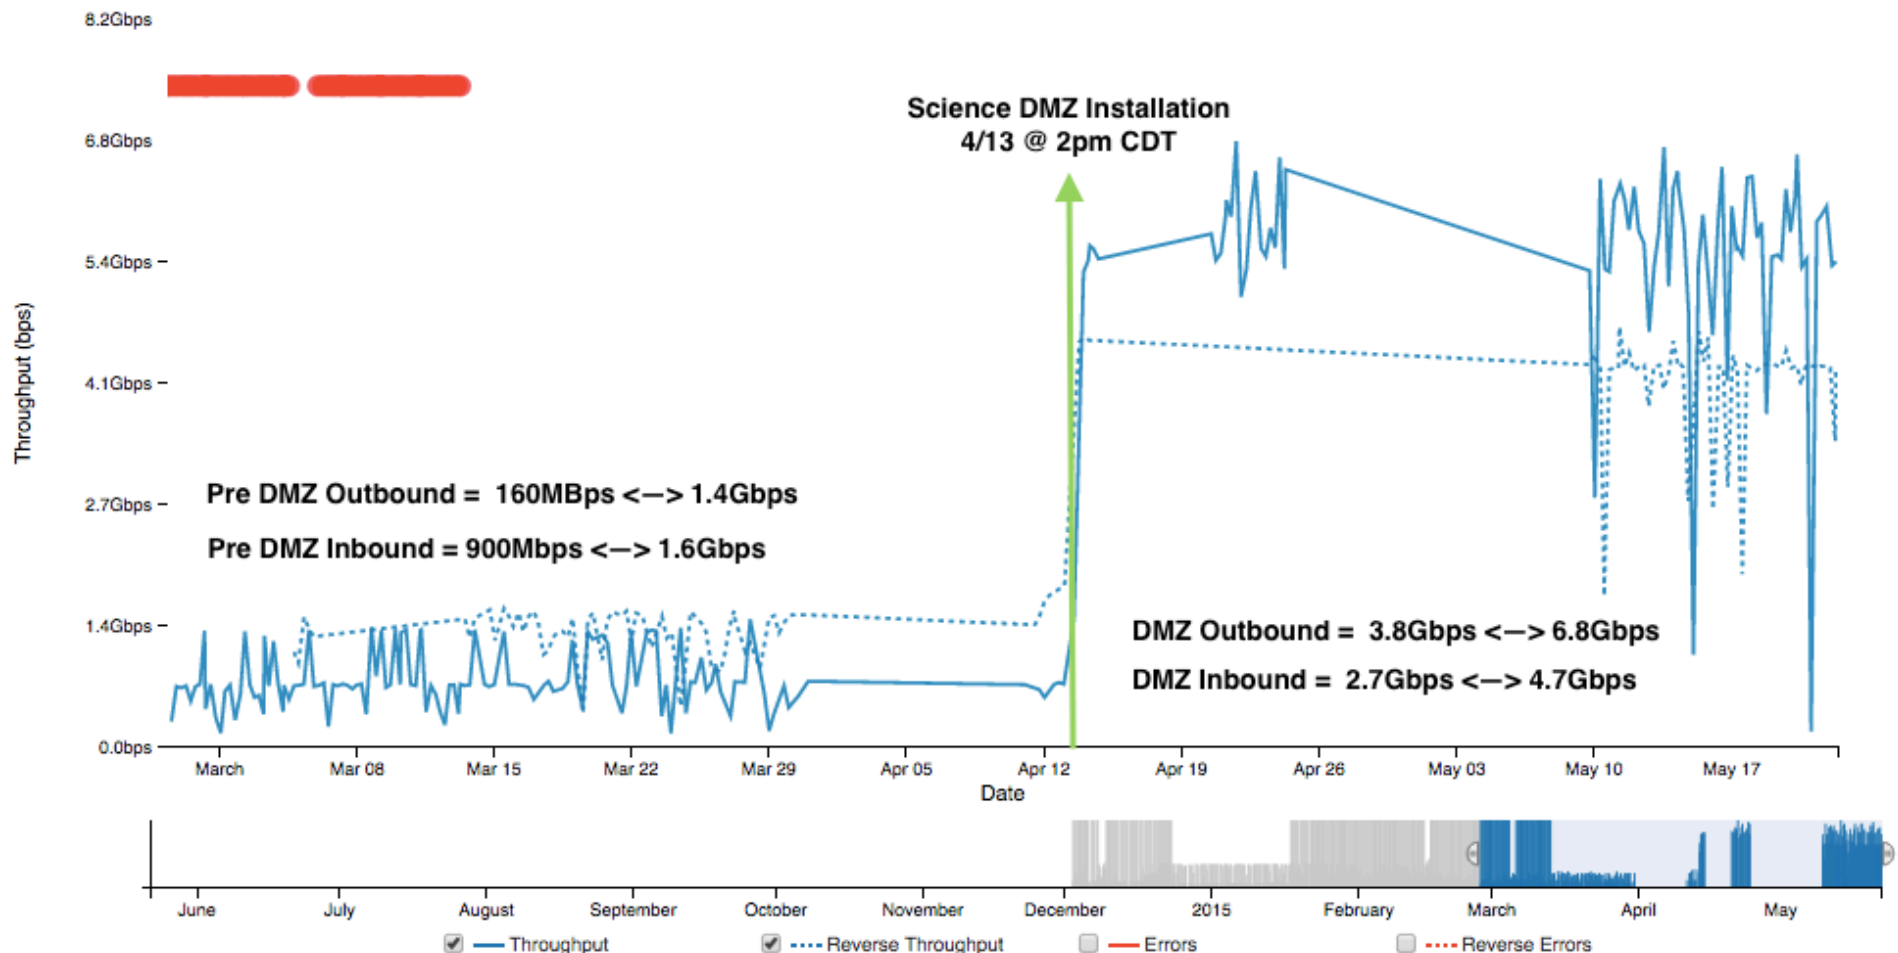
[Link to this chart](#)

**~40ms RTT from ANL to ESnet ELPA**

Zoom: 1d 3d 1w 1m 1y

Thu May 22 10:04:36 2014 -- Fri May 22 10:04:36 2015

[Previous 1y](#)



# perfSONAR Results

Source	Destination
192.5.180.130 - 192.5.180.130 Capacity: Unknown MTU: Unknown	-- newy-pt1.es.net - 198.124.238.54 [traceroute] Capacity: 10G MTU: 1500
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	-- newy-pt1.es.net - 198.124.238.54 Capacity: 10G MTU: 1500

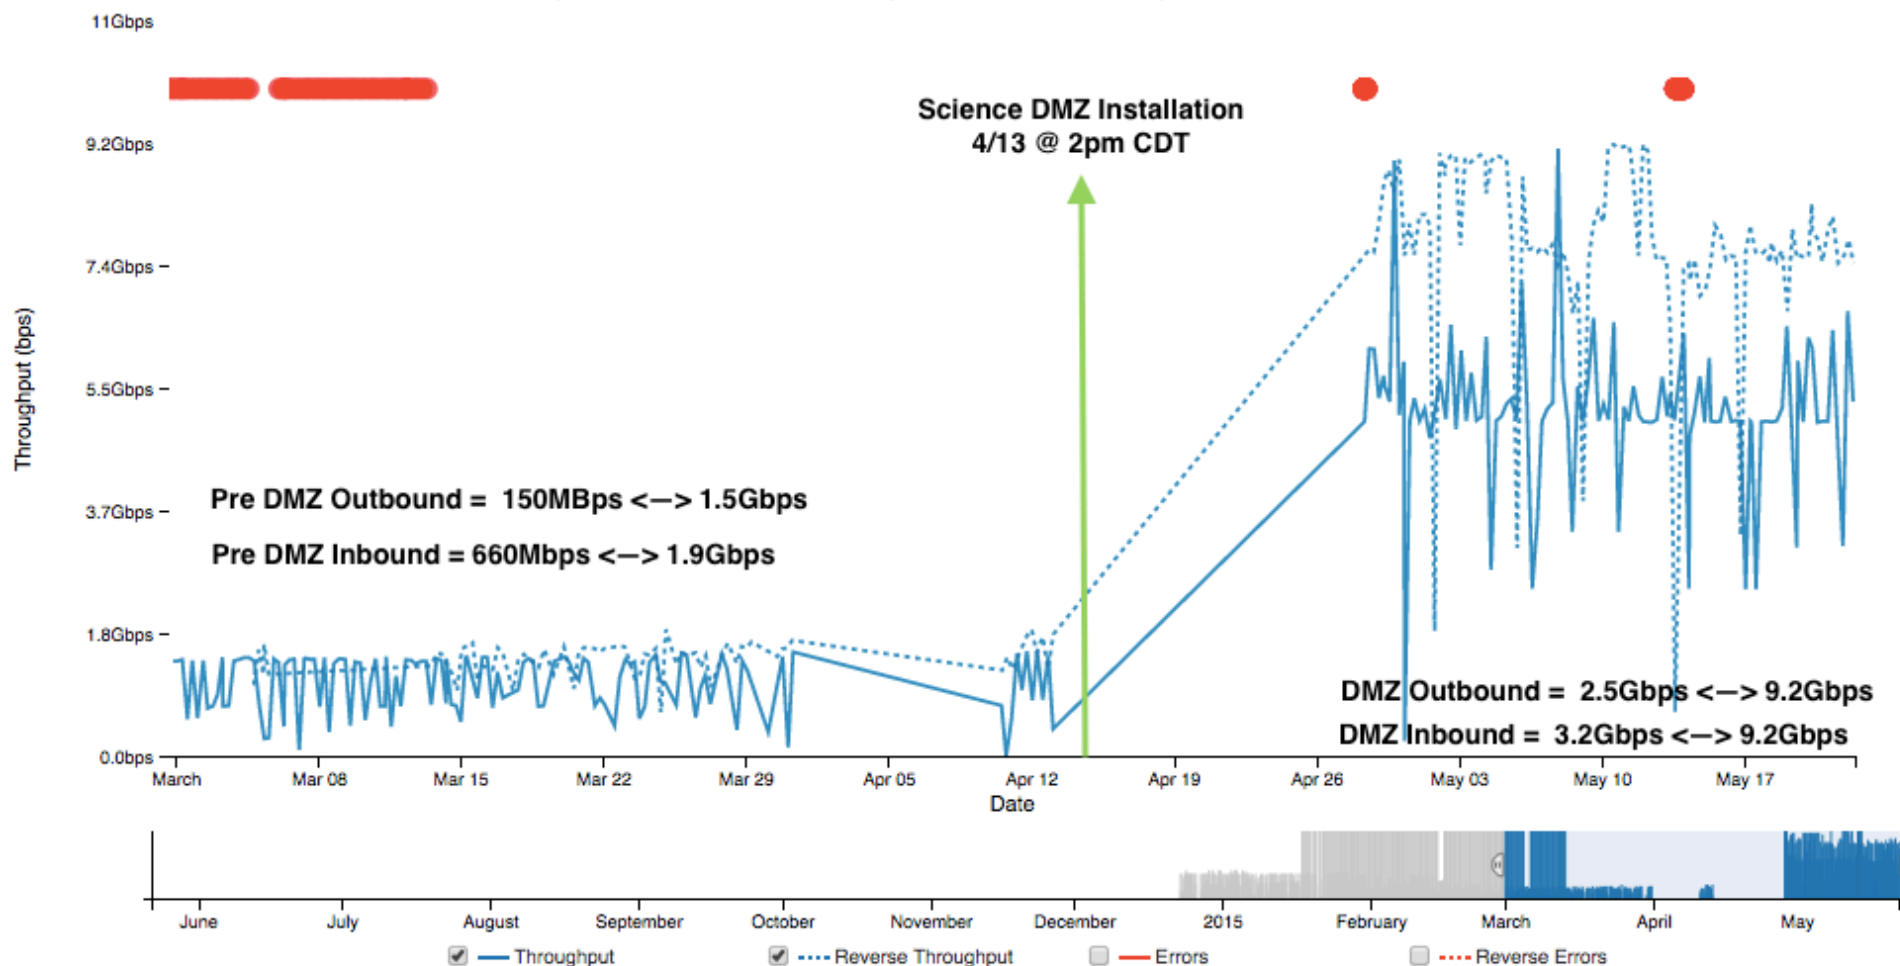
[Link to this chart](#)

**~27ms RTT from ANL to ESnet NEWY**

Zoom: 1d 3d 1w 1m 1y

Thu May 22 10:04:46 2014 -- Fri May 22 10:04:46 2015

Previous 1y



# perfSONAR Results

Source	Destination
192.5.180.130 - 192.5.180.130 Capacity: Unknown MTU: Unknown	-- sacramento1.es.net - 198.129.254.38 [traceroute] Capacity: 10G MTU: 9000
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	-- sacramento1.es.net - 198.129.254.38 Capacity: 10G MTU: 9000

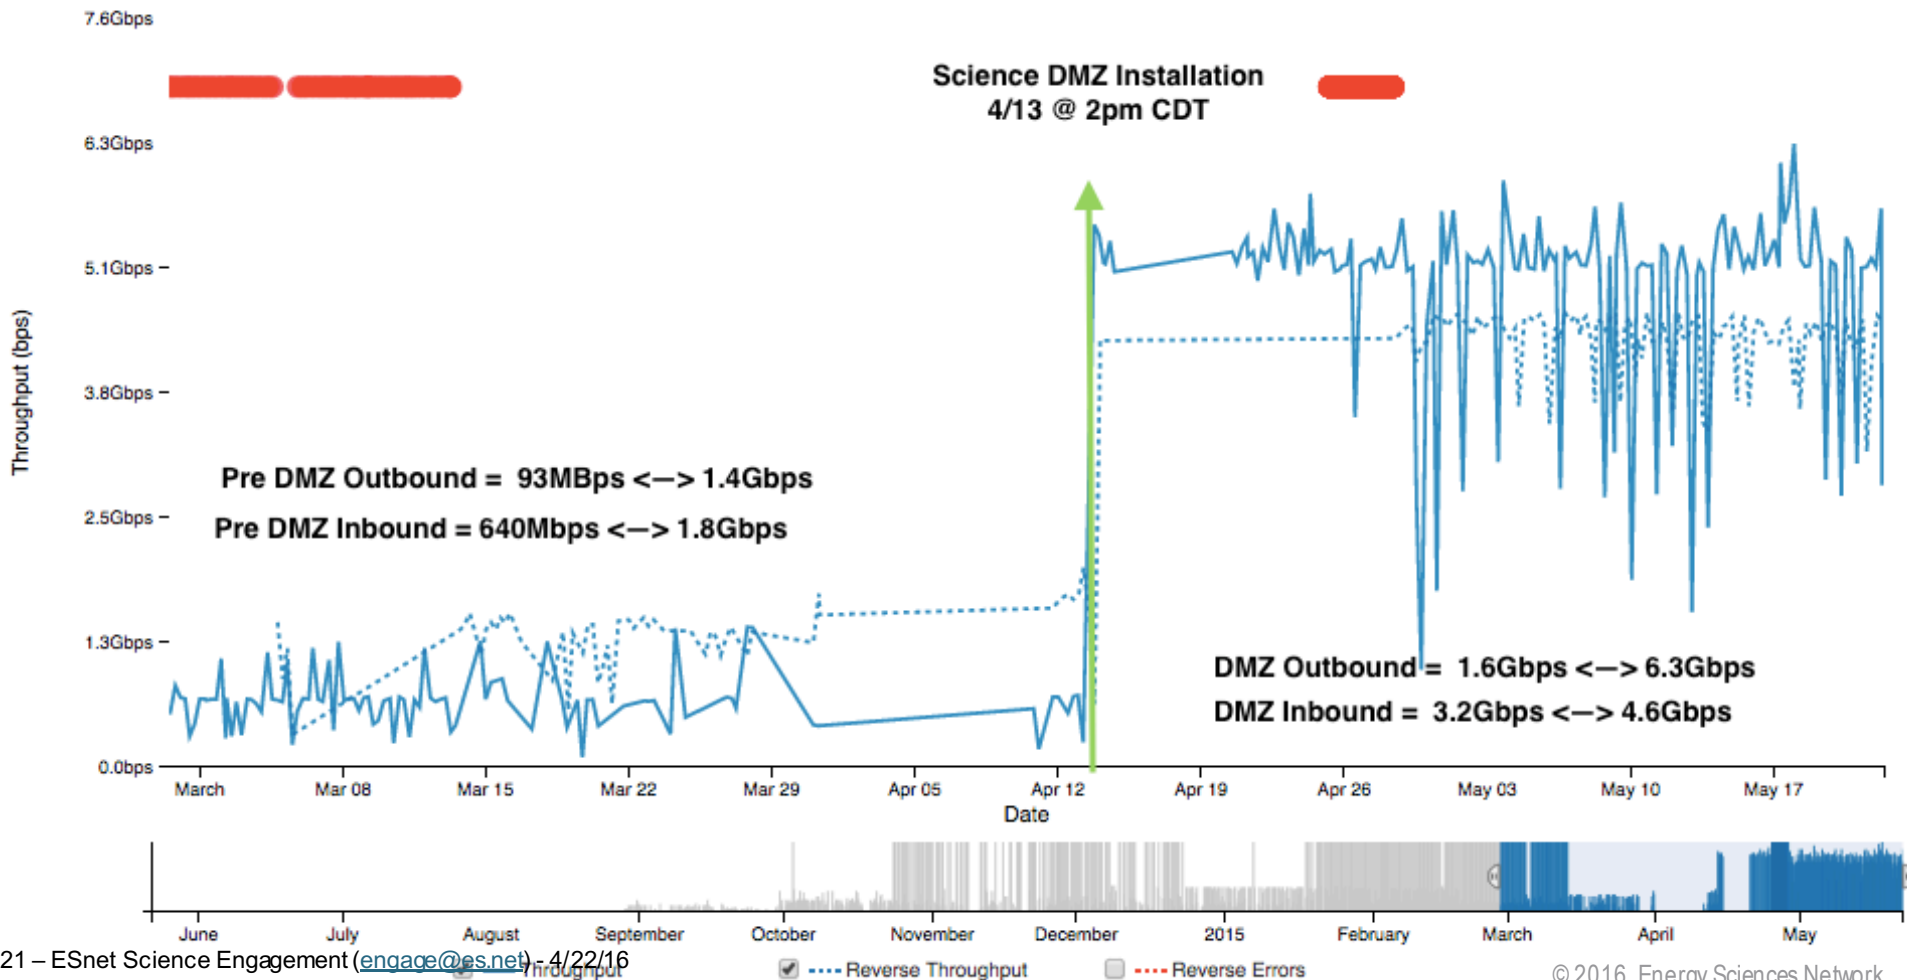
[Link to this chart](#)

**~43ms RTT from ANL to ESnet SACR**

Zoom: 1d 3d 1w 1m 1y

Thu May 22 10:04:15 2014 -- Fri May 22 10:04:15 2015

[Previous 1y](#)



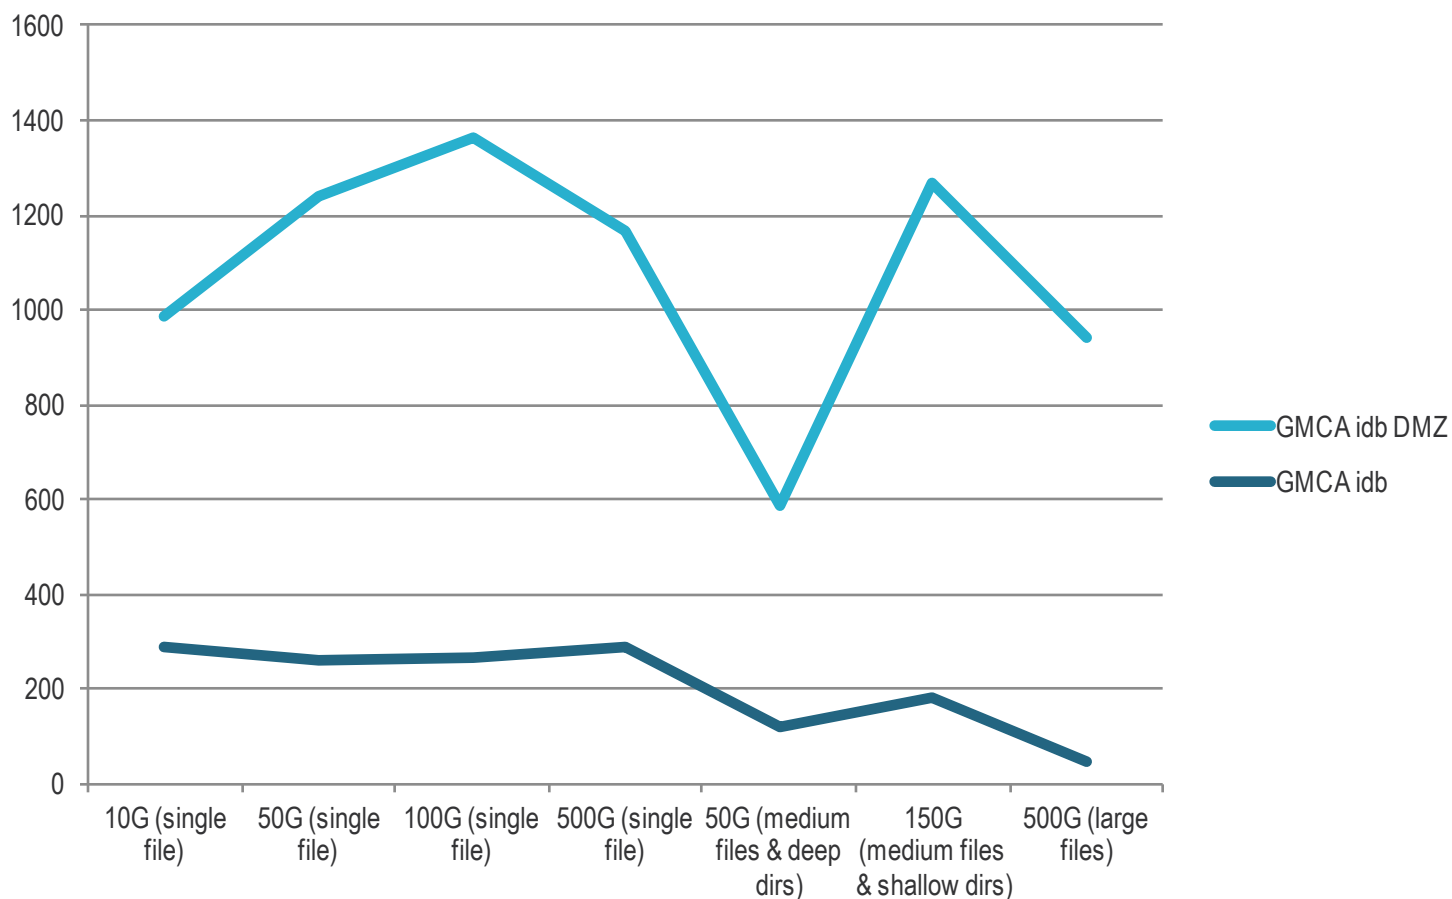
# perfSONAR is One Metric – what about GridFTP?

- Setup of the DTN Architecture:
  - 2 Servers
  - Shared filesystem
  - Each server has 2 interfaces/routes:
    - Science DMZ routes (plumbed to specific external locations via ANL, ESnet/MREN) and avoiding slower path
    - Traditional routes over slower path
  - **Goal**: show that with the correct network the **\*SAME\*** underlying infrastructure can perform as well as the perfSONAR testing



# Testing ESnet CERN (long RTT) to GMCA


- Scale is Mbps



- Across the board improvement of around 5x-20x
- Exercise helped to tune filesystem performance.



# Testing GMCA to ESnet CERN (long RTT)

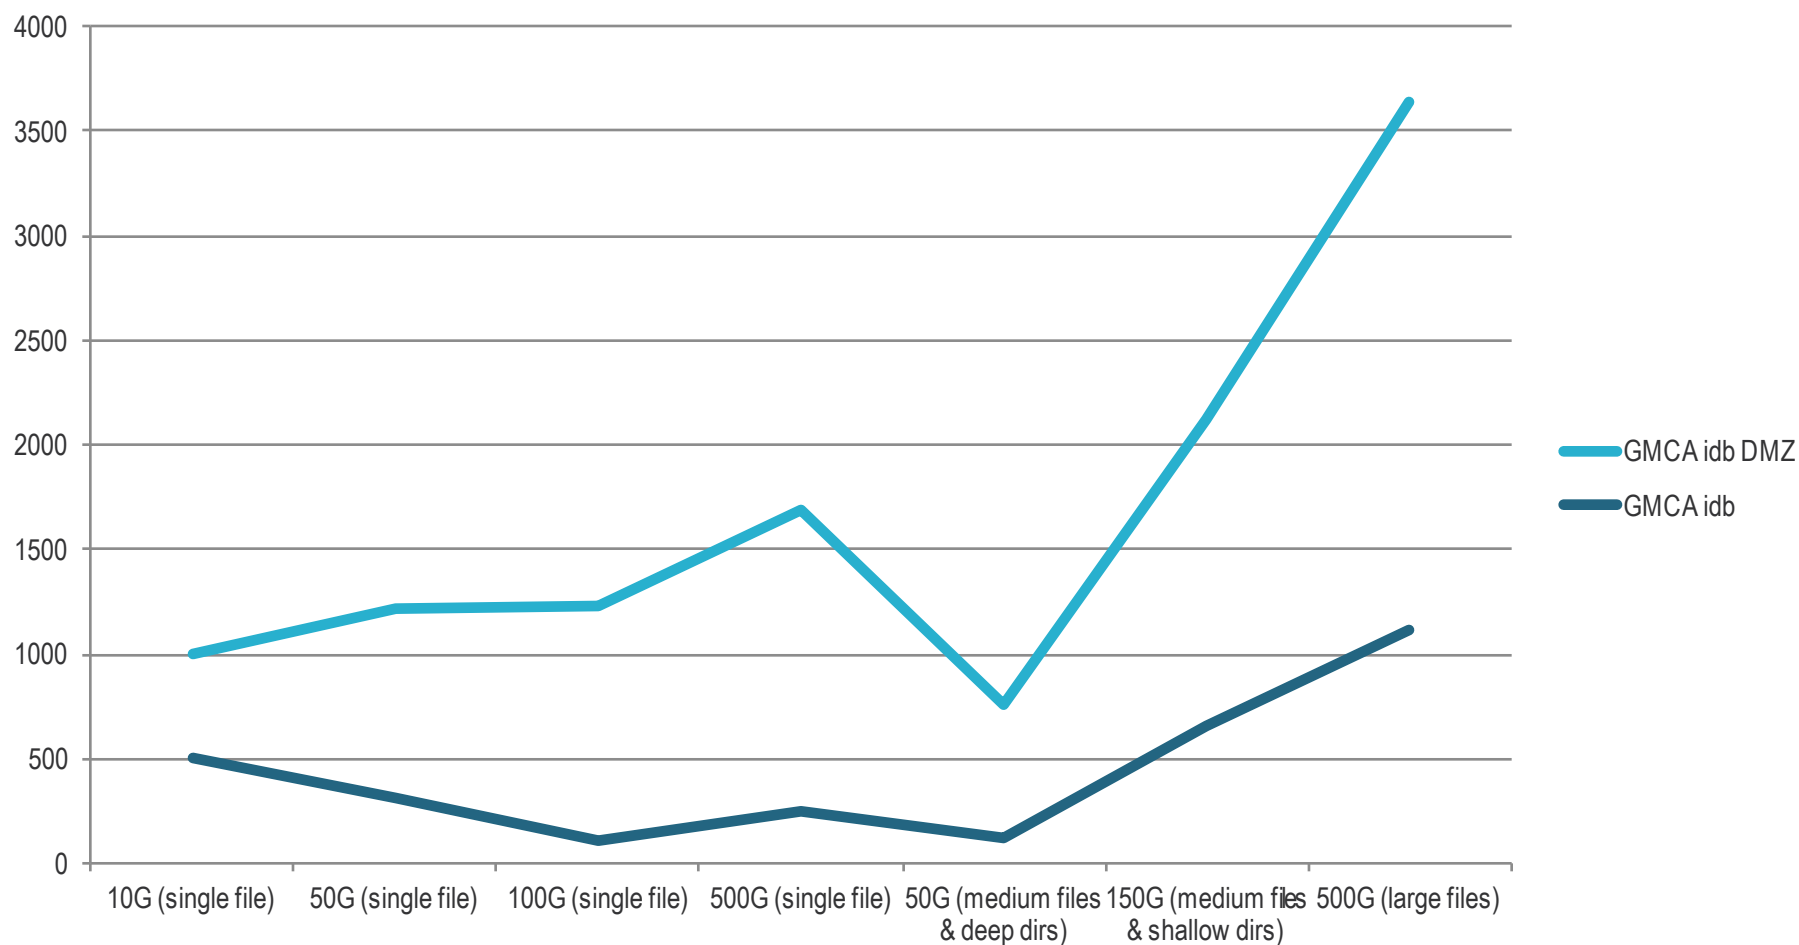


- Metadata operations (e.g. deeply nested directories) coupled with smaller files remains hard to crank up





# Testing GMCA to ESnet LBL (medium RTT)

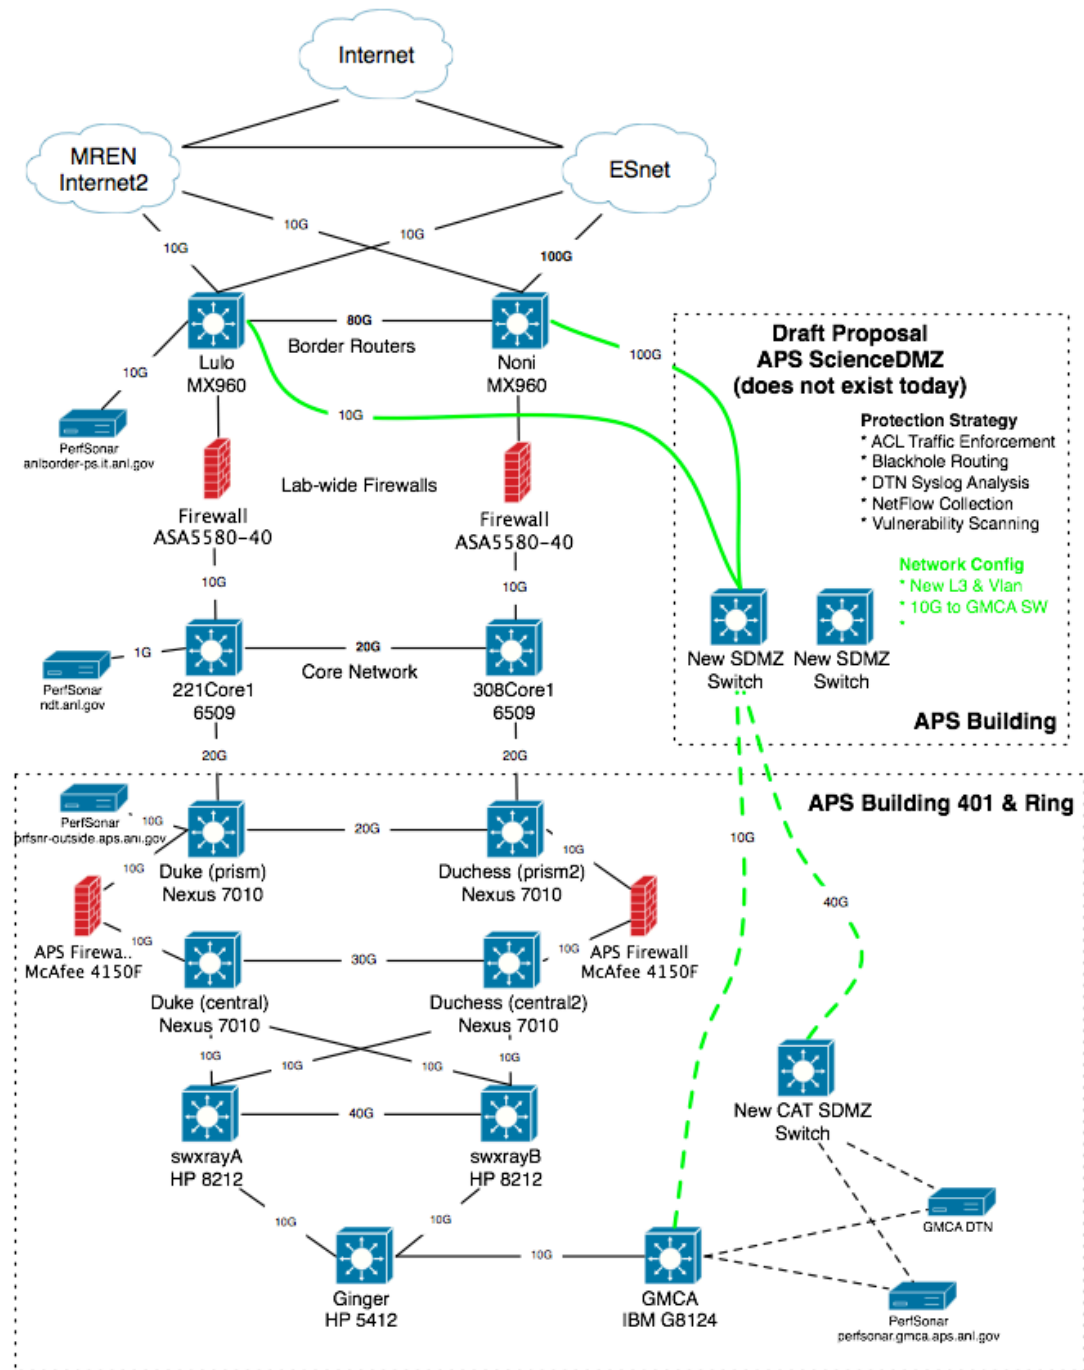


- Metadata operations (e.g. deeply nested directories) coupled with smaller files remains hard to crank up



# Future

- Wider use @ APS
  - Solution now goes directly to a specific beamline – how to support the entire facility?
- Wider use @ ANL
  - Service for other research groups (e.g. ALCF, ARM, etc.)
  - Pool of DTN resources w/ Globus, instead of each group manning their own (allows to back up to communal storage)
- Defining Policy
  - ACLs – ports exposed for things like perfSONAR, Globus. Shut off for things that don't need it. Gray area is defining what is and is not science.
  - Who gets on, who doesn't
  - Managing routing @ the border to best utilize the available WAN links





# ESnet

ENERGY SCIENCES NETWORK

# Science DMZ Implementation at the Advanced Photon Source

Jason Zurawski – [zurawski@es.net](mailto:zurawski@es.net)

Science Engagement Engineer, ESnet

Lawrence Berkeley National Laboratory

GlobusWorld 2016

April 22<sup>nd</sup> 2016



U.S. DEPARTMENT OF  
**ENERGY**

Office of Science

