



ESnet

ENERGY SCIENCES NETWORK

Improving Scientific Outcomes at the APS with a Science DMZ

Jason Zurawski – zurawski@es.net

Science Engagement Engineer, ESnet

Lawrence Berkeley National Laboratory

GlobusWorld 2015

April 15th, 2015



U.S. DEPARTMENT OF
ENERGY
Office of Science



Outline

- Background
- Current Network
- Pilot Network
- Future Work



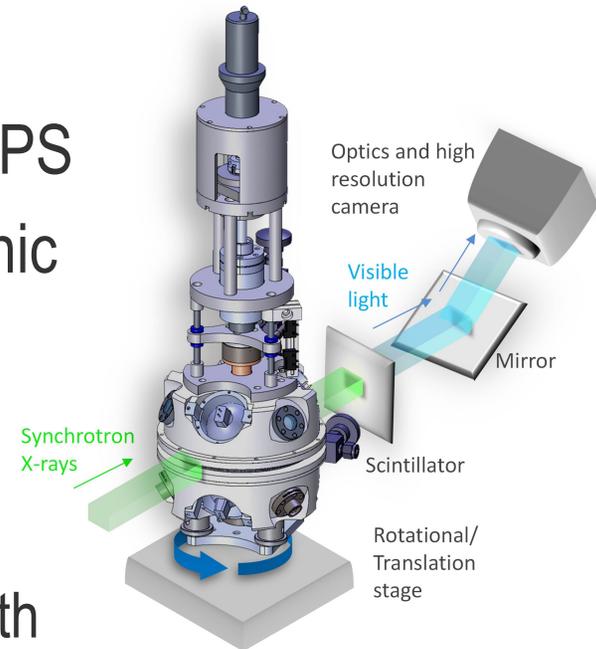
Why Build A Science DMZ Though?

- What we know about scientific network use:
 - Machine size decreasing, accuracy increasing
 - HPC resources more widely available – and potentially distributed from where the scientists are
 - WAN networking speeds now at 100G, MAN approaching, LAN as well
- Value Proposition:
 - If scientists can't use the network to the fullest potential due to local policy constraints or bottlenecks – they will find a way to get their done outside of what is available.
- Without a Science DMZ, this stuff is all hard



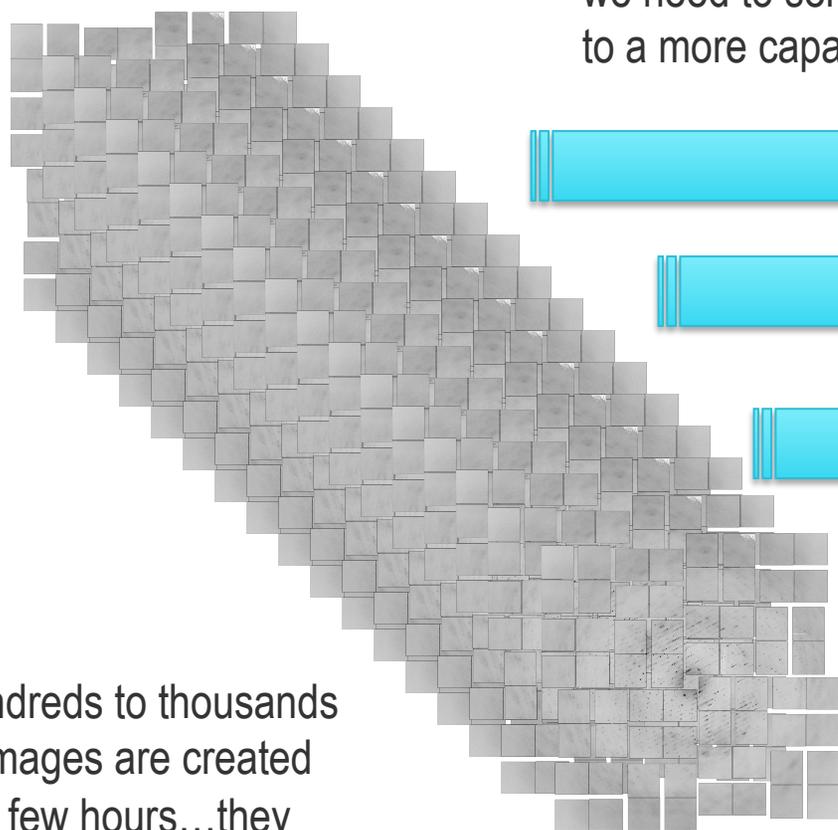
GM/CA

- 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’

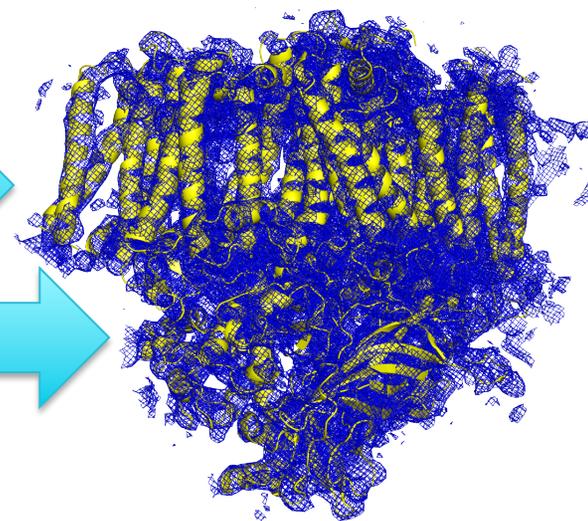
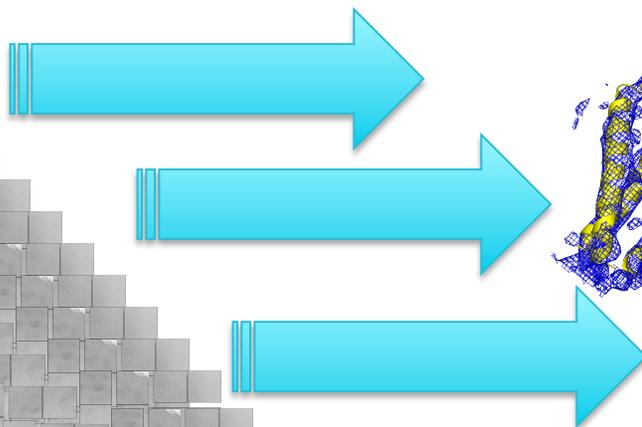


E Pluribus Unum

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

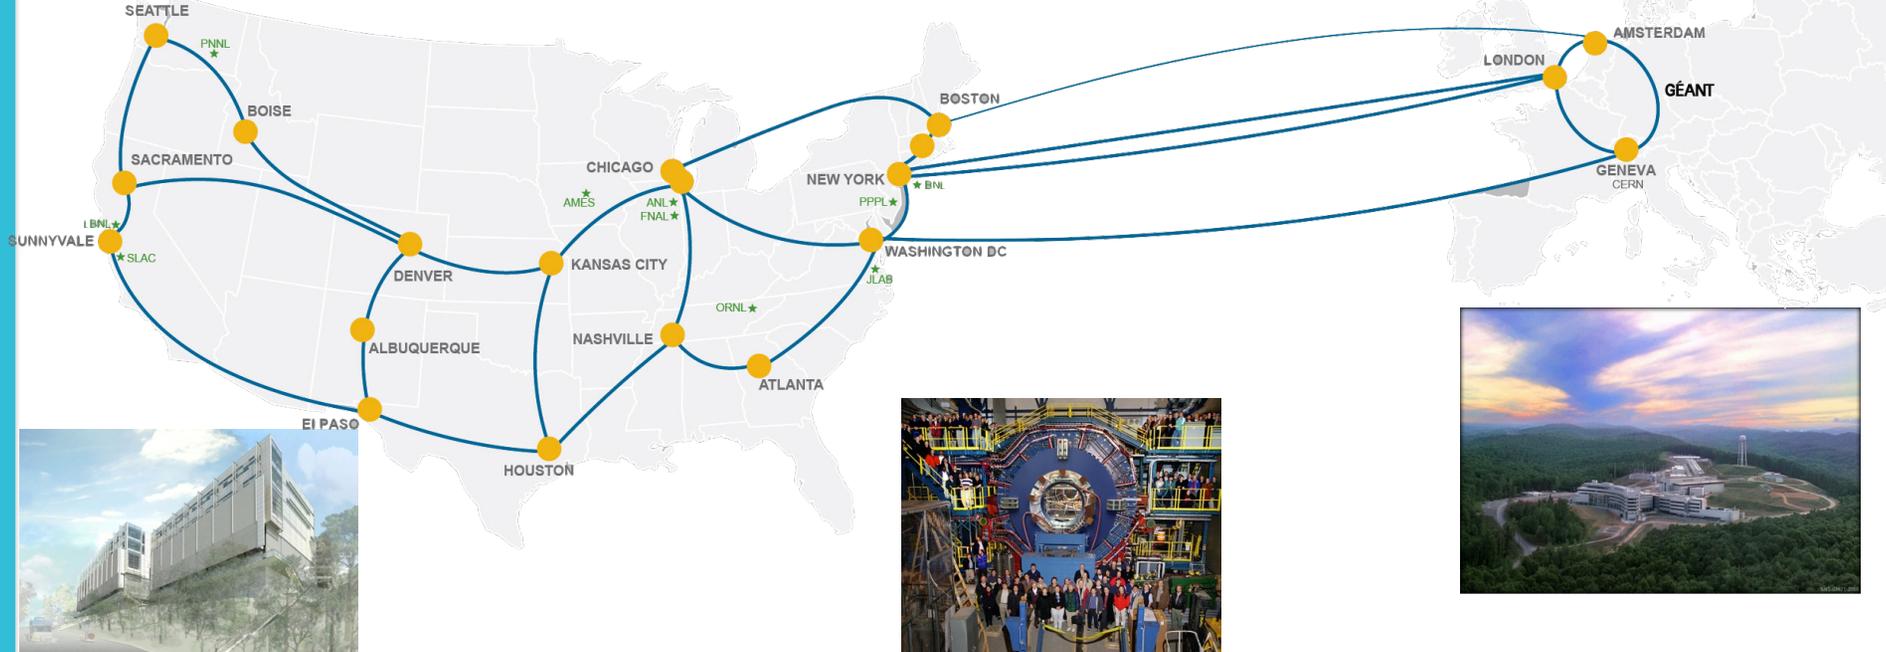
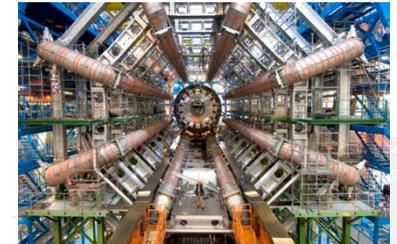


Hundreds to thousands of images are created in a few hours...they can range in size from MB to TB



After processing on a supercomputer, models are created.

Network as Infrastructure *Instrument*



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

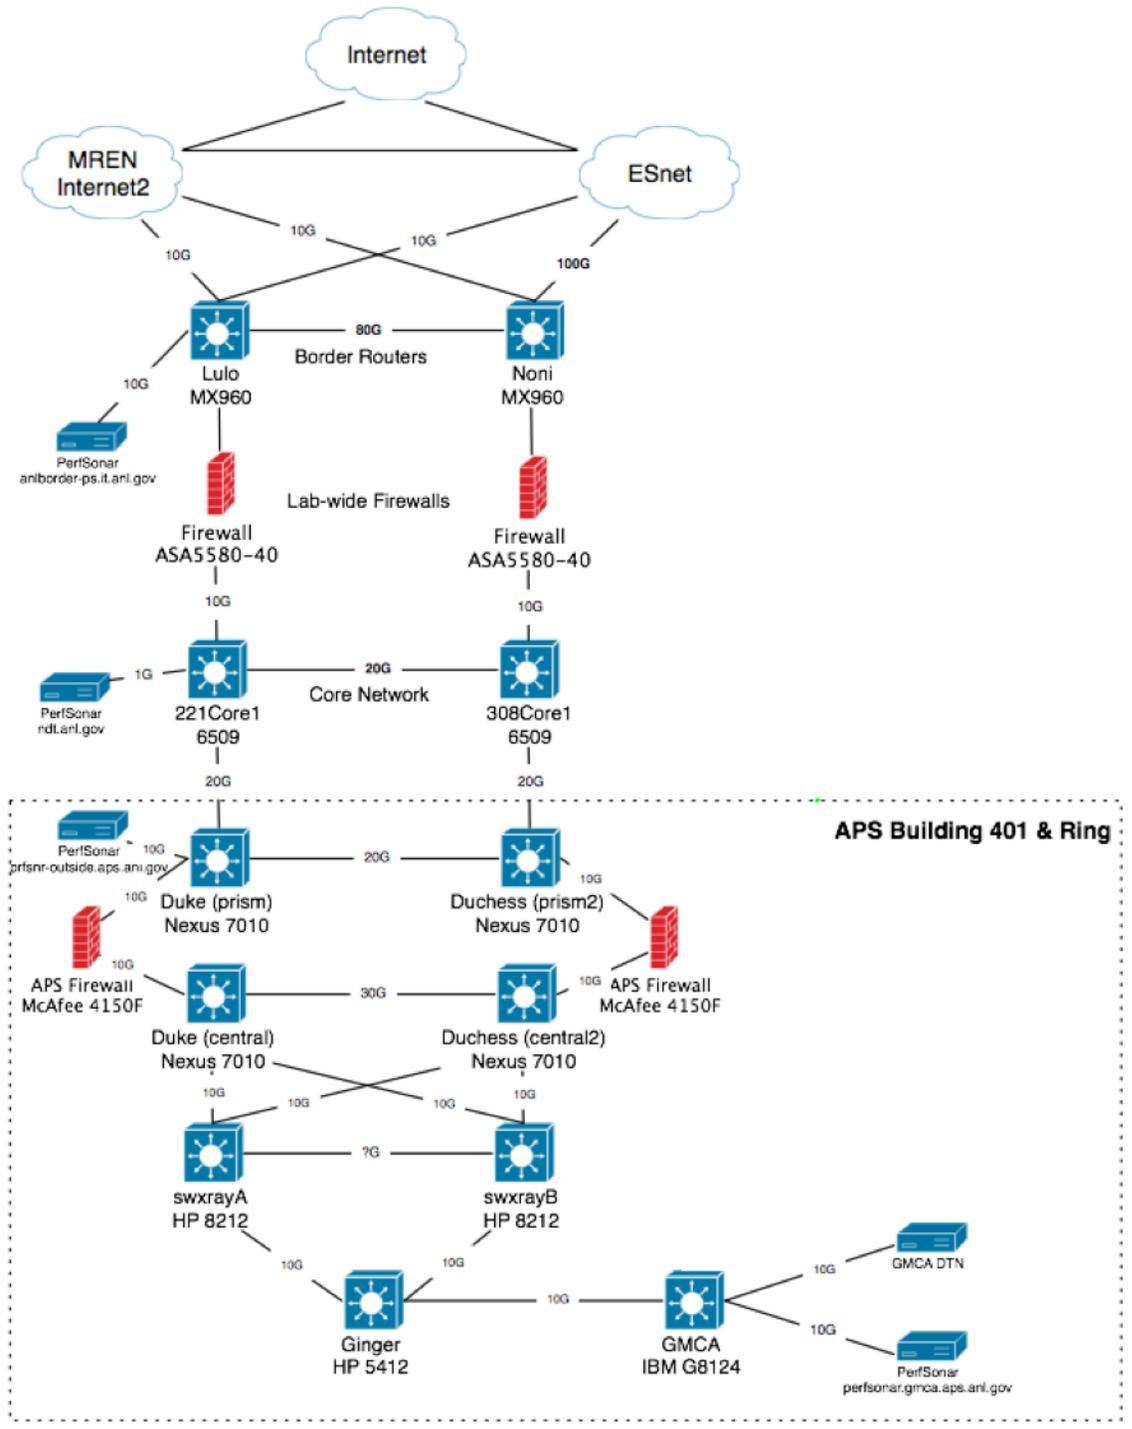


Outline

- Background
- **Current Network**
- Pilot Network
- Future Work



Current Design



ANL Border to ESnet New York (perfSONAR Test)

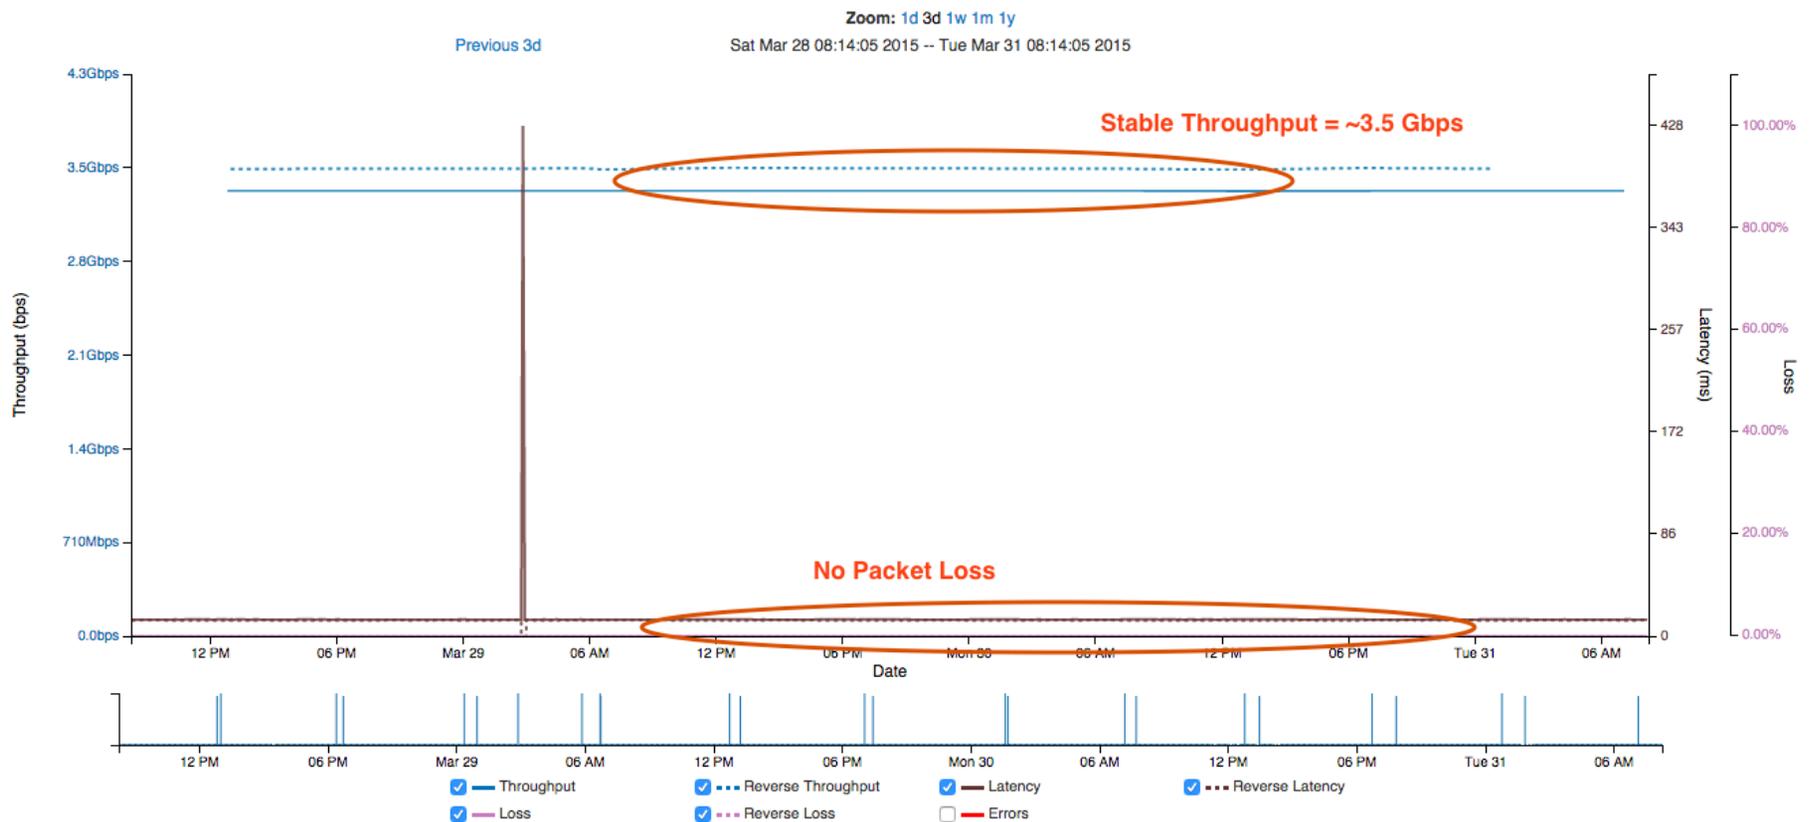
Source
 anlborder-ps.it.anl.gov - 130.202.222.58
 Capacity: 10G MTU: 9000

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

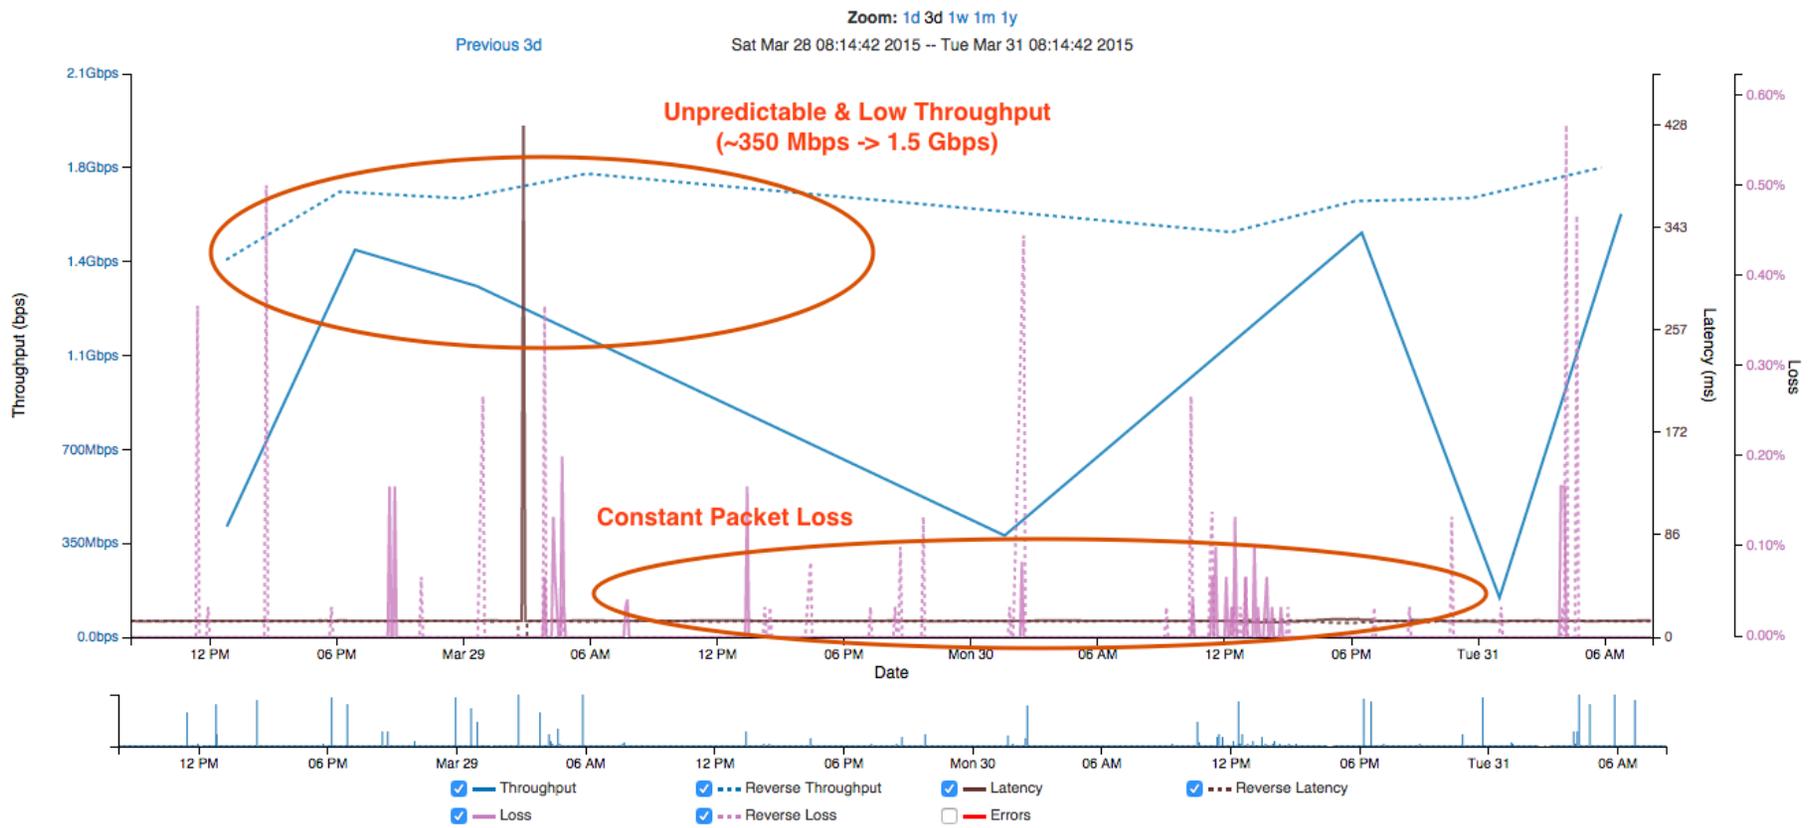
Destination
 newy-pt1.es.net - 198.124.238.54
 Capacity: 10G MTU: 9000

Source
 perfsonar.gmca.aps.anl.gov - 164.54.103.141
 Capacity: 10G MTU: 1500

Destination
 newy-owamp.es.net - 198.124.252.158
 Capacity: 1.0G MTU: 1500

[Link to this chart](#)

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 ANL Border (perfSONAR Test)

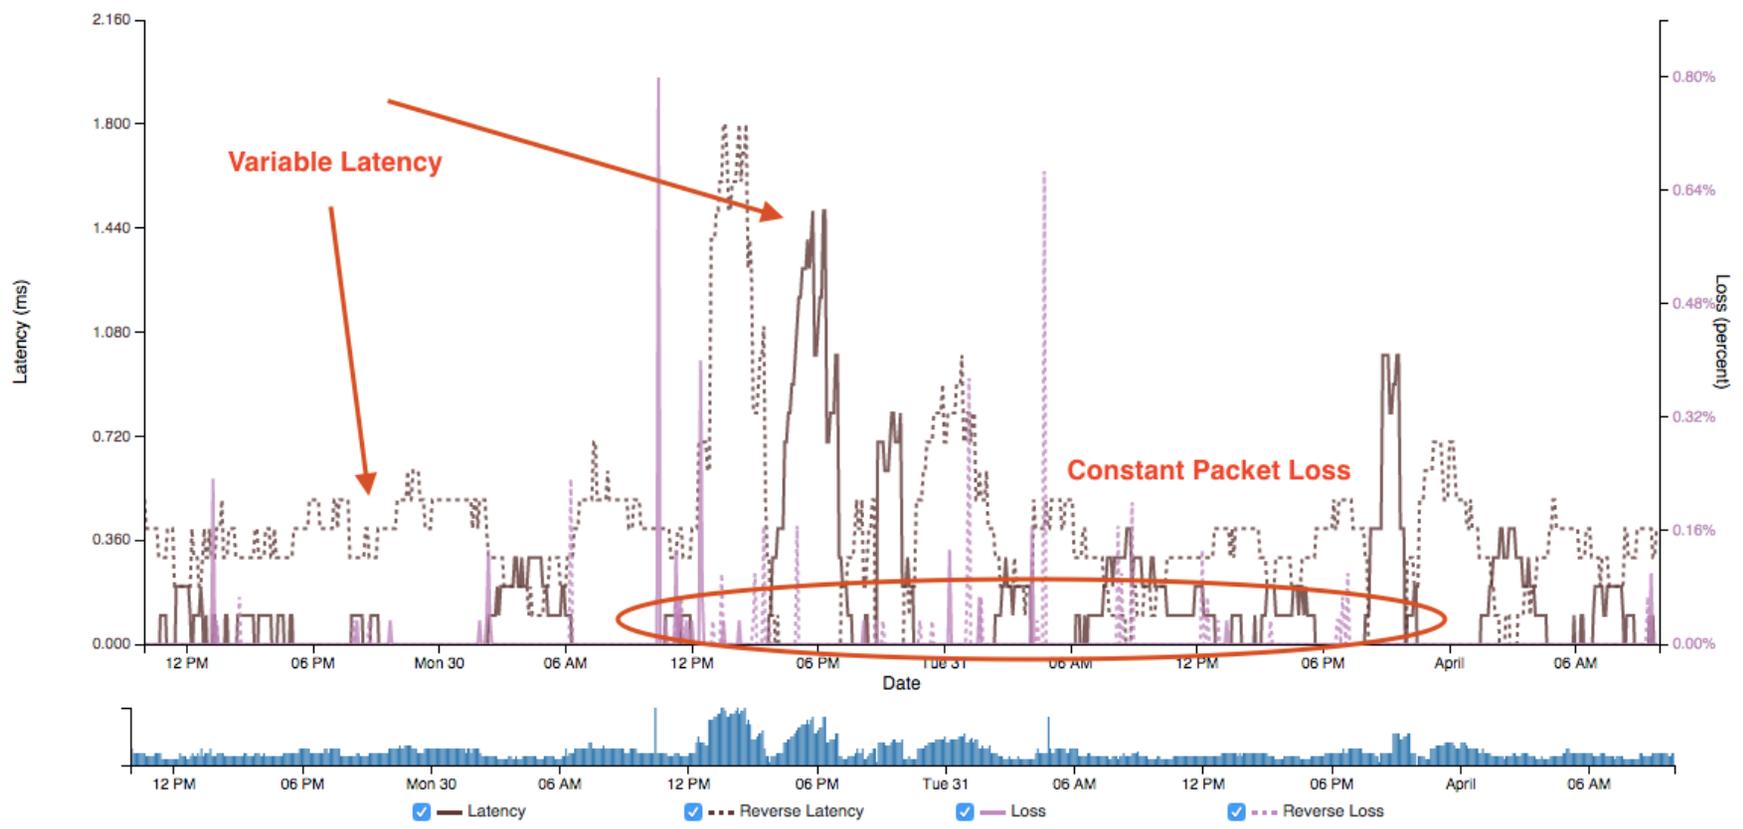
Source: **perfsonar.gmca.aps.anl.gov** - 164.54.103.141 -- Destination: **anlborder-ps.it.anl.gov** - 130.202.222.58
Capacity: Unknown MTU: Unknown

[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.

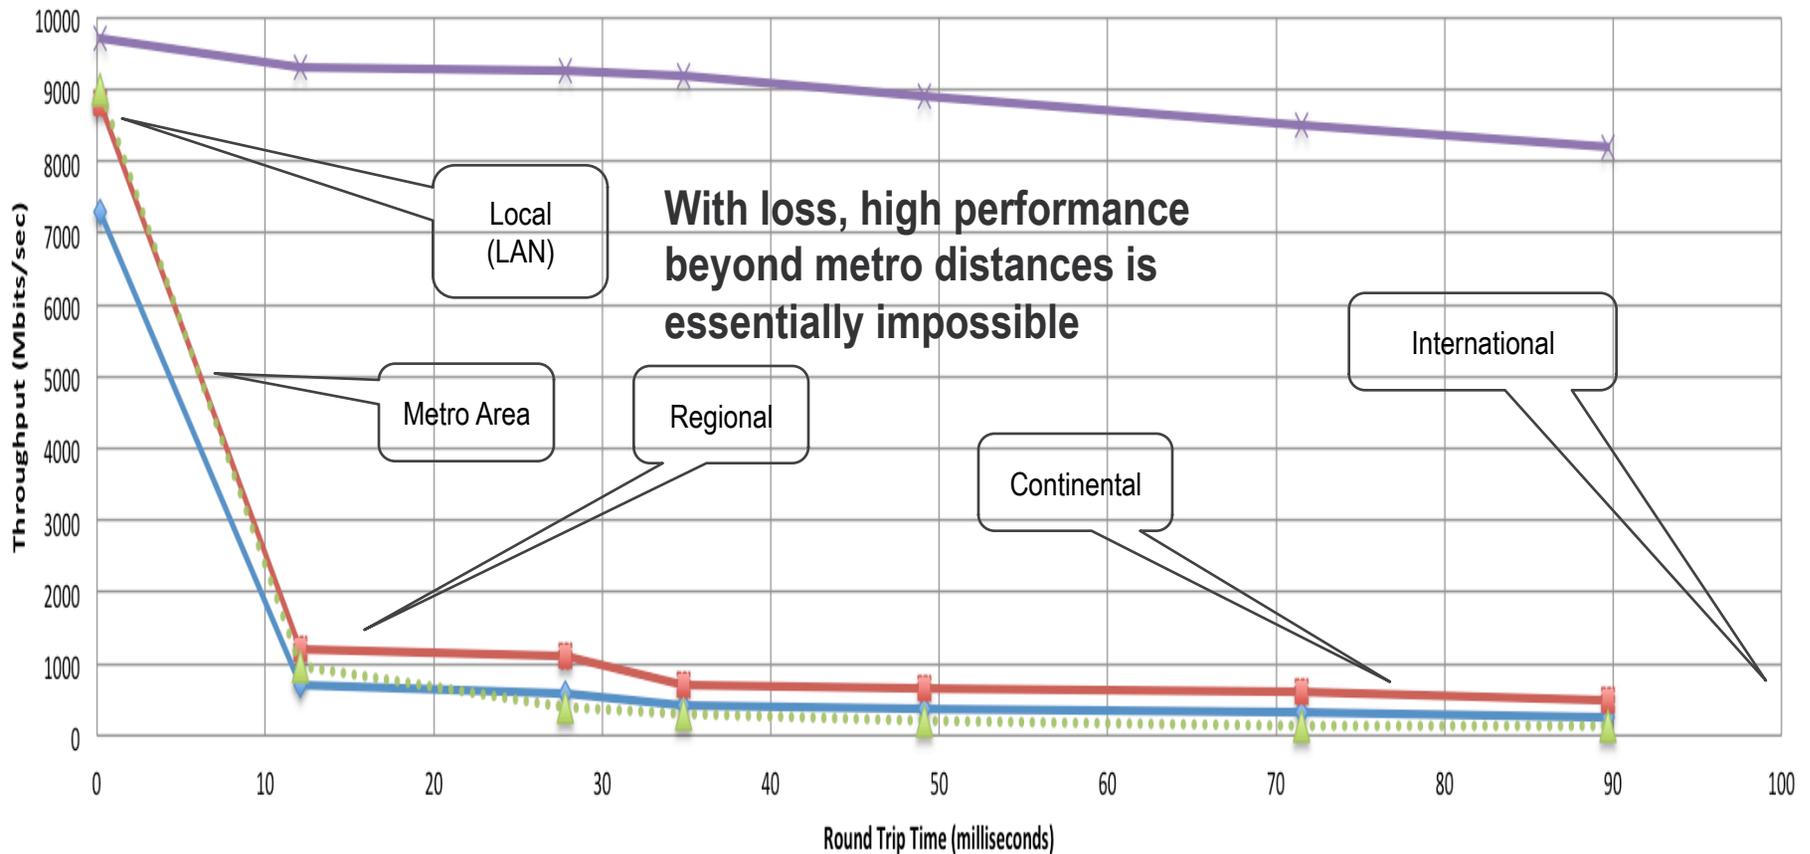
Zoom: 1d 3d 1w 1m 1y

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)

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

Globus Results

- GM/CA Downloading from:
 - ESnet ANL (~2ms)
 - Average: 285 Mbps
 - Max: 360 Mbps
 - ESnet BNL (~30ms)
 - Average: 26 Mbps
 - Max: 28 Mbps
 - ESnet LBL (~50ms)
 - Average: 16 Mbps
 - Max: 17 Mbps
- GM/CA Downloading from:
 - ESnet CERN (~110ms)
 - Average: 7 Mbps
 - Max: 8 Mbps
 - NERSC HPSS (~50ms)
 - Average: 127 Mbps
 - Max: 134 Mbps
- GM/CA Uploading to:
 - NERSC HPSS (~50ms)
 - Average: 112 Mbps
 - Max: 142 Mbps



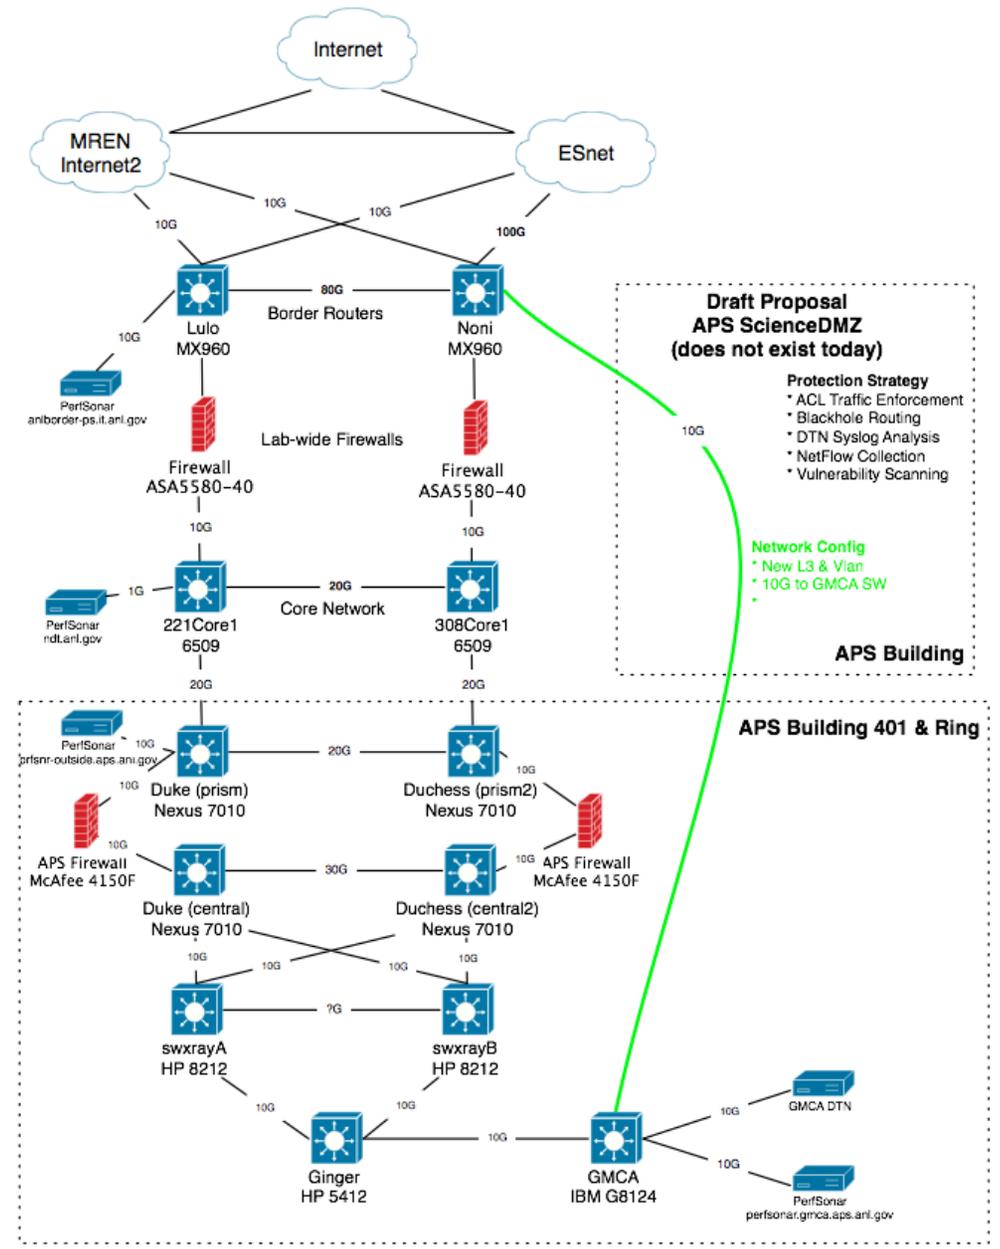
Outline

- Background
- Current Network
- Pilot Network
- Future Work



Pilot Network

- Plan:
 - Use campus and building fiber resources for 2nd 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)

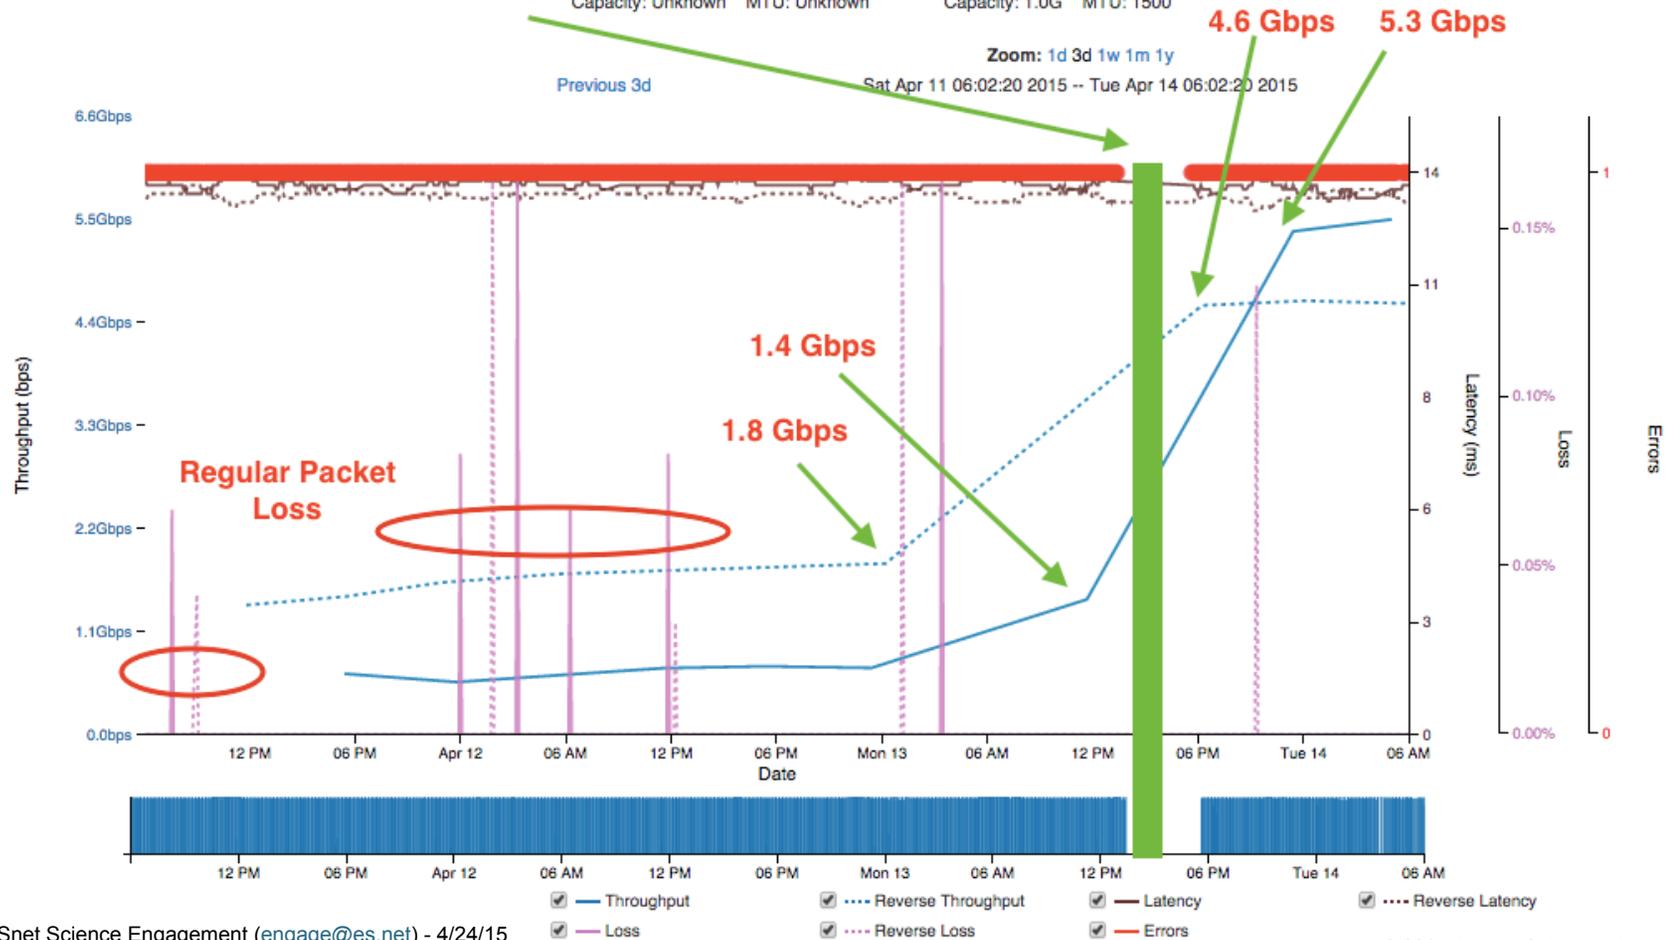


Results (as of 2pm CDT on April 13th)

Source	Destination
192.5.180.130 - 192.5.180.130 Capacity: 10G MTU: 1500	elpa-pt1.es.net - 198.129.254.82 Capacity: 10G MTU: 9000
192.5.180.130 - 192.5.180.130 Capacity: 10G MTU: 1500	hous-owamp.es.net - 198.129.254.66 [traceroute] Capacity: 1.0G MTU: 1500
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	elpa-pt1.es.net - 198.129.254.82 Capacity: 10G MTU: 9000
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	hous-owamp.es.net - 198.129.254.66 Capacity: 1.0G MTU: 1500

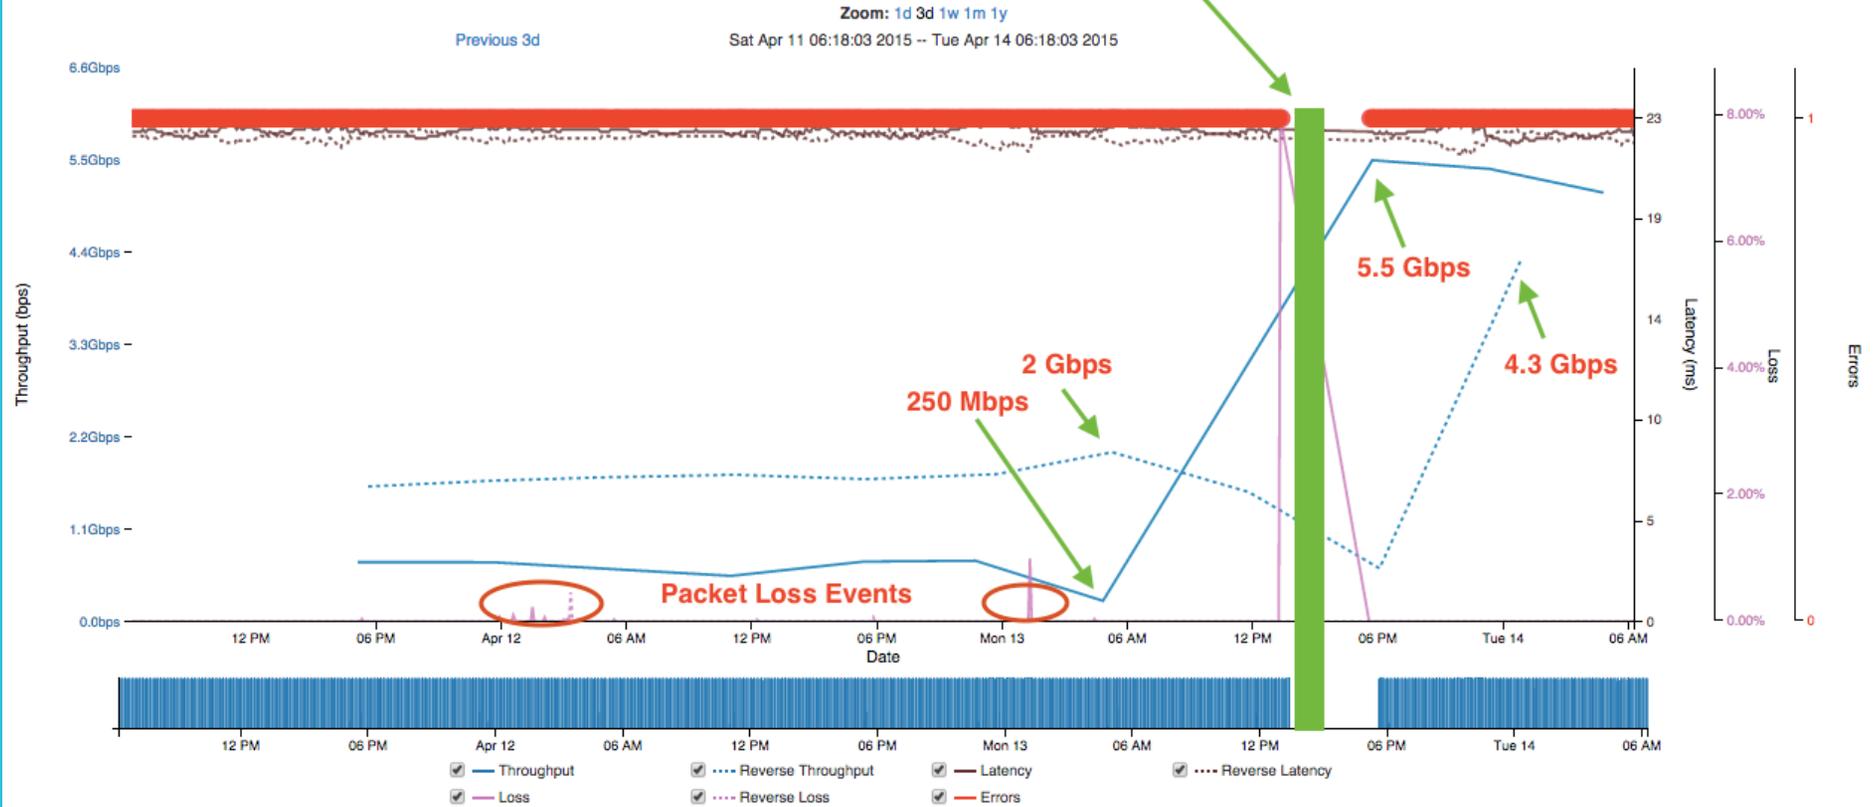
[Link to this chart](#)

**DMZ Upgrade
(~2pm CDT 4/13)**



Results (as of 2pm CDT on April 13th)

Source	Destination	Link to this chart
192.5.180.130 - 192.5.180.130 Capacity: 10G MTU: 1500	-- sacr-pt1.es.net - 198.129.254.38 Capacity: 10G MTU: 9000	Link to this chart
192.5.180.130 - 192.5.180.130 Capacity: 10G MTU: 1500	-- sacr-owamp.es.net - 198.129.254.18 [traceroute] Capacity: 1.0G MTU: 1500	
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	-- sacr-pt1.es.net - 198.129.254.38 Capacity: 10G MTU: 9000	
164.54.103.141 - 164.54.103.141 Capacity: Unknown MTU: Unknown	-- sacr-owamp.es.net - 198.129.254.18 Capacity: 1.0G MTU: 1500	



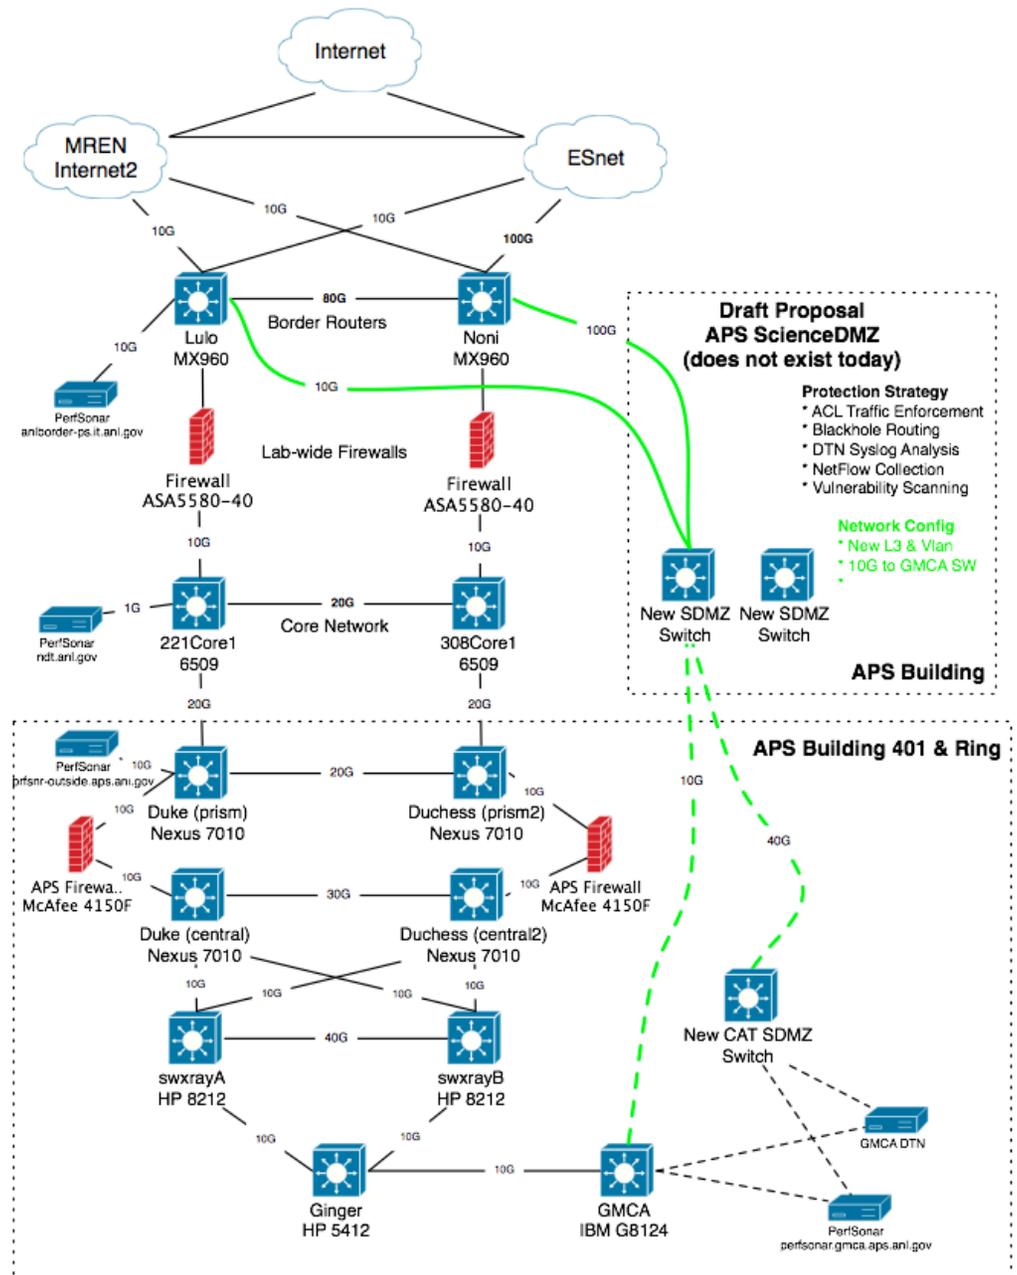
Outline

- Background
- Current Network
- Pilot Network
- **Future Work**



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



Participants & Thanks

- Argonne National Laboratory Networking
 - Corey Hall
 - Linda Winkler
 - Corby Schmitz
- Advanced Photon Source Networking
 - David Leibfritz
 - Mary Westbrook
- GM/CA Beamline
 - Sergey Stepanov





ESnet

ENERGY SCIENCES NETWORK

Improving Scientific Outcomes at the APS with a Science DMZ

Jason Zurawski – zurawski@es.net

Science Engagement Engineer, ESnet

Lawrence Berkeley National Laboratory

GlobusWorld 2015

April 15th, 2015



U.S. DEPARTMENT OF
ENERGY
Office of Science

