



Give me a ping, Vasili.

One ping only, please.

Random comments on USCMS perfSONAR status

Garhan Attebury  
University of Nebraska-Lincoln

# So we've deployed these “things”

- sites don't actively go hunting for problems...  
(or just don't have time to do anything about them)
- we can't just set it up and walk away?
- who's responsible for watching this stuff?
- networks are hard (and don't fix themselves)

# Changes in the past year

- More modern-ish (CentOS 6 x86\_64)
- Current stable version = 3.3.2 (3.2.x deprecated last week)
- 10GbE is the new 1GbE
- Mesh configurations (!)

`/opt/perfsonar_ps/mesh_config/etc/agent_configuration.conf`

<https://grid-deployment.web.cern.ch/grid-deployment/wlwg-ops/perfsonar/conf/central/testdefs/jsons/tests-us-cms.json> (owamp ~5hr, bwctl ~6hr or 12hr)

<https://grid-deployment.web.cern.ch/grid-deployment/wlwg-ops/perfsonar/conf/central/testdefs/jsons/tests-wlwg-all.json> (traceroutes, pinger, bwctl ~7d)

# Changes in the past year

- More modern-ish (CentOS 6 x86\_64)
- Current stable version = 3.3.2 (3.2.x deprecated last week)
- 10GbE is the new 1GbE
- **Mesh configurations (!)**

`/opt/perfsonar_ps/mesh_config/etc/agent_configuration.conf`

<https://grid-deployment.web.cern.ch/grid-deployment/wlwg-ops/perfsonar/conf/central/testdefs/jsons/tests-us-cms.json> (owamp ~5hr, bwctl ~6hr or 12hr)

<https://grid-deployment.web.cern.ch/grid-deployment/wlwg-ops/perfsonar/conf/central/testdefs/jsons/tests-wlwg-all.json> (traceroutes, pinger, bwctl ~7d)

# While we're at it ...

- Do the basic setup



pS-Performance Node For Purdue University In West Lafayette , IN , US

Host Information	
Organization Name	Purdue University
City, State, Country	West Lafayette, IN, US
Zip Code	47906
Latitude, Longitude	40.2525, 86.5516
Administrator Name	
Administrator Email	

- If you've upgraded from the start, tidy things up  
`ps faux | grep -i knoppix`

mannperf2.itns.purdue.edu (128.211.143.4)	psonar3.fnal.gov (131.225.205.23)	Yes	TCP	30	0.636	4.517	Select ▾
mannperf2.itns.purdue.edu (128.211.143.4)	perfonar02.hep.wisc.edu (144.92.180.76)	Yes	TCP	30	0.362	0.915	Select ▾
mannperf2.itns.purdue.edu (128.211.143.4)	psonar3.fnal.gov (131.225.205.23)	Yes	TCP	20	0.642	4.857	Select ▾

- Watch the disk usage / ensure cleanup scripts work
- Put instances near your data, and your admins!

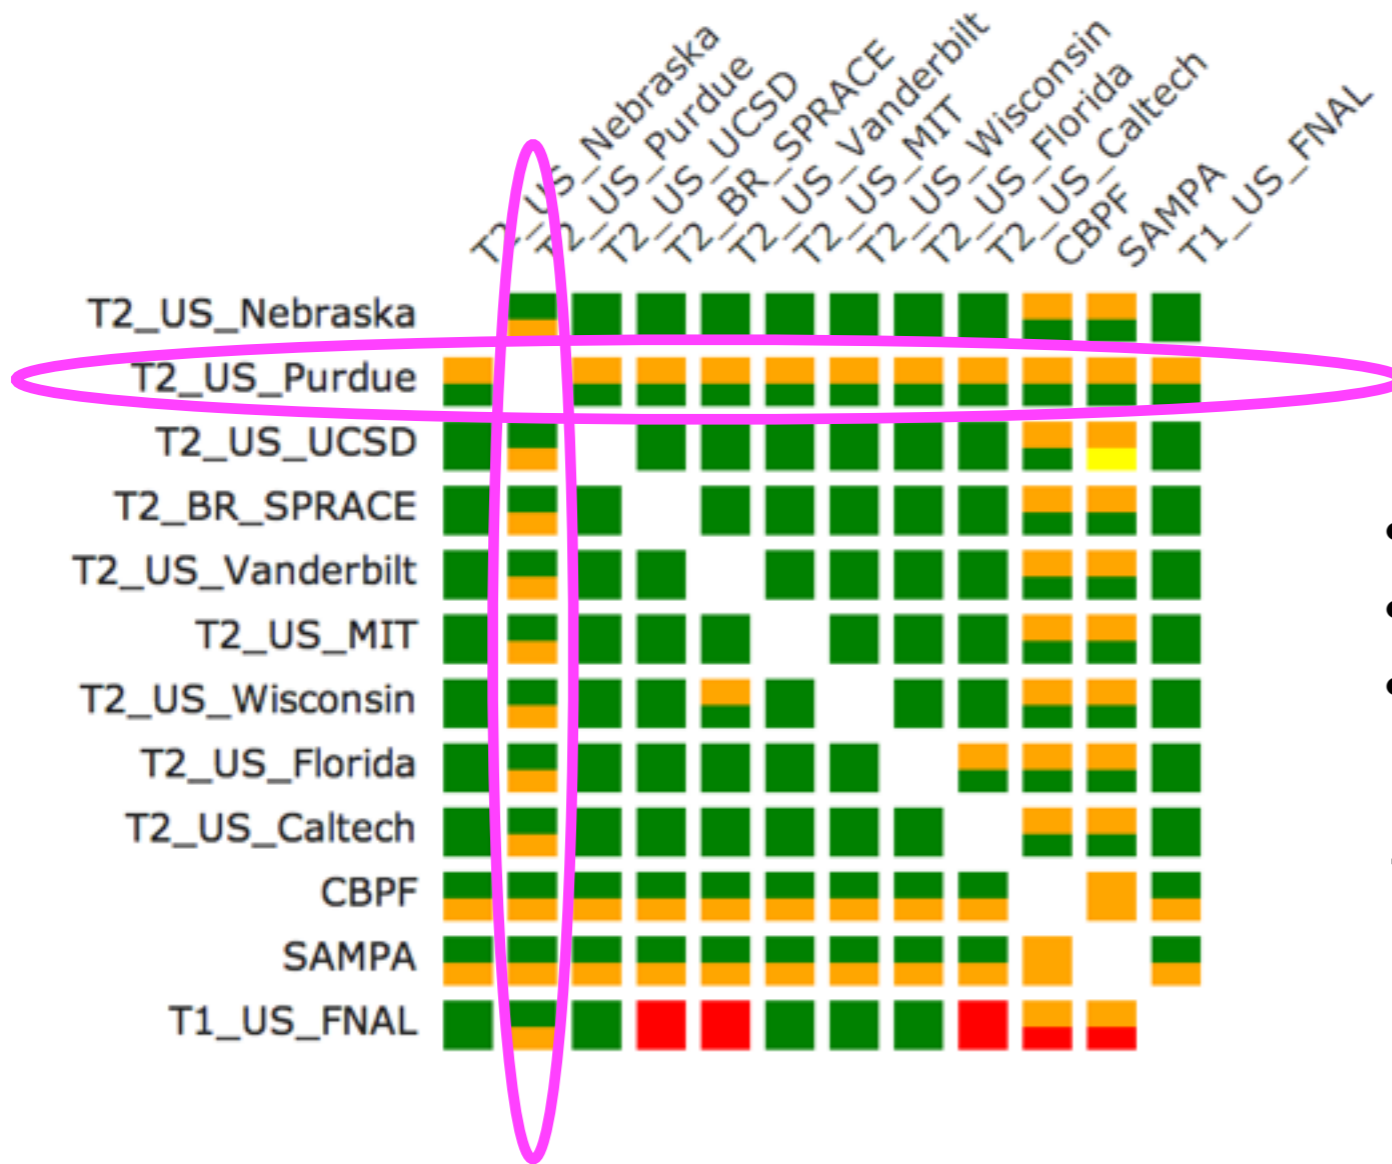
# What to watch...

- <http://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=US%20CMS%20T2%20Sites>
- The hosts!
  - Services get stuck / fail to start
  - Disks full, NTP failing, iptables ...



# US CMS T2 Sites - USCMS Cloud OWAMP Mesh Test

■ Loss rate is  $\leq 0$     
 ■ Loss rate is  $\geq 0$     
 ■ Loss rate is  $\geq 0.01$     
 ■ Unable to retrieve data    
 ■ Check has not yet run




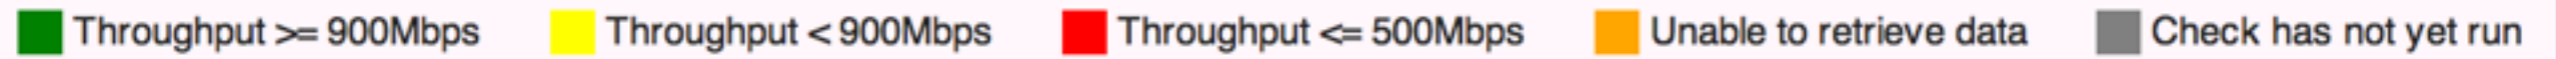
No data for Purdue across the board

- OWAMP service running
- Regular testing enabled
- OWAMP graph page won't load

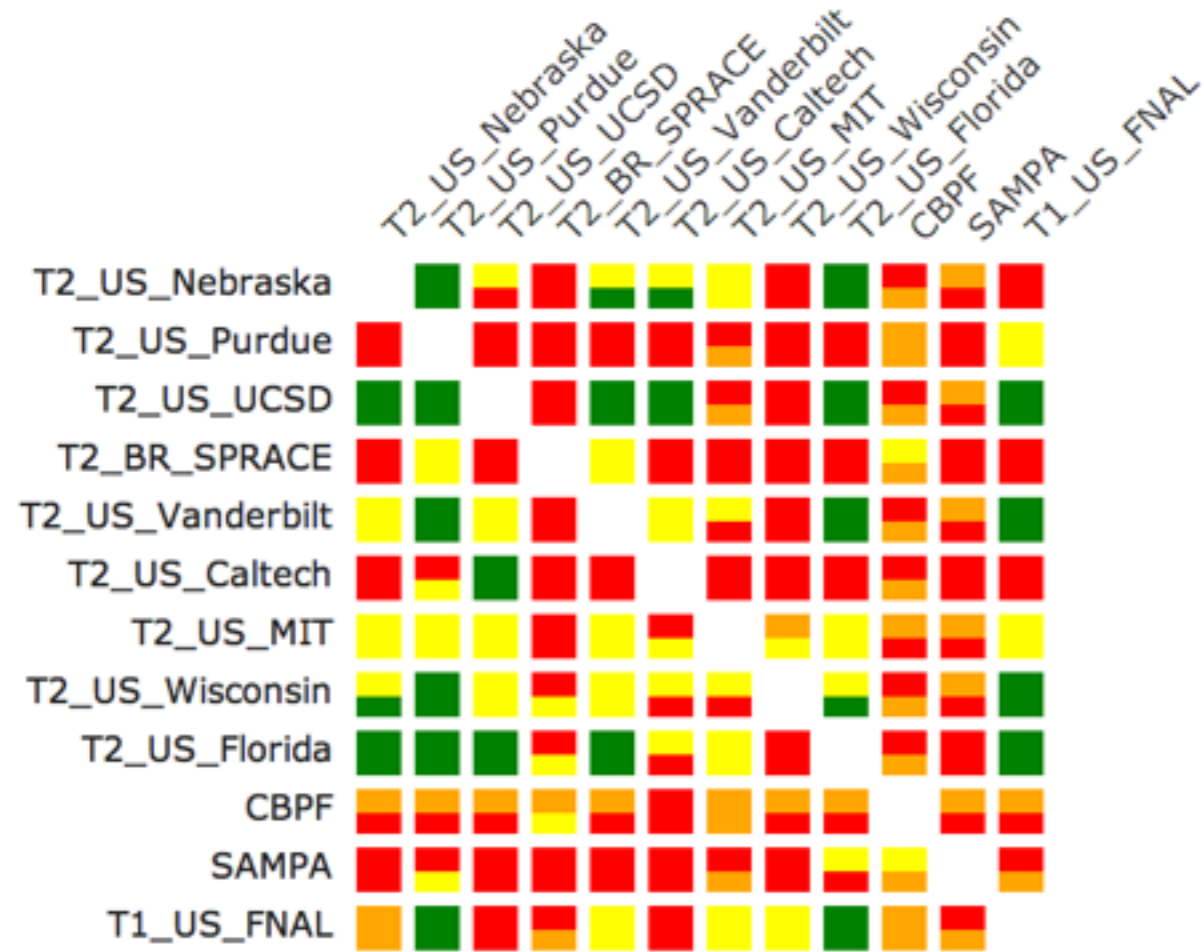

... probably just stuck / disk full



# US CMS T2 Sites - USCMS Cloud BWCTL Mesh Test



# US CMS T2 Sites - USCMS Cloud BWCTL Mesh Test




“That looks horrible”

“It’s so ... RED!”

“OH HH THE HUMANITY!”

# US CMS T2 Sites - USCMS Cloud BWCTL Mesh Test

■ Throughput  $\geq$  900Mbps   
 ■ Throughput  $<$  900Mbps   
 ■ Throughput  $\leq$  500Mbps   
 ■ Unable to retrieve data   
 ■ Check has not yet run

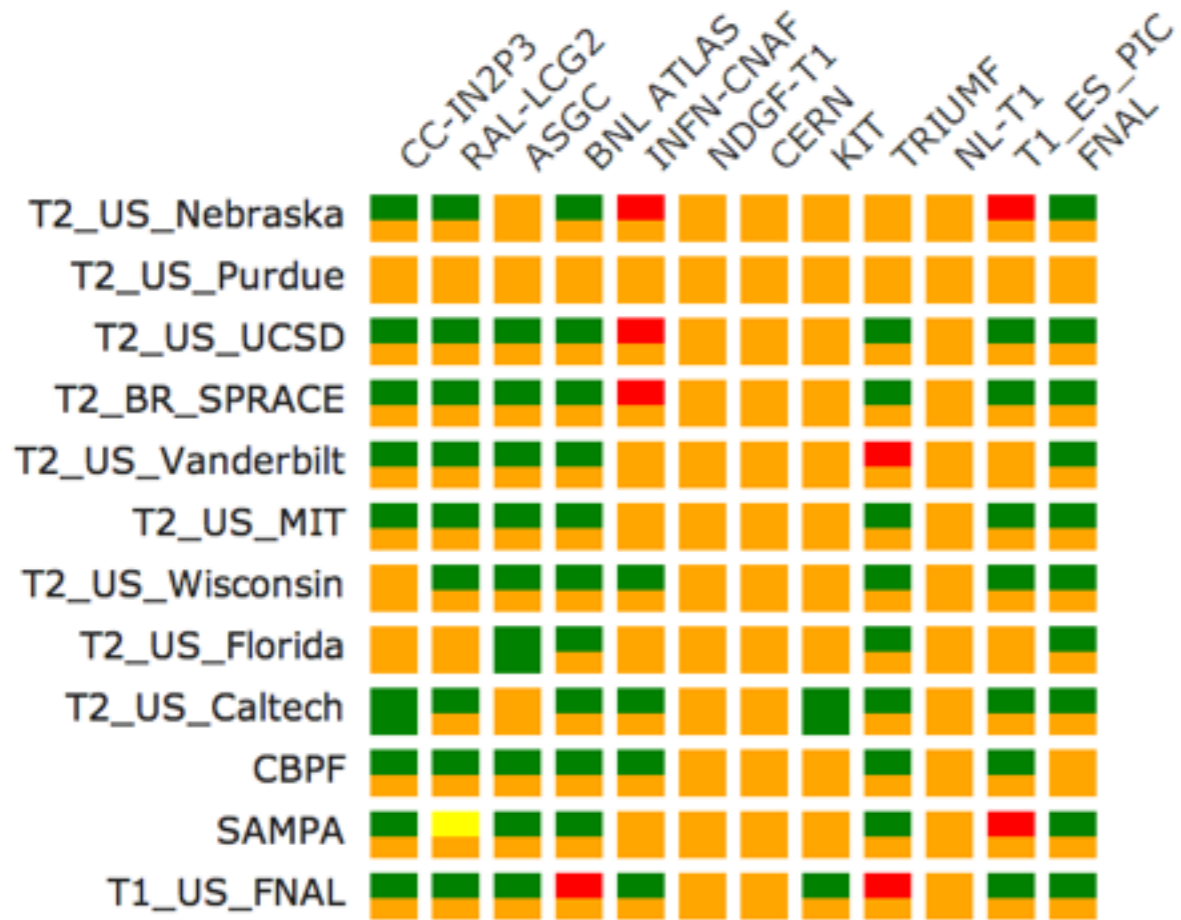


“We pay a lot of \$\$\$ for these 10Gb links...”




Try for  $>$  1Gb at least... even if in ‘optimal’ conditions

# USCMS Intercloud Tests




OWAMP




BWCTL









NOT A STEREOGRAM

# So, some problems then...

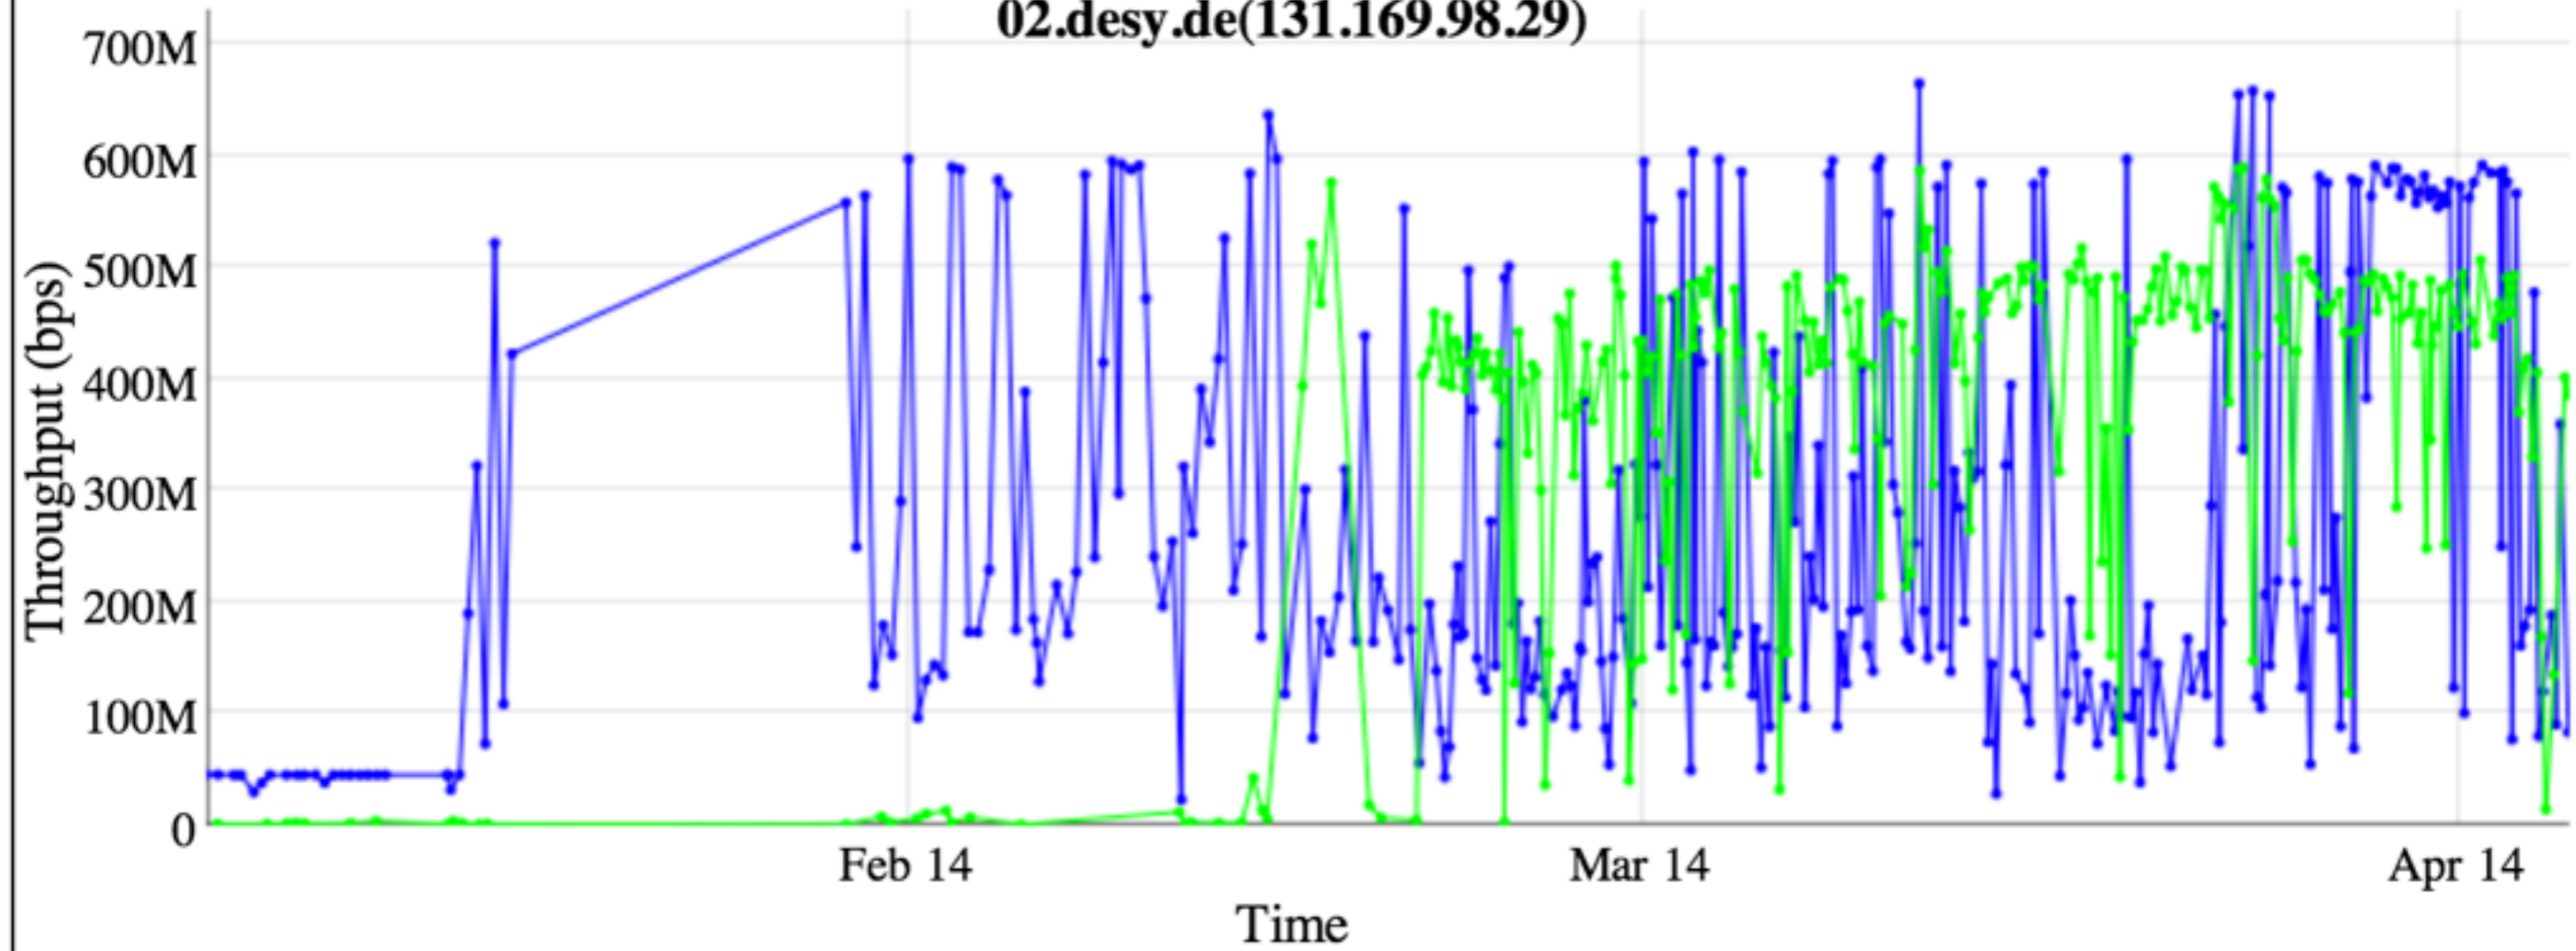
- ... but are they actually problems?
- TCP, UDP, parallel streams ohh my!
- pS is still 'new' in most networks
- network admins need to know the tool exist
- ... and then they need to care

# Say you do care...

- Get curious
  - `bwctl -c hcc-ps02.unl.edu -s perfsonar-ps-02.desy.de -T iperf -P 1 -i 5 -f m -t 30`
  - `owping -c 1000 perfsonar01.hep.wisc.edu`
  - traceroute / tracepath / MTU insanity
- Tools are always improving - learn them
- Find and interact with the right humans!



Throughput test between Source: hcc-ps02.unl.edu(129.93.239.163) -- Destination: perfsonar-ps-02.desy.de(131.169.98.29)

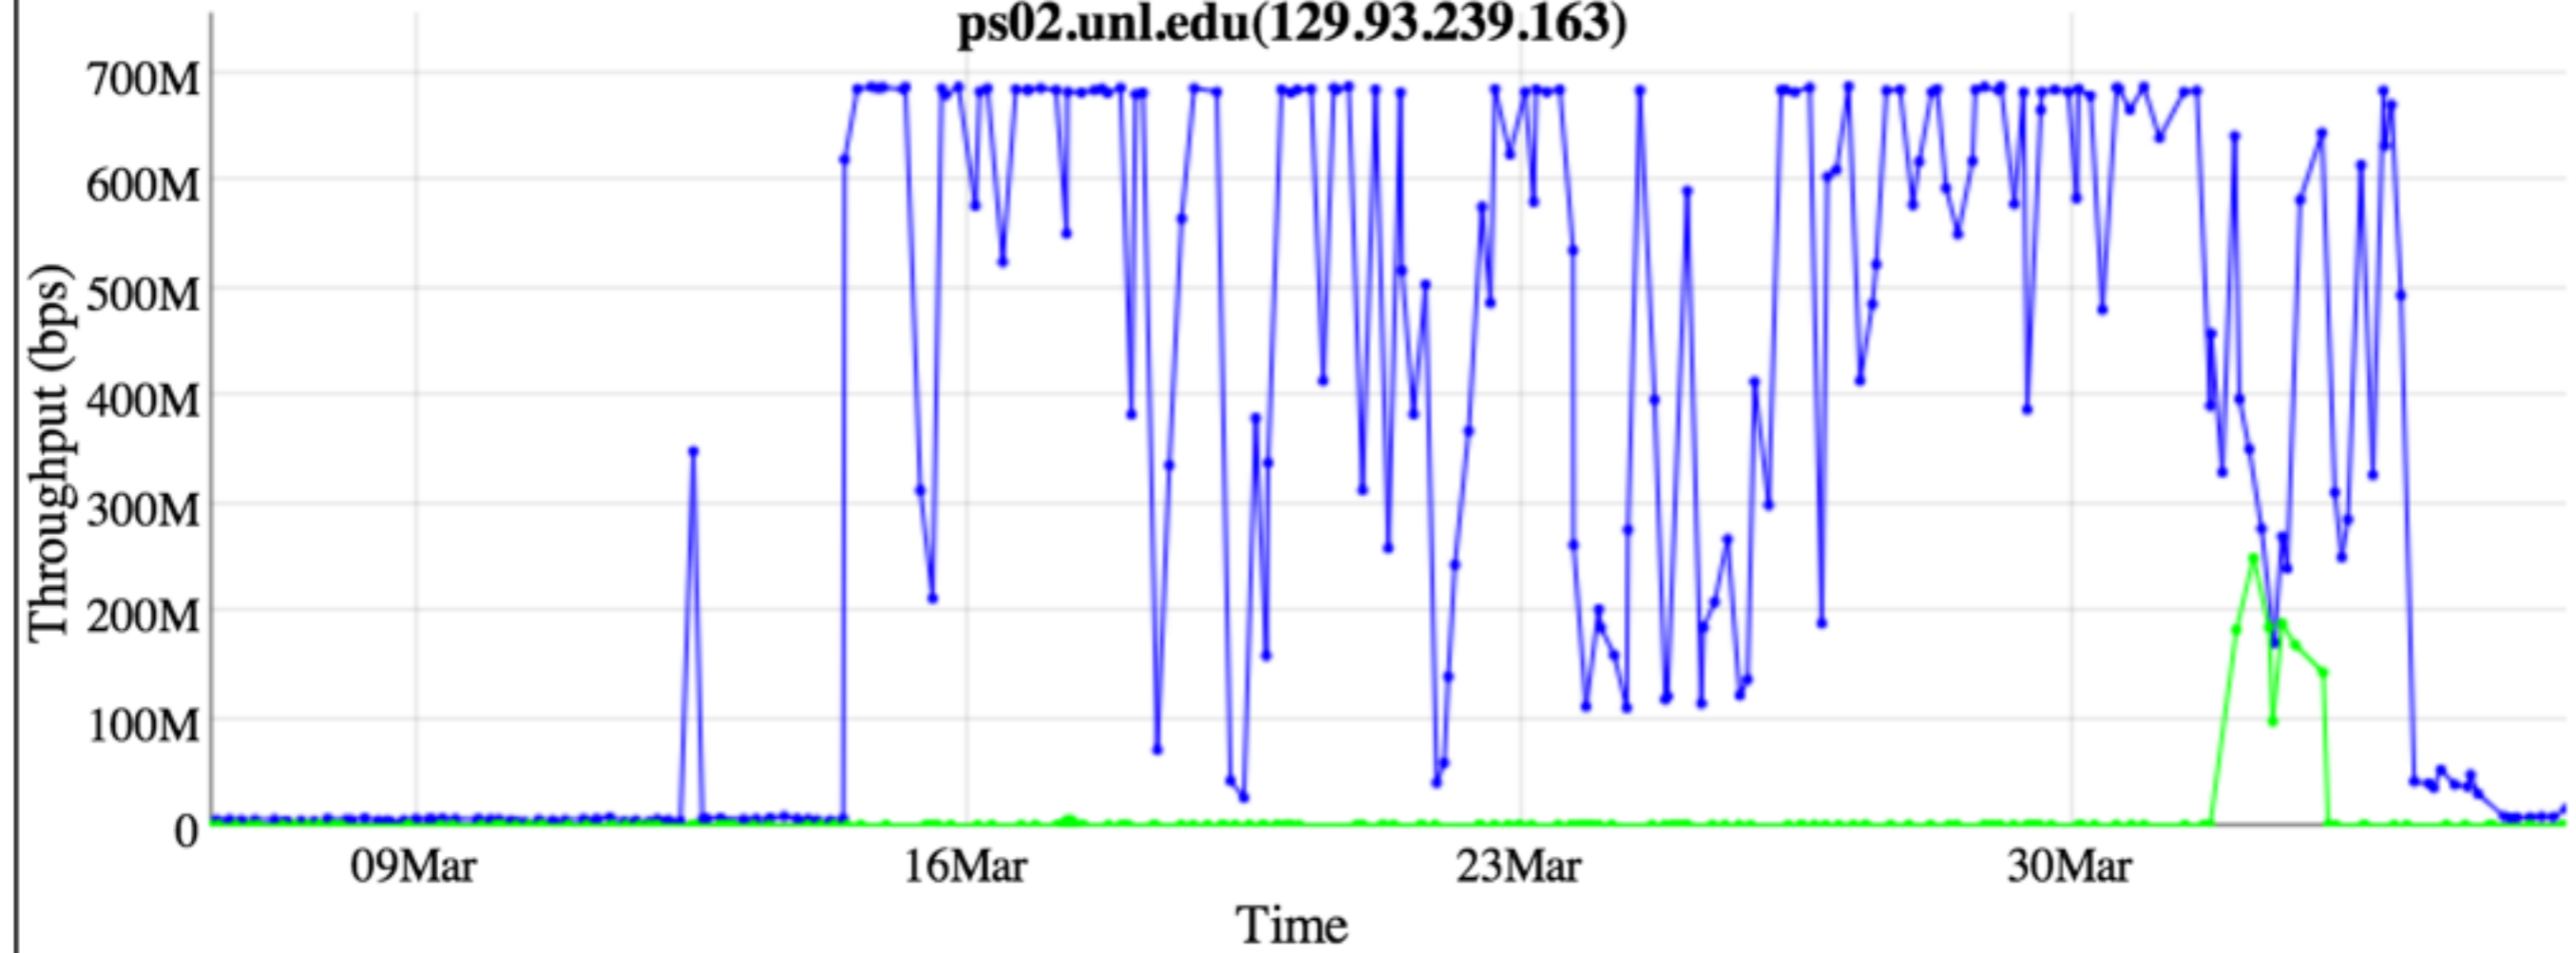


UNL + DESY issue via LHCONE

Took ~year to 'resolve'

Lots of moving parts + people involved

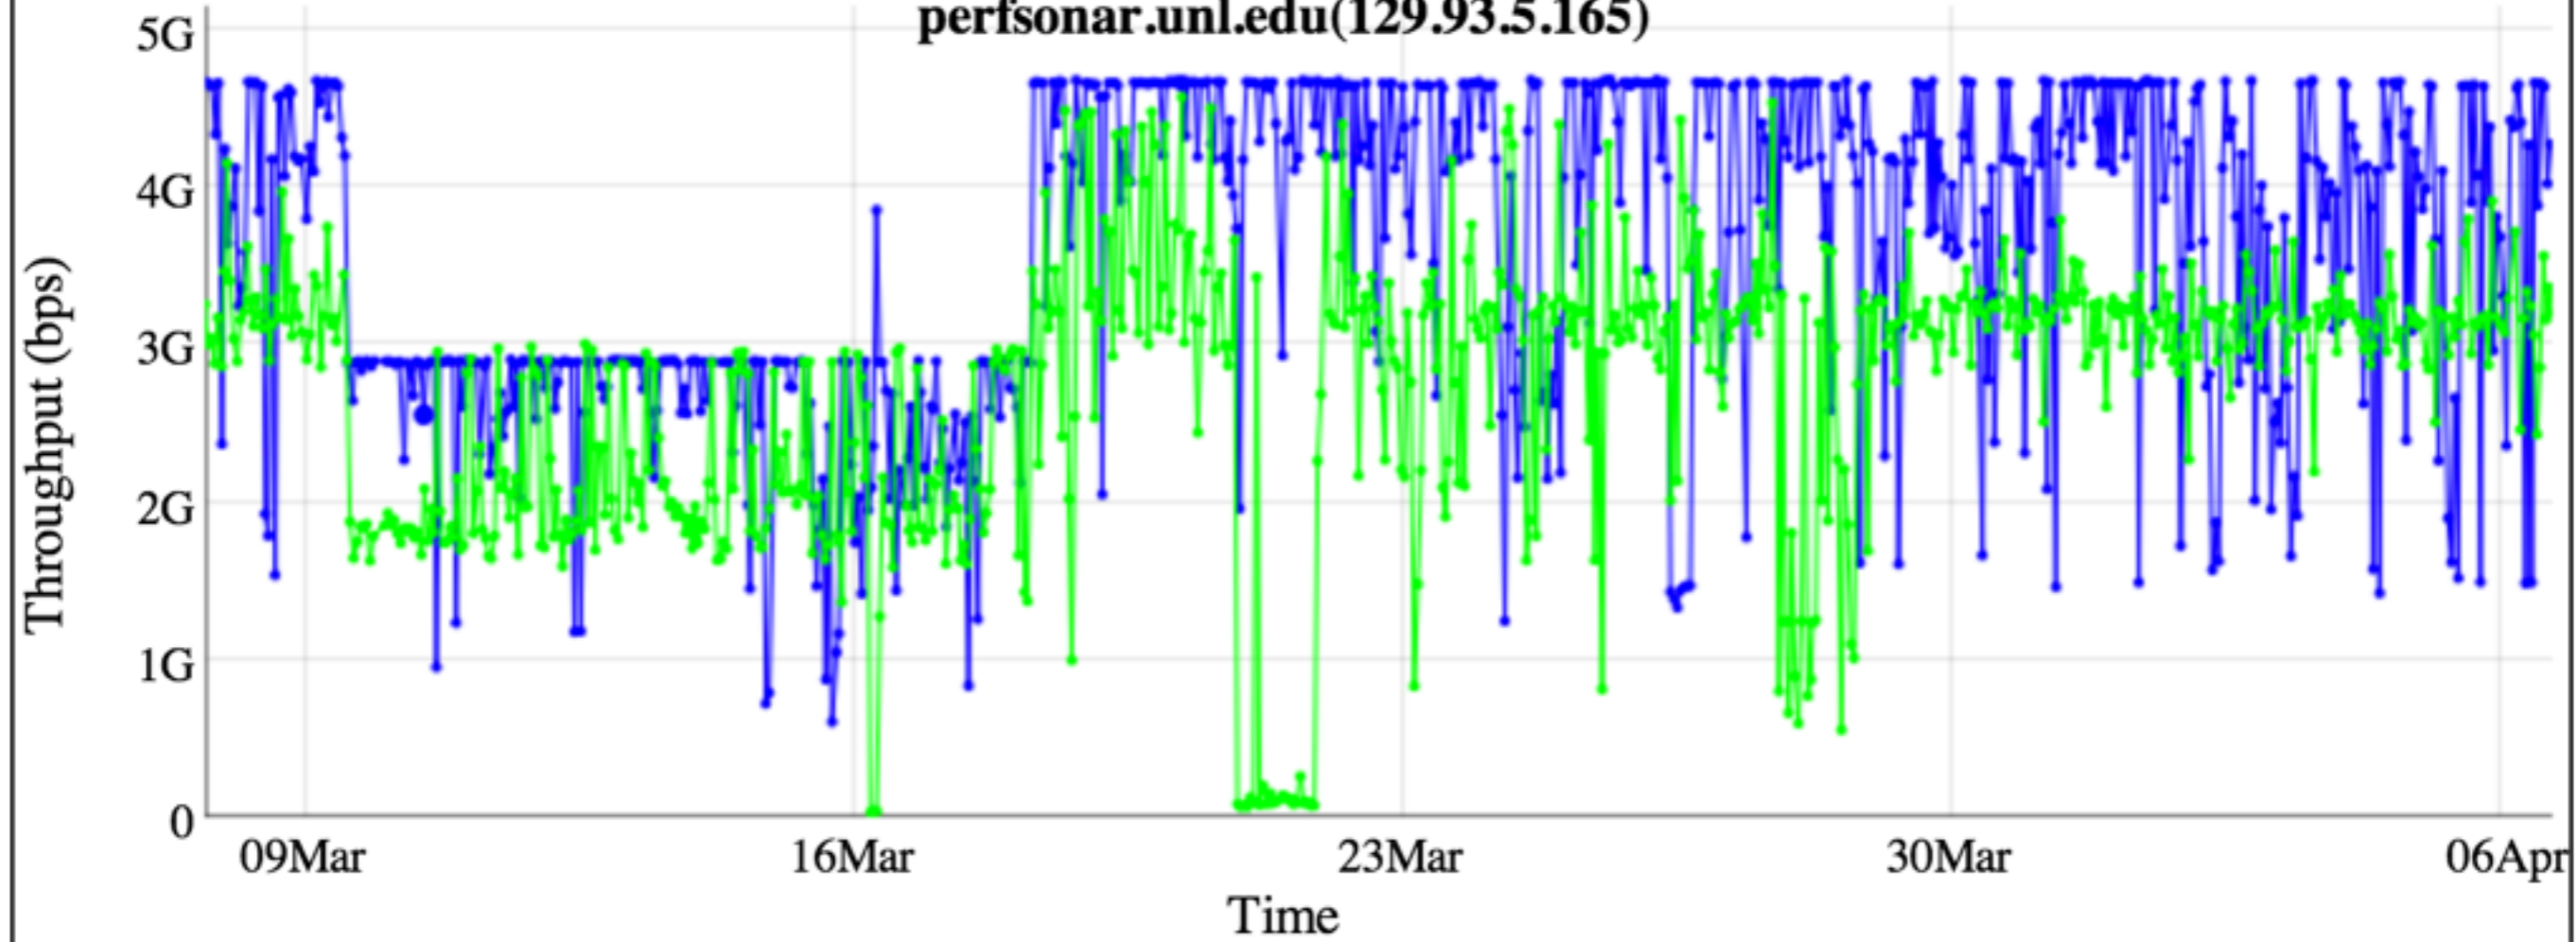
Throughput test between Source: perfsonar-ps.cern.ch(128.142.223.236) -- Destination: hcc-ps02.unl.edu(129.93.239.163)



What is this I don't even ...  
Suspiciously close to ANA100 testing?  
Haven't opened as new thread (yet)

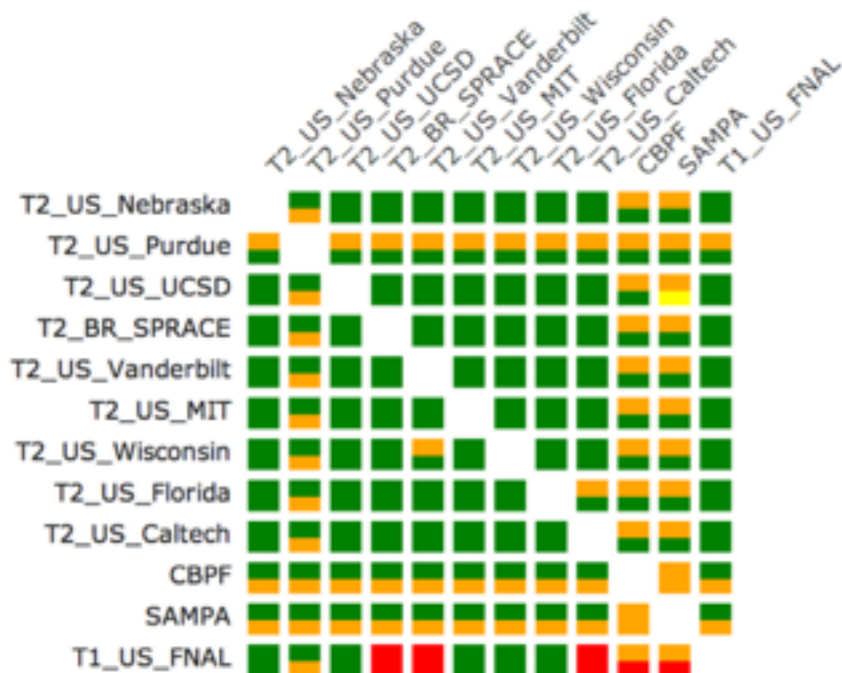
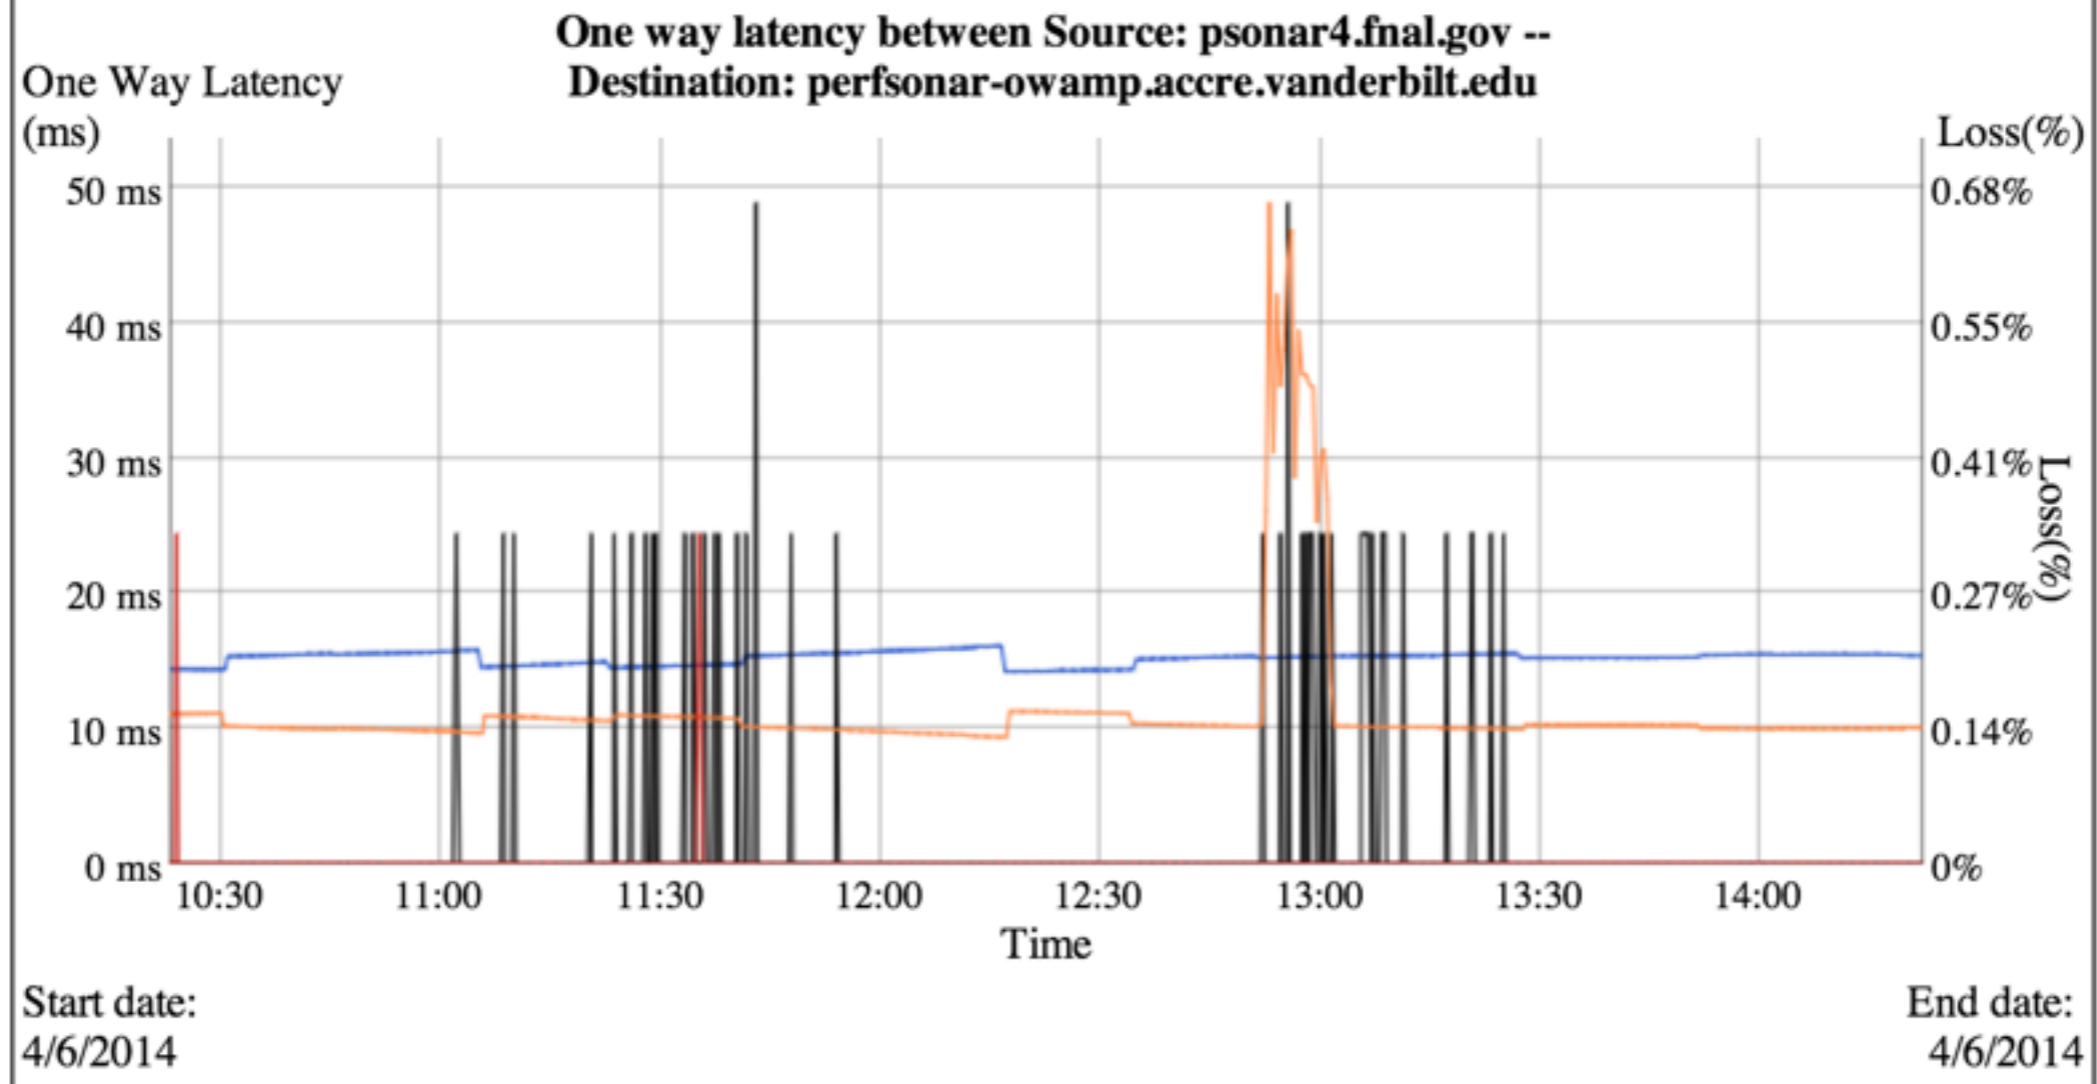
(not even sure this is what real transfers see)

Throughput test between Source: komatsu.chtc.wisc.edu(128.104.100.67) -- Destination: perfsonar.unl.edu(129.93.5.165)



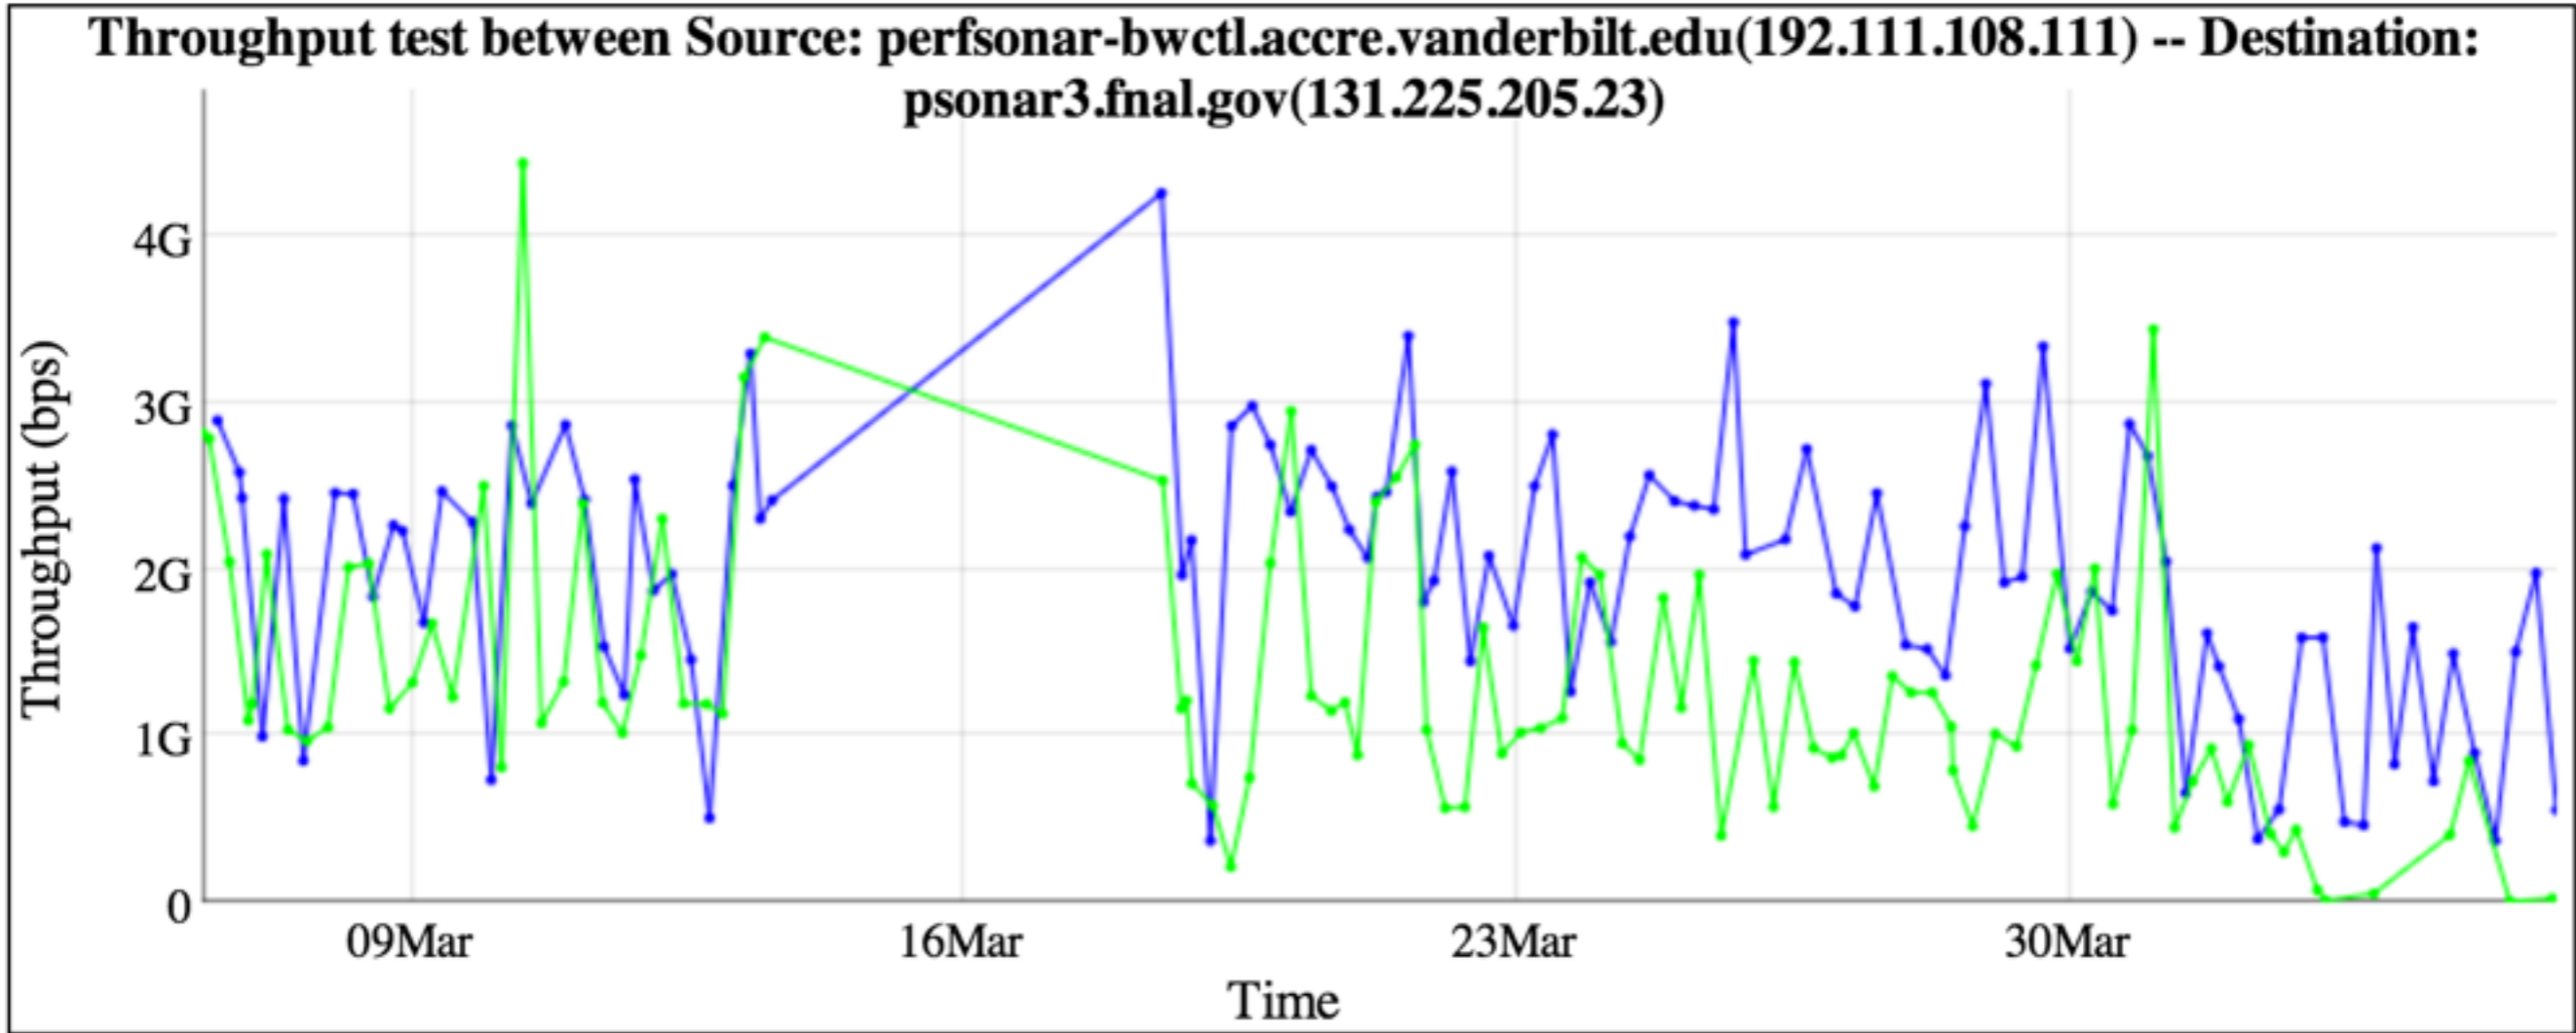
If a butterfly flaps its wings...

Not easy to correlate changes

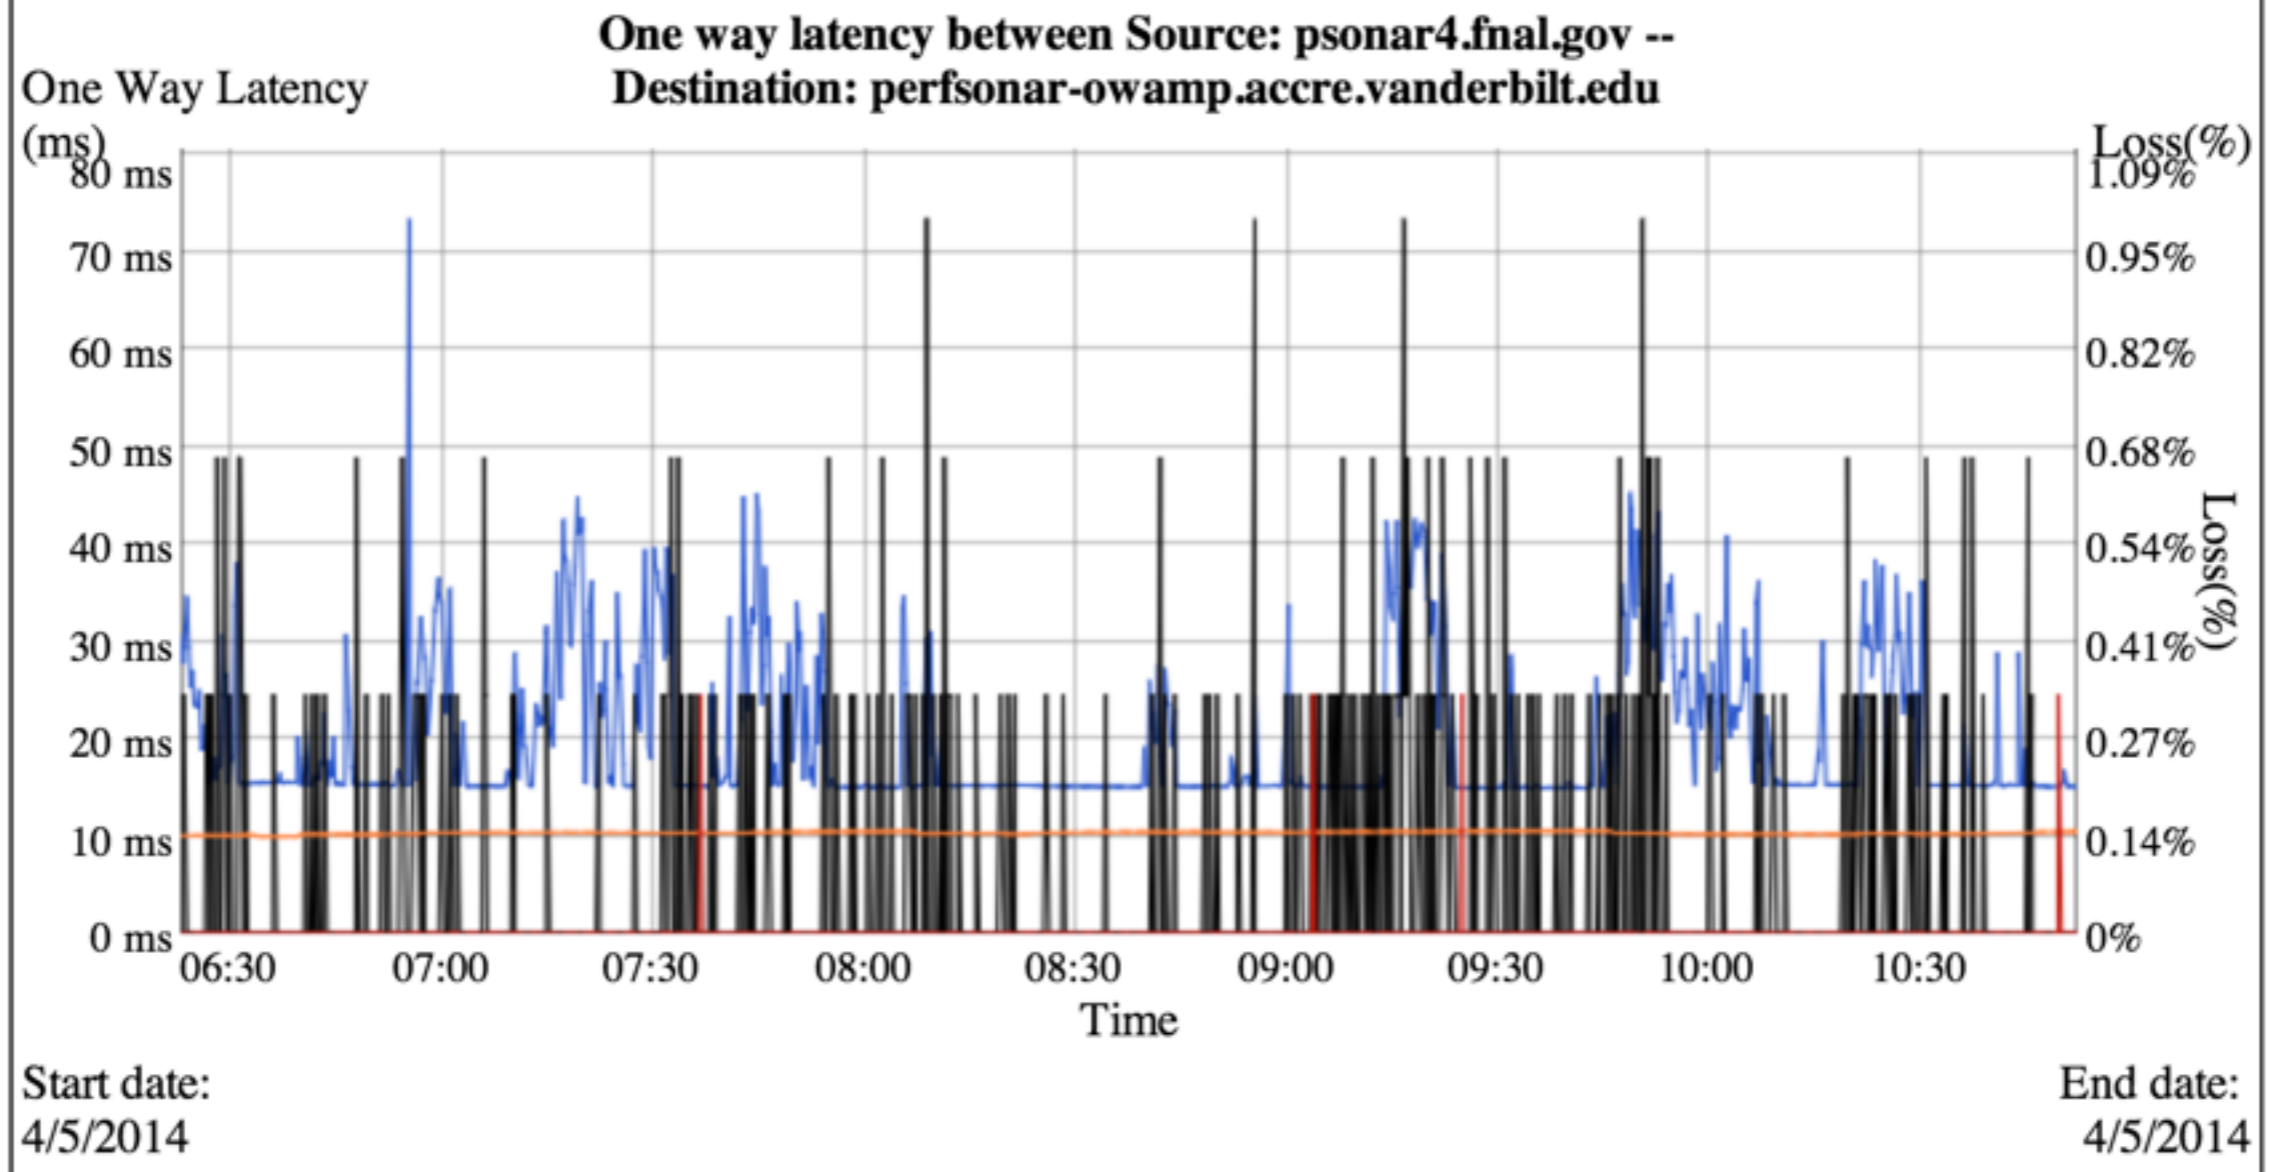


Vanderbilt / FNAL packet loss?



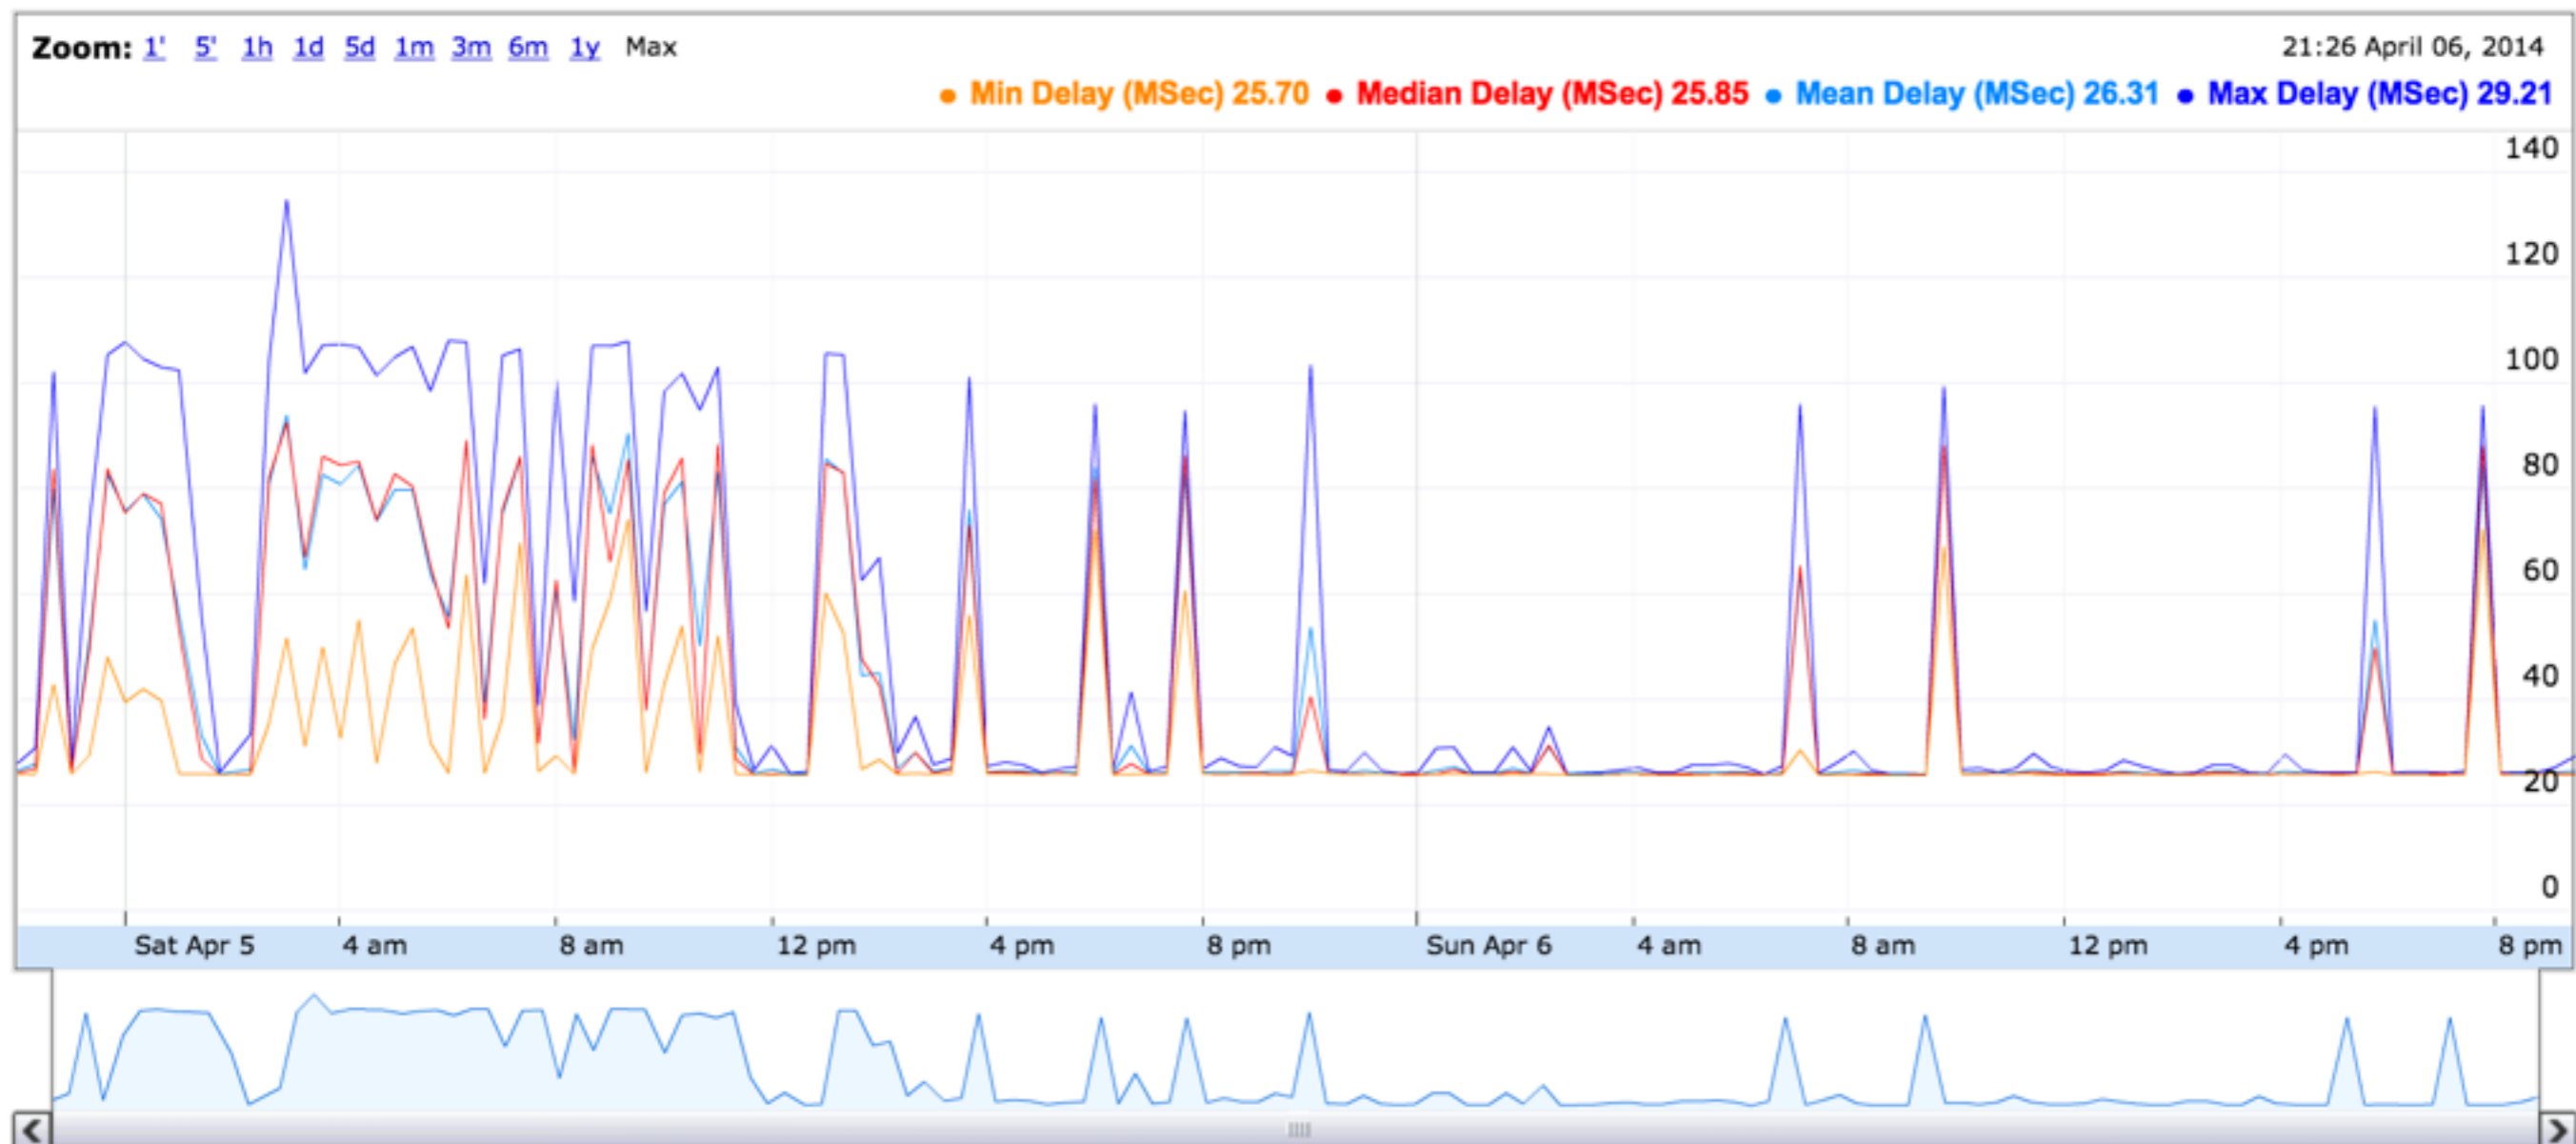


Looks like a degradation in recent days

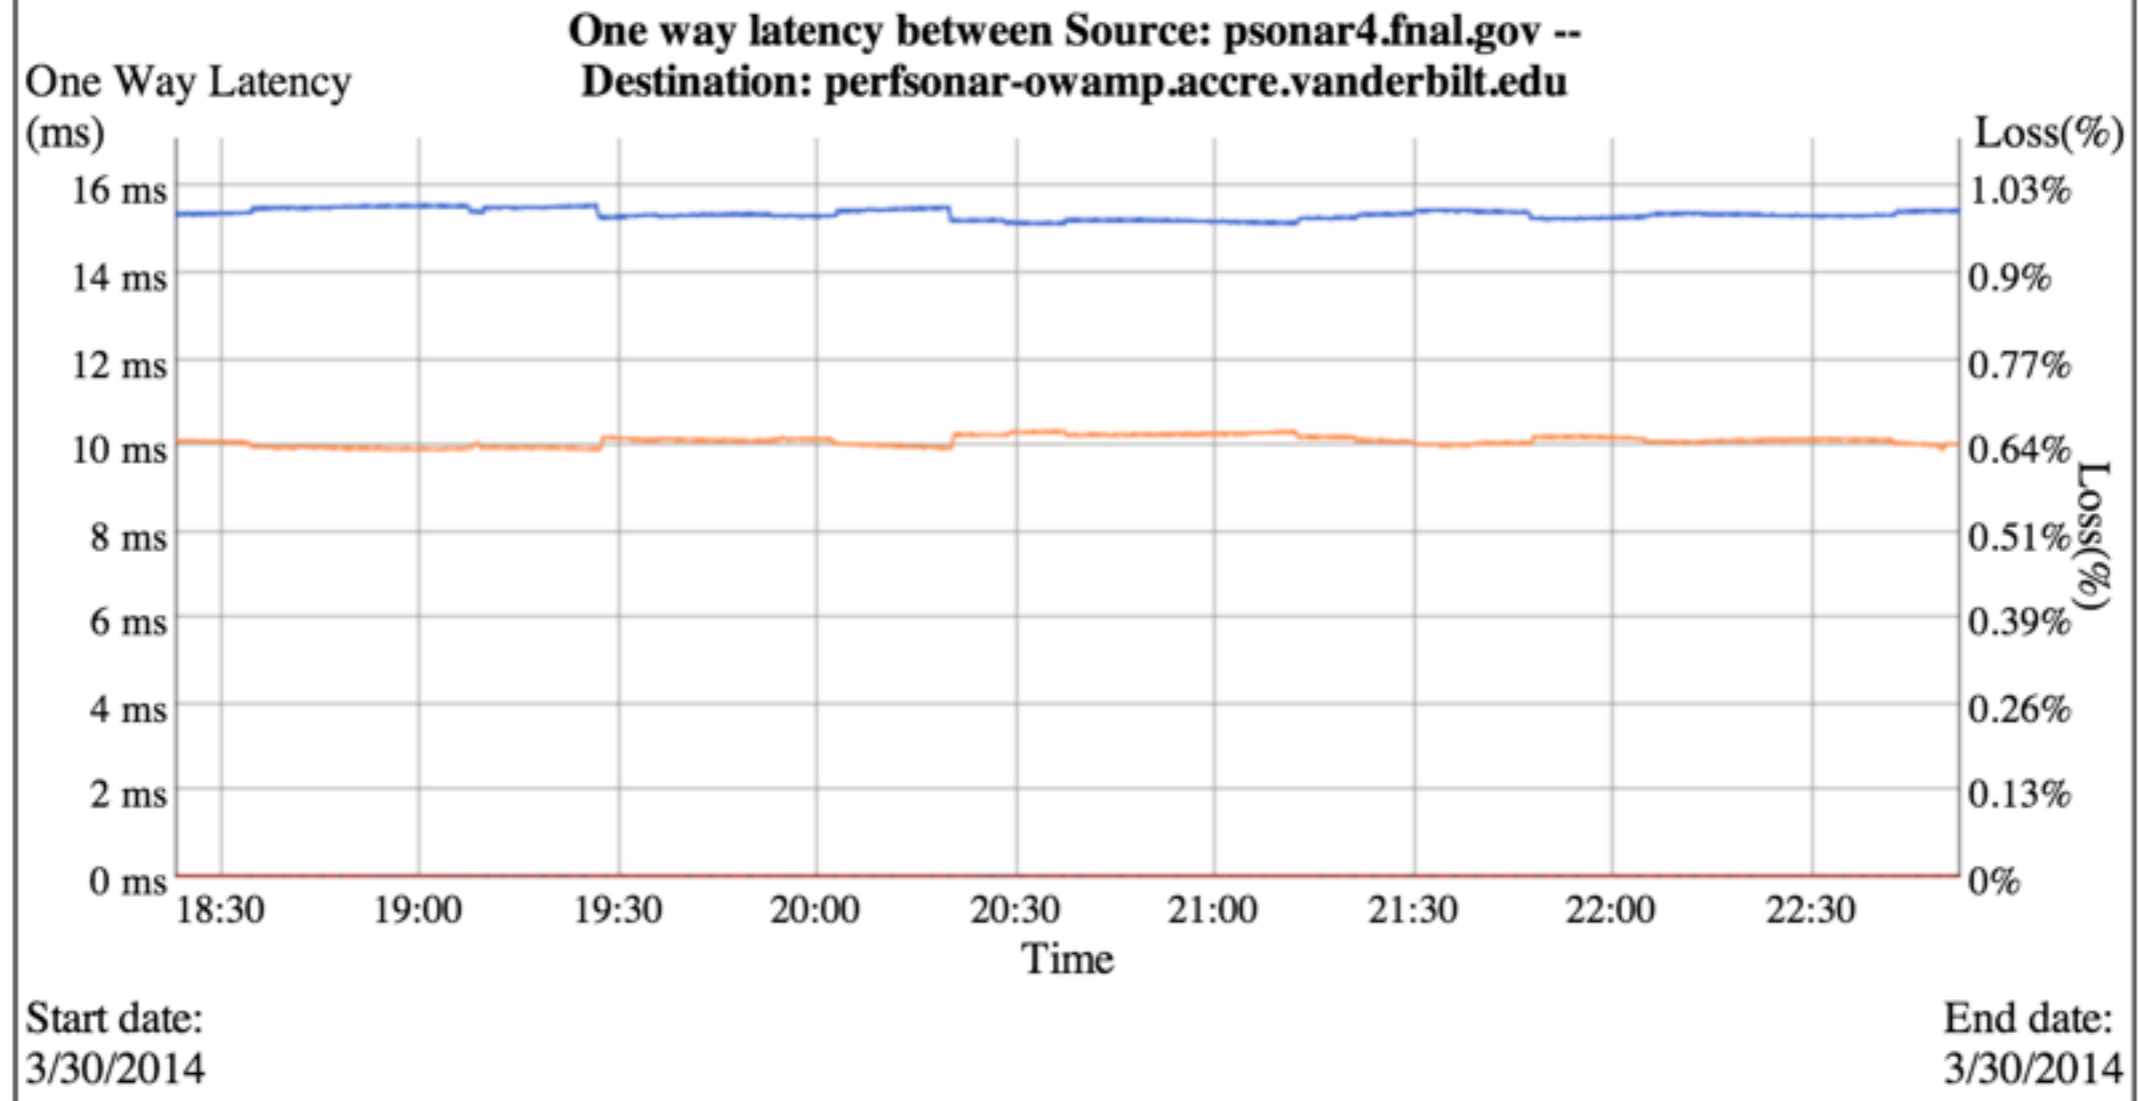


Packet loss + TCP + distance = yeah, we know...

Source: perfsonar-owamp.accre.vanderbilt.edu (192.111.108.112) -- Destination: psonar2.fnal.gov (131.225.205.141)



pingER of the same



A few days prior, the latency is calm  
and the throughput great

(bad pun, I know)



# 15 + n minutes later...

- Watch the dashboard

<http://maddash.ag1t2.org/maddash-webui/index.cgi?dashboard=US%20CMS%20T2%20sites>

- Internet2 perfsonar-user mailing list

<https://lists.internet2.edu/sympa/subscribe/perfsonar-user>

- Throughput mailing list

<https://lists.bn1.gov/mailman/listinfo> (Throughput-1)

- North American Throughput Meeting

Bi-weekly, Tuesdays, 3PM EST, occasional overlaps  
Meeting info is announced on above Throughput-1 list

- Know your network contacts (and buy them beer)

Campus, Regional, NREN, and of course other CMS sites!