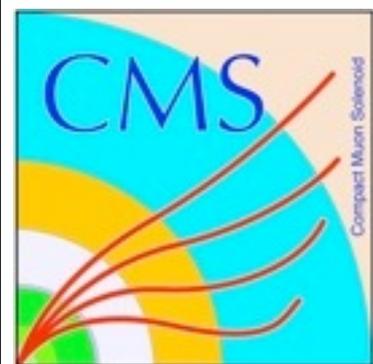


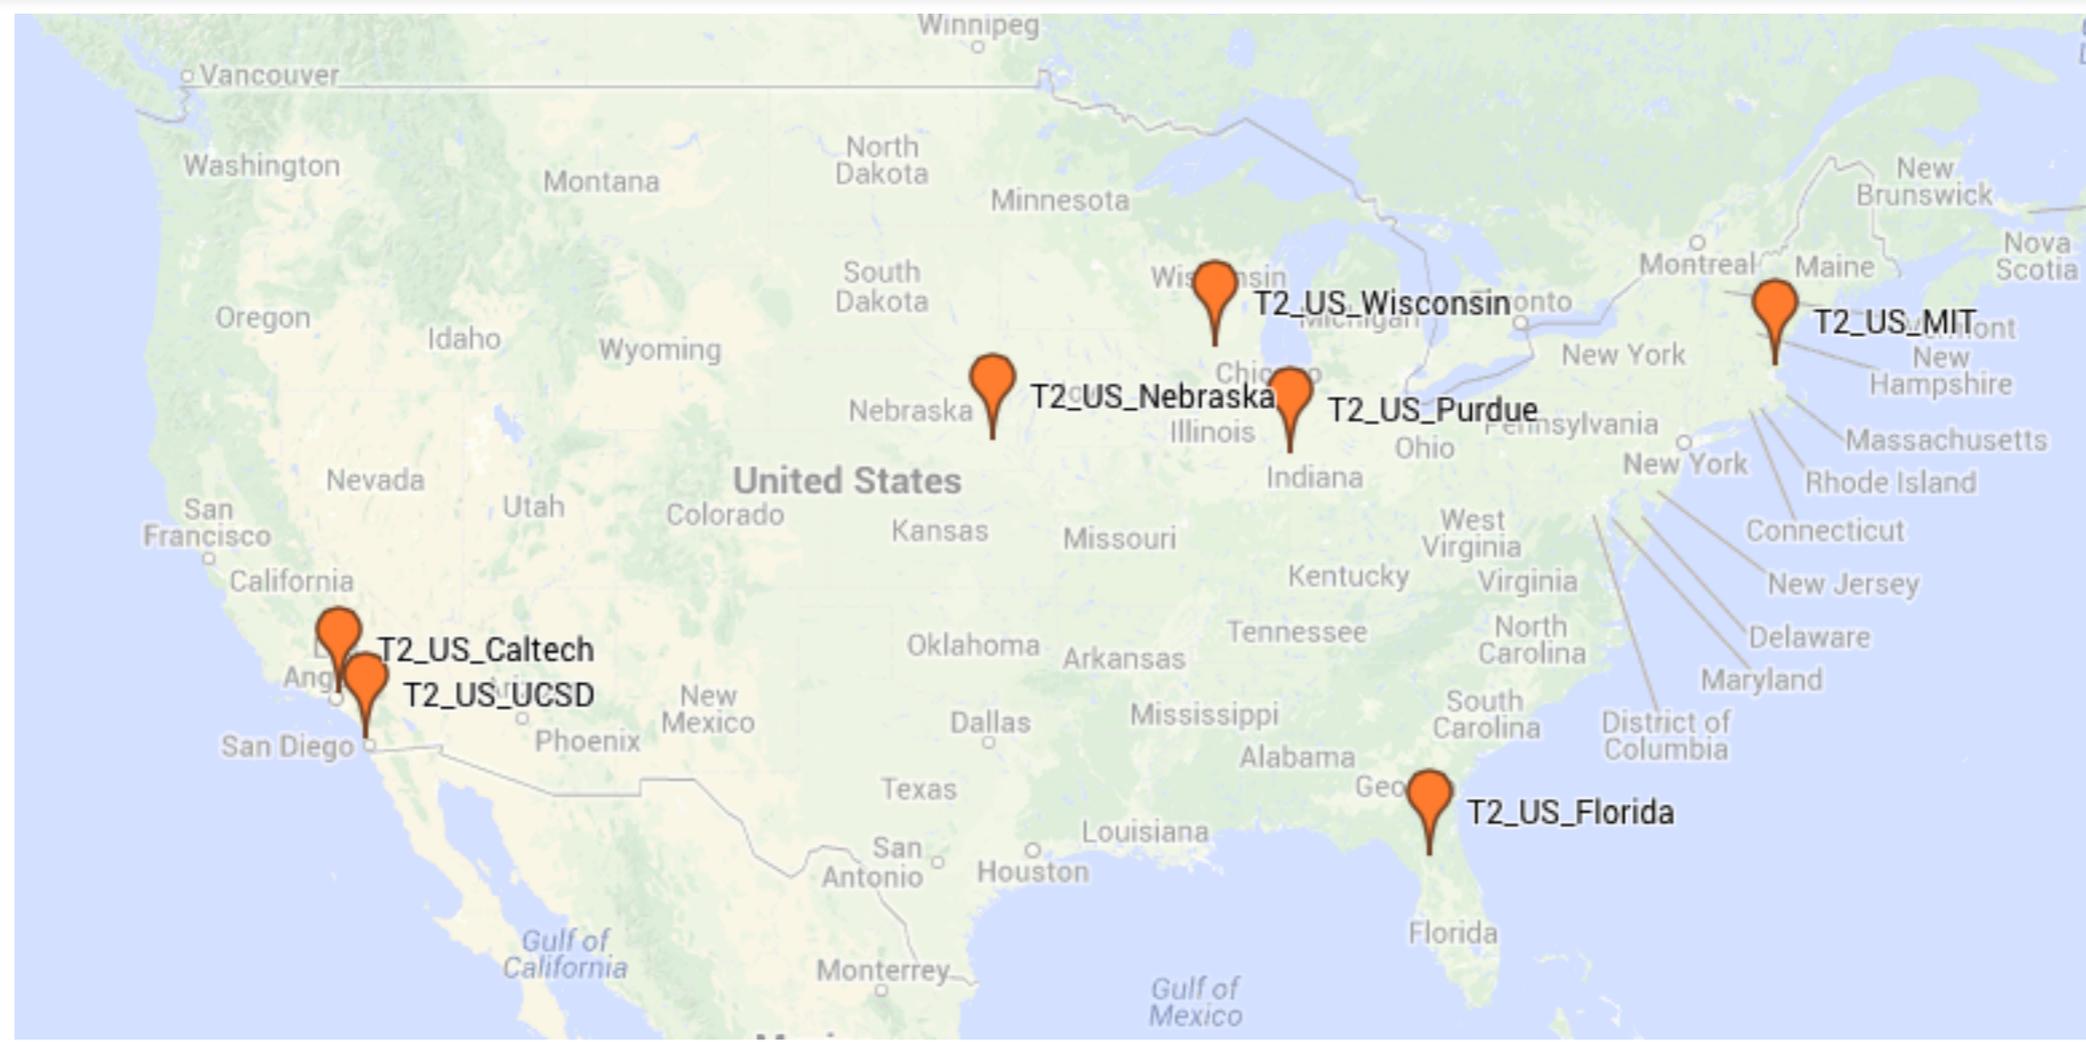
US CMS Tier-2: Comments on LHCONE

Ken Bloom

10 February 2014

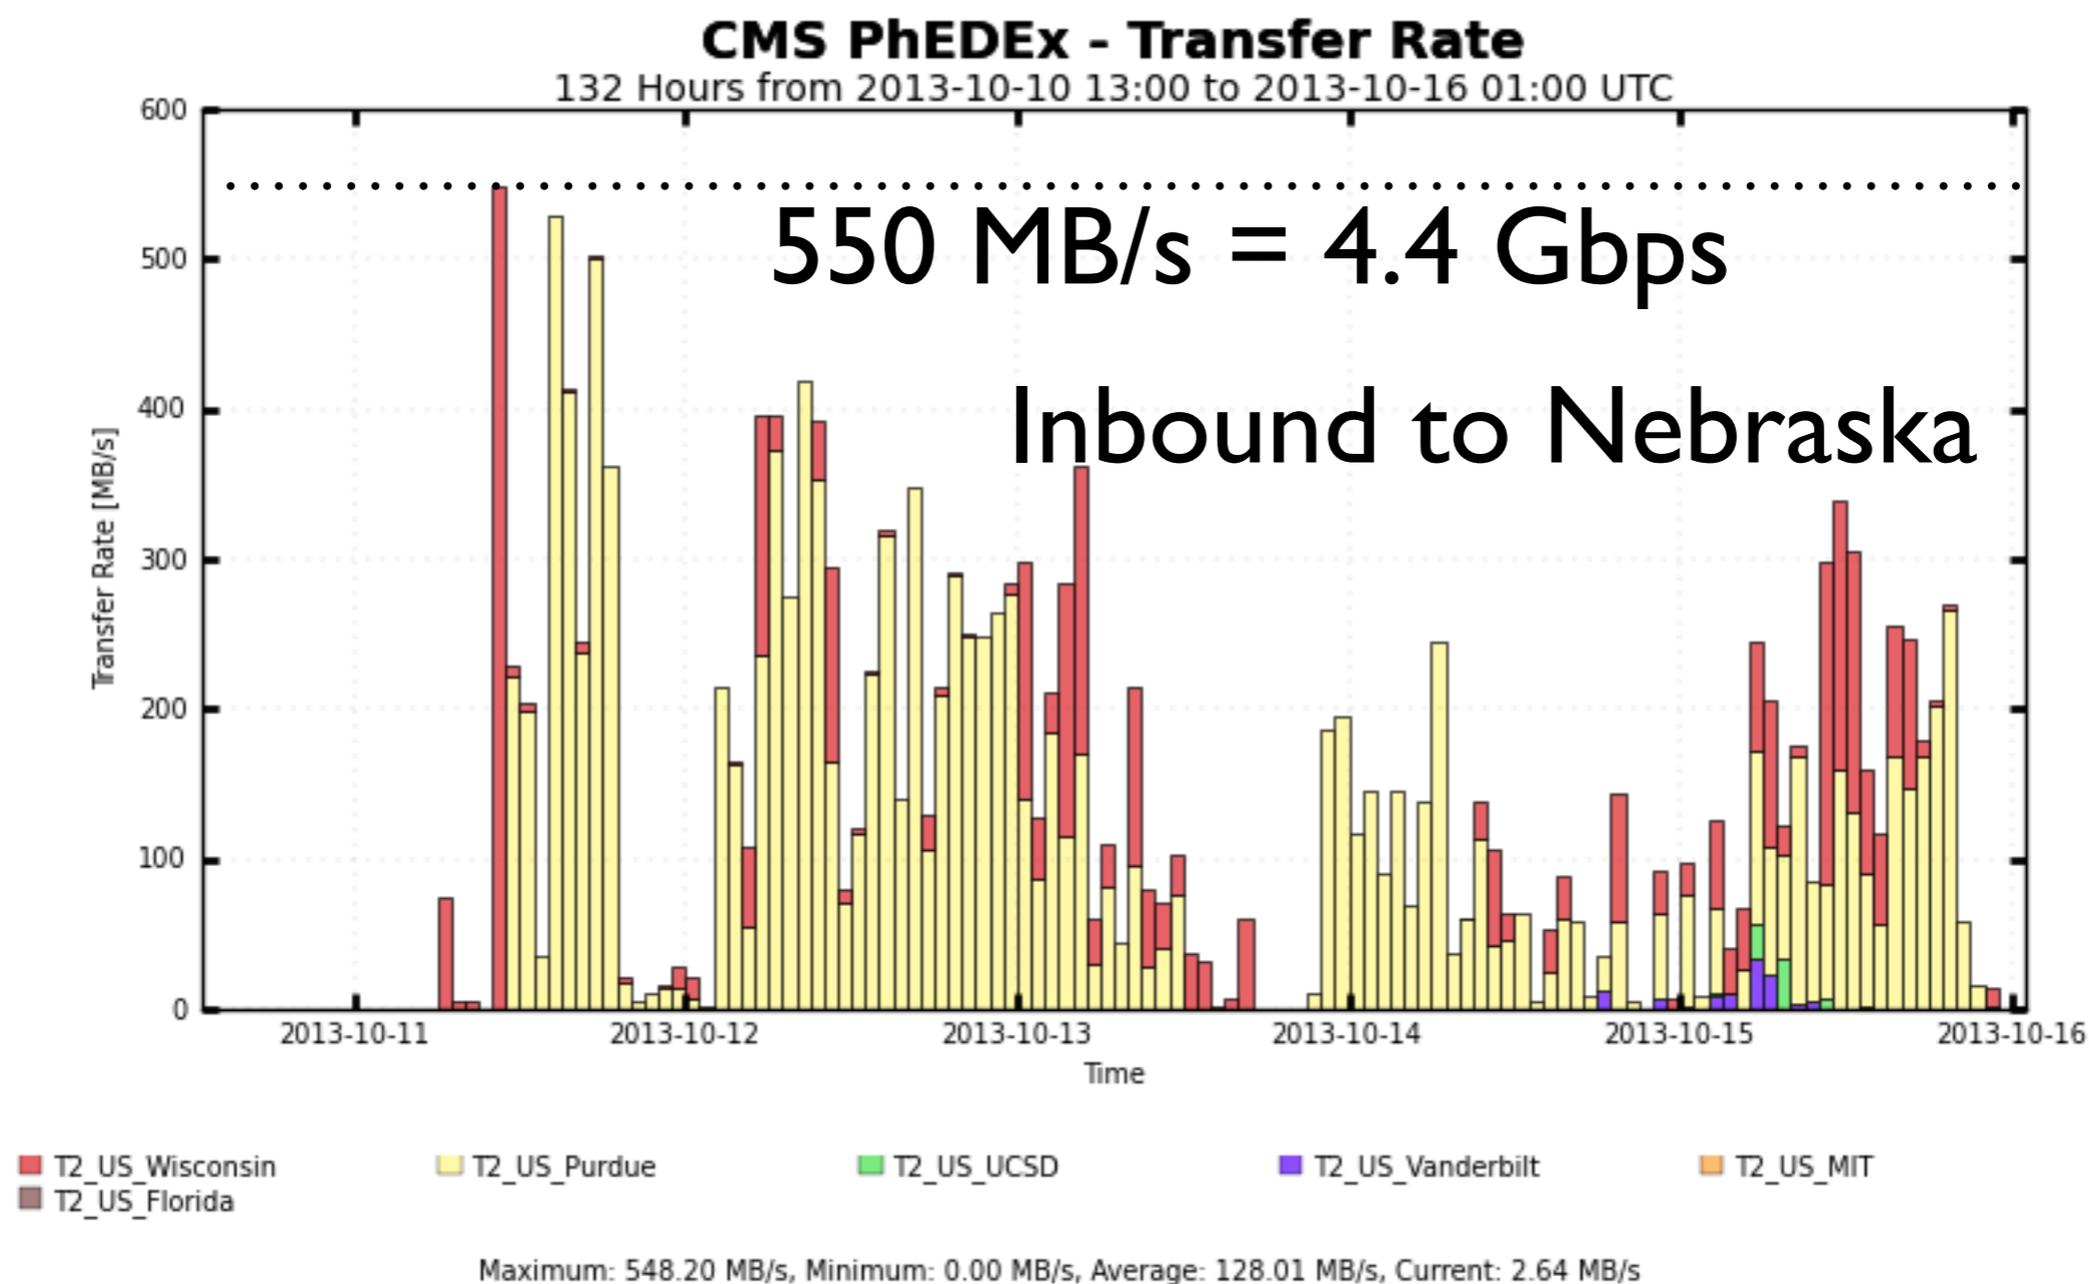
LHCONE/LHCOPN Workshop





- ▶ US CMS operates seven T2 sites for CMS
- ▶ Processing for simulations and user analysis, storage for analysis
- ▶ These sites have substantial resources: in aggregate, ~30K batch slots potentially available, 10 PB disk
- ▶ All have at least 10 Gbps WAN, with plans to expand

- ▶ Recent example of data movement within the US last year
- ▶ Only transfer data when needed, but when we do, it works well



Site	Upgrade plan	LHCONE
Caltech	100 Gbit by March 2014	Yes
Florida	100 Gbit available	Planning to
MIT		
Nebraska	100 Gbit in March 2014	Yes
Purdue	100 Gbit available	No plan
UCSD	100 Gbit in August 2014	“Depends”
Wisconsin	40 Gbit by Summer 2014	No plan

- ▶ Note: 1000 T2 batch slots can analyze 2.4 Gbit/s of CMS data
- ▶ Needless to say, given the effort and expense needed to upgrade the campus network infrastructure, we want to make the best use of it for scientific productivity

- ▶ (Thanks to G. Attebury (Nebraska) for this timeline!)
- ▶ 11/2012: Nebraska peers with LHCONE at an Internet2 router in Kansas. Observe > 400 Mbit/s to/from transatlantic sites.
- ▶ 9/5/13: Rates from IN2P3, INFN, DESY drop to < 10 MB/s
- ▶ 28/5/13: Start investigating. Opened a ticket with GEANT, although confusion over underlying cause led it to be closed later.
- ▶ Investigations with traceroute, perfSONAR, assisted by S. McKee, J. Zurawski (ESnet), D. Finkelson (I2)
 - ▶ But no understanding of how LHCONE handles trouble reports, or even how to contact LHCONE
 - ▶ Apparently lhcone-operations@cern.ch is not the preferred venue for trouble reports

- ▶ 27/6/13: As a result of more email, ticket #54202 is opened with GEANT. It remains open today. Was noted that there were packet ordering issues that might be the cause.
- ▶ Following months: more traceroutes requested and sent, more tests, no changes
- ▶ 17/10/13: Another mail to lhcone-operations, some discussion at Monday LHCONE call with renewed “interest”
- ▶ 19/11/13: M. O'Connor (ESnet) asks for more traceroutes, as ESnet observed some changes in the network. But changes to traceroutes seemed insignificant.
- ▶ Phone call with GEANT NOC for live troubleshooting, assisted by T. Khan. UDP performance great in both directions, problem was only with TCP. Packet reordering still seems like the culprit. Could do further testing when I2 could provide a test endpoint...

- ▶ 24/1/14: MO'C asks for the I2 ticket number for the problem. Apparently there is one, but we don't know what it is. Perhaps the matter will be revitalized.
- ▶ Still no improvement to the throughput from the EU sites, eight months after the issue was originally raised.
 - ▶ Note: NRENs have been cooperative throughout this process
 - ▶ Also note: Vanderbilt is currently having a very similar experience
- ▶ What might a site conclude from this?
 - ▶ LHCONE has no working model of operations
 - ▶ No ticket system, no operations center, no clear points of contact
 - ▶ Issue resolution lacks transparency
 - ▶ Not allowed to see the tickets, can't tell if there is any movement
 - ▶ No entity is taking responsibility to solve problems observed and confirmed by multiple parties
 - ▶ R&E networks do not appear to consider LHCONE a priority

- ▶ From the perspective of network users who have the current set of LHC use cases, we do not see a need for LHCONE and virtual circuits within the United States, where an abundance of bandwidth and routed circuits are working very well for us
 - ▶ Virtual and/or on-demand circuits add complexity we don't need
 - ▶ R&E networks appear to have enough to do without having to debug additional complexity
 - ▶ Also concerned about additional funding/operations costs
- ▶ The LHCONE approach may provide benefits for networking outside the US, but current state of operations support does not make it possible
- ▶ On the basis of this, a focus on maintaining bandwidth and connectivity via “science DMZ’s” to all sites where LHC science is being done seems preferable