T2_US_Nebraska
Site Report

Garhan Attebury
HEPiX Spring 2015

Holland Computing Center

- University of Nebraska-Lincoln
  … and UNO, UNK, UNMC
- ~25k cores in four clusters
- USCMS Tier2 “red”
  T2_US_Nebraska
  5.4k cores, 2.2PB storage
- CMS Tier3 for UNL, KU, KSU
- Involvement with the OSG
- “… that flat place from the Fall 2014 meeting …”

T2_US_Nebraska @ HCC

- “Typical” USCMS Tier2 site
  - CE: HTCondorCE, ~5,392 slots @ ~57.7k HS06
  - SE: Bestman2 + gridftp-hdfs / xrootd-hdfs, ~2.2PB
  - 100Gb connectivity to ESnet / I2 / LHCONE
    - Mix of 1Gb and 10Gb connectivity to hosts
    - (almost) all dual-stack with IPv6

T2_US_Nebraska CE/SE Hardware

Qty	Type	CPU	RAM	Net	Job Slots	HDFS
16	Dell R730xd	E5-2640v3	128GB	10Gb	32	43.5TB
28	Huawei RH1288 V2	E5-2650v2	128GB	1Gb	32	-
48	Dell R720xd	E5-2660	128GB	10Gb	32	43.5TB
36	Dell R410	X5650	64GB	1Gb	24	9.8TB
60	Dell R510	E5620	32GB	1Gb	16	21.7TB
12	Dell R410	E5520	24GB	1Gb	12	6.9TB
20	Dell R710	E5520	24GB	1Gb	12	10.5TB
30	Sun X2200	Opteron 2354	24GB	1Gb	8	-
250					**5,392 slots**	**4.7PB (raw)**

- Job slots dependent on RAM (most nodes = # of cores/threads)
- HDFS storage is 2+ replicated, usable is more ~2.2PB
- Out of warranty: ~34% storage, ~29% slots, ~23% HS06
- Estimated 57.7k HS06

!Storage/Worker Tier2 Hardware

- 9x 10Gb transfer servers (GridFTP + xrootd, IPv6) (3x also act as OSG squid caches for HCC)

- 4x old R710 + 3x new R730 VM hosts (gatekeepers, GUMS, SRM, xrootd.unl.edu, hcc-crabserver.unl.edu, development nodes, etc…)

- ~20x other physical boxes (LVS redirectors, CMS specific squid, HDFS, ELK stack, Tier3 login/NFS, hcc-cvmfs.unl.edu, etc…)

Other resources at HCC

- Other CMS resources
  - Tier3 for UNL, KU, KSU
  - xrootd.unl.edu, hcc-crabserver.unl.edu, hcc-cvmfs.unl.edu
    - Various OSG related grid infrastructure
- HPC Clusters: Crane, Tusker, Sandhills
  lustre+IB, OSG opportunistic via HTCondorCE
- Crane/Tusker account towards T2_US_Nebraska
- Some OpenStack contraption…

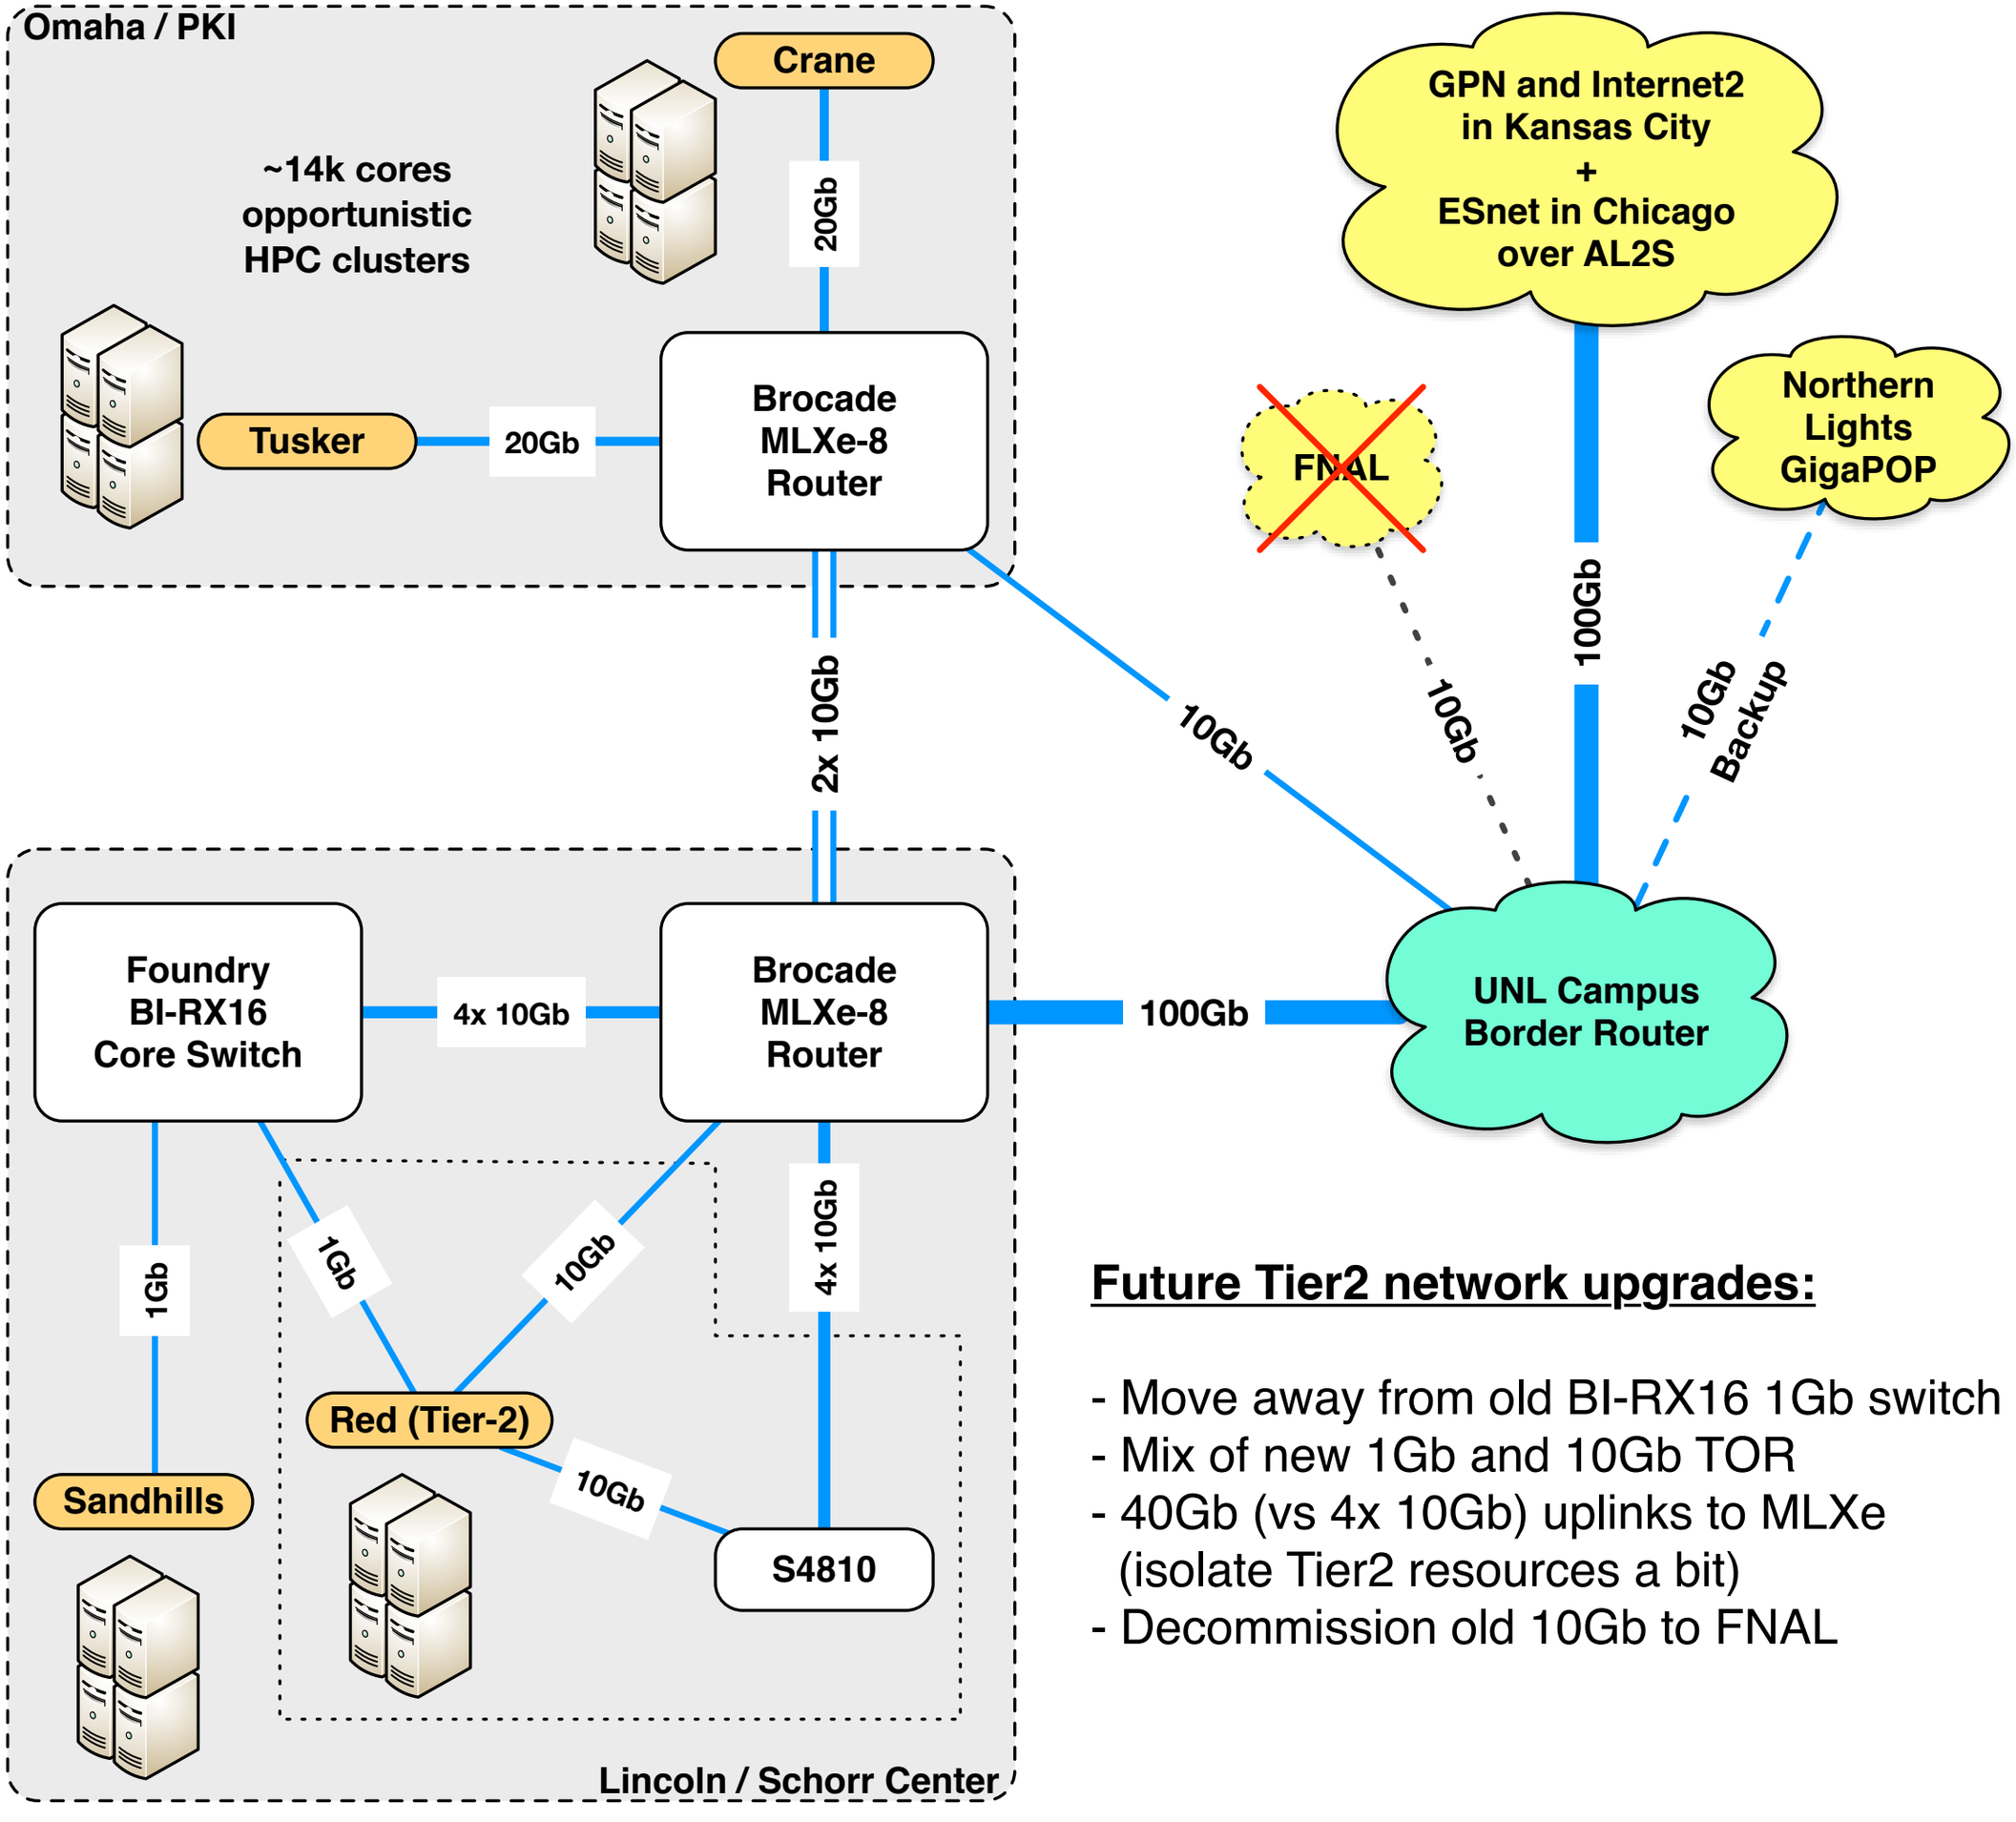

Future Tier2 network upgrades:

- Move away from old BI-RX16 1Gb switch
- Mix of new 1Gb and 10Gb TOR
- 40Gb (vs 4x 10Gb) uplinks to MLXe
  (isolate Tier2 resources a bit)
- Decommission old 10Gb to FNAL

Source
hcc-ps02.unl.edu - 129.93.239.163
Capacity: Unknown MTU: Unknown

--

Destination
perfsonar.ultralight.org - 192.84.86.121 [traceroute]
Capacity: Unknown MTU: Unknown

Link to this chart

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

Previous 1w

Wed Mar 11 08:09:12 2015 -- Wed Mar 18 08:09:12 2015

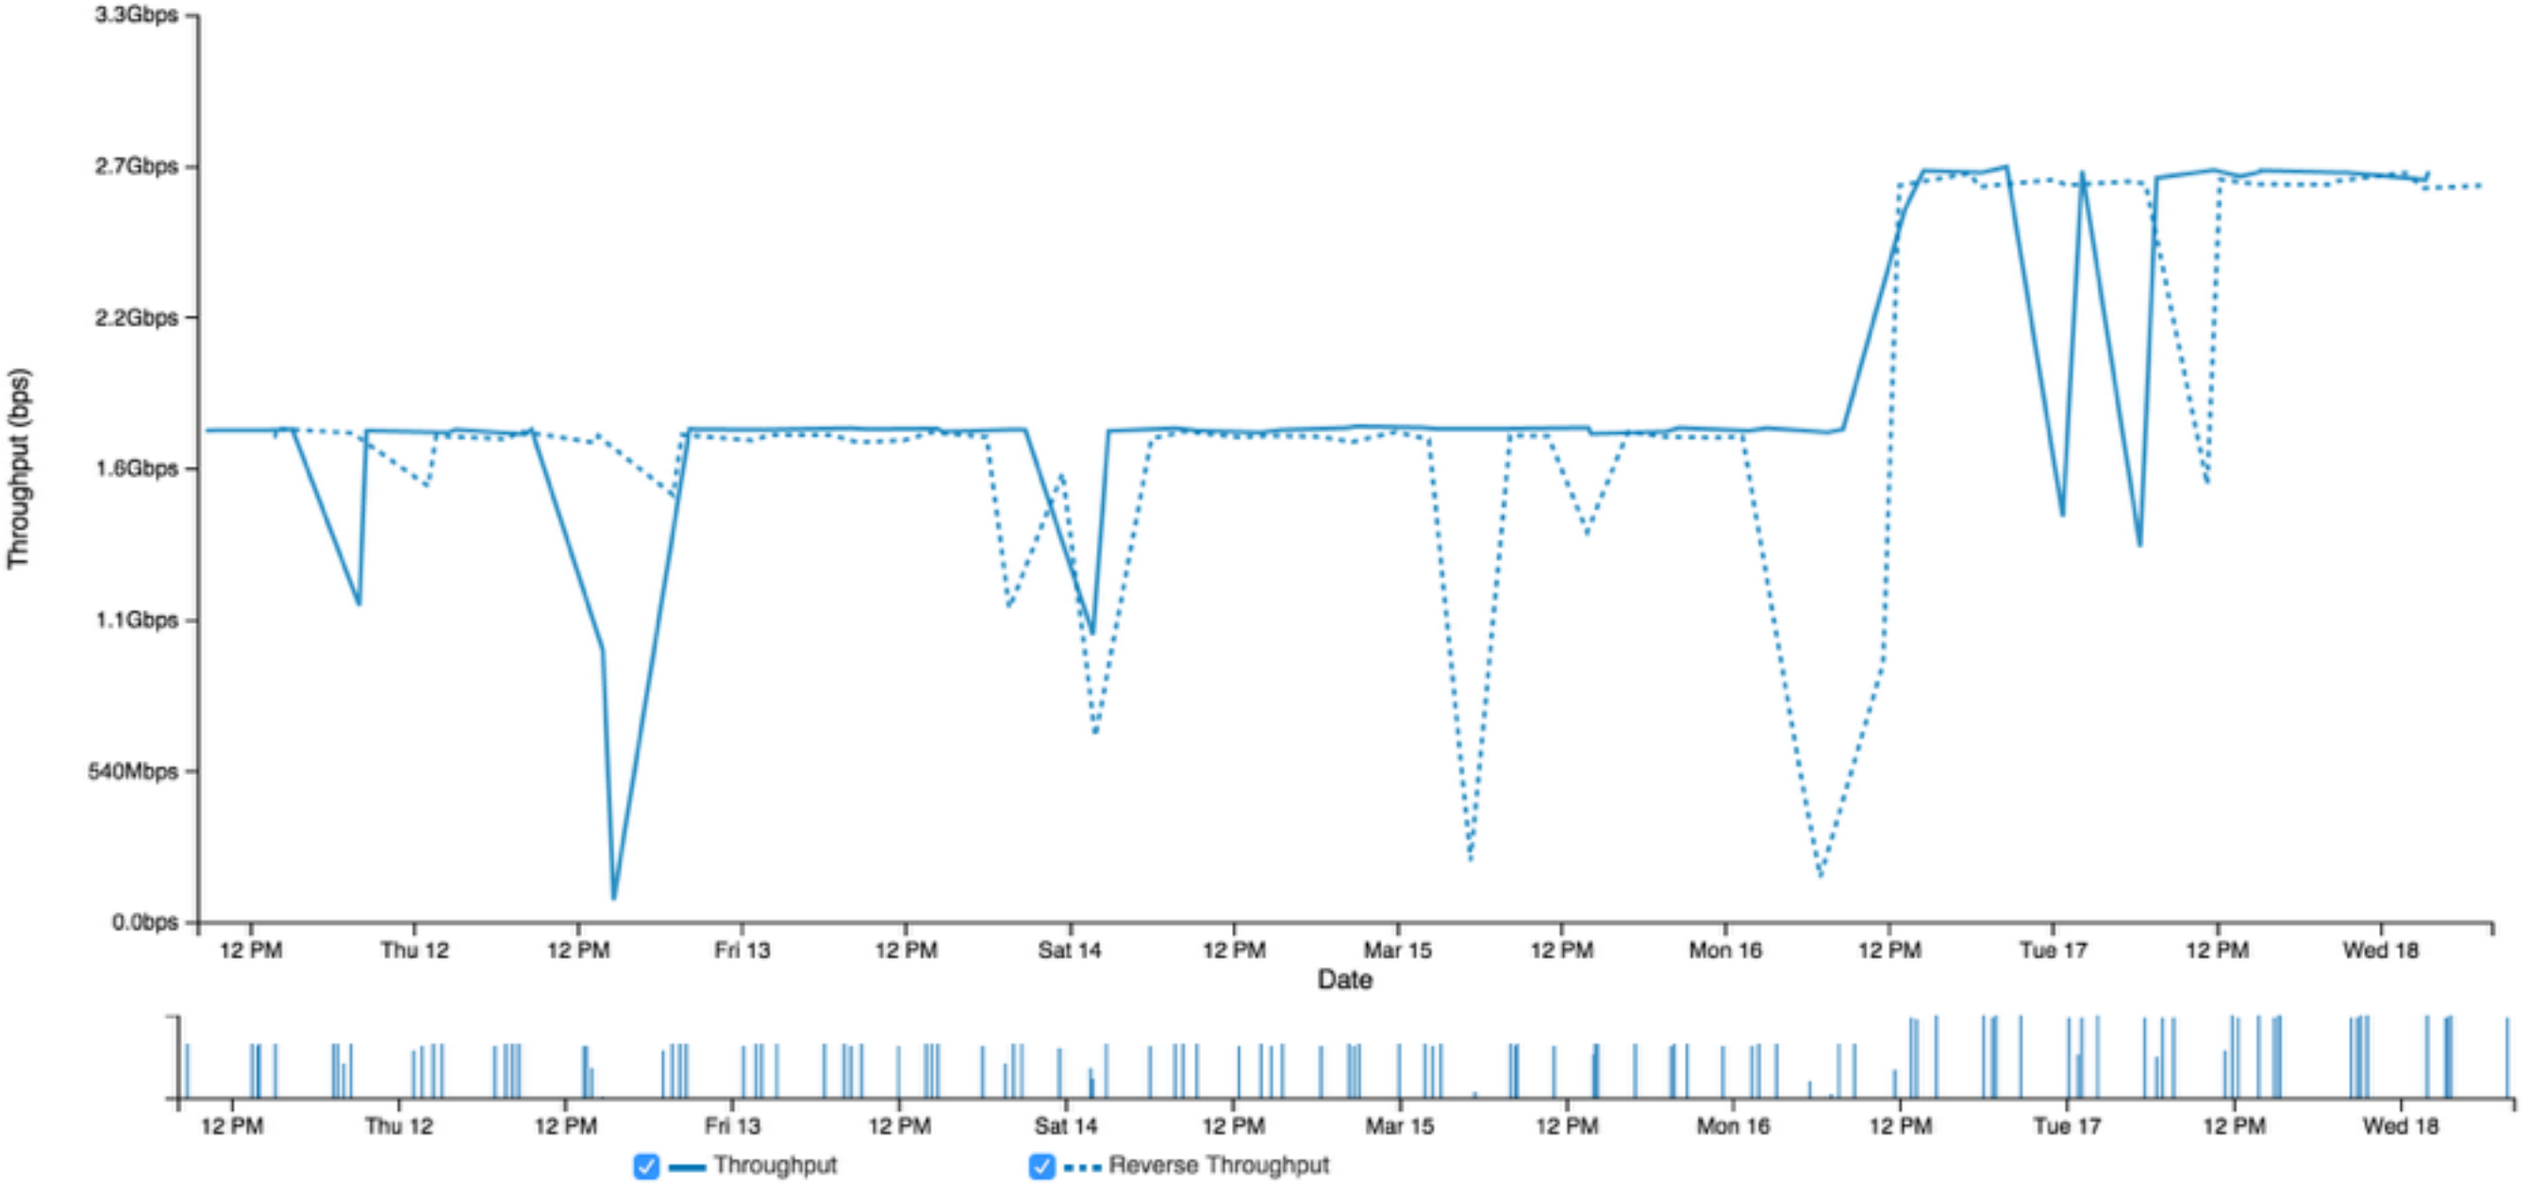

Other networking notes

- DNS … QIP … Many facepalm moments
  - “play nice with campus” vs “get things done”
- IPv6 support in OSG/grid software
  - Mostly dual-stack across the entire Tier2
  - Squids? Bestman?
- Old core switch —> some updates needed
- Overall, networking in good shape

Future Plans

- Maintenance mode for Tier2 - budget cuts?
- Pledge levels for CPU / Storage
  “We need more blinkenlights!
  … but really we’ll be just fine.”
- OpenFlow in production at HCC
  (at least for the Tier2)
- OpenStack

Future Plans

- OpenStack
  (because now we have to)

  Controllers (2x R710)
  Network (2x R710)
  Storage (5x Sun X4275 w/20TB)
  Compute (8x Sun X2200, 8 core, 32GB)

- $630k == Many, many unknowns
  (anyone looking for a job?)

Future wishes…

- 07 Bestman2 
- Squid that does IPv6 … without effort
- No really, EL7 this time, I mean it (!)

Questions?

Garhan Attebury
garhan.attebury@unl.edu