NET2

- Status of procurement, deployment, production operations for remainder of Columbia CA purchases
- Capacity comparison to 2017 & 2018 pledges
- CPU-hours by activity (Production/Analysis; single/multi-core)
- LOCALGROUPDISK capacity and usage
- Planned retirements in 2017
- HTCondorCE deployment status
- Local Site Mover
- Bestman sites: SRM --> Gridftp strategy
- ADC operational issues
- Networking: current BW capacity, peering, usage and upgrade plans
- Plans for IPv6
- Plans for SL7

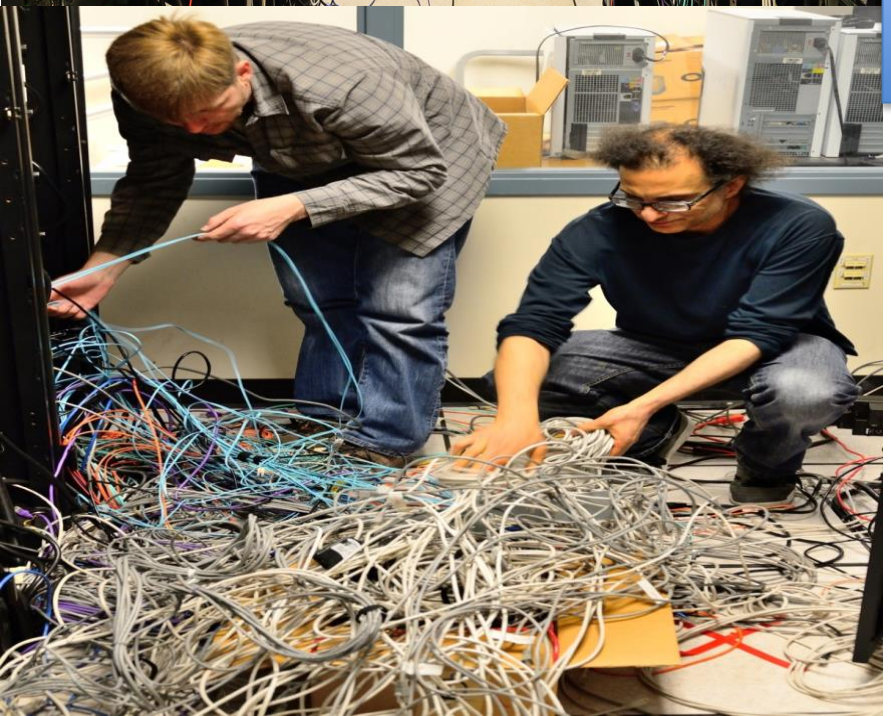

Daily Busiest Hour for the Past 15 Months

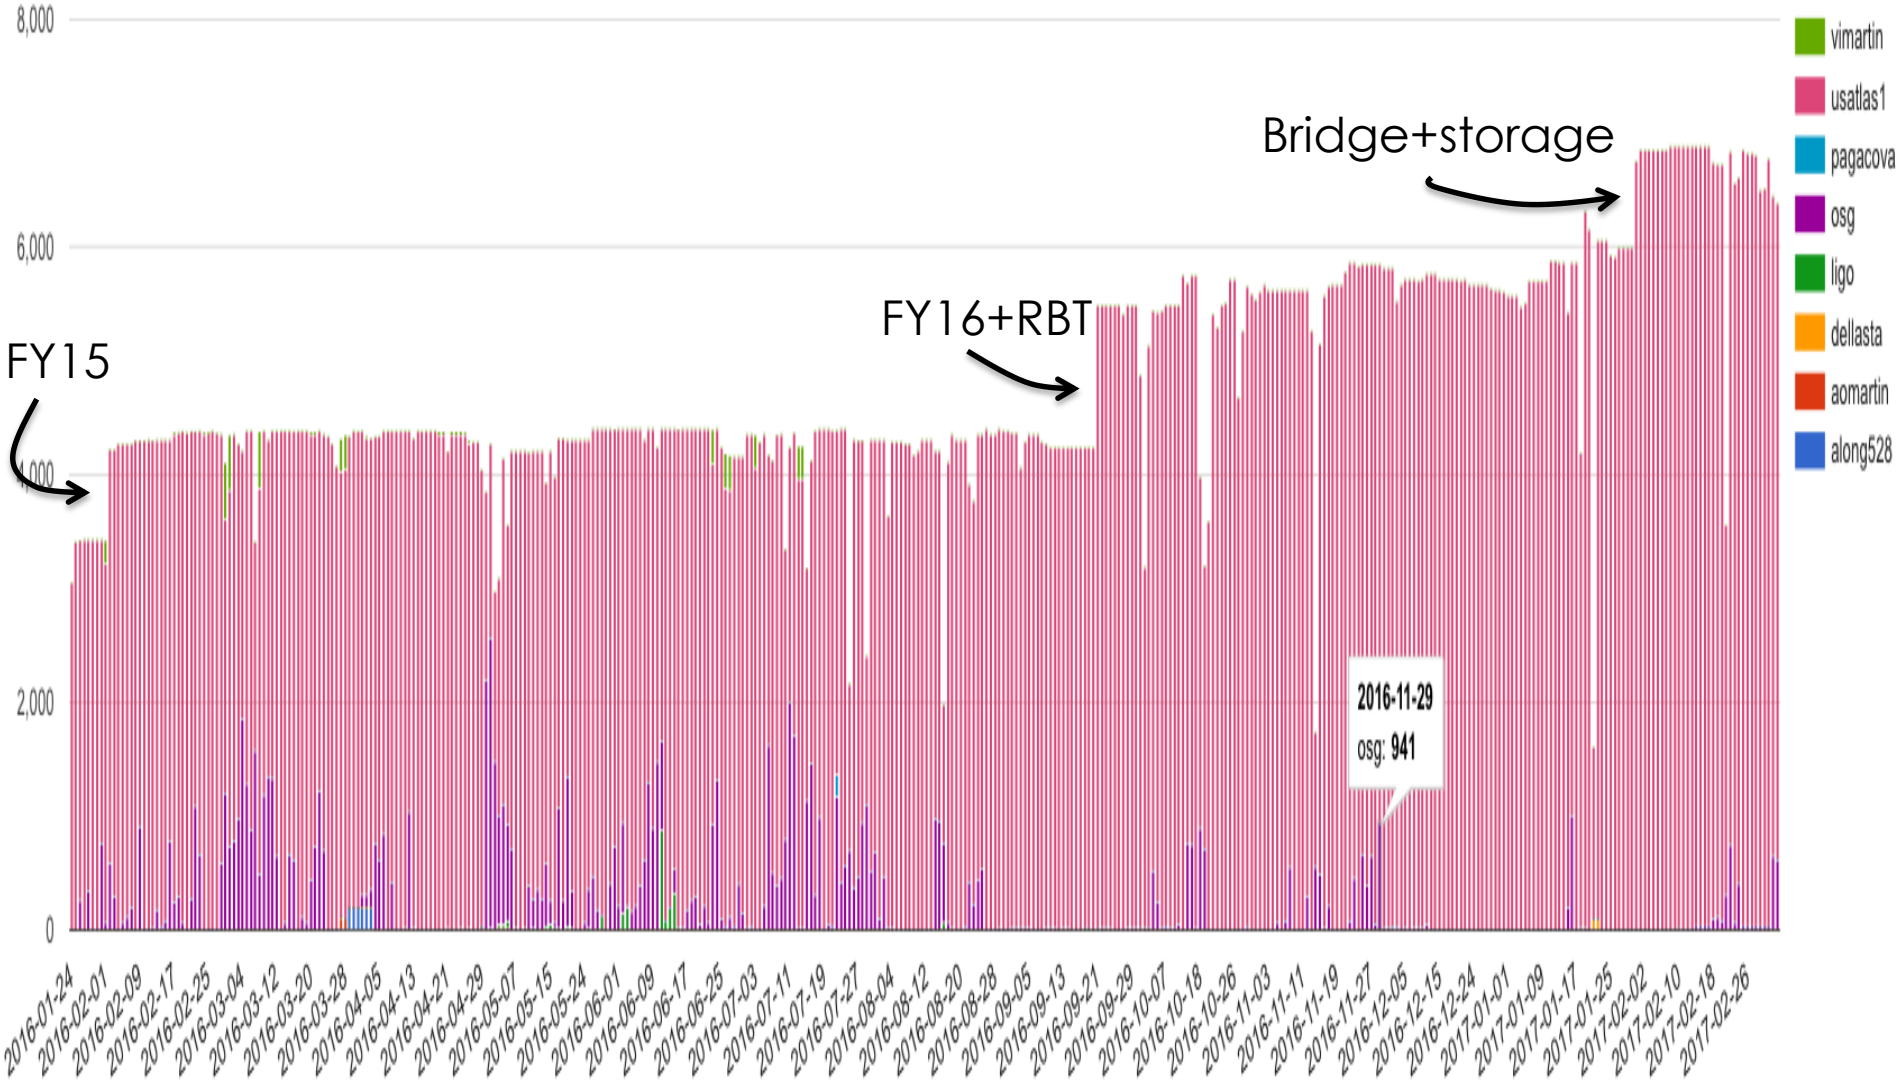

We let OSG fill in the gaps during ATLAS production lulls

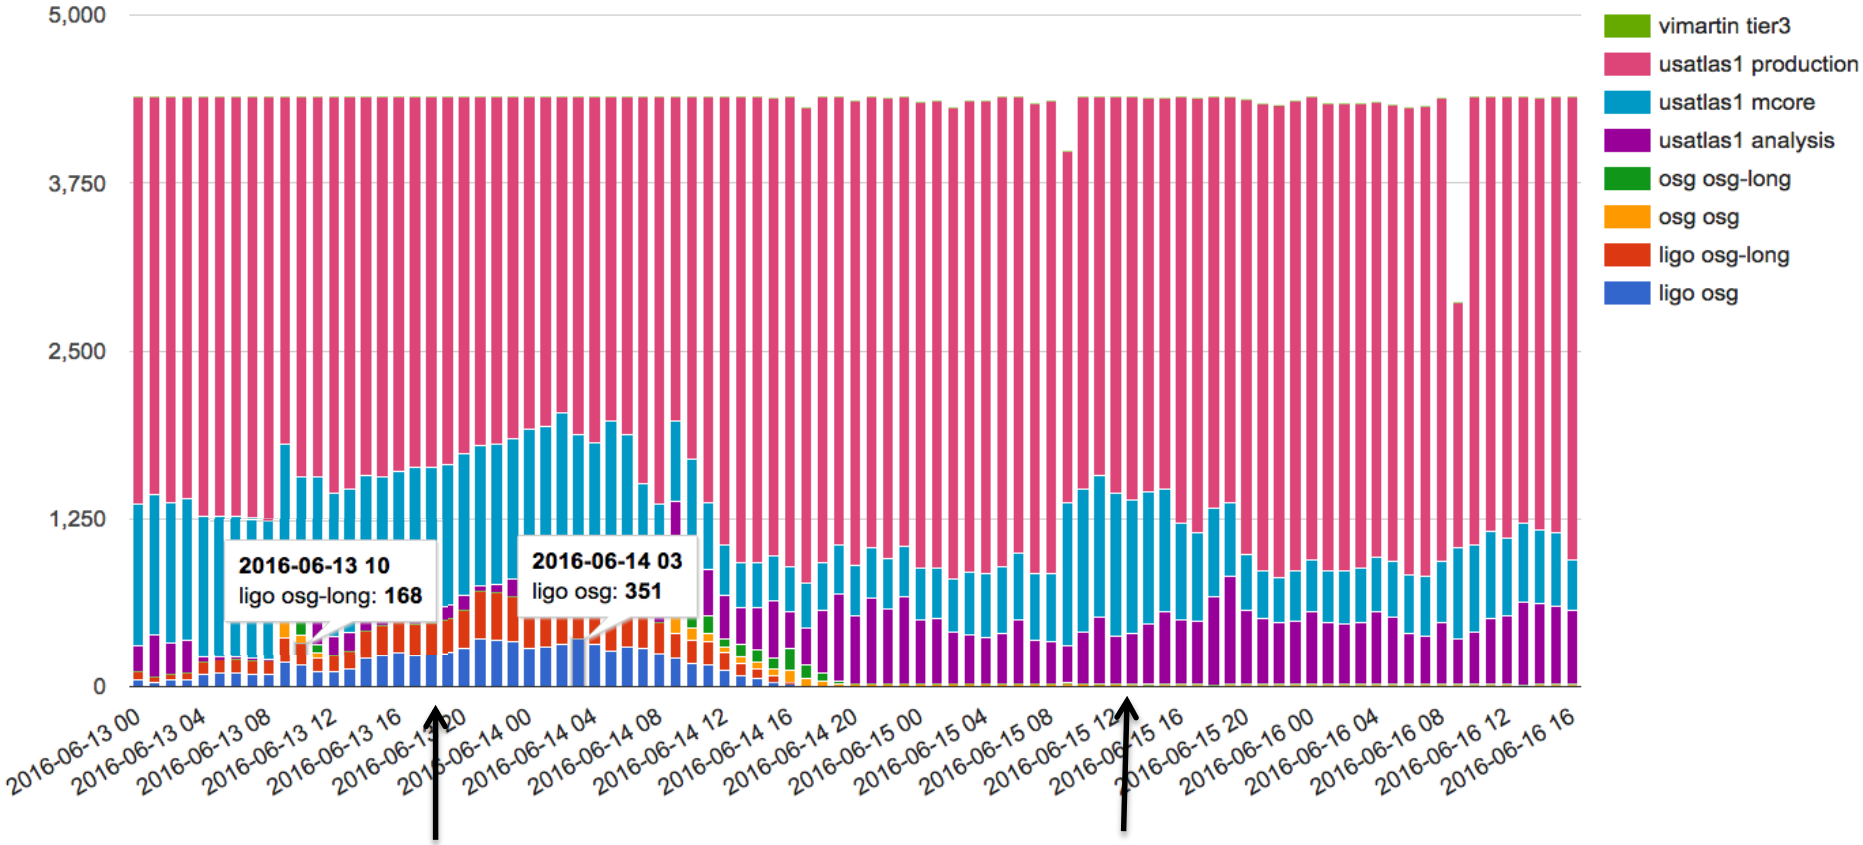

LIGO computing at one of the US ATLAS Tier2s peaks at about 700 cores on the day before the press conference

LIGO press conference announcing 2d black hole merger

All FY16, RBT, Bridge and Storage funds spent, installed and on-line

NET2									Comments
HU_ATLAS_Tier2	Resource = HU_ATLAS_Tier2								
2007	Quad core Intel E5520		2.27	1,536	9.84	15,114	HT	12	
2013	Intel(R) Xeon(R) CPU X5650		2.67	480	9.97	4,786	HT	24	
		Summary	Totals	1,536		19,900			
				HS06 average	9.87				
BU_ATLAS_Tier2o	Resource = BU_ATLAS_Tier2								
2012	Intel(R) Xeon(R) CPU X5650		2.67	1,248	9.97	12,443	HT	24	
2015	Intel(R) Xeon(R) CPU E5-2650 v3 @ 2.30GHz		2.3	1,000	10.44	10,440	HT	8	
2,012	Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz		2.6	1,280	10.66	13,644	HT	32	
2016	Intel(R) Xeon(R) CPU E5-2660 v3 @ 2.60GHz		2.6	960	11.8	11,328	HT	40	
2012	Quad core Intel E5620		2.4	32	19	608		8	
2016	Intel Xeon E5-2680 v4 2.4GHz		2.4	392	11.57	4,535	HT	56	
2016	Intel Xeon E5-2640 v4 2.4GHz		2.4	1,160	11.57	13,421	HT	40	FY16+RBT
2017	Intel Xeon E5-2640 v4 2.4GHz		2.4	920	11.57	10,644	HT	40	Bridge funding
		Summary	Totals	6,992		77,064			
				HS06 average	11.02				
NET2 TOTAL		Summary		9,008		96,964			
				Total HS06 average	10.76				

9000 cores, 97k hs06, 5.5 PB

2017 Pledged capacities at the US ATLAS Facilities (updated from October 2016 C-RSG Report		
	CPU [HS06]	DISK [TB]
Site	2017 Pledge	2017 Pledge
Tier-1	211830	15640
AGLT2	57500	4242
MWT2	86250	6363
NET2	57500	4242
SWT2	57500	4242
WT2	0	0
Total	470580	34730
T2 Total	258750	19090

GPFS Main Storage

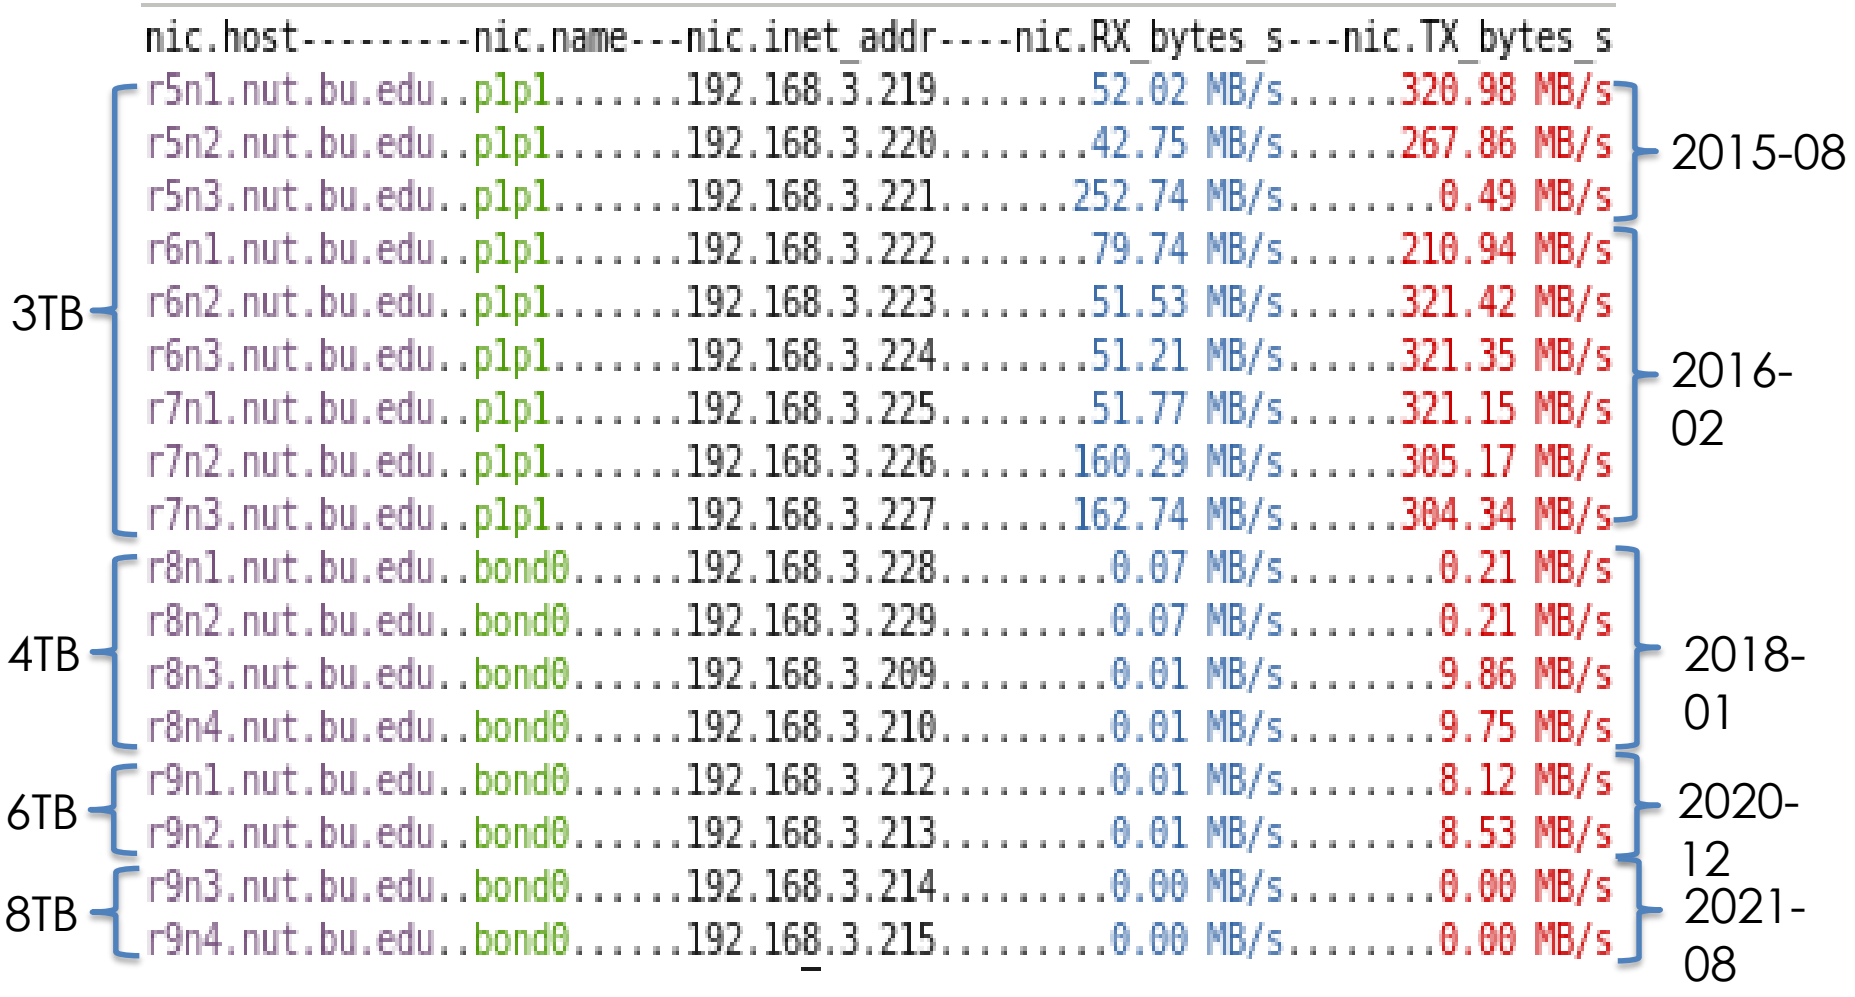

290 Gb/s Total Bandwidth
6.0 PB total useable storage
0.5 PB about to retire, 1.0 PB additional in 0.5-1.0 years

LSM

We have a custom python code written to Marco Mambelli's original specifications (still in effect, we understand). It

1. Lets us control cp, scp, direct, <other> for different nodes (BU, Harvard, MOC)
2. Allows us to control where the LSM traffic goes, e.g. 20G dedicated link from BU to HU.
3. Allows us to control burstiness of read/writes to GPFS, which can cause problems otherwise.
4. Allows us to efficiently compute Adler checksums on GPFS servers, storing in extended attributes.

Plan to add S3 and/or librados I/O to NESE.

SRM Migration

- Plan to migrate main storage to NESE
- Most likely use Object Store (prob. S3)
- FTS transfers either via Gridftp or FTS S3 endpoint directly
- Convert LSM to use S3
- Xrootd TBD with Wei

NET2 Current Networking

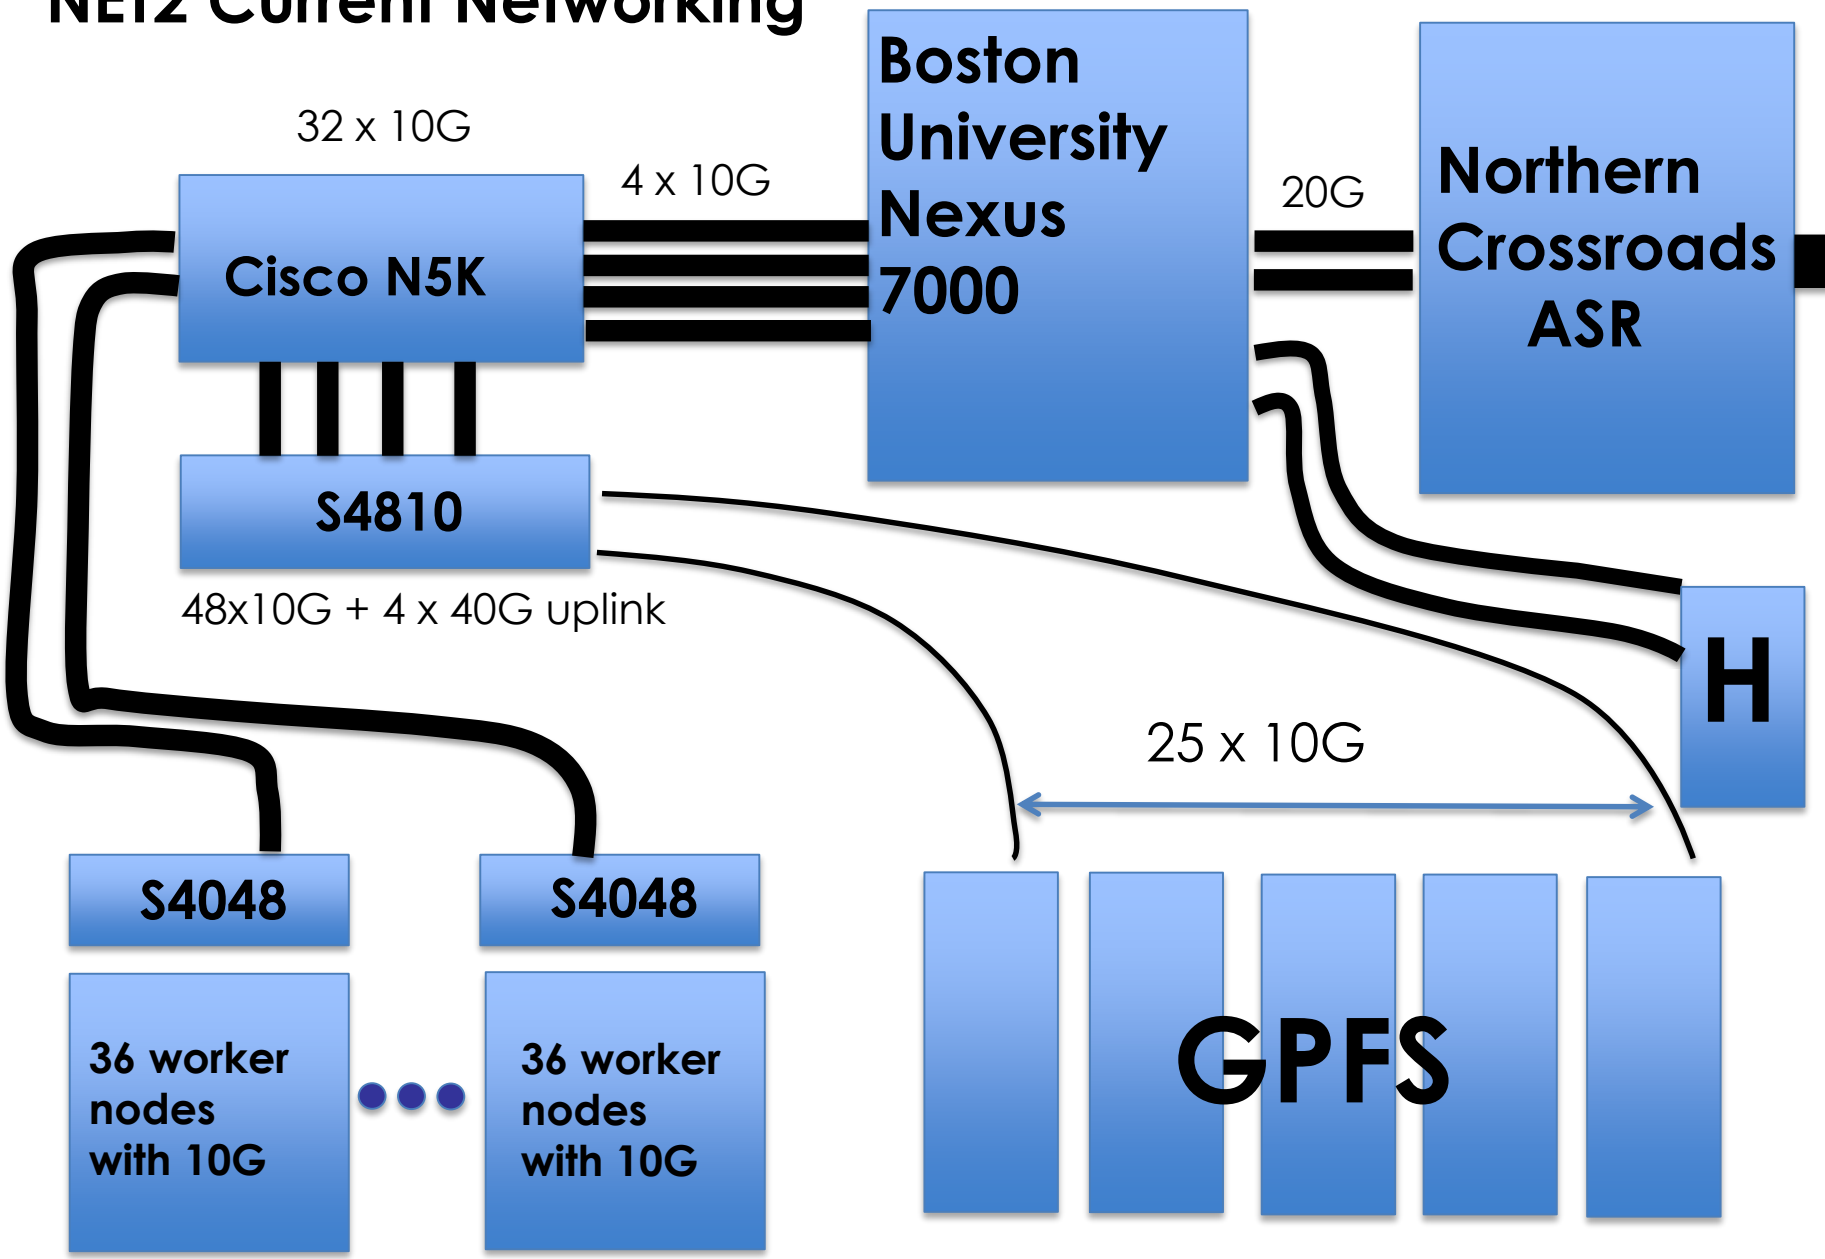

Incoming Gridftp Transfers Resume at 20G following switch to new Cisco Equipment at NoX

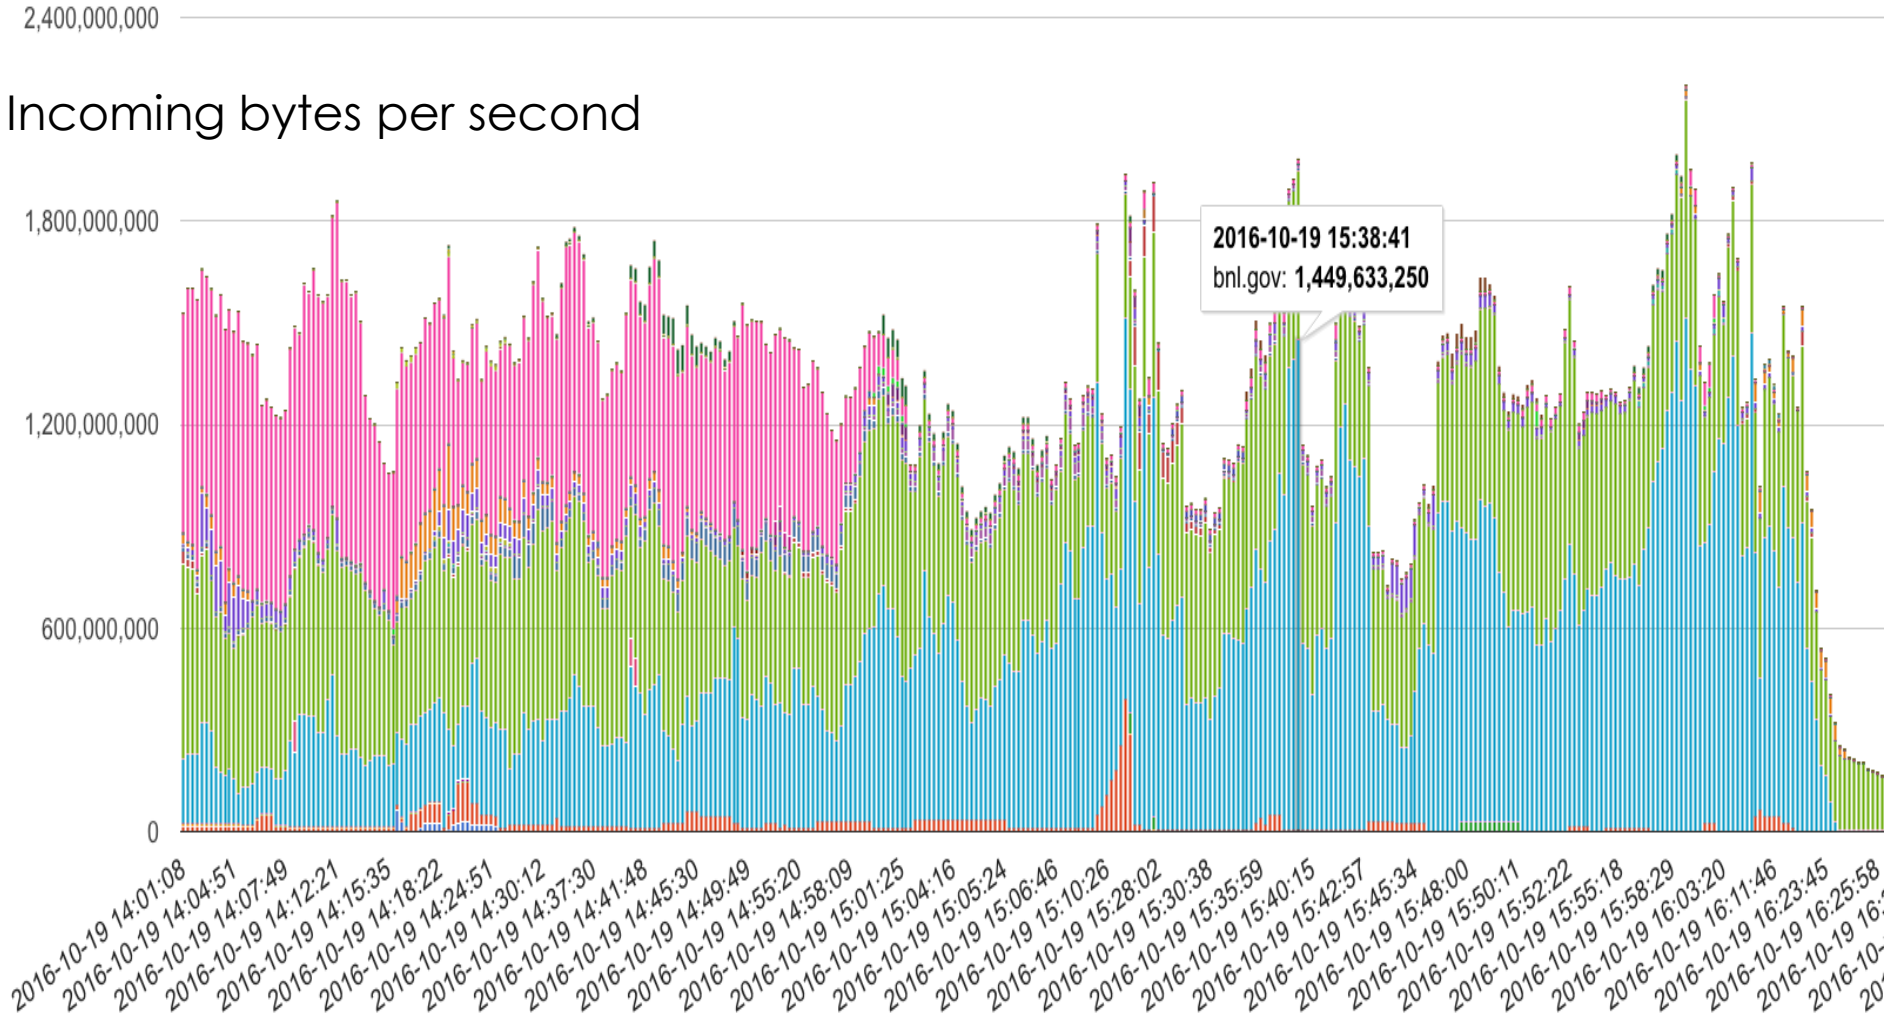

Joined LHCONE
Haven't seen >12 Gb/s since…needs testing
No saturation problems so far

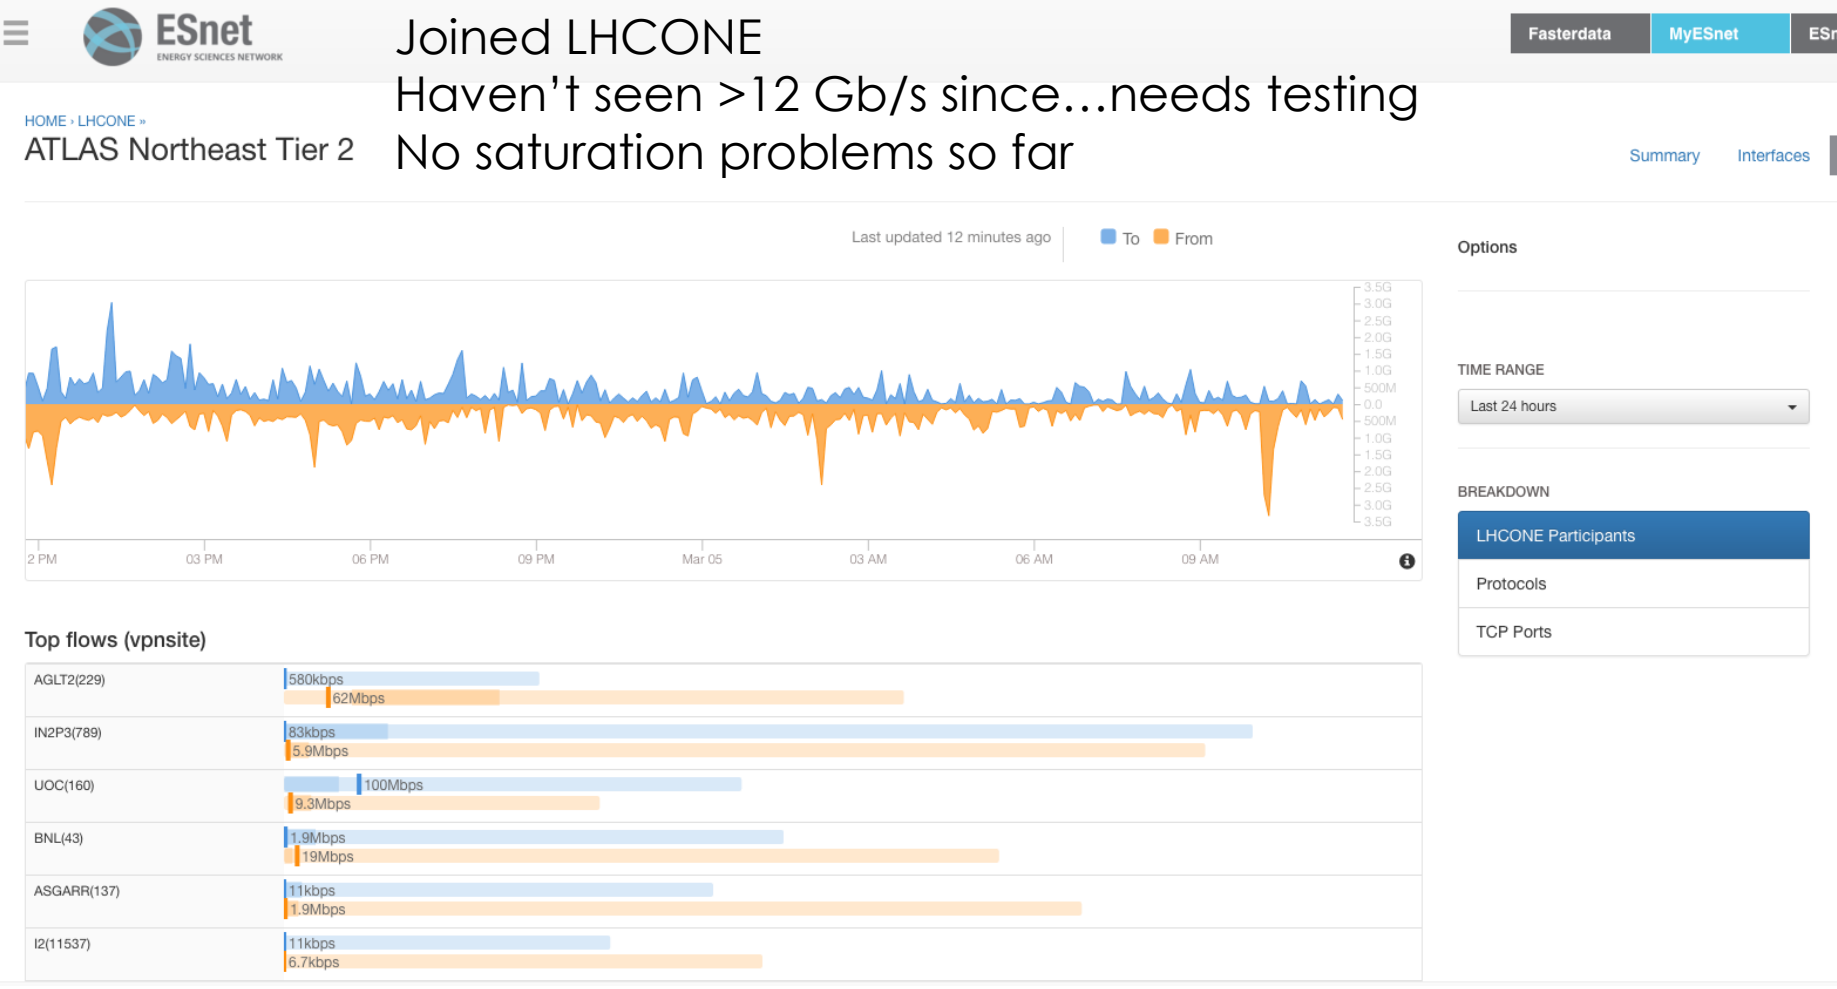

Networking Plan:

- Take maximum advantage of NESE project
- Network via NESE via Harvard 100Gb/s to NoX/internet2 and LHCONE
- Upgrade to multi x 40Gb/s WAN avoiding Cisco gear
- Initially need ~4 x 40G from NET2 to NESE fabric
- NESE networking being planned for IPv6. This is connected with re-thinking networking on the MGHPCC floor.

```
srm.spacetoken--------srm.totalSize---srm.usedSize---srm.unusedSize
ATLASDATADISK.............3328.00 TB.....3082.29 TB........245.71 TB
ATLASGROUPDISK............140.00 TB.......87.28 TB..........52.72 TB
ATLASHOTDISK...............0.00 TB........0.00 TB...........0.00 TB
ATLASLOCALGROUPDISK.......450.00 TB......138.59 TB........311.41 TB
ATLASPRODDISK..............2.00 TB........0.00 TB...........2.00 TB
ATLASSCRATCHDISK..........110.00 TB.......27.49 TB..........82.51 TB
ATLASUNALLOCATED..........270.00 TB........0.00 TB........270.00 TB
ATLASUNPOWERED.............0.00 TB........0.00 TB...........0.00 TB
ATLASUSERDISK.............200.00 TB.......75.77 TB........124.23 TB
@,home,table,table
```

LOCALGROUPDISK

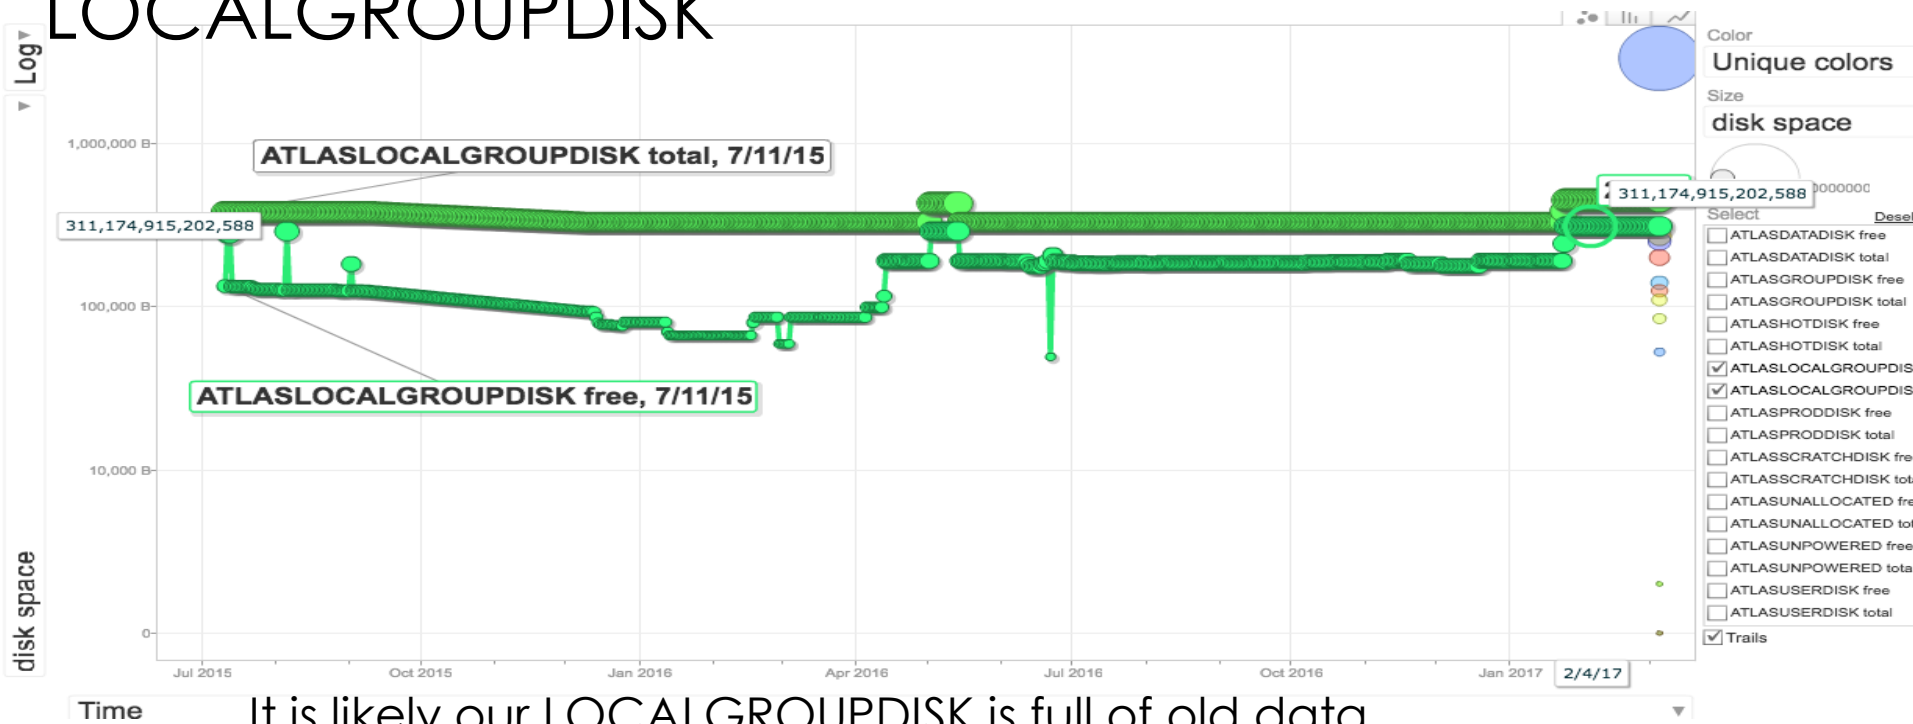

It is likely our LOCALGROUPDISK is full of old data that nobody cares about anymore.

HTCONDOR-CE Status: Both BU and HU converted

HU AuditLog History: 2017-03-05 22:20:39 UTC

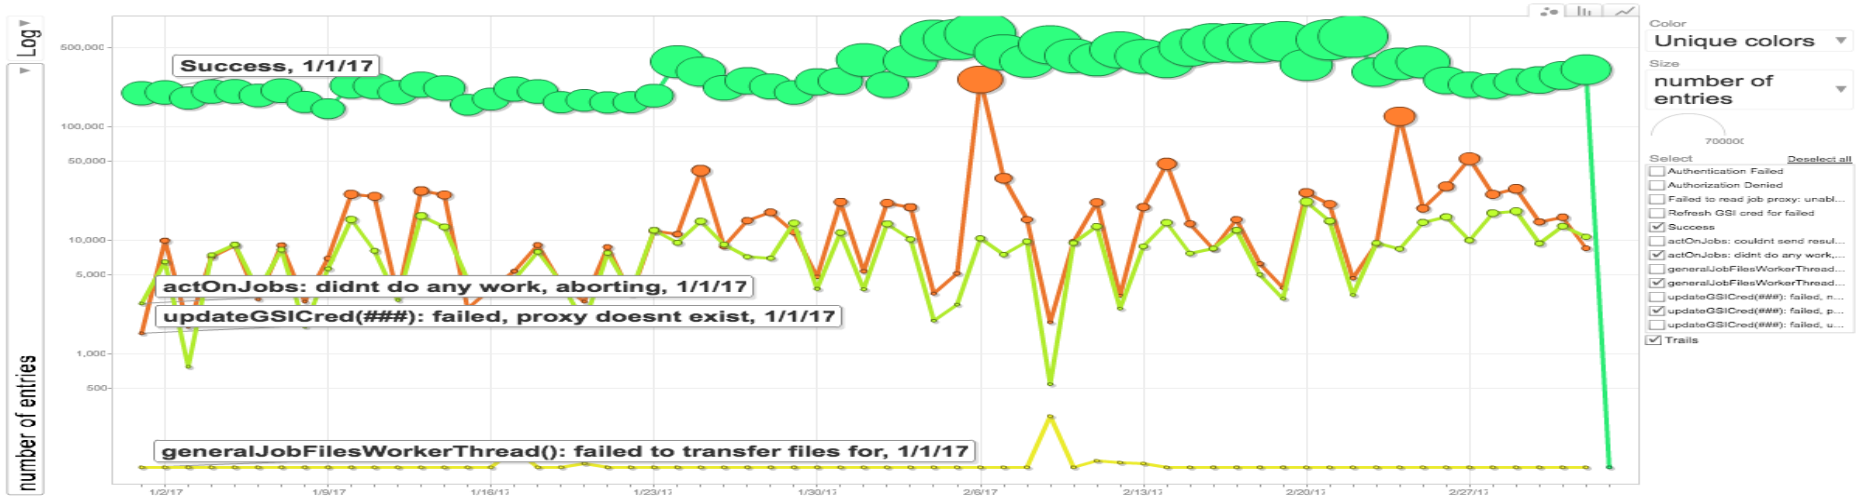

BU AuditLog History: 2017-03-05 22:15:55 UTC

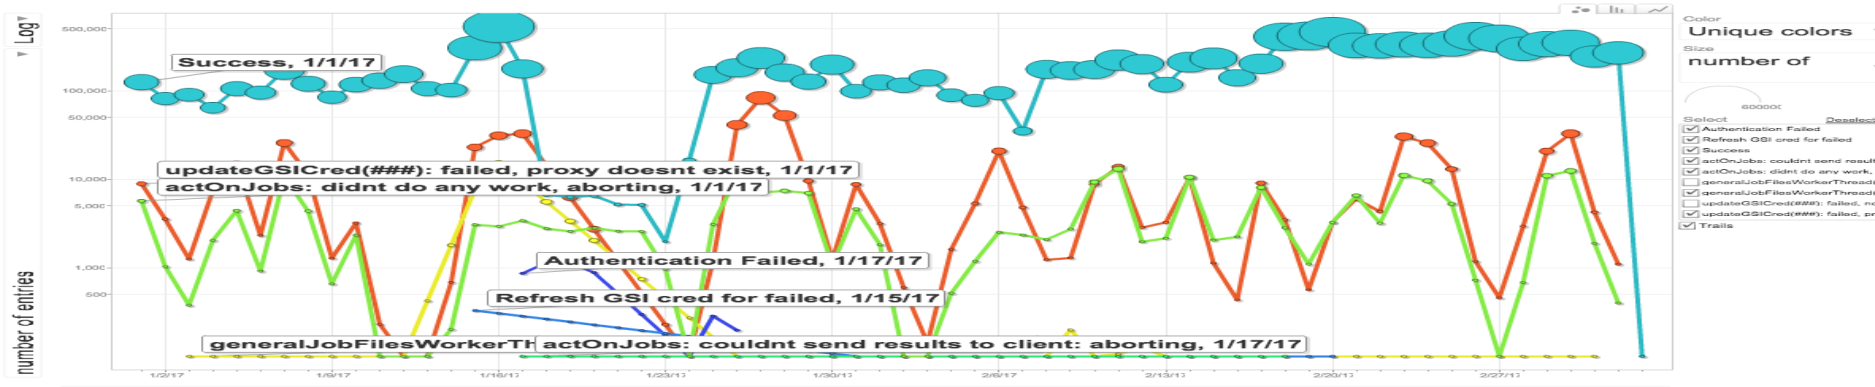

Plan for SL7 Migration

1. Wait for someone else to try first.
2. Test that GPFS still works.
3. Migrate as before coordinating with other T2s.
4. We have done this before going from RH3->4->5->6. We don't anticipate problems.

Current Worker Nodes

linux.model	linux.vendor	nc	hs06.hs06
Intel(R) Xeon(R) CPU E5-2640 v4 @ 2.40GHz	GenuineIntel	960	11049.6
Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz	GenuineIntel	1280	13452.8
Intel(R) Xeon(R) CPU E5-2650 v3 @ 2.30GHz	GenuineIntel	1000	10450.0
Intel(R) Xeon(R) CPU E5-2660 v3 @ 2.60GHz	GenuineIntel	1060	12508.0
Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	GenuineIntel	168	1933.7
Intel(R) Xeon(R) CPU E5520 @ 2.27GHz	GenuineIntel	1904	18069.0
Intel(R) Xeon(R) CPU E5620 @ 2.40GHz	GenuineIntel	24	240.0
Intel(R) Xeon(R) CPU X5650 @ 2.67GHz	GenuineIntel	1680	14599.2

307 nodes
8076 cores
82 k hs06

Same for NET2/Harvard

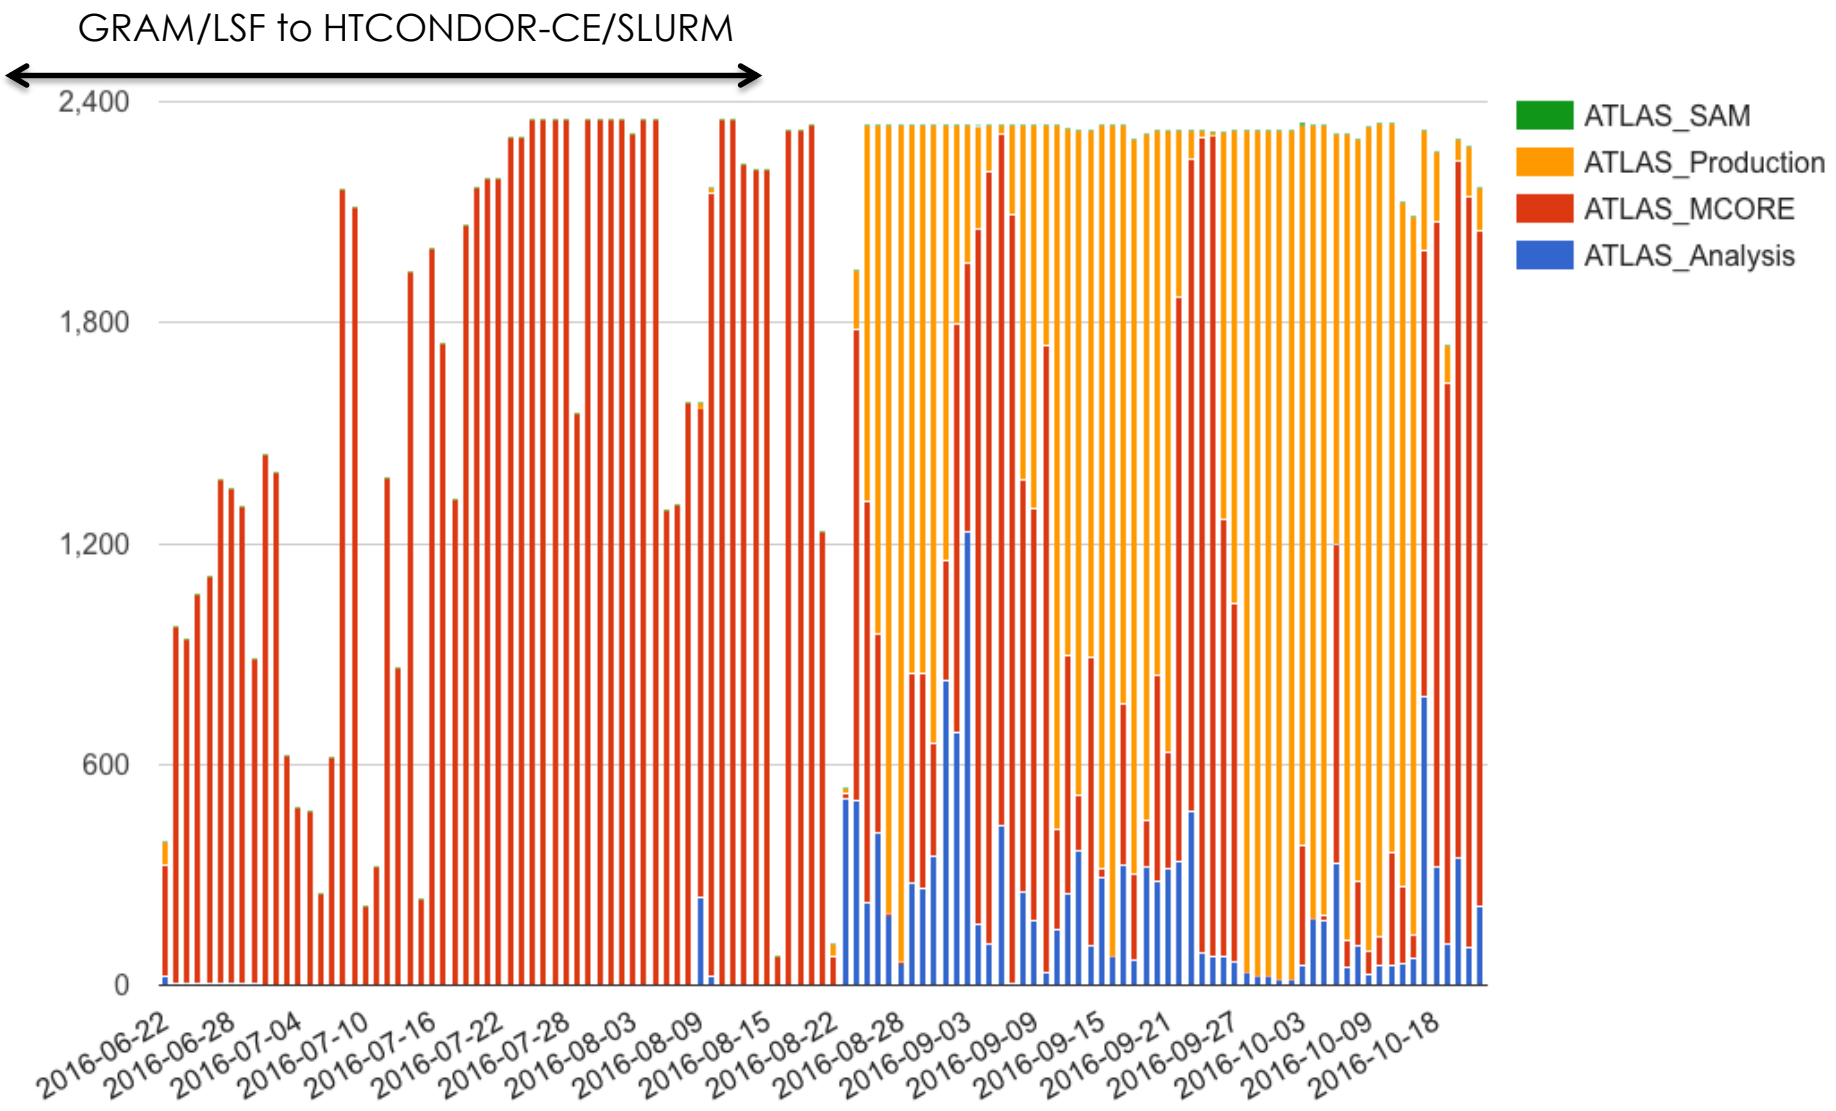

One Entire Year of GGUS Tickets!

Last Update	Subject	
2016-10-17	Transfers to NET2 DATADISK are failing w...	Aborted attempt to move WAN to new NoX routers
2016-10-20	transfer errors for transfers from CA-MC...	McGill networking problem in Canadian WAN.
2016-09-09	Open webdav frontend	Request to install webdav. – not a site issue.
2016-07-01	CA-MCGILL-CLUMEQ <-> NET2 data tra...	McGill networking problem in Canadian WAN.
2016-06-10	US-NET2: DESTINATION OVERWRITE srm-ifce ...	McGill networking problem in Canadian WAN.
2016-04-11	Deletion errors at US NET2_DATADISK and ...	Burst of SRM attempts to delete non-existing files caused Bestman to clog up.
2016-04-18	NET2 high failure rate for data transfer...	Similar GPFS migration problem. Fixed in 1 hour.
2016-03-07	NET2_DATADISK: two stuck files and ~600 ...	GPFS migration slowdown and DDM errors. Fixed in 5 hours.
2016-03-01	US NET2 : DATADISK, LOCALGROUPDISK, USER...	SRM problem caused deletion errors. Fixed in 6 hours.
2016-02-09	Issues with storage dumps for consistenc...	Trivial problem with requested storage dumps fixed in 10 minutes.
2015-12-23	BU_ATLAS_Tier2_MCORE: Production tasks f...	One task has low efficiency due to WAN saturation due to NET2 being used as a "nucleus" site.
2015-12-28	NET2 'Stager did not answer within the r...	SRM problem. Fixed in 10 minutes.
2015-12-21	US NET2: serveral transfer staging and d...	GPFS experiment during low PanDA production gone wrong. 7 hours to recover.
2015-12-23	ATLAS request- storage consistency check...	Request for storage dumps – not a site problem.
2015-10-26	NET2_PRODDISK deletion errors	SRM deletion problem, fixed in ~1 hour