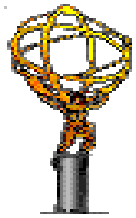

BROOKHAVEN NATIONAL LABORATORY

BNL Site Report

ATLAS Sites Jamboree, CERN

18 - 20 January, 2017

Xin Zhao, Hironori Ito
Brookhaven National Laboratory, USA

Outline

- General Status
- Object Store and its Integration with dCache (Hiro)
- Tier1 Network Migration
- Increased Bandwidth for LHCONE/LHCOPN
- Production Jobs on Local (Opportunistic) Tier3 queue
- Staging Test directly from HPSS tapes
- Running ATLAS jobs on OSG Opportunistic Sites

General Status

- BNL CSI (Computational Science Initiative)
  - Centralizing scientific computing at BNL
  - RACF/Tier1 => SDCC (Scientific Data and Computing Center)
- ATLAS Computing Facility (Tier1)
  - Running fine overall
  - Pledges fulfilled for 2016
    - CPU capability
      - Added 100 Dell R430 systems, total ~ 18k condor job slots
    - dCache:
      - disk storage 14PB (pledge 11.5PB)
    - HPSS:
      - ~22PB of ATLAS data on tapes (pledge 27PB)
  - Now move to some highlighted topics …

Object Store and Integration with dCache

- We have been running two independent **TEST** instances of Ceph storage by using the retired dCache storages and other retired servers.
  - One of them, which have previously used as the main S3 storage for event service, has been retired completely to make the space for newly retired storage.
    - At the time of retirement, there exist more than 30M objects (data) in the S3 storage. It seems that (a) the deletion service is not working and/or (b) the data is not by any means temporary and/or (c) the data is not being used.
  - The other one, previously called a test instance, is now the primary and only S3 service currently. But, due to various re-organization of backend storage servers, it is currently at less than ½ capacity and performance.
  - The new instance of Ceph storage is currently being installed.
- S3 storage is not the only use of Ceph storage at BNL.
  - Ceph librados: dCache storage pool, Gridftp, XROOTd
  - CephFS is being used for the cloud storage, Gridftp, XROOTd

S3 Performace via Event Services

- Doug, Taylor and Wen has been the primary persons who are closely looking at the behavior of S3 storage through the event service jobs at HPC sites (NERSC and ANL) as well as opportunistic event service jobs in the grid.
- Doug and Taylor has reported the performance issues seen in Event service jobs at HPC sites. They also has done a simulated tests to evaluate the observed S3 client behavior.
  - https://drive.google.com/open?id=1U00Zl11gjd_0yoYWFnkZ3Aj3ElHLSjLRgMR7I1H2SMw
  - https://drive.google.com/open?id=1qta2BQst1dsLb0wzgXvqT7O2NMQQVKzgtKWfy9DM6f8
  - https://drive.google.com/open?id=1ZuUK0rnrlTYgiT7Z0_15IBVeCHHzwqkd2X7vc0kPIZQ
- ADC has asked BNL to do more tests to study the S3 performance systematically by looking at the following parameters.
  - The number of clients, RTT, data size, the number of buckets, the deletion rate
  - The study is currently under way and still on-going.

Tier1 Network Migration

- Move Tier1 (ACF) network out of BNL Campus Network

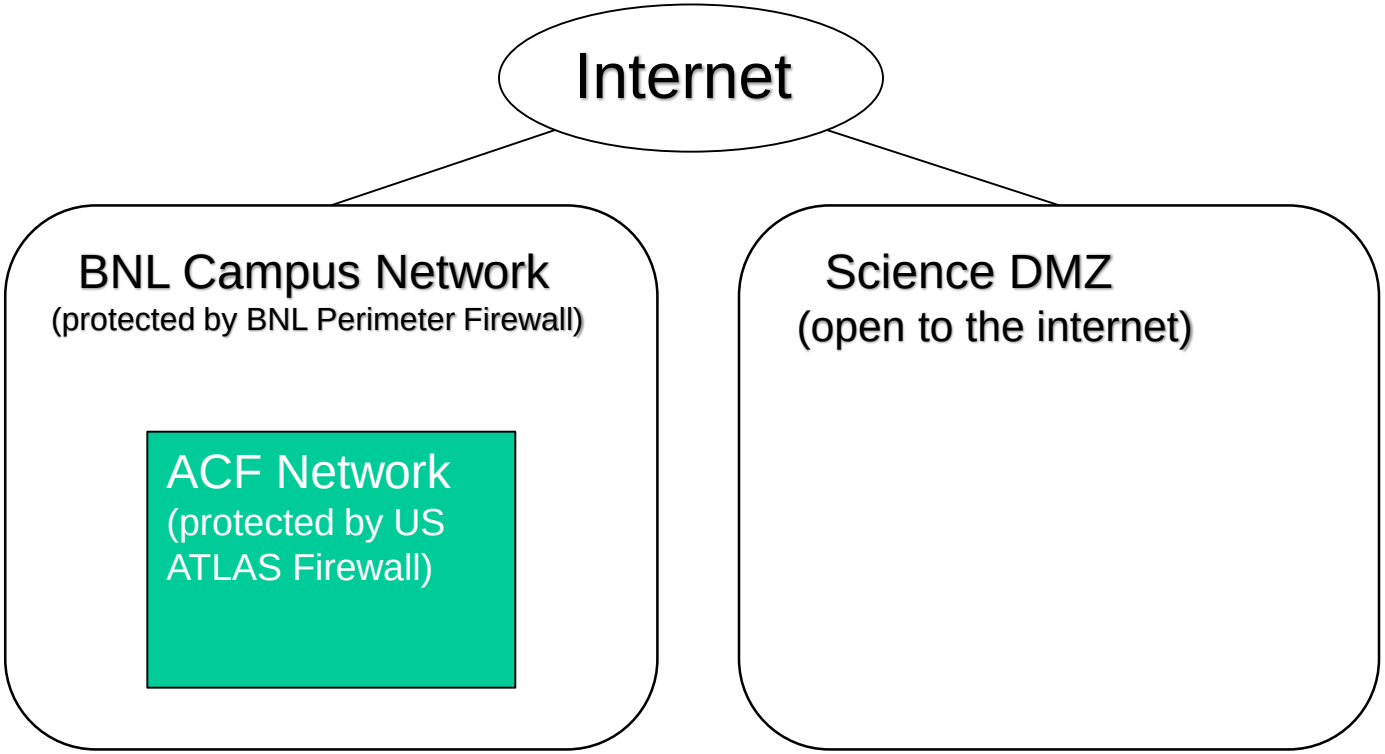

Tier1 Network Migration

- Why the move?
  - Allow IPv6 rollout to ATLAS Tier1 facility
  - Separate high bandwidth firewalled scientific internet traffic from BNL campus general purpose internet traffic
  - Isolate Tier1 from the BNL campus network, to allow the Tier1 facility to benefit from cybersecurity rules that govern “scientific” traffic.
- Schedule
  - January 30th is the cut-over day, downtime may be scheduled
- Effect to users, after the migration
  - Transparent to users/jobs coming from outside of BNL
  - Minor changes in ways of accessing interactive nodes for local ATLAS (Tier3) users

Increased Bandwidth for LHCONE/LHCOPN

- BNL network (perimeter/science DMZ) topology
  - One primary (100G) for ALL traffic, one backup (100G)

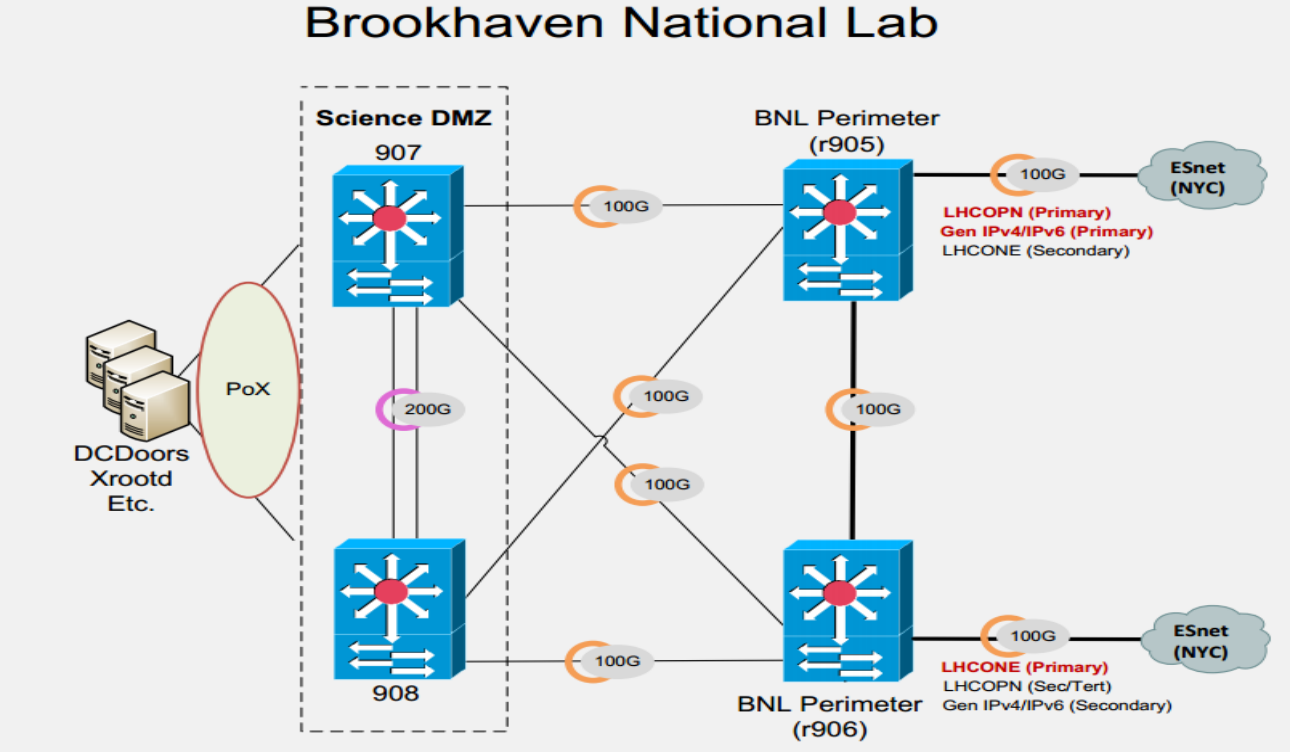

Increased Bandwidth for LHCONE/LHCOPN

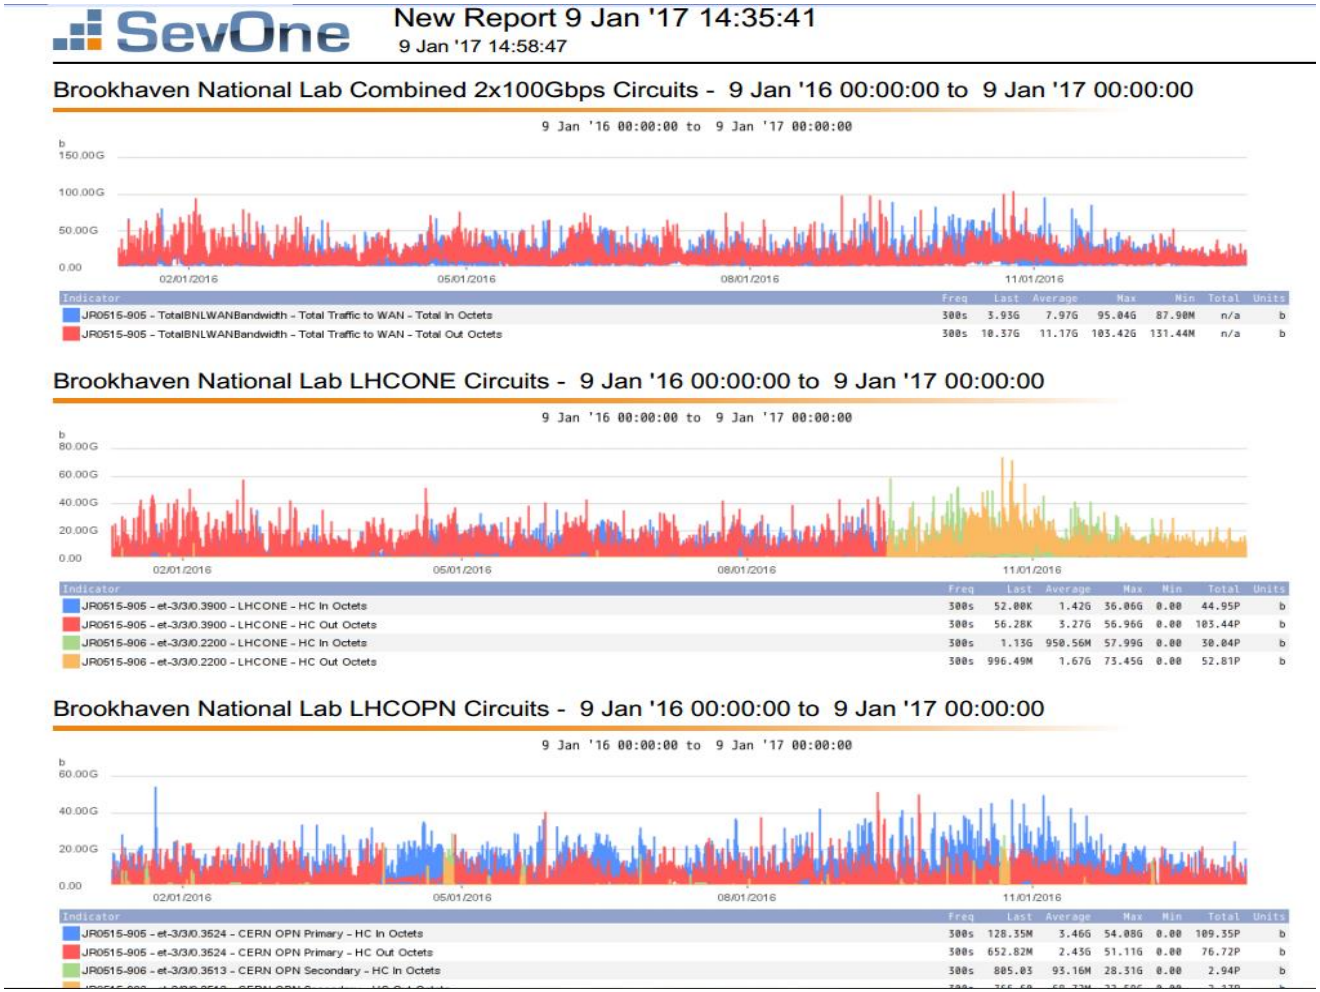

- **LHCONE+LHCOPN saturated the primary circuit at 100Gbps, for the first time, around the end of August, 2016**
- The old backup 100G connection is now the primary circuit for LHCONE
- The two circuits back up each other.
- **Effectively, total bandwidth to ESnet is doubled**

Production Jobs on Local (Opportunistic) Tier3 Queue

- BNL local Tier3 resources
  - ~2k job slots, local users jobs preempt others
  - single core jobs only
    - Preemption doesn't work well with Partitionable slots in condor (in contact with HTCondor developers)
- Backfill of Production Jobs
  - PanDA queue: BNL_LOCAL
  - High failure rate (>50%) for regular single core jobs, due to preemption
  - Single core ES jobs are ideal
    - Implemented and tested by ADC successfully
    - Need more ! --- many idle CPU-hours

Production Jobs on Local (Opportunistic) Tier3 Queue

Spotty Backfill

Tier-3 Backfill Idle/Running

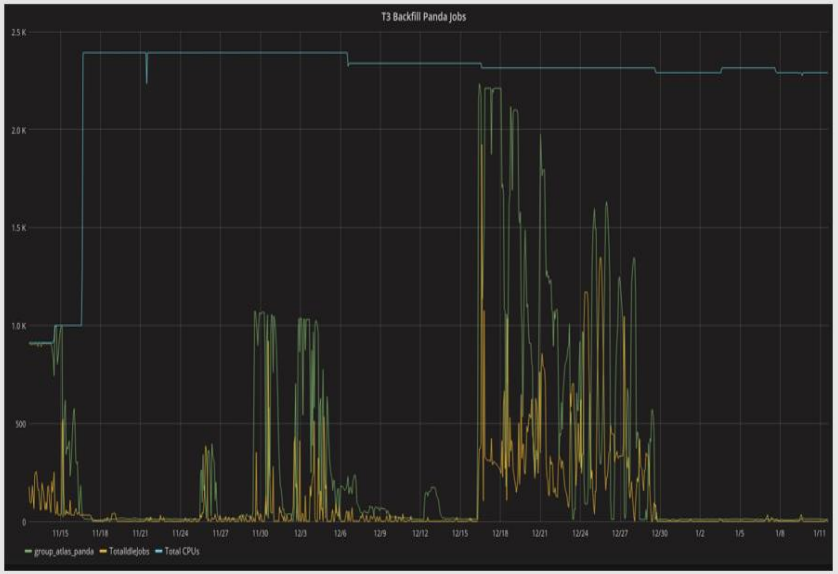

T3 Occupancy

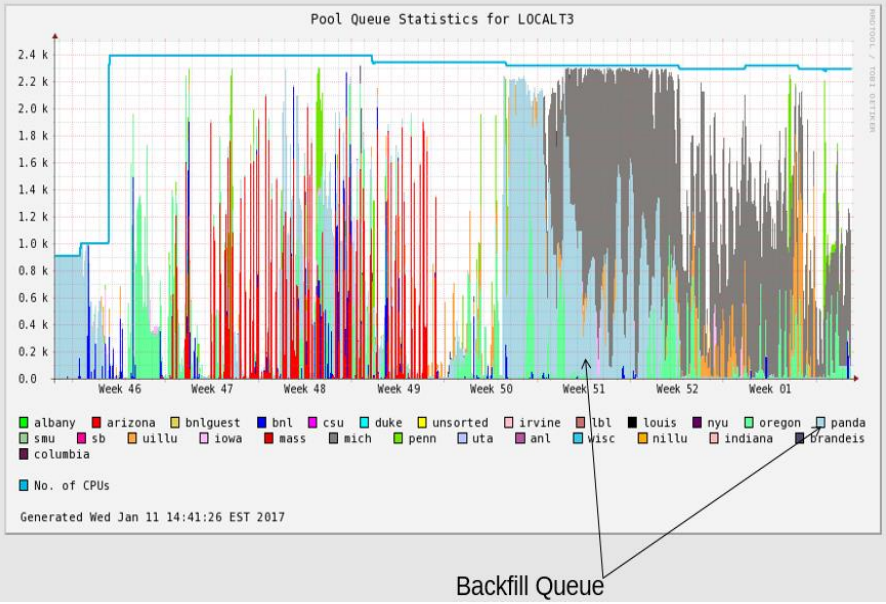

Staging Test from Tapes

- Staging Test (Jan 10-14):
  - replicate 150TB AODs from DATATAPE to DATADISK
  - ~1500 new reqs added, per hour
  - Transfer rate : not constant, average at 385MB/s

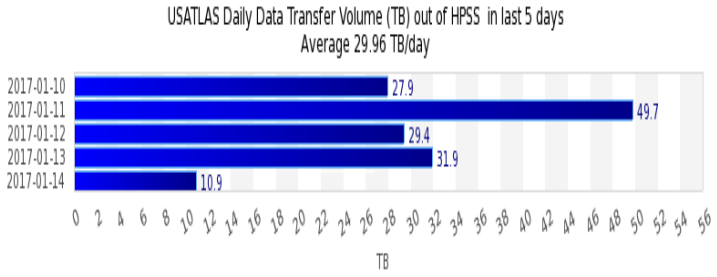

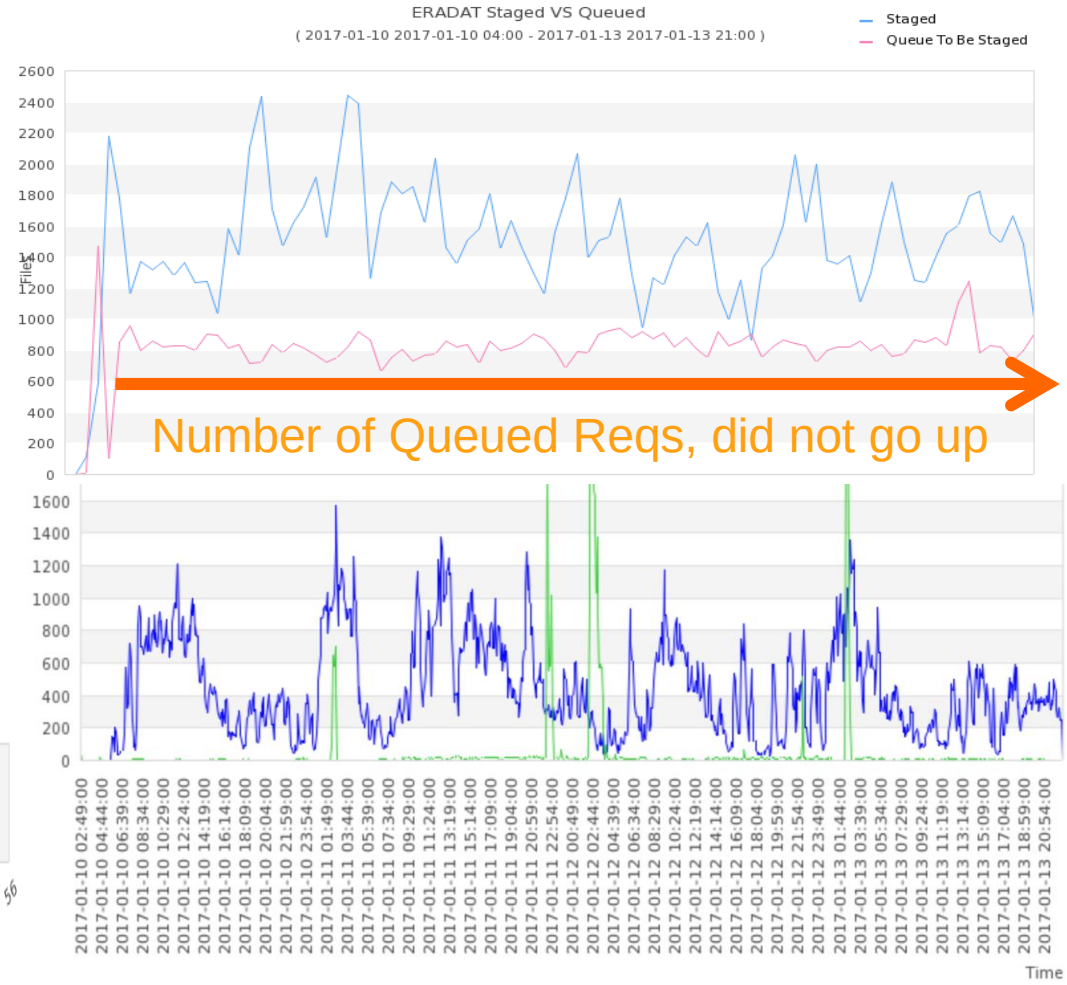

Staging Test from Tapes

- Improvements ?
  1. Increase File size
  2. Bulk request

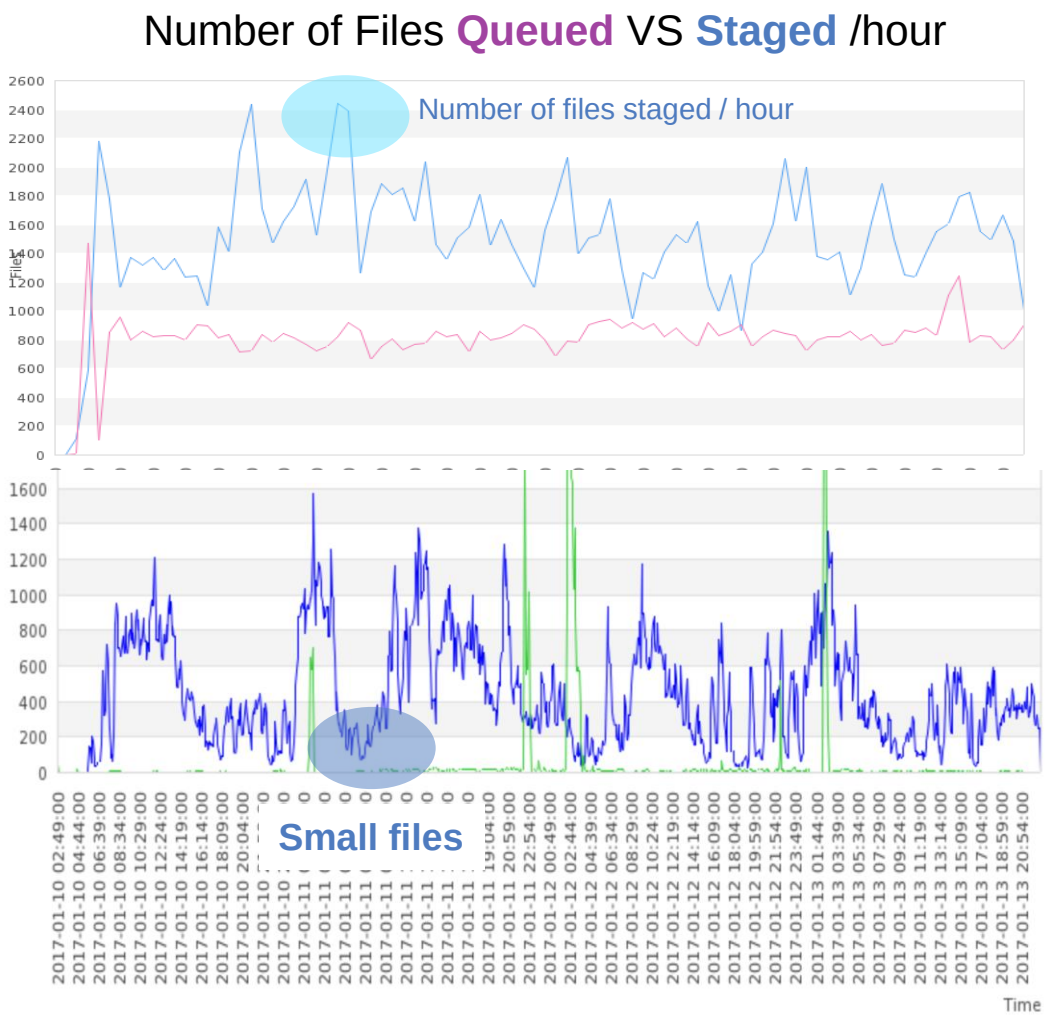

Staging Test from Tapes

- Improvements ?
  1. Increase File size
  2. Bulk request:
     BNL tape system optimizer reduces tape remounts

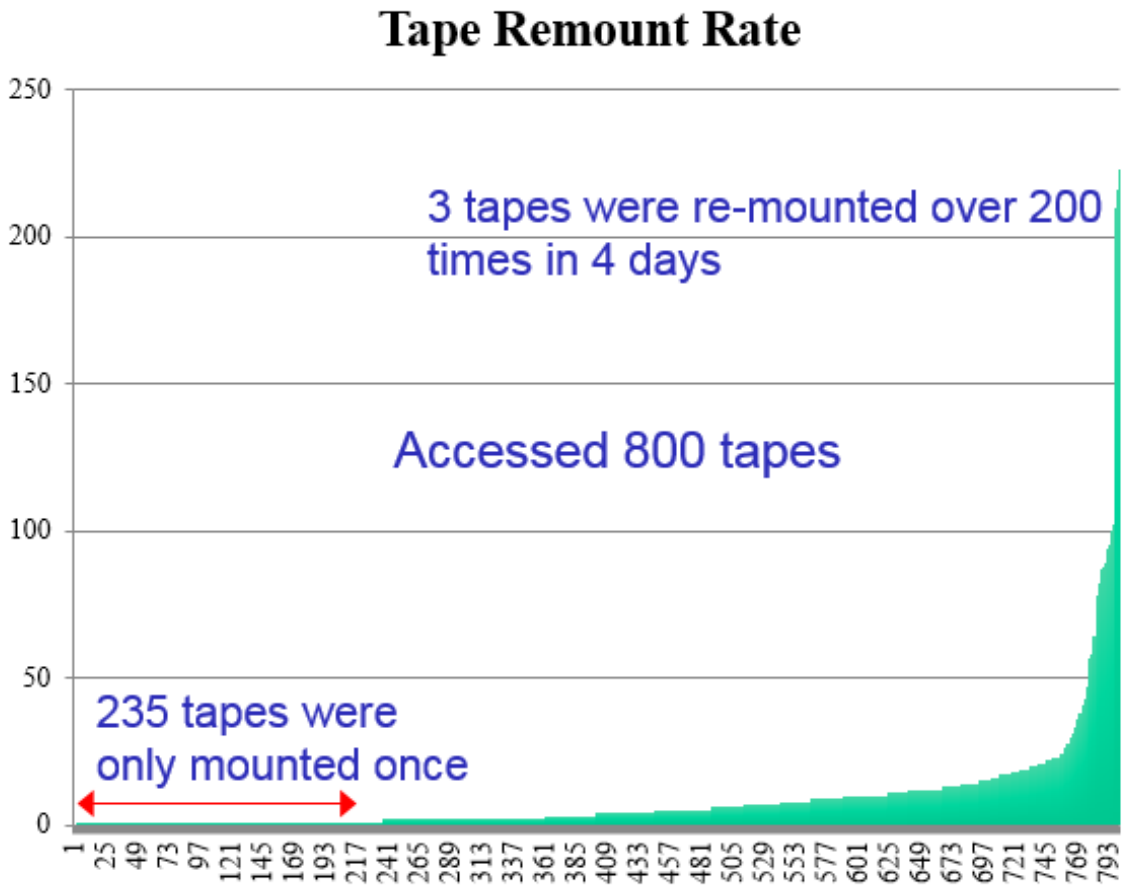

Staging Test from Tapes

Use Case Study of Tape System Performance: BNL STAR Experiment

Sep 2016, STAR submitted ~245,000 files

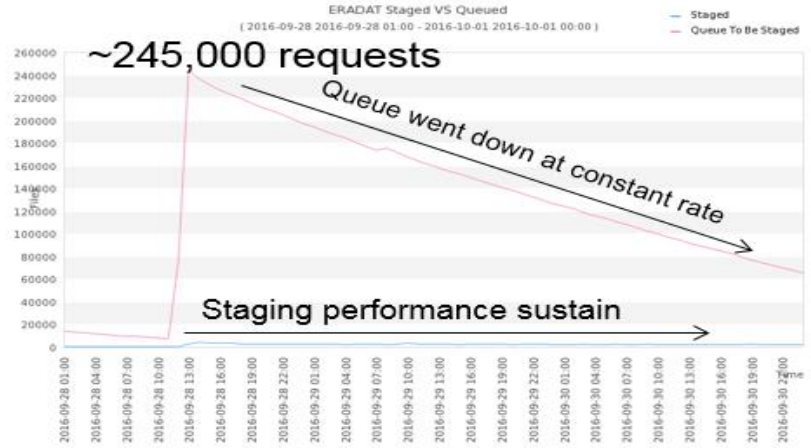

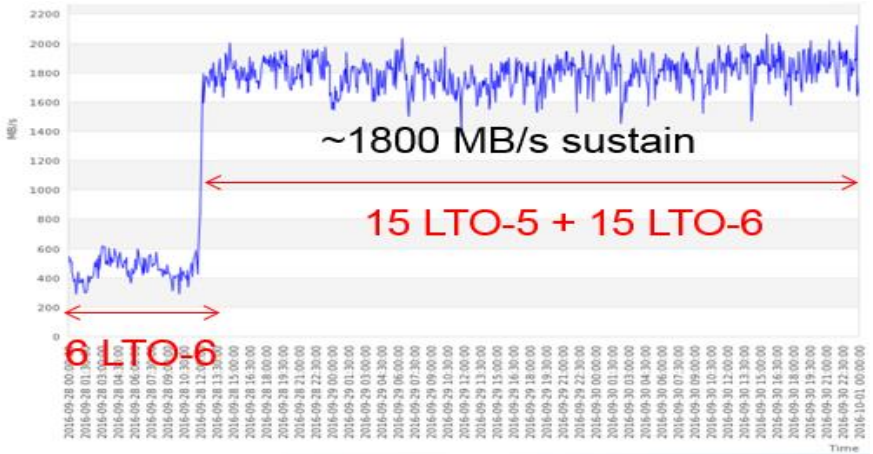

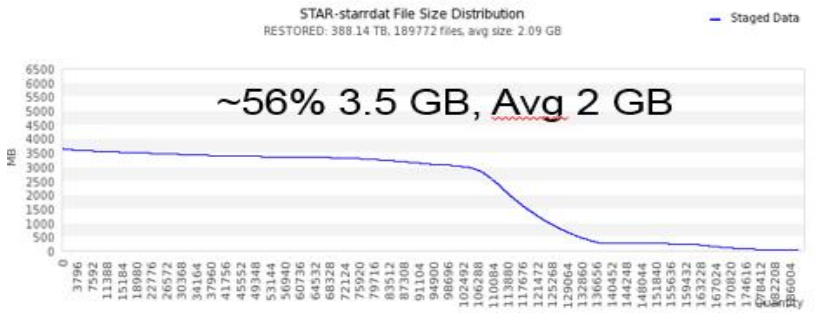

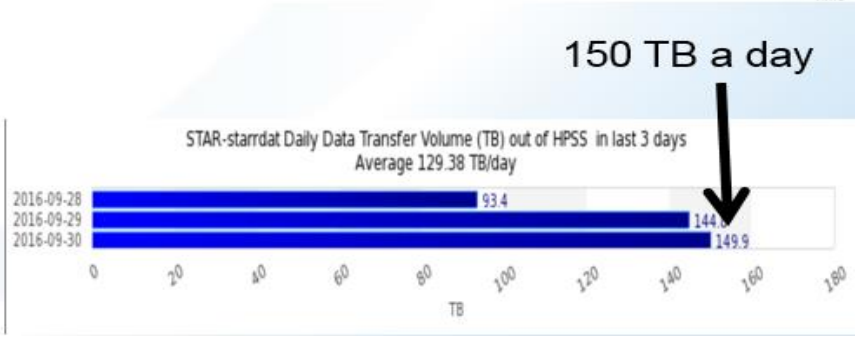

ATLAS Jobs on OSG Opportunistic Sites

- Each OSG Opportunistic site has its own PanDA queue and AGIS entries
  - Initially one PanDA queue for all OSG opportunistic sites : proved to be too course-grained and difficult to troubleshoot
  - Required working out AGIS hierarchy for opportunistic sites in order to separate pledged vs. non-pledged resources for accounting (thanks to Alexey Anisenkov)
  - Most sites are CMS-owned, some non-LHC sites as well, so we get slots during CMS lulls.

ATLAS_OPP_OSG-CIT_CMS_T2	online
ATLAS_OPP_OSG-CLEMSON_PALMETTO	online
ATLAS_OPP_OSG-FNAL_GPCE01	online
ATLAS_OPP_OSG-FNAL_GPCE02	online
ATLAS_OPP_OSG-GLOW	online
ATLAS_OPP_OSG-HYAK	online
ATLAS_OPP_OSG-MIT_CMS	test
ATLAS_OPP_OSG-NYSGRID_CORNELL_NYS1	online
ATLAS_OPP_OSG-SU_OG	online
ATLAS_OPP_OSG-USCMS-FNAL-WC1-CE3	online
ATLAS_OPP_OSG_ES-PURDUE	online
ATLAS_OPP_OSG_ES-UCONN	online
ATLAS_OPP_OSG_MCORE-UCSDT2	online

ATLAS Jobs on OSG Opportunistic Sites

- Usage is volatile, but peak simultaneous jobs can be significant

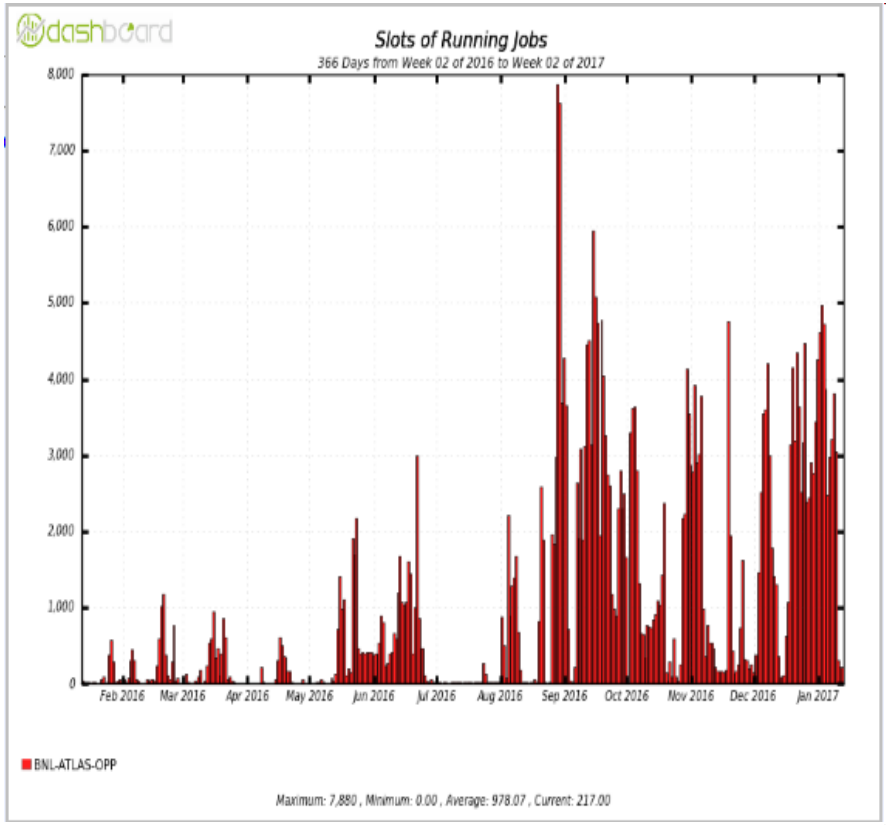

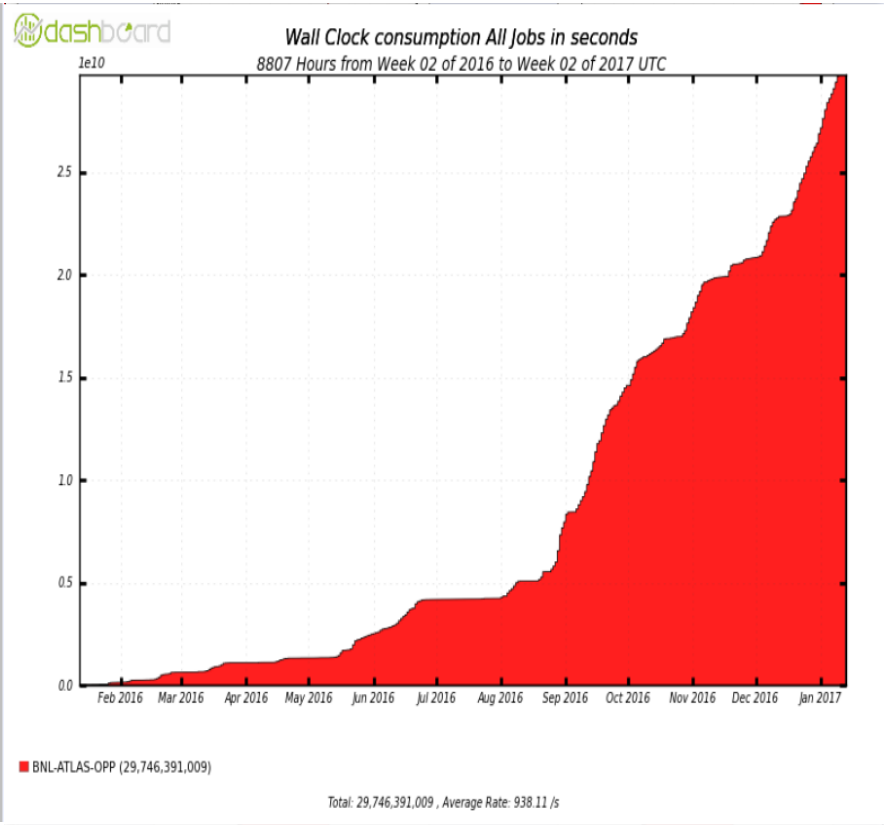

Questions ?

SDCC Current Configuration

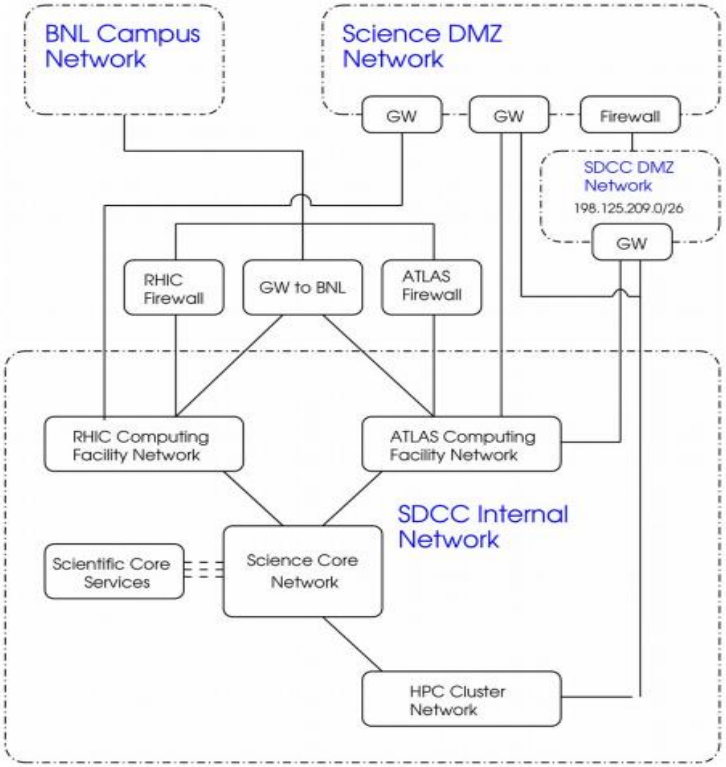

Configure After Jan 30

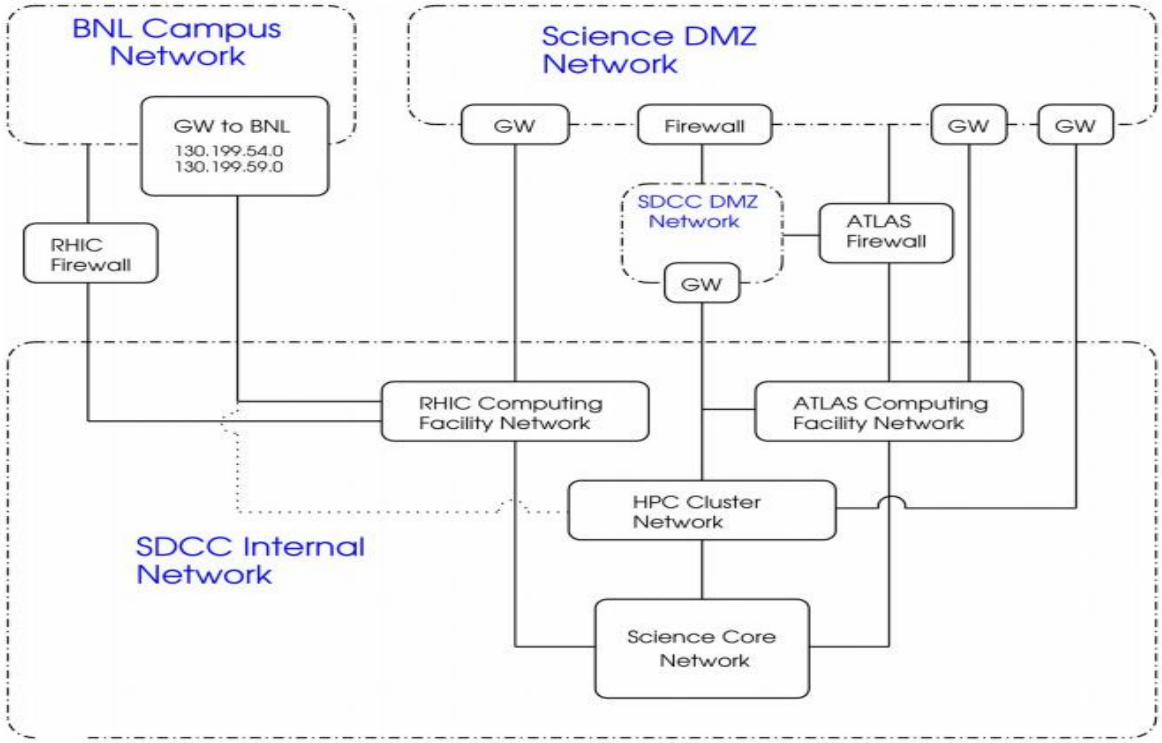

Backup Slides

Date	New Req	Staged	Failed	Data Volume	Average MB/s
01/11/2017 00:00:00	1605	1633	0	585,396,061,953	155.08
01/11/2017 01:00:00	2007	1951	0	760,356,521,658	201.43
01/11/2017 02:00:00	1791	1855	0	3,062,424,578,459	811.27
01/11/2017 03:00:00	1982	2033	0	3,888,970,095,398	1030.23
01/11/2017 04:00:00	1699	1710	0	3,796,246,265,582	1005.66
01/11/2017 05:00:00	2517	2482	0	3,003,709,885,109	795.71
01/11/2017 06:00:00	2251	2241	0	762,490,212,934	201.99
01/11/2017 07:00:00	2081	2015	0	648,813,882,979	171.88
01/11/2017 08:00:00	1228	1215	0	440,803,029,680	116.77
01/11/2017 09:00:00	2297	2420	0	972,648,422,679	257.66
01/11/2017 10:00:00	2229	2202	0	1,672,576,231,070	443.08
01/11/2017 11:00:00	1981	1964	0	2,662,066,393,059	705.21
01/11/2017 12:00:00	1919	1947	0	2,565,759,484,005	679.69
01/11/2017 13:00:00	1747	1721	0	4,062,789,294,705	1076.27
01/11/2017 14:00:00	2176	2194	0	3,186,628,257,301	844.17
01/11/2017 15:00:00	1845	1777	0	2,370,421,391,742	627.95
01/11/2017 16:00:00	1326	1375	0	3,358,896,586,460	889.80
01/11/2017 17:00:00	1628	1599	0	2,618,689,730,919	693.72
01/11/2017 18:00:00	1557	1692	0	2,920,748,610,089	773.73
01/11/2017 19:00:00	2040	1883	0	2,192,466,095,603	580.81
01/11/2017 20:00:00	1446	1525	0	1,448,261,789,917	383.66
01/11/2017 21:00:00	1734	1705	0	3,202,032,394,601	848.25
01/11/2017 22:00:00	1539	1507	0	1,991,100,218,575	527.46
01/11/2017 23:00:00	1429	1368	2	1,238,955,364,515	328.21

Jan 11 Sample log

Backup Slides

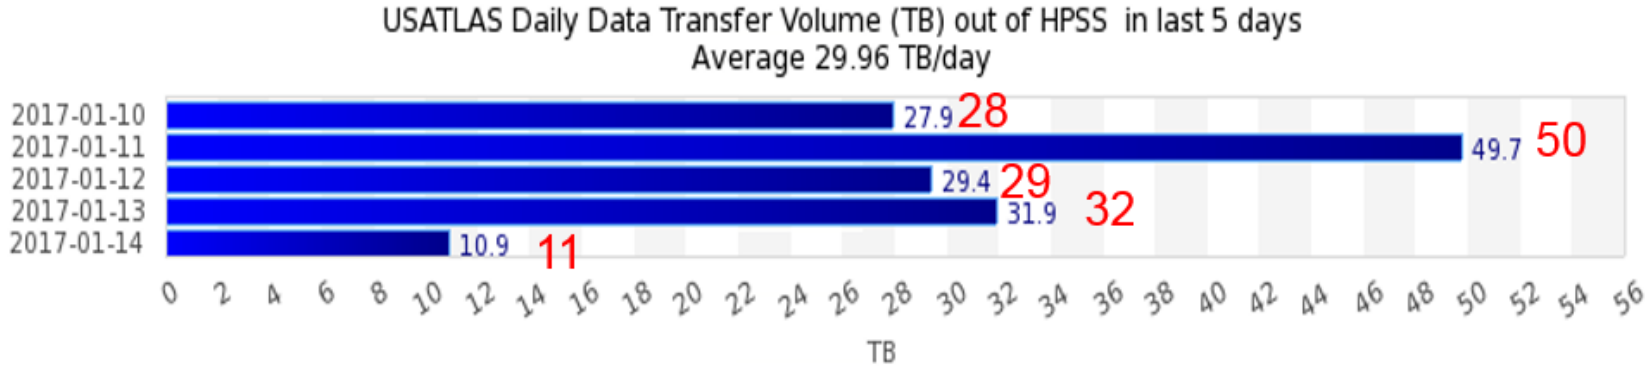

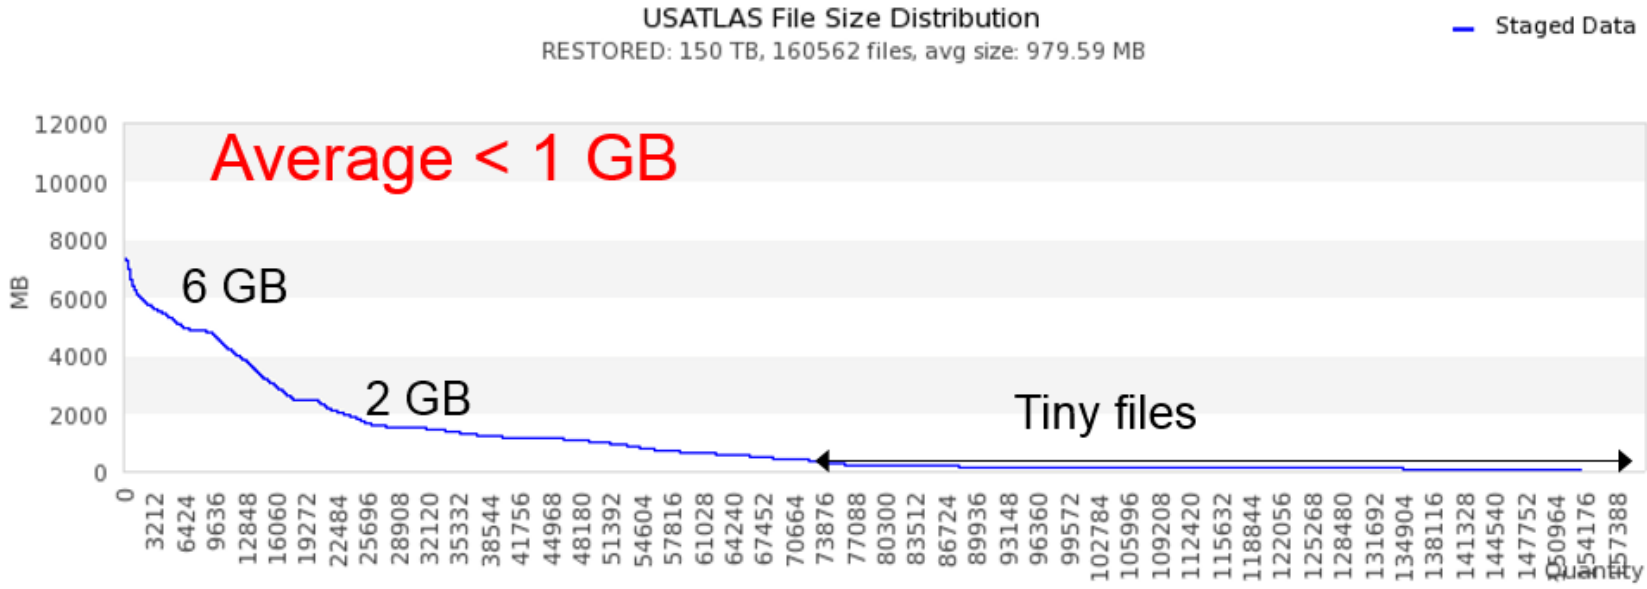

Backup Slides

Efficient Access to Massive Amounts of Tape-Resident Data

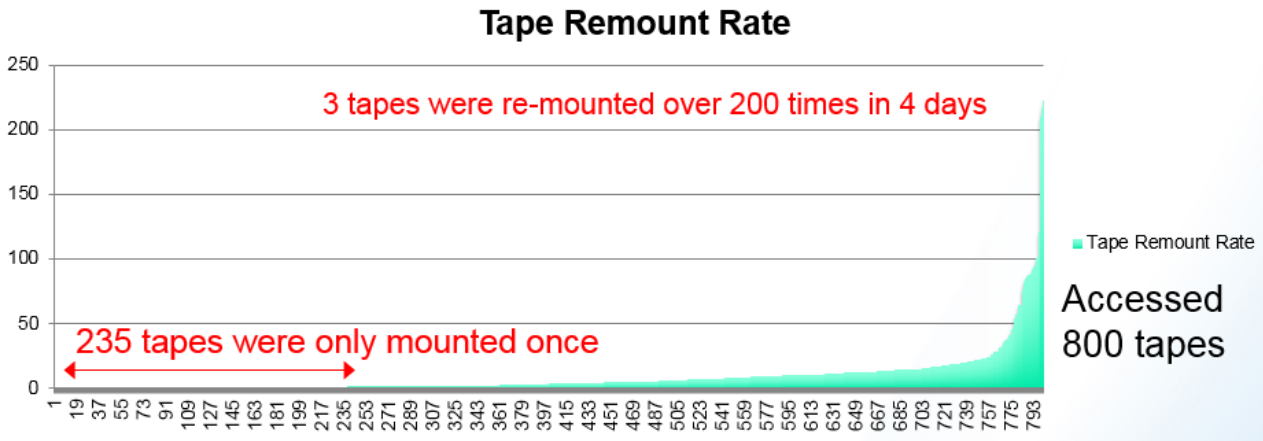

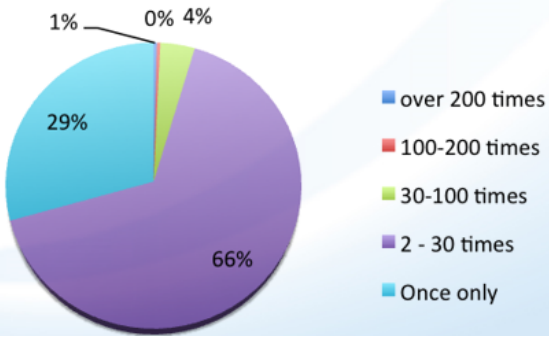

Need to reduce the re-mounts

Brookhaven Science Associates