U.S. ATLAS Computing Facilities Overview

Michael Ernst

Brookhaven National Laboratory

U.S. ATLAS Tier-2 & Tier-3 Meeting

LIGO Observatory, Livingston

March 3rd, 2009

HEP office of high energy physics

BROOKHAVEN NATIONAL LABORATORY

Facility Organization

Facility divided into two principal lines

- Production and Analysis Coordination & Operation
- Computing Deployment, Integration and Operation

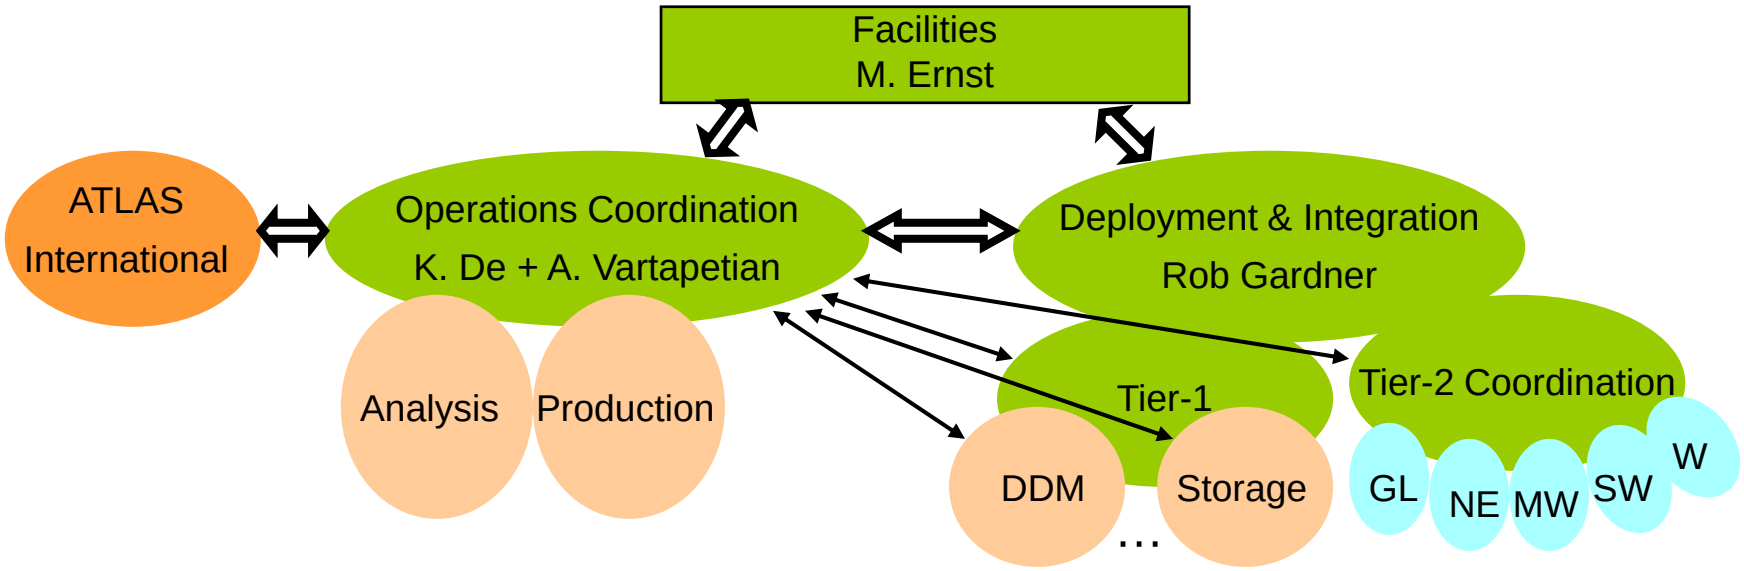

Securing the Facilities' Readiness

- Towards ATLAS Milestones
  - A Computing Integration Program is in place which aims at building the Integrated Virtual Computing Facility that we need to support LHC Data Analysis for the ATLAS Community
    - Coordinated by Rob Gardner (UC) and Michael Ernst (BNL)
    - Conducting weekly Computing Integration and Operation meetings
      - **Chaired by Rob Gardner**
      - **Excellent continuing participation by Tier-2 PIs, site administrators, production and analysis coordinators, technical experts, and Rich Carlson!**
    - Organizing ~ quarterly F-to-F meetings w/ Tier-2 and Tier-3 admins
      - **March '07: UCSD, jointly with Open Science Grid (OSG) AH Meeting**
      - **June '07: Indiana University, ~40 participants**
      - **November '07: SLAC, ~50 participants**
      - **March '08: RENCI, jointly with Open Science Grid (OSG) AH Meeting**
      - **May '08: University of Michigan, ~40 Participants**
      - **September '08: BNL**
      - **March '09: LIGO Observatory, jointly with Open Science Grid (OSG) All-Hands Meeting**
    - Mature set of processes for Facility Integration
      - **Well structured WBS, Tracking progress**
    - Tasks organized into quarterly phases
      - **Quarterly reporting of progress to the Facility project (WBS line item)**
  - Much more details about the Integration Program in Rob's talk

U.S. ATLAS Facility Capacities

- 2008 and 2009 Pledges for CPU and Disk (P08, P09), and currently installed capacities

Site	CPU (SI2K)	CPU P09 (SI2K)	CPU P08 (SI2K)	Disk (TB)	Disk P09 (TB)	Disk P08 (TB)
T1	5,451,738	7,337,000	4844000	3100	5,822	3136
NET2	1,031,068	1049000	665000	170	445	244
SWT2	2,341,050	1734000	1386000	260	328	256
MWT2	3,478,432	978000	1112000	520	358	282
AGLT2	2,003,088	1406000	965000	670	542	322
WT2	1,171,568	1202000	820000	220	413	462
USATLAS FACILITY	15,476,944	13,706,000	9,792,000	4,940	7,908	4,702
USATLAS TIER2	10,025,206	6,369,000	4,948,000	1,840	2,086	1,566

The currently installed CPU capacity allows to simulate 12M fast simulated or 500k full simulated events per day

T1	U.S. ATLAS Tier-1 Center at BNL
NET2	North East Tier-2 Center at Boston University and Harvard University
SWT2	Southwest Tier-2 Center at University of Texas at Arlington and Oklahoma University
MWT2	Midwest Tier-2 Center at University of Chicago and Indiana University
AGLT2	ATLAS Great Lakes Tier-2 Center at University of Michigan and Michigan State University
WT2	Western Tier-2 Center at SLAC

ATLAS Resource Requests for 2009

- Given small differences for 2009 between current model and published (07/2007) requests for Tier-1 and Tier-2 resources ATLAS prefers no change at this stage
  - Official requests for 2009 are summarized in the table below

	CPU (MSI2k)	Disk (PB)	Tape (PB)
Tier-0	*7.6*	*0.7*	*8.6*
CAF	*5.8*	*3.3*	*1.1*
Total CERN	**13.4**	**4.0**	**9.7**
Sum of Tier-1s	**28.4**	**20.9**	**15.8**
Sum of Tier-2s	**27.0**	**13.3**	**–**

  - ATLAS resource requirements strategy document available at
    - http://cdsweb.cern.ch/record/1131818/files/LHCC-G-141.pdf
  - Resources above approved at RRB meeting on 11/17/2008
    - http://lcg.web.cern.ch/LCG/planning/phase2_resources/WLCGResources2008-2013%20171108.pdf
  - ATLAS is further refining the Computing Model, and the 2009 and 2010 requirements
    - Accelerator Schedule and Experiment Run Scenarios were discussed at the Chamonix Workshop
    - Planning for a long physics run from October 2009 until October 2010 (70 days of run time)
  - U.S. ATLAS is ~23% of the total

Modified Equipment Deployment Schedule

- Following the discussion in the Management Board on Jan 13th, the following is what we agreed would be the process now for updating the resource procurement planning for 2009 and 2010 (Ian Bird, January 16th, 2009).
  1) It now appears unlikely that we will see collisions for physics much before September this year, assuming that according to the presently understood schedule the accelerator complex restarts in the summer once the repairs are complete. However, the experiments continue to run with cosmic data, as well as simulation productions in preparation for data taking and analysis. Although significant portions of the 2009 resources have already been procured and installed at many sites, we nevertheless feel that for this year it is reasonable to relax the requirement to have the full 2009 resources commissioned by April, and to push this back to July for CPU and the first part of disk commitment, and later in the year for the remainder of the disk.
     a. This raises the opportunity to get next generation equipment in some cases where the procurement process allows it;
     b. In many cases, changing the process at this late date is not possible as commitments have already been made (as sites were requested to do given the information previously available);
  2) The LHC workshop to be held in Chamonix on Feb 2-6 will discuss the details of the accelerator schedule for 2009 and 2010, taking into account updated input from the experiments. It is expected to also provide information on likely running scenarios including energy, and timescales for heavy ion running.
  3) Once there is a better understanding of the expectations for the accelerator and experiments for 2009 and 2010 coming from the Chamonix workshop we can then update the experiment requirements and thus the resource procurement plans for 2009 and 2010 (we should now treat these together). The timescale on which this planning should be updated is the following:
     a. Feb 6: input from Chamonix workshop,
     b. Feb 16: discuss the resource planning with the LHCC referees during the mini review of WLCG. Since this is only 10 days after the workshop we probably will not have the final numbers available;
     c. First week of April: meet with experiment management and jointly plan the presentation to the RRB of an updated resource schedule for 2009/10; inform C-RSG of the result of the discussions;
     d. April 28: presentation to the Computing RRB for agreement of the updated plan
  4) All WLCG sites will be informed of changes in the anticipated schedule through the Management Board.

Modified Equipment Deployment Schedule

- Following the discussion in the Management Board on Jan 13th, the following is what we agreed would be the process now for updating the resource procurement planning for 2009 and 2010 (Ian Bird, January 16th, 2009).
  1) It now appears unlikely that we will see collisions for physics much before September this year, assuming that according to the presently understood schedule the accelerator complex restarts in the summer once the repairs are complete. However, the experiments continue to run with cosmic data, as well as simulation productions in preparation for data taking and analysis. Although significant portions of the 2009 resources have already been procured and installed at many sites, we nevertheless feel that for this year it is reasonable to relax the requirement to have the full 2009 resources commissioned by April, and to push this back to July for CPU and the first part of disk commitment, and later in the year for the remainder of the disk.
     a. This raises the opportunity to get next generation equipment in some cases where the procurement process allows it;
     b. In many cases, changing the process at this late date is not possible as commitments have already been made (as sites were requested to do given the information previously available);
  2) The LHC workshop to be held in Chamonix on Feb 2-6 will discuss the details of the accelerator running.

LHC machine will be running in "Physics mode" from ~November 2009 - ~October 2010

  3) Once there is a better understanding of the expectations for the accelerator and experiments for 2009 and 2010 coming from the Chamonix workshop we can then update the experiment requirements and thus the resource procurement plans for 2009 and 2010 (we should now treat these together). The timescale on which this planning should be updated is the following:
     a. Feb 6: input from Chamonix workshop,
     b. Feb 16: discuss the resource planning with the LHCC referees during the mini review of WLCG. Since this is only 10 days after the workshop we probably will not have the final numbers available;
     c. First week of April: meet with experiment management and jointly plan the presentation to the RRB of an updated resource schedule for 2009/10; inform C-RSG of the result of the discussions;
     d. April 28: presentation to the Computing RRB for agreement of the updated plan
  4) All WLCG sites will be informed of changes in the anticipated schedule through the Management Board.

Projected Tier-2 Capacity Development

- "Flat-Flat" Funding may result in difficulties in 2012 and beyond
  - Counting on Moore's Law to compensate rising salaries
  - Inexpensive Storage Components are heavily taxing on personnel
  - Replenishment of outdated equipment may become a problem
    - Space, Power, Cooling a potential problem if old equipment cannot be replaced

CPU Pledge (kSI2k)

Site	2009	2010	2011	2012	2013
T2 sum	6369	7681	19948	19948	27370
NET2	1049	1592	1966	3990	5474
SWT2	1734	1966	2514	3990	5474
MWT2	978	1262	1785	3990	5474
AGLT2	1406	1670	2032	3990	5474
WT2	1202	1191	1685	3990	5474

Disk Pledge (TB)

Site	2009	2010	2011	2012	2013
T2 sum	2086	3067	4481	9255	12888
NET2	445	727	1024	1851	2578
SWT2	328	650	1103	1851	2578
MWT2	358	362	1103	1851	2578
AGLT2	542	709	914	1851	2578
WT2	413	619	928	1851	2578
Tier-2 Funding (Ak$s)	3000	3000	3000	3000	3000

Walltime Days in U.S. Cloud in 2008 – Successful Production Jobs (Simulation)

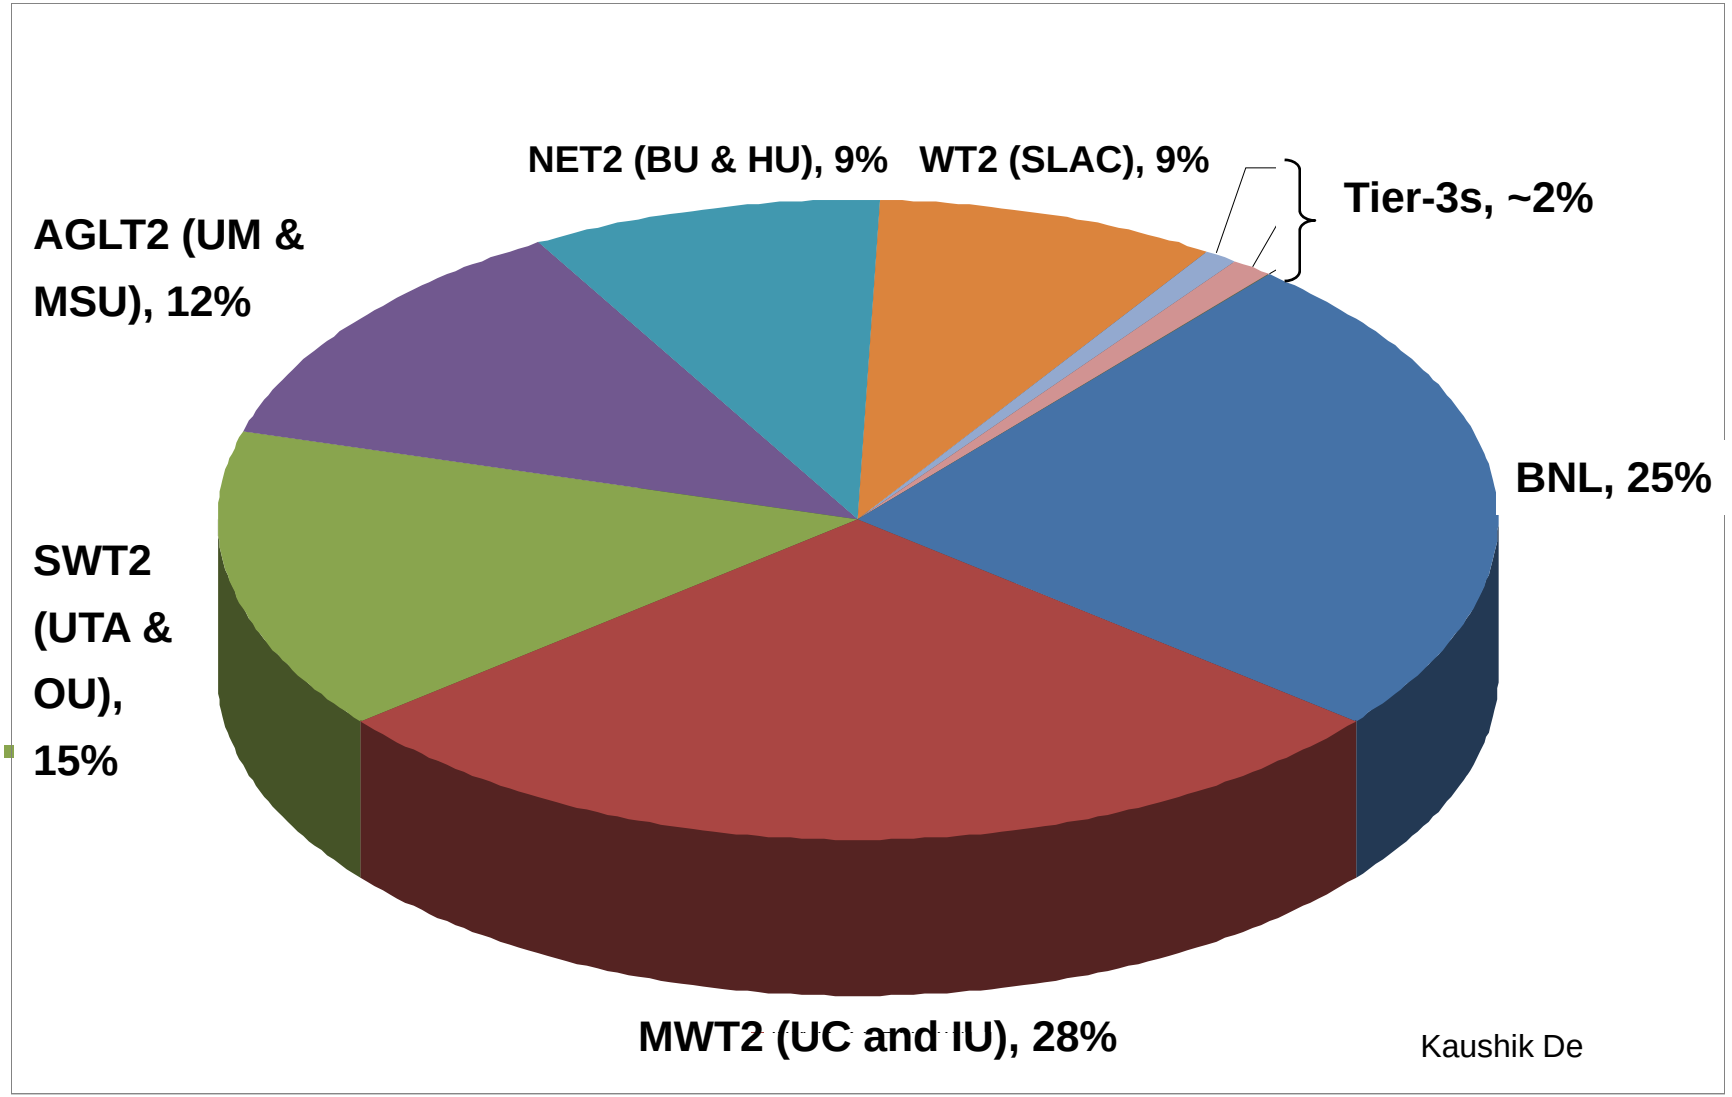

Common Job Errors at U.S. Sites – From PanDA Monitoring

All errors in 2008

(Note: >85% of all jobs Completed w/o Error)

Lost Heartbeat (site/scheduler/…) 21%

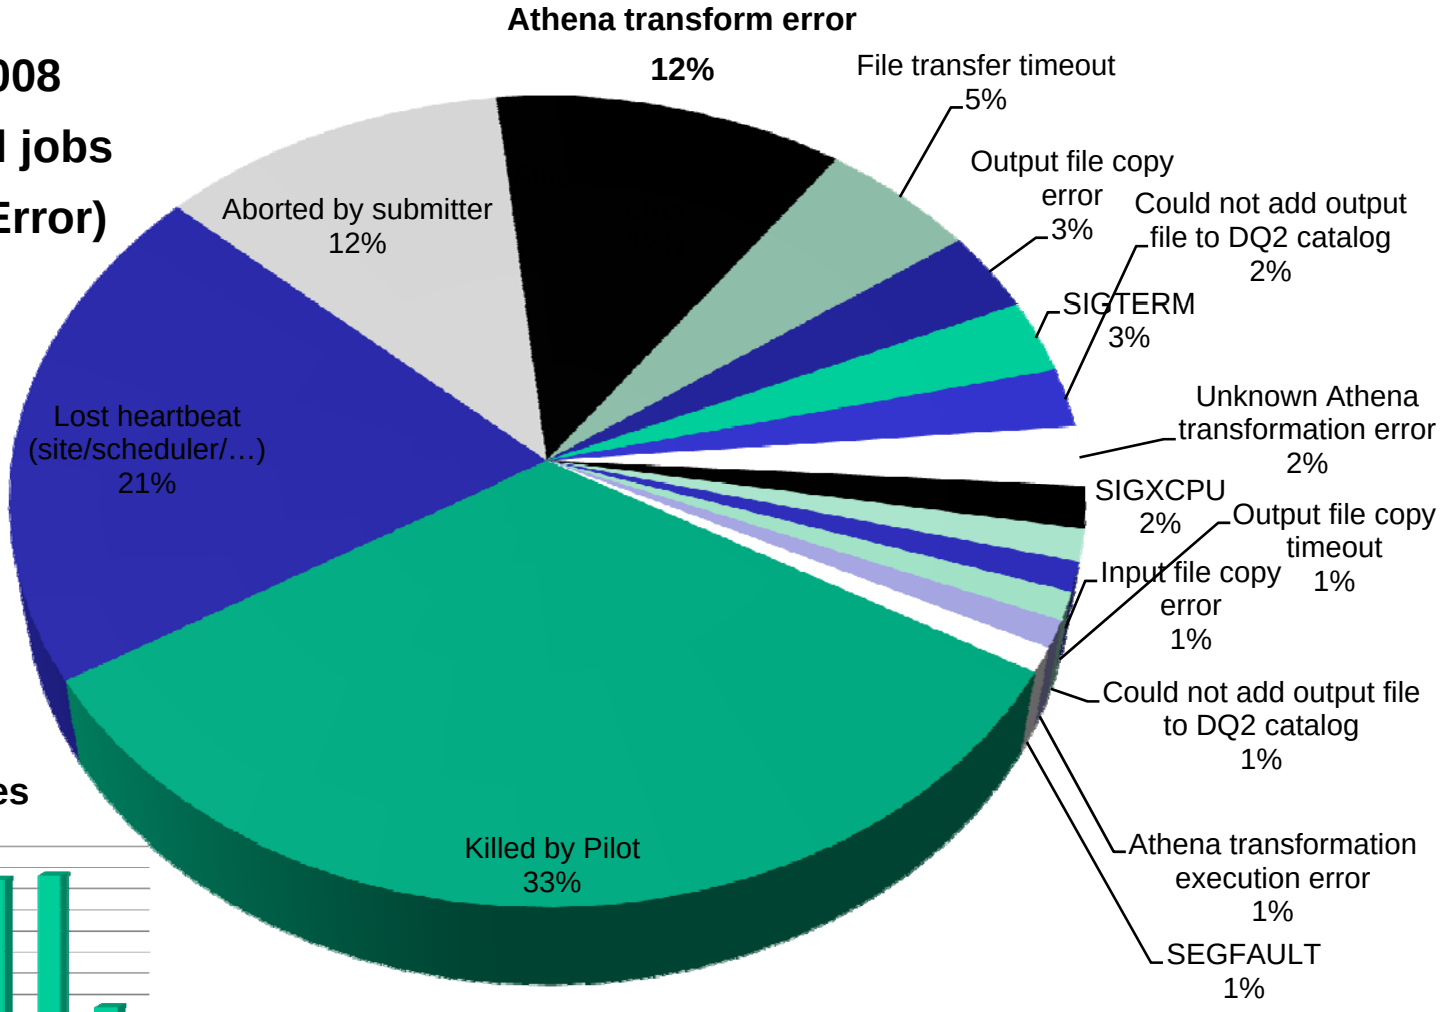

Efficiency at U.S. Sites

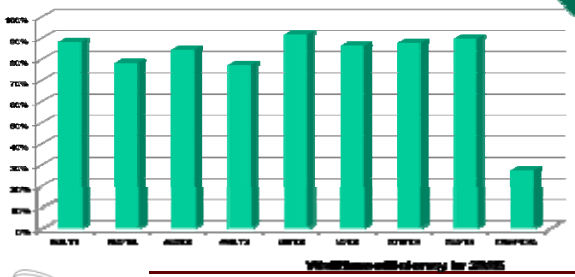

Kaushik De

Reprocessing

- All ‘good’ cosmic and single-beam data collected by ATLAS in 2008 was reprocessed in December
  - Massive ATLAS-wide exercise
    - >500 TB RAW data distributed to Tier-1 centers according to MoU shares
  - 10 different detector and performance streams, >250M events
  - Multiple formats of output data, to be available widely for analysis
- 42% of the reprocessing jobs successfully done in US cloud
  - More than the U.S. share (23%)
    - US was asked by ATLAS to process additional shares from DE, ES, NL, TW to complete reprocessing in time
    - Not because of over-provisioning of computing resources
    - Facilities (Tier-1 and Tier-2’s) working well
- US (alone in ATLAS) allows reprocessing at Tier 2’s
  - Almost 20% of all December reprocessing was done at US Tier 2’s
  - Automated staging of input data by Production System/PanDAMover
  - We plan to maintain this option (other large T2’s in EGEE may join)
  - Very successful exercise

December 2008 Reprocessing - Worldwide

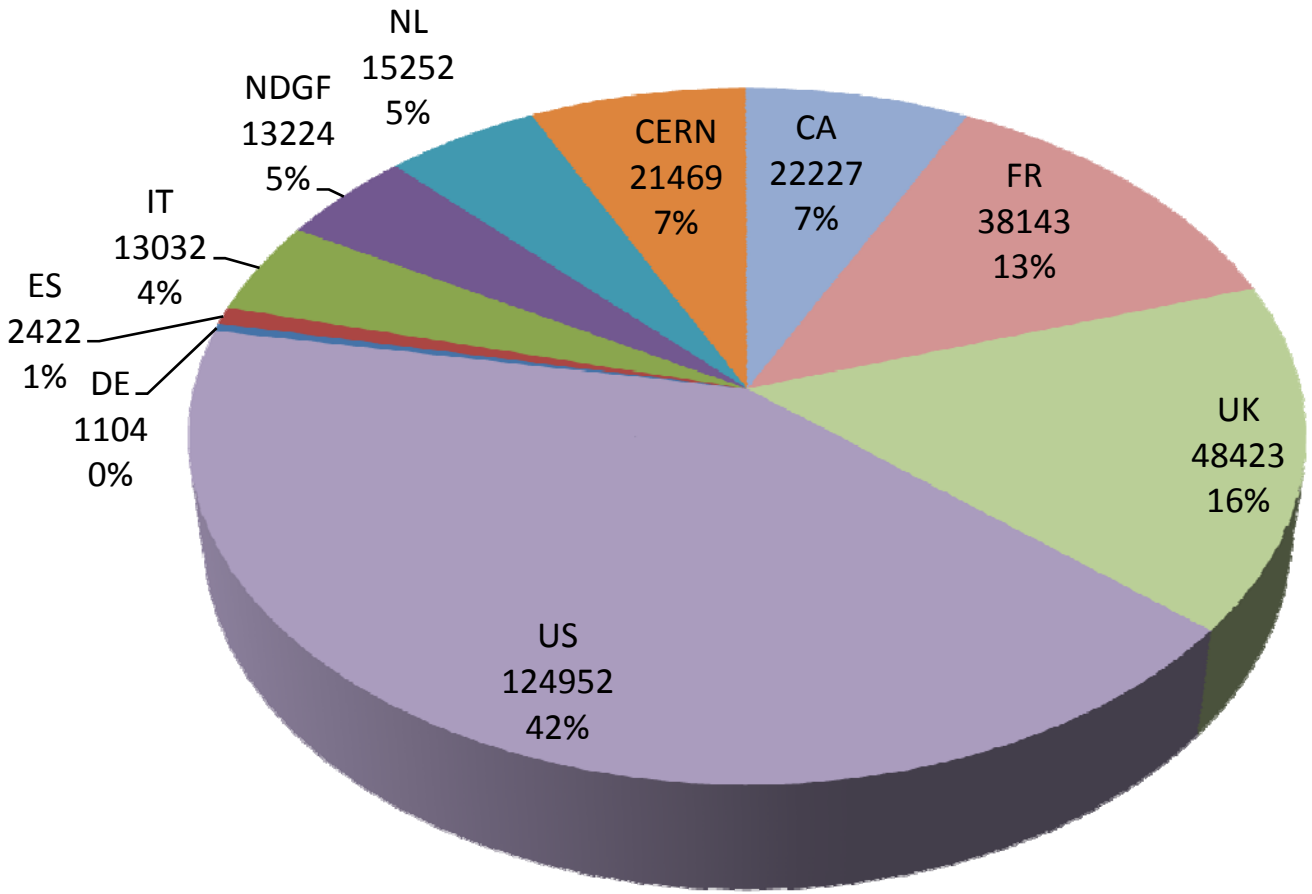

Number of Successful Jobs

Kaushik De

December 2008 Reprocessing - US

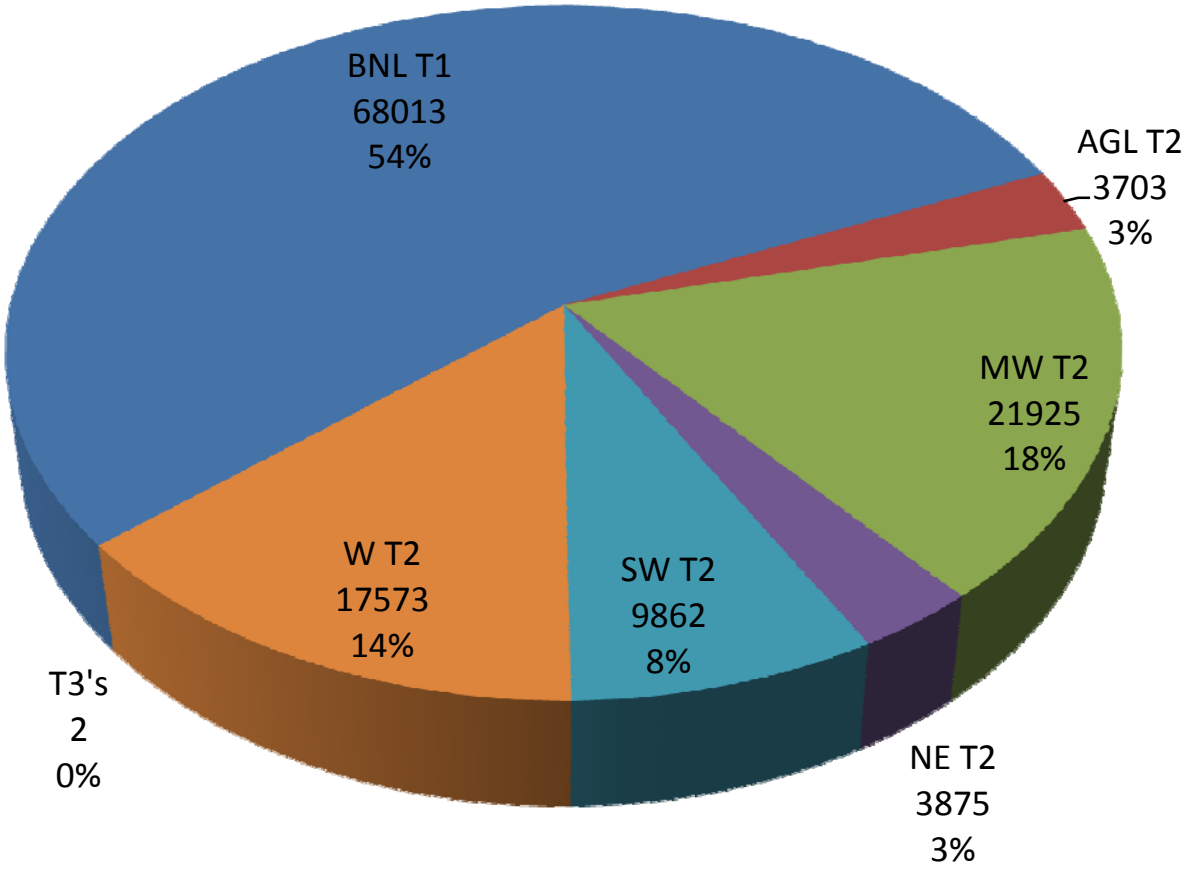

Successful Jobs per Site

Kaushik De

Tier-0/1 ⇔ Tier-1 ⇔ Tier-2 Connectivity

- For Tier-1 <=> Tier-2 Connectivity we are in the process of moving from ESnet/Internet2 based "best effort" IP Services to Circuit based Network Paths
  - To circumvent potential bottlenecks in ESnet Backbone (IP Services)
- All Tier-2 Sites but UTA are connected at 10 Gbps
- ESnet added 2 * 10 Gbps links to BNL WAN Connectivity
  - 10 Gbps will be used to increase Tier-0/1 <=> US Tier-1 B/W to 20 Gbps
  - 10 Gbps will be used to provide dedicated connectivity between BNL and major peering points in the U.S. relevant to Tier-2 connectivity
    - StarLight (MWT2, AGLT2, OU, possibly UTA) and MANLAN (BU, via Internet2)
    - BNL ⇔ UC circuit delivered by ESnet, waiting for UC equipment to arrive
    - BNL ⇔ BU to be delivered soon (04/01) by Internet2 & ESnet
    - Next will be AGLT2 (following the same engineering model we are using for UC)
- Connection has now the desired redundancy and diversity between NY and BNL

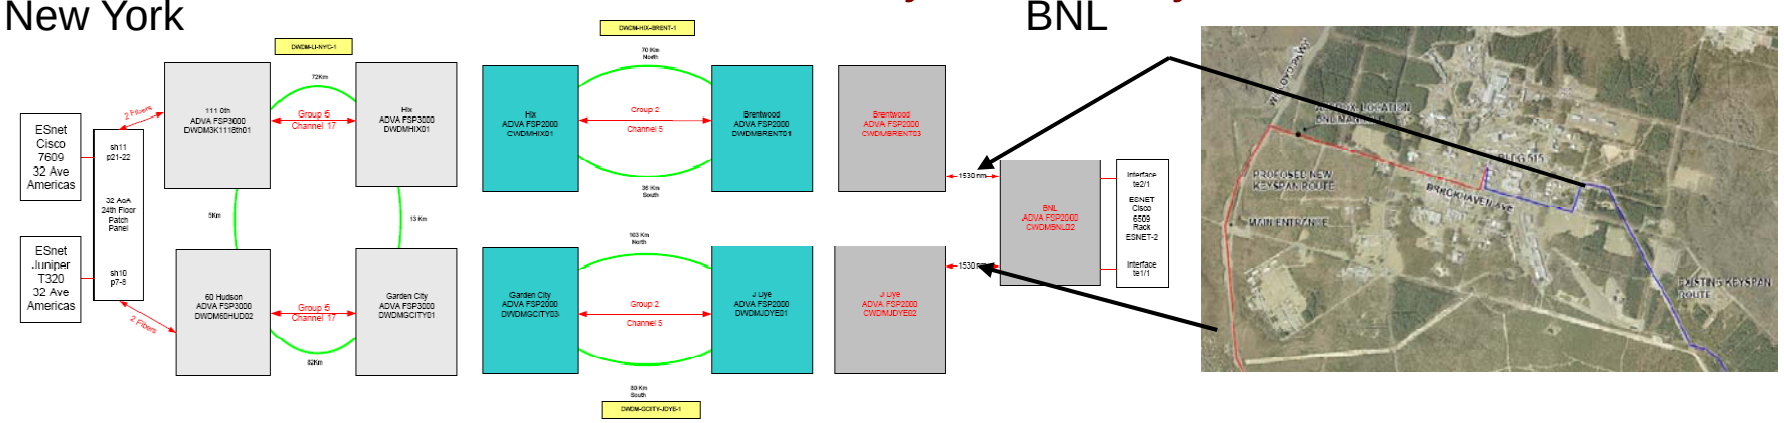

U.S. ATLAS Tier-1 Ingest Rate from CERN

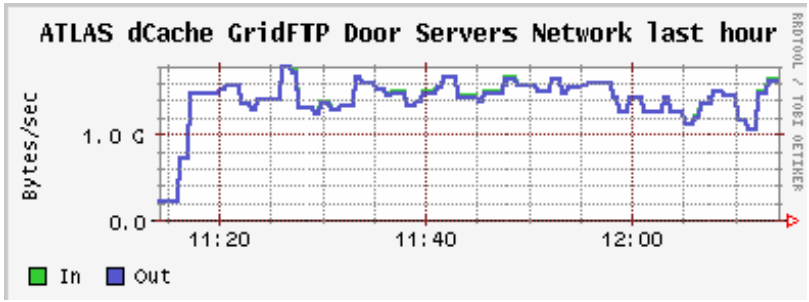

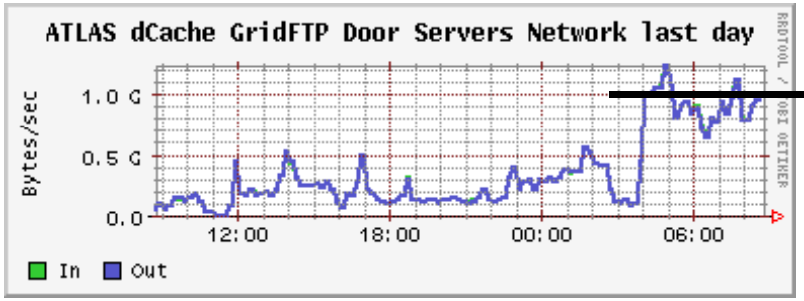

1 GB/s

Fully saturating the 10 Gbps link in a test …

… and in production on 10/29/2008

MB/s

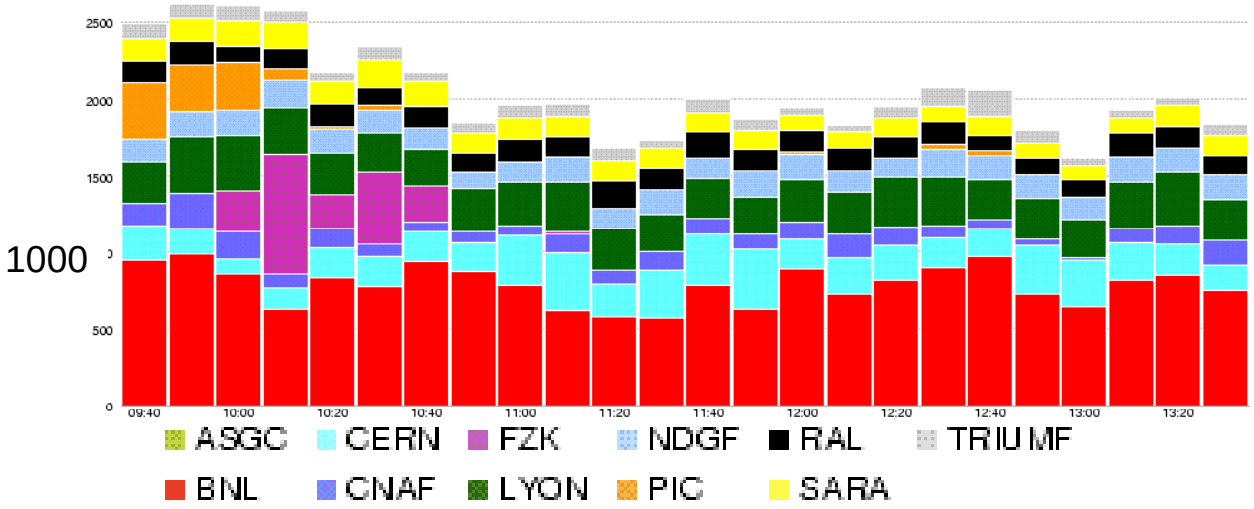

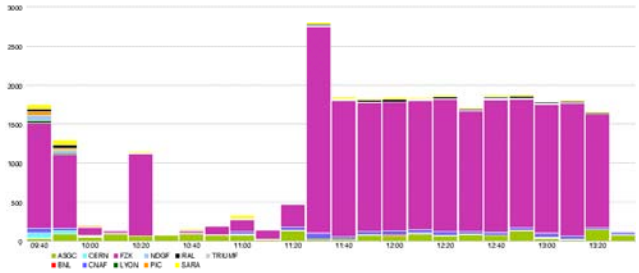

Export data rate to all ATLAS Tier-1 Centers on October 29th

Errors

Throughput (MB/s, daily average)

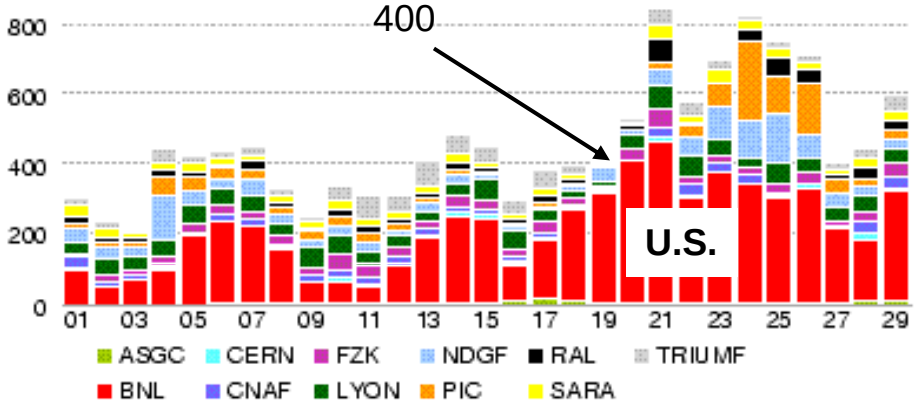

Data Transferred (GBytes/day)

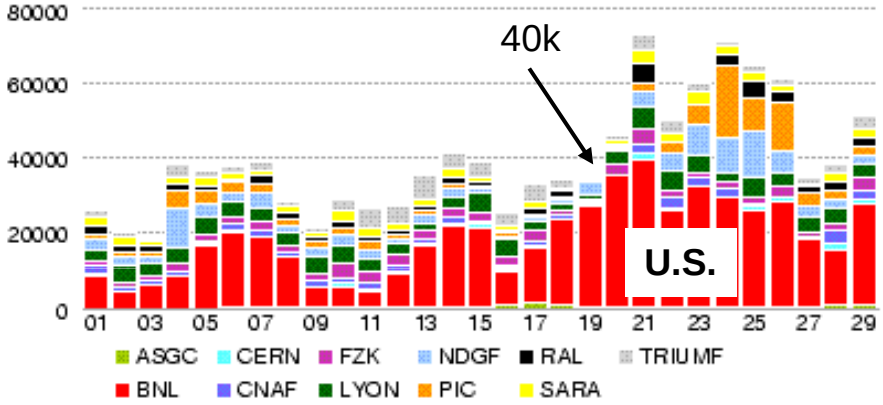

Completed File Transfers per day

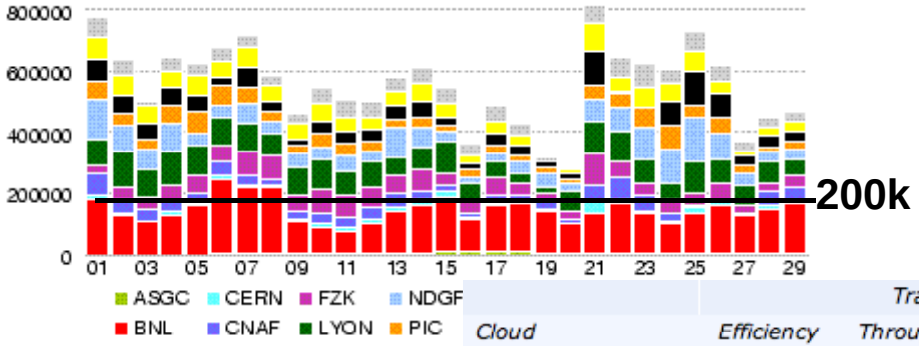

Total Number Transfer Errors per day

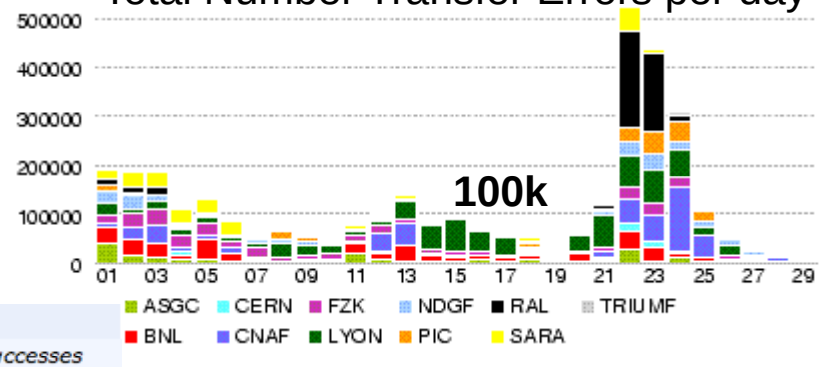

	Transfers		
Cloud	*Efficiency*	*Throughput*	*Successes*
ASGC	29%	3 MB/s	66243
BNL	91%	209 MB/s	4144181
CERN	84%	7 MB/s	314019
CNAF	66%	16 MB/s	955978
FZK	80%	25 MB/s	1479592
LYON	75%	42 MB/s	2282913
NDGF	88%	42 MB/s	1772825
PIC	83%	37 MB/s	1127813
RAL	74%	20 MB/s	1301170
SARA	83%	22 MB/s	1400394
TRIUMF	98%	31 MB/s	1297612

Data Replication within the U.S. Cloud in November 2008

- U.S. ATLAS Tier-1 => 5 Tier-2s
- 209 MB/s average, 91% Efficiency
- 4 M files total, 200k files per day
- Up to 40 TB/day

ATLAS DDM/SE-Stress-Test (Jan 2009)
10M Files Distributed Tier-1⇔Tier-1 + CERN

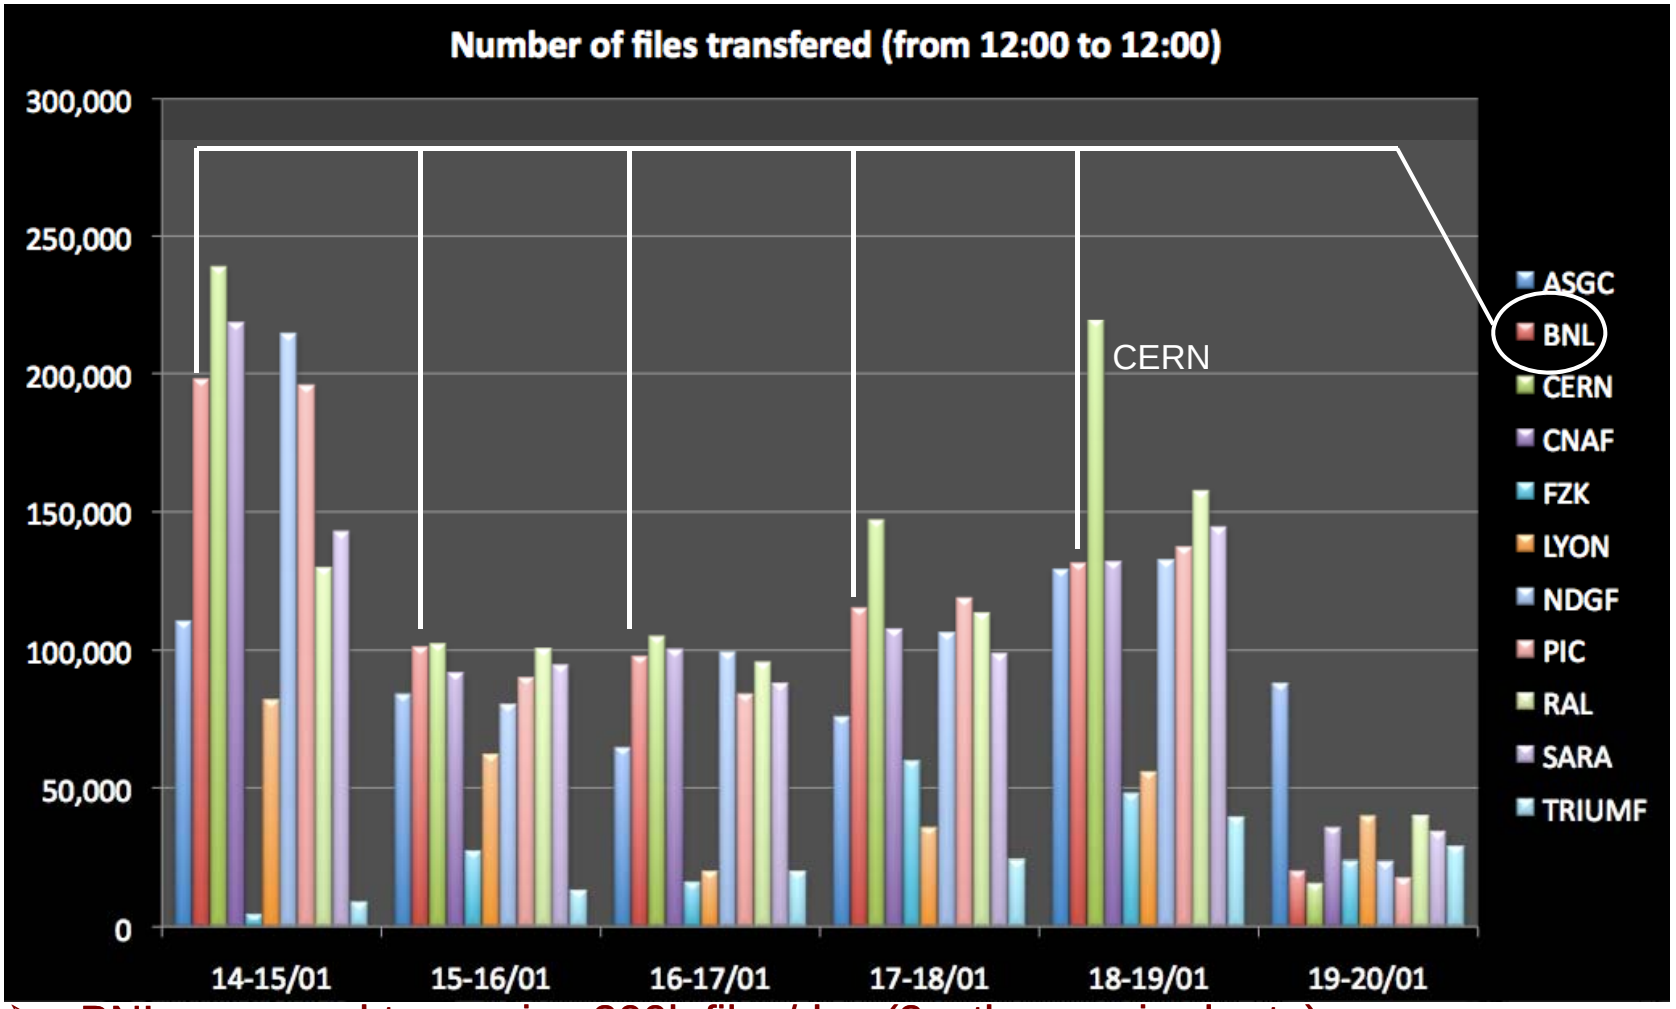

- BNL managed to receive 200k files/day (2 x the required rate)
  - Important in case U.S. has to catch up after maintenance or distribution problems

Facilities Accomplishments - 2008

- Fabric Resources
  - The Tier-1 center and each Tier-2 is deploying processing and storage resources according to a profile which meets WLCG pledge requirements and those of the U.S. ATLAS physics community
    - Almost achieved – Currently adding 1PB at Tier-1 to complete '08 WLCG pledges, Tier-2s ~OK
- Physics exercises
  - Full participation in FDR-1, FDR-2, M* (ATLAS Milestone), CCRC'08 exercises and cosmic data taking & analysis, and reprocessing at the Tier-1 and all Tier-2s
  - Support for Physics Analysis Workshops (Jamborees)
    - Achieved with good results. Ready for cosmic data taking and challenges in Spring 2009
- Data throughput benchmarks
  - Sustained 200-400 MB/s Tier-1 to every Tier-2 site, Tier-1 aggregate at 1GB/s
    - Partially achieved – Working with Internet2 & ESnet on Network Performance Optimization
- Storage Resource Manager (SRM) v2 functionality at all Storage Elements
  - SRM/dCache and Bestman/xrootd implementations
  - Supported by OSG (VDT packaging and Storage Group)
  - ATLAS defined space tokens
    - All sites completed configuration in early May, right before the start of CCRC'08
- Migration from LRC (Local Replica Catalog, U.S. only) to ATLAS-wide LFC (LCG File Catalog, now part of OSG/VDT distribution)
  - Improves efficiency of Distributed Data Management throughout ATLAS Computing Facility
    - All 10 sites completed migration by early December 2008
- U.S.ATLAS Facilities fully based on OSG Grid Middleware (OSG 1.0)
- Network monitoring infrastructure (perfSonar) deployed at all sites, needs more work
- Implemented Managed ATLAS S/W Release installations at all sites
- WLCG
  - Accounting statistics forwarded via OSG middleware - Achieved
  - Site-level monitoring with OSG-RSV, forwarded to WLCG-SAM/Gridview - Achieved

U.S. ATLAS benefits from OSG

- Serves as integration and delivery point for core middleware components including compute and storage elements (VDT)
- Cyber Security operations support within OSG and across Grids (e.g. WLCG) in case of security incidents
- Cyber Security infrastructure including site-level authorization service, operational service for updating certificates and revocation lists
- Service availability monitoring of critical site infrastructure services, i.e. Computing and Storage Elements (RSV), and forwarding of results to WLCG
- Site level accounting services and forwarding of accumulated results to WLCG
- Consolidation of Grid client utilities incl. incorporation of LCG client suite, resolving Globus library inconsistencies
- LCG File Catalog (LFC) server and client packaging – needed in support of the ATLAS global Distributed Data Management system (DDM)
- BestMan and xrootd: SRM and file system support for Tier 2 and Tier 3 facilities
- dCache and BestMan/xrootd packaging with VDT, and support through OSG-Storage
- Essential support for basic Middleware services at Tier-3 sites
- Support for integration and extension of security services in the PanDA workload management system and the GUMS grid identity mapping service, for compliance with OSG security policies and requirements
- Integration testbed for new releases of the OSG software, pre-production deployment testing with Panda
- Excellent response from OSG to specific U.S. ATLAS requirements

Drivers for 2009

The 2007/2008 drivers were primarily designed to

- Demonstrate the functionality of the computing model elements and to improve the stability of the services while scaling to nominal performance level
  - To be continued in 2009
- Participation in global ATLAS exercises
  - Cosmic Data Taking and Data Processing
  - Encouraging ATLAS to run Analysis Challenges
- Goal is to have end-to-end data collection, distribution, and access
  - Includes DAQ and Tier-0 elements, Tier-1 and Tier-2 data replication, data management and operations
  - Application services
    - Reprocessing at the Tier-1 Center
      - **Working jointly with ATLAS International Distributed Computing Operations**
      - **Running large scale exercises based on Cosmic Data**
    - Analysis at the Tier-2 Centers, and beyond …
      - **Infrastructure in place since early 2008**
      - **Data relevant to Analysis available at Tier-2's**
      - **Understand and (help) implement end user analysis facility components**
- Continuous operation with increasing functionality and activity level
  - Required integrated reliability over a month is >90% for Tier-2 centers
    - No more than 3 days per month should be lost due to unscheduled outages

Analysis – From HLT to D3PD

- Now that (we believe) we understand Facility Services up to the Tier-2 Level we are working on Facility related End User Analysis issues
  - Focus of this Meeting!

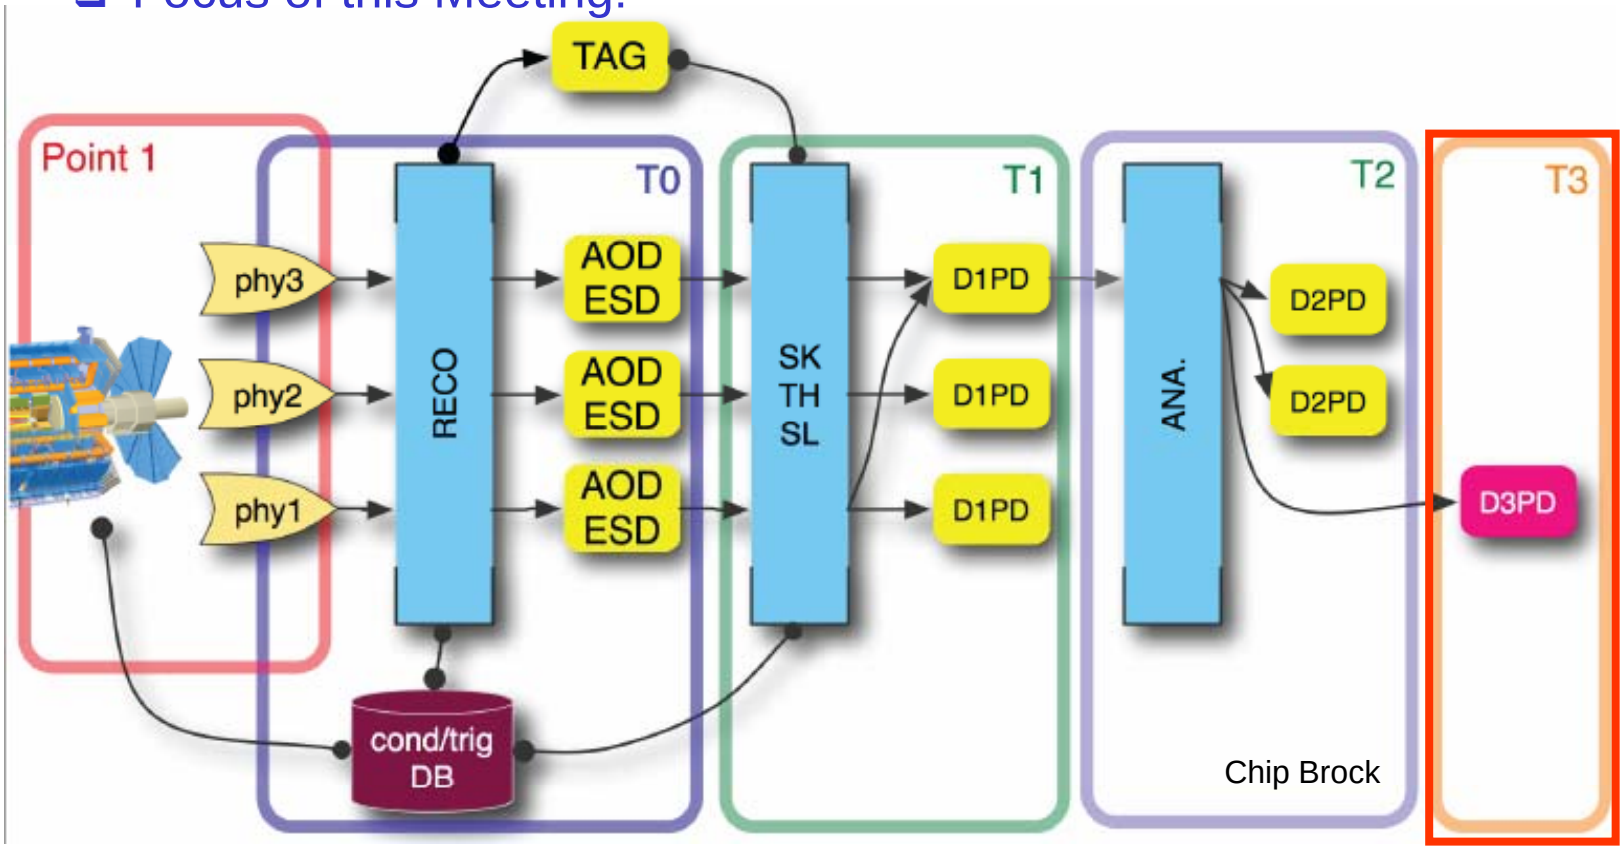

What is working well - a selection

- Integration with
  - OSG as our Middleware Platform
  - ATLAS Distributed Computing Operations & Contribution to Production
- Excellent Teamwork within Facility Group
  - Excellent level of (distributed) expertise and spirit
  - People (YOU !!!) volunteer, are ready at all times to pick up the slack, which is key to letting us progress
  - Good response in case of service disruptions
- Task Forces focusing on specific issues

Migration to SRM-capable Storage Elements	DONE
Migration from LRC to LFC	DONE
Throughput Initiative and E-2-E Monitoring (perfSONAR)	In Progress
Analysis queue performance	NEW
Data and Storage Management issues	NEW

- DDM/DQ2 Performance & Stability Improvements

What needs improvement - a selection (1/2)

- Site Stability
  - Some sites are suffering from frequent component failures
    - Most problems due to SE issues
      - **Are our technology choices right?**
    - Sometimes caused by site infrastructure (e.g. AFS)
    - While this is ~tolerable with centrally managed Production it will become a real issue with End-User Analysis
- Analysis support
  - Analysis Queues are configured - but hardly used
    - From the little usage we can't tell whether we're ready for prime time
      - **"What has not been stress tested doesn't work !!!"**
    - End User Analysis
      - **Facilities still lacking requirements estimates (can User Forums help?)**
    - Facilities need this information to "right-size" setup
      - **Tuning done according to well defined (production) job profile so far**
      - **Observed user jobs at BNL leading to component failures (i.e. SE, LRC)**
      - **Almost no way to protect the components - best we can do is instrument & react as fast as possible**
      - **Can the "framework" protect us from these cases?**
    - Need representative exercises & continuing workload involving ALL instances
      - **If ATLAS doesn't step up we have to organize challenges in the U.S.**

What needs improvement - a selection (2/2)

- Coordination of Production with Facilities
  - Site admins would benefit from knowing what is running at their sites (simulation vs reprocessing vs …), and what the required resources are (i.e. storage space)
- Interaction with NSF to secure Tier-2 funding in a timely manner
- Miscellaneous
  - Monitoring
    - PanDA Mover based transfers hard to debug - even for experts
    - …
  - …

Summary

- The facilities, the Tier-1 and the Tier-2's, have performed well in ATLAS computer system commissioning and specific exercises
  - An Integration Program is in place to ensure readiness in view of the steep ramp-up
  - The Tier-2's provide resources for User Analysis
  - Excellent contribution of U.S ATLAS Tier-2 Sites to high volume production (event simulation, reprocessing) and analysis in 2008
- Space, Power & Cooling additions at the Tier-1 center are well underway
  - Funding in 2010/2011 for power and cooling infrastructure unclear
- Overall, the Integrated Facility in the U.S. has demonstrated excellent capabilities in centrally managed production, however, more work is needed to prove our readiness for LHC data analysis

Additional Material

Pledged Capacities at the Tier-1 and the Tier-2s

US-ATLAS Tier1	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	2560	4844	7337	12785	18193	21132	32200
Disk (Tbytes)	1100	3136	5822	11637	16509	16965	27600
Tape (Tbytes)	603	1715	3277	6286	9820	15085	18630
Nominal WAN (Mbits/sec)	9952	19904	29856	39808	39808	39308	39808
USA, Northeast ATLAS T2	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	394	665	1049	1592	1968	3990	5474
Disk (Tbytes)	103	244	445	727	1024	1851	2578
USA, Southwest ATLAS T2	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	998	1386	1734	1966	2514	3990	5474
Disk (Tbytes)	143	256	328	650	1103	1851	2578
USA, Midwest ATLAS T2	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	826	1112	978	1262	1785	3990	5474
Disk (Tbytes)	213	282	358	362	512	1851	2578
USA, Great Lakes ATLAS T2	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	581	965	1406	1670	2032	3990	5474
Disk (Tbytes)	155	322	542	709	914	1851	2578
USA, SLAC ATLAS T2	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	550	820	1202	1191	1685	3990	5474
Disk (Tbytes)	228	462	794	1034	1462	1851	2578
US ATLAS FACILITY	2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K)	5909	9792	13706	20446	28175	41080	59570
Disk (Tbytes)	1942	4702	8289	15119	21524	26220	40488

Tier-1 Wide Area Network

- BNL's current WAN bandwidth provided by ESnet is 20 Gbps
  - 10 Gpbs best effort IP (shared by entire lab) and 10 Gbps Lightpath to CERN (ATLAS)
  - 10 Gpbs will be added for ATLAS Tier-0 / Tier-1 connectivity this month
  - Excellent technical support from BNL/ITD Networking and ESnet
- Connection has now the desired redundancy and diversity between NY and BNL