US ATLAS Computing Operations

Kaushik De

University of Texas At Arlington

U.S. ATLAS Tier 2/Tier 3 Workshop, UTA

November 10, 2009

Overview

- We expect the LHC to start in a few weeks
- ATLAS is ready – after 15 years of preparations
- As soon as collisions start, the focus will be on physics
- The distributed computing infrastructure must perform
  - US facilities are required to provide about one quarter of ATLAS computing (though historically we have often provided one third)
  - US primarily responsible for PanDA software used ATAS wide
  - We have done many readiness exercises during the past couple of years – with excellent success, learning from each exercise
  - But the stress on the system will be far greater when data arrives
  - We have to adapt quickly to circumstances, as they arise

Facilities Organization

- See Michael Ernst's talk for overview
- Integration program covered in Rob Gardner's talk
- Operations activity started 1.5 years ago

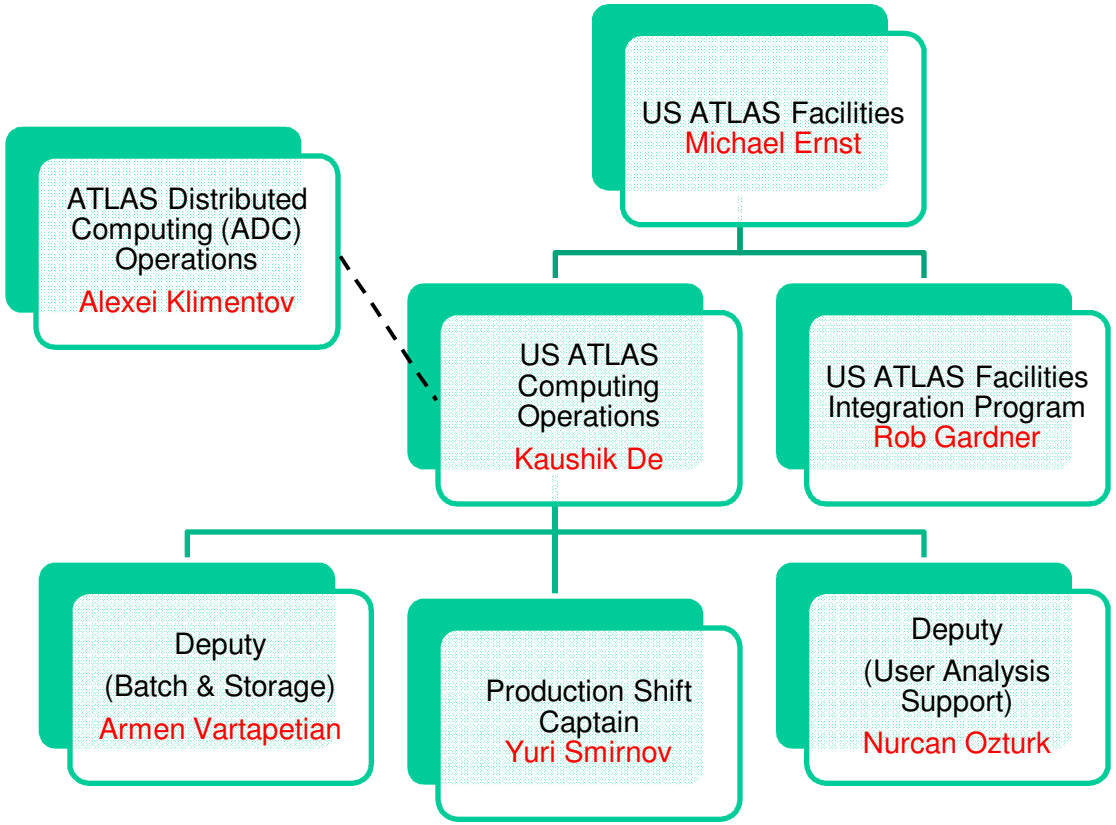

Operations Checklist

- Data production – MC, reprocessing
- Data management – storage, distribution
- User analysis
- All three common areas rely on smooth site operations

US Production - Steady

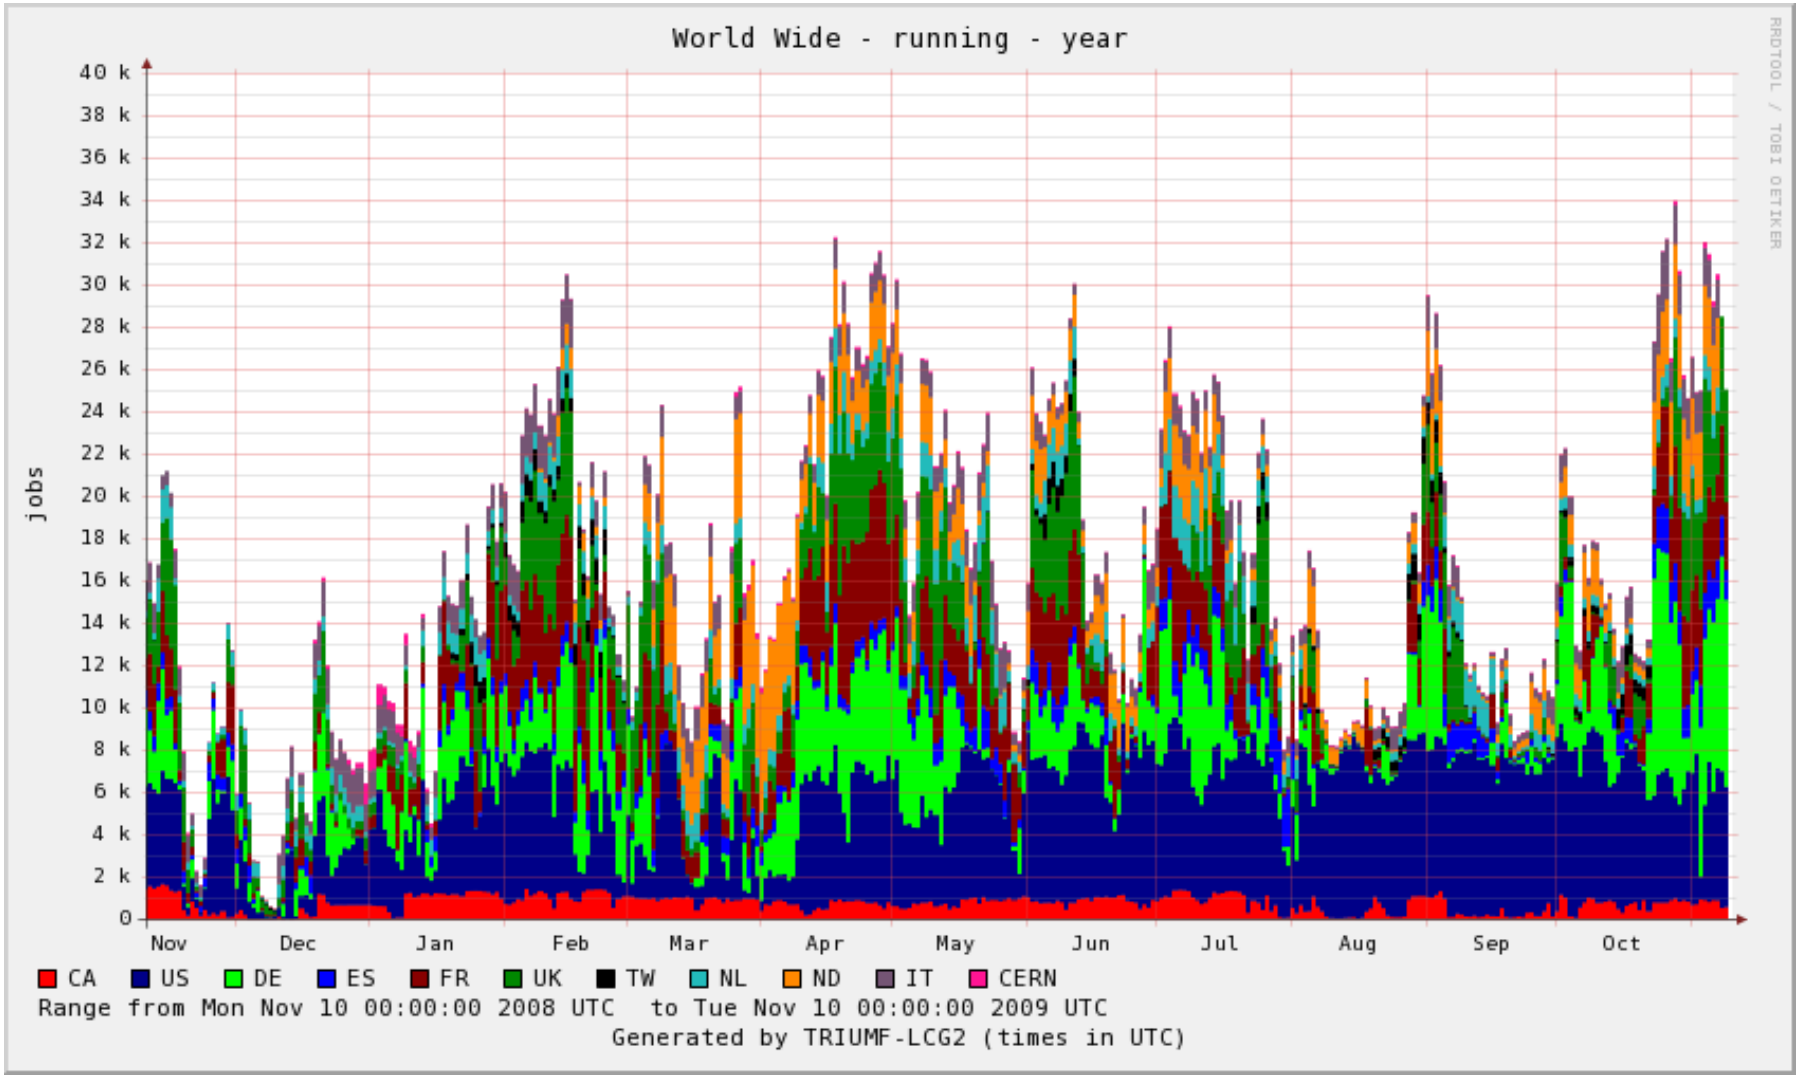

U.S. Production Shares

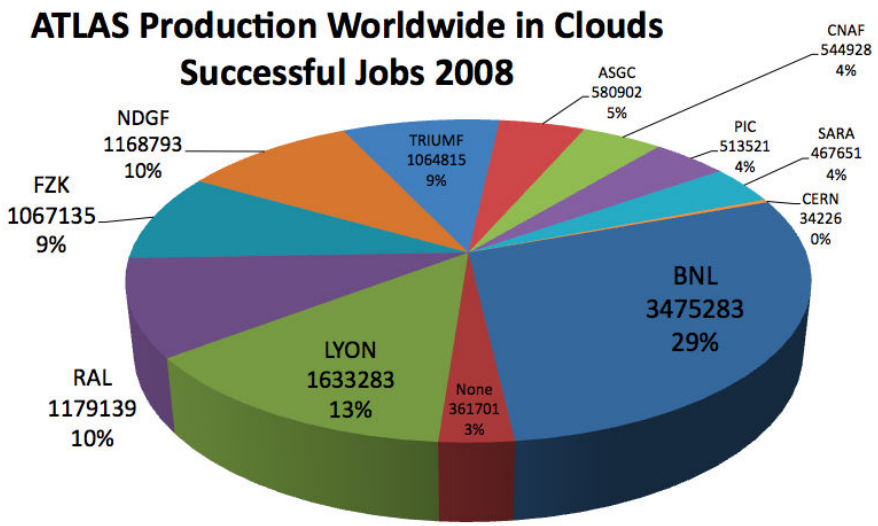

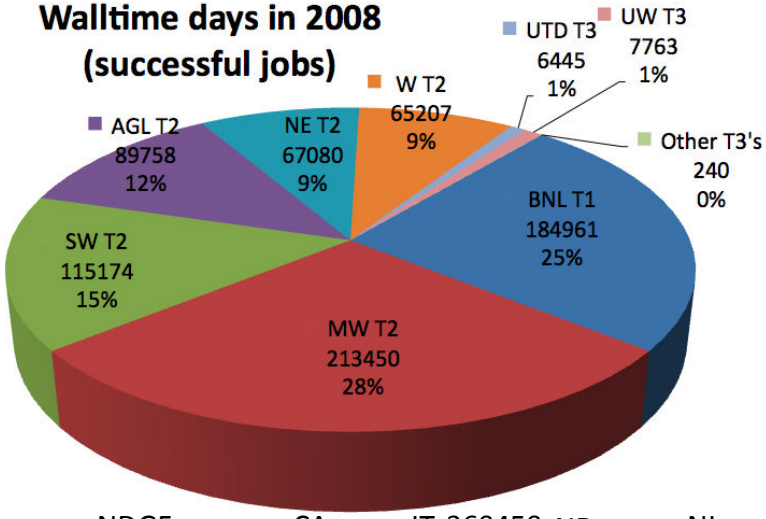

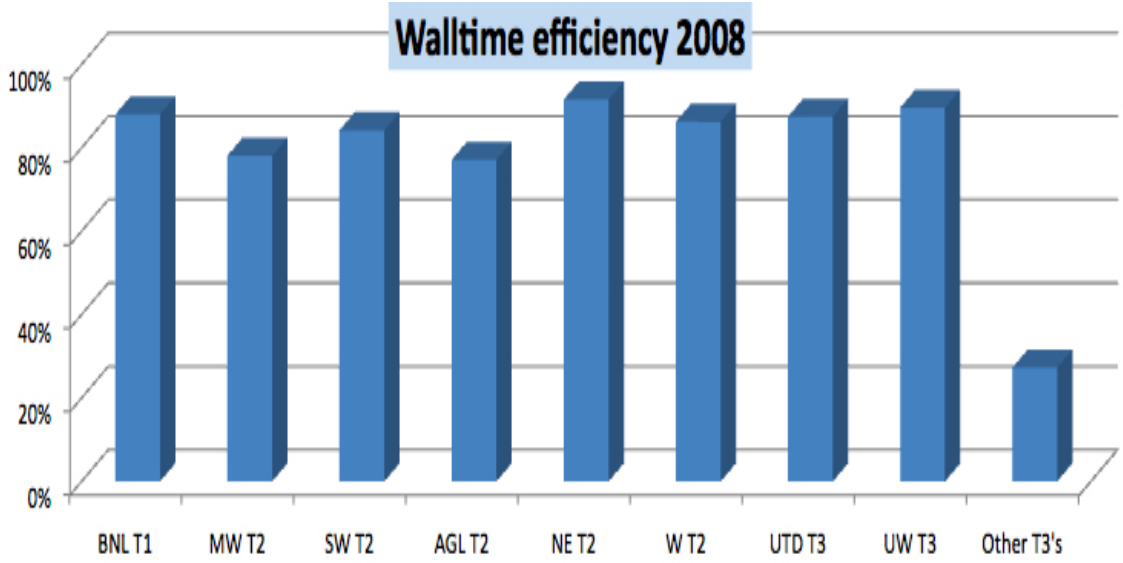

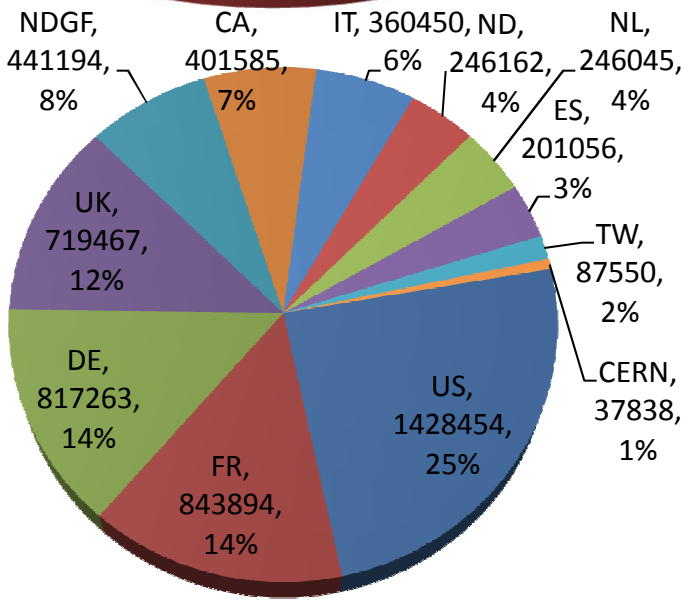

US production Q1, 2009

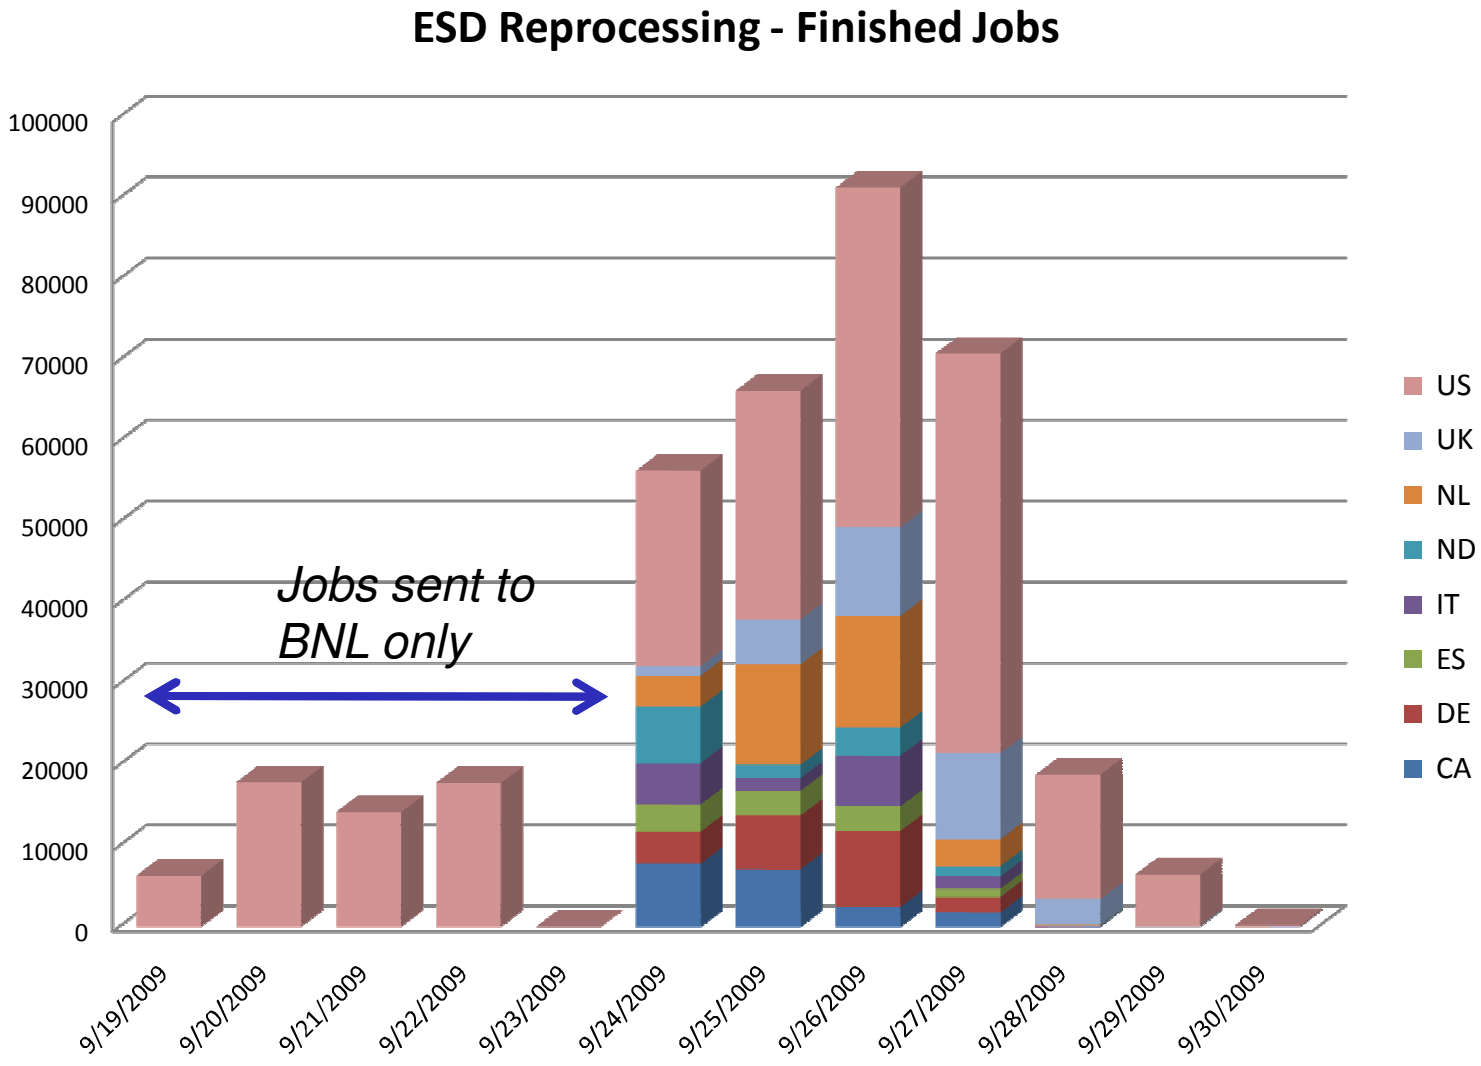
ESD Reprocessing - Finished Jobs
100000
90000
80000
70000
60000
50000
40000
30000
20000
10000
0
Jobs sent to BNL only
9/19/2009
9/20/2009
9/21/2009
9/22/2009
9/23/2009
9/24/2009
9/25/2009
9/26/2009
9/27/2009
9/28/2009
9/29/2009
9/30/2009
US
UK
NL
ND
IT
ES
DE
CA

ESD Reprocessing by Cloud

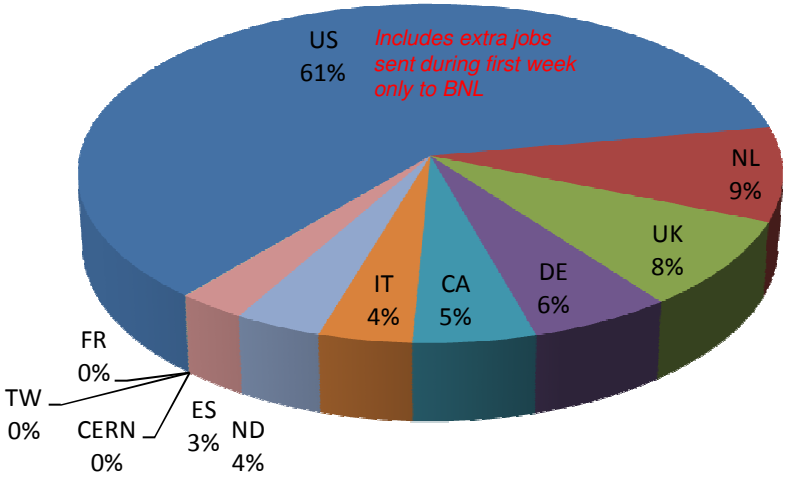

ESD Reprocessing US SIte s

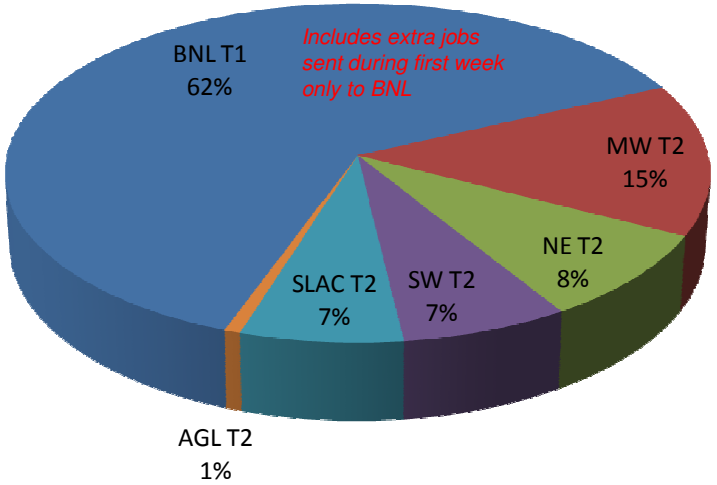

ESD Reprocessing Jobs Sep. 2009

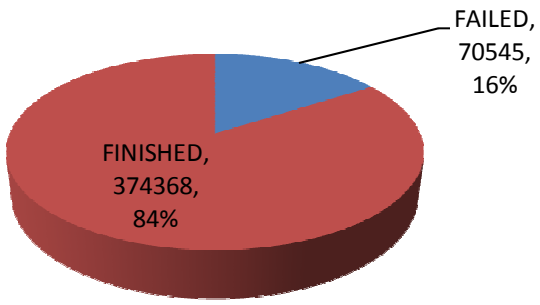

Job Error Summary From all Sites

ESD Reprocessing Errors

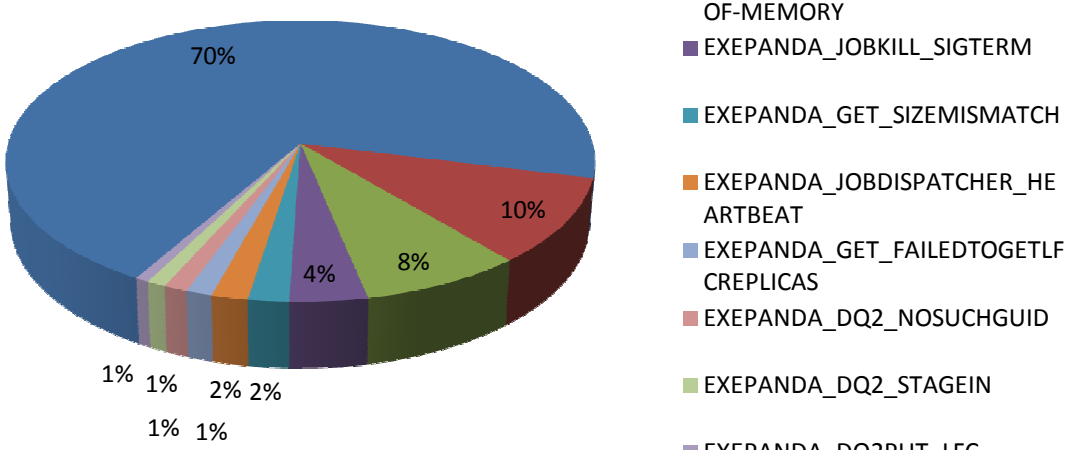

Top 10 errors only

ESD Reprocessing Job Error Rates (Failed/Successful)

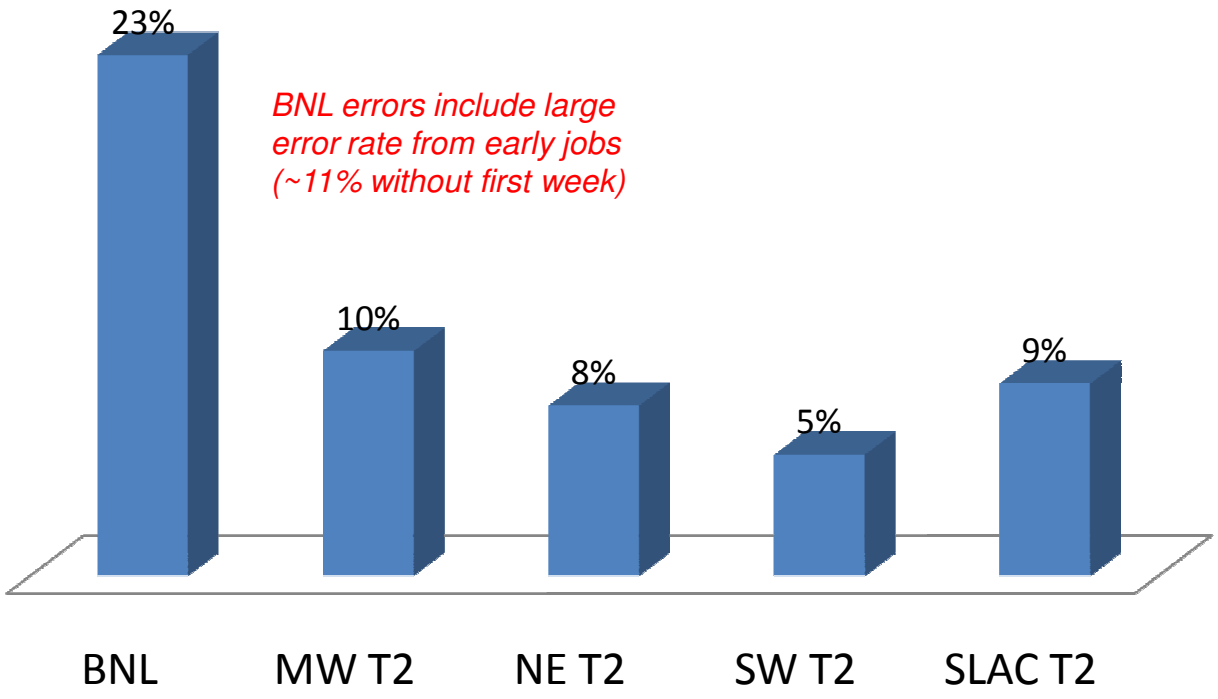

Errors at BNL T1

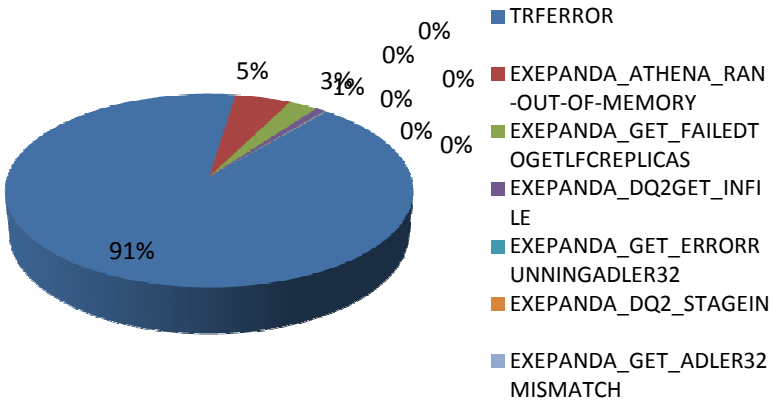

Errors at MW T2

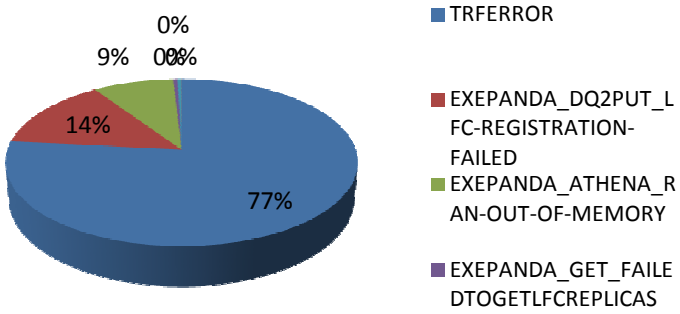

Errors at NE T2

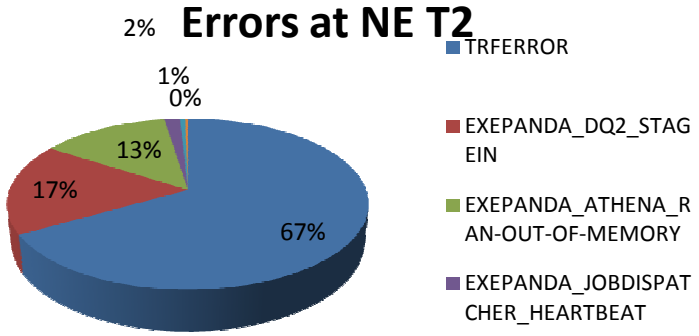

Errors at SW T2

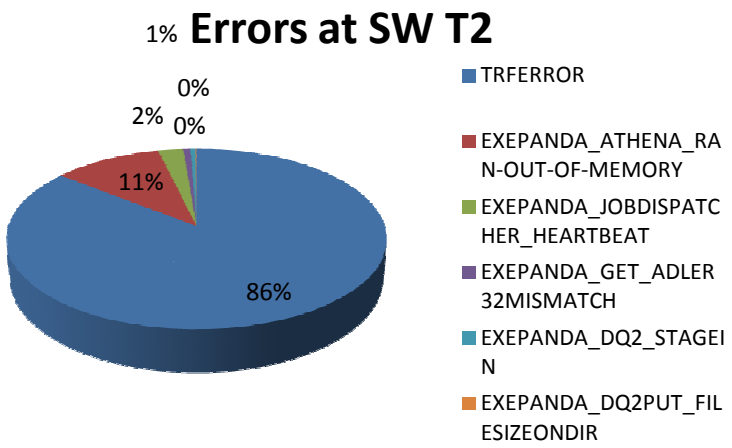

Errors at SLAC T2

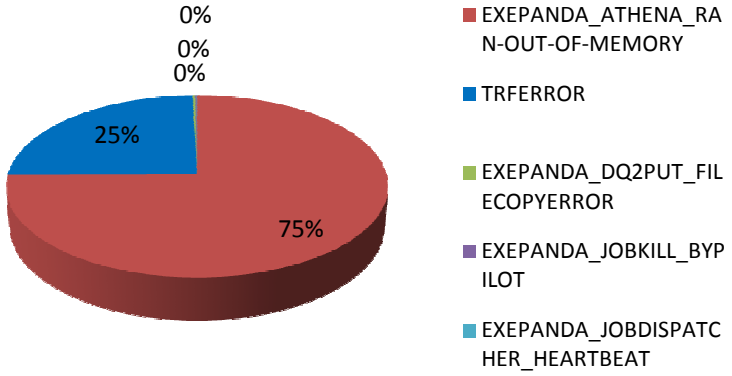

ADCoS (ADC Operations Shifts)

- ADCoS combined shifts started January 28th, 2008
  - Coordinated by K. De and Xavier Espinal (PIC/IFAE)
- ADCoS Goals
  - World-wide (distributed/remote) shifts
  - To monitor all ATLAS distributed computing resources
  - To provide Quality of Service (QoS) for all data processing
- Organization
  - Senior/Trainee: 2 day shifts, Expert: 7 day shifts
  - Three shift times (in CERN time zone):
    - ASIA/Pacific: 0h - 8h
    - EU-ME: 8h - 16h
    - Americas: 16h - 24h
- U.S. shift team
  - In operation long before ADCoS was started
  - Yuri Smirnov (captain), Mark Sosebee, Wensheng Deng, Barry Spurlock, Armen Vartapetian, Rupam Das

Round-the-clock shifts (data taking)

Phase III: 5 Shifters on Duty (+ Trainees) - 24h coverage

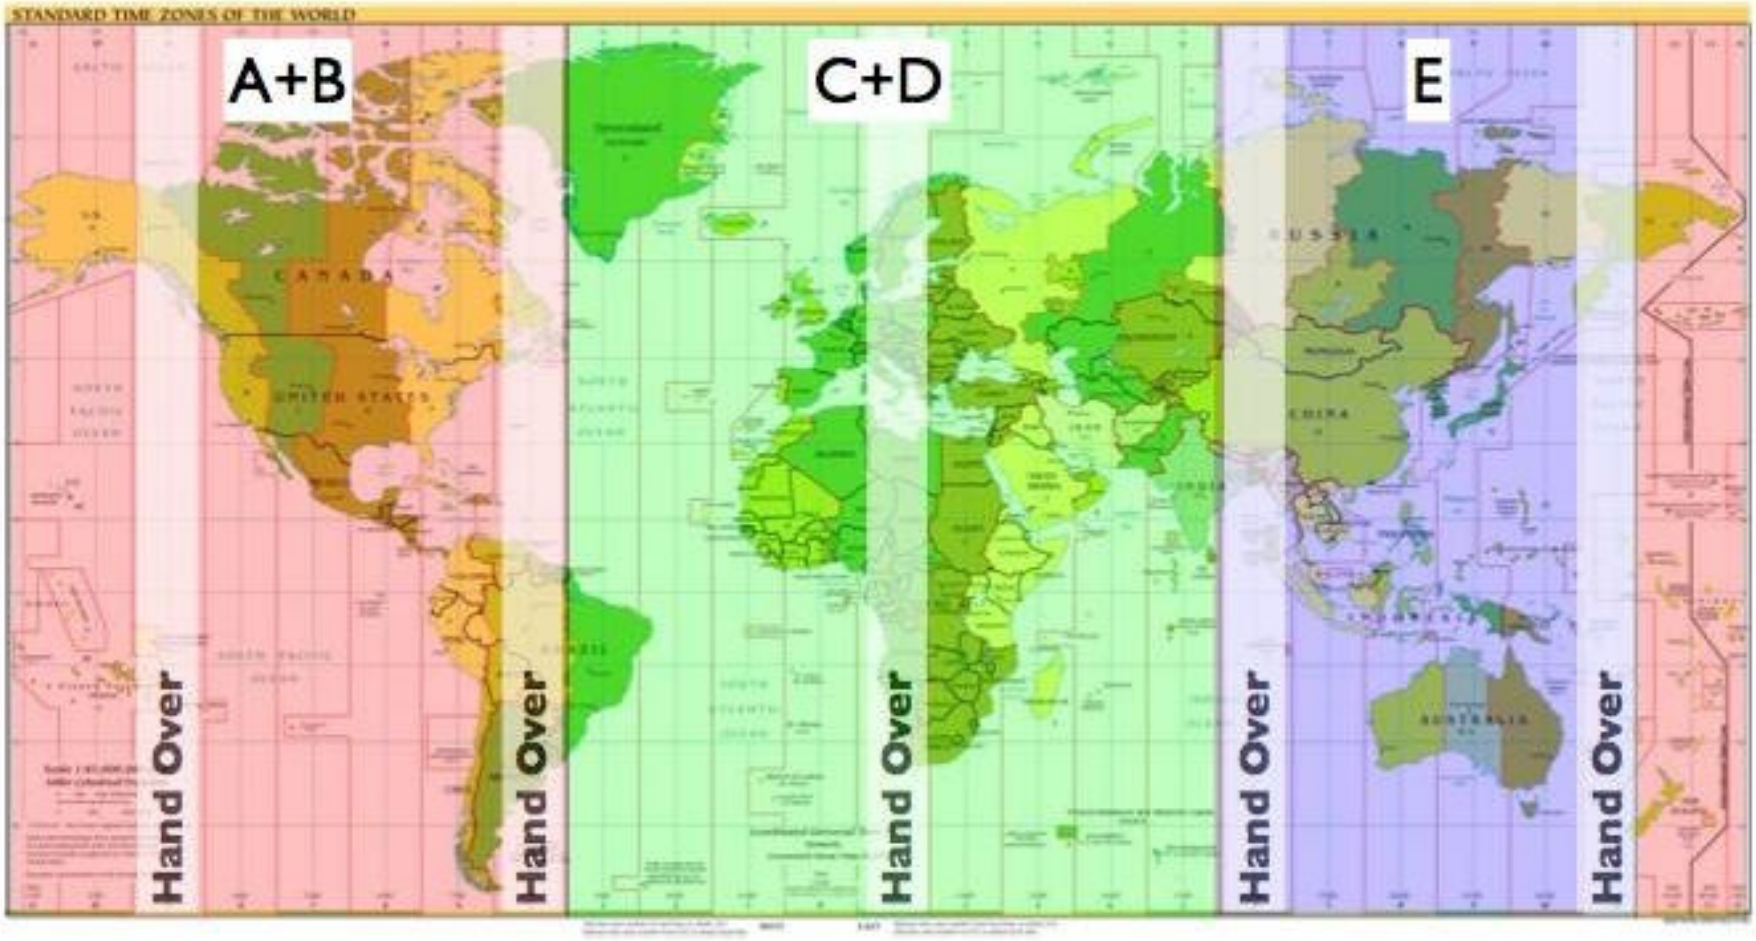

Storage Issues

- US will have ~10 PB by Q1 2010
  - Already have 4-5 PB deployed
  - Fast ramp-up needed
- Space token management
  - Each site must provide 6-10 different storage partitions (tokens)
  - This is quite labor intensive – ADC trying to automate
  - Need to decide soon about group data placement and policies
  - Good management of space tokens is essential to physics analysis

Storage Tokens (~mid Sep09)

Site	HOTDISK	DATADISK	MCDISK	PRODDISK	USERDISK	SCRATCH DISK	GROUP DISK	LOCAL GROUP DISK
BNL	5 TB	767/935 TB	953/1078	No	0/16	3/16	4/23	2/8
AGLT2	No	43/110 TB	127/138	17/23	16/18	0/21	1/17	No
MWT2-UC	No	52/100 TB	127/181	27/40	12/30	0/2	2/10	No
NET2	No	Total 173/245 TB	Yes	Yes	Yes	Yes	Yes	No
SLACT2	Yes	Total 203/230 TB	Yes	Yes	Yes	No	Yes	No
SWT2	No	Total 191/229 TB	Yes	Yes	Yes	No	Yes	No

File Transfers - Steady

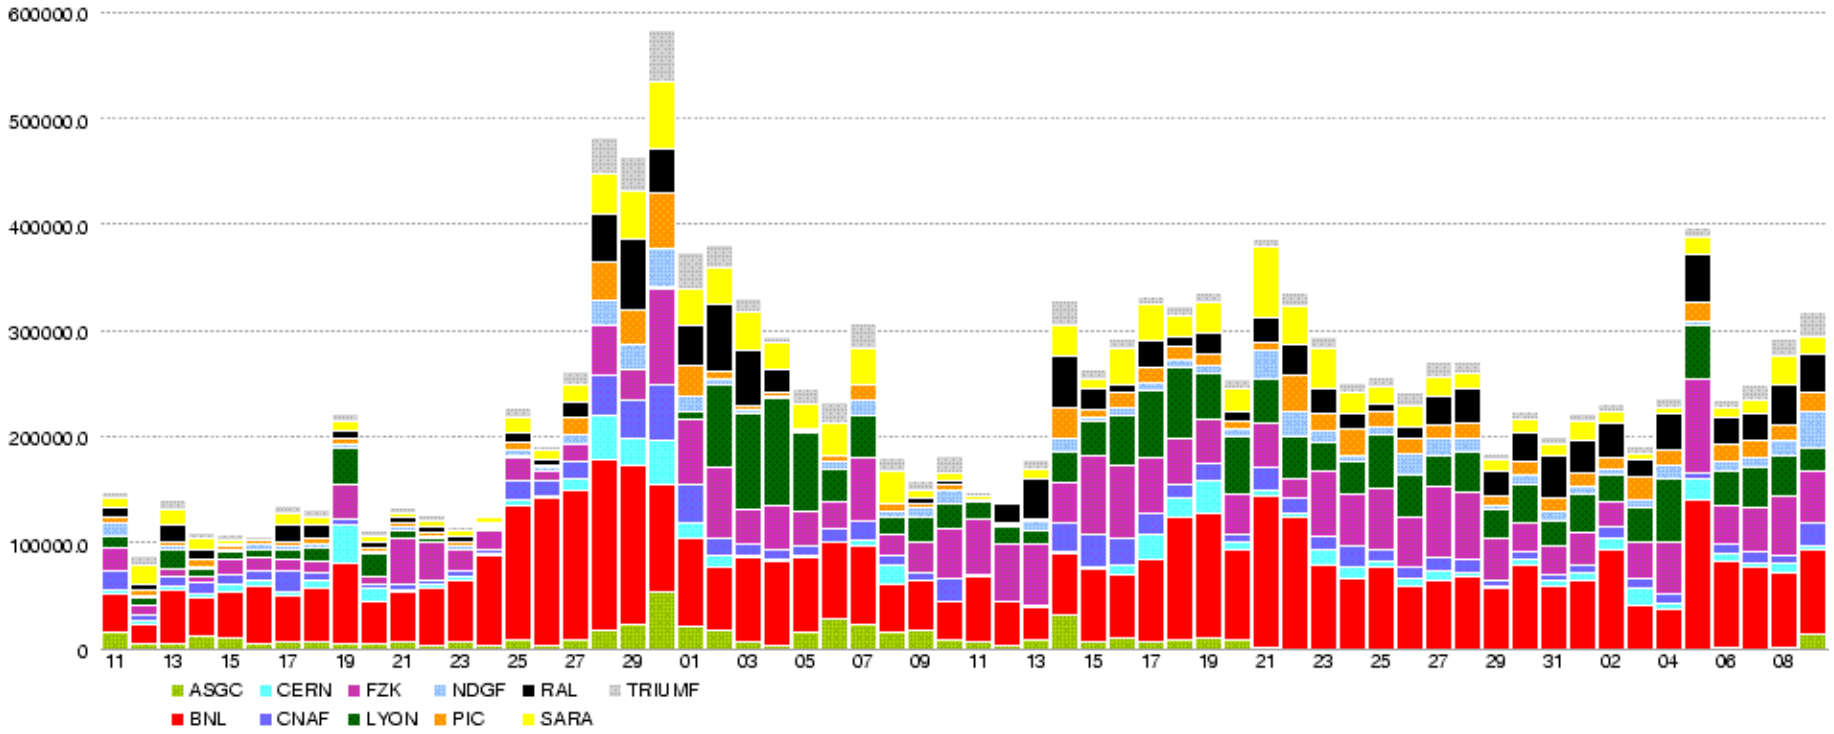

User Analysis

- U.S. ATLAS has excellent track record of supporting users
  - US is the most active cloud in ATLAS for user analysis
  - Analysis sites are in continuous and heavy use for >2 years
  - We have regularly scaled up resources to match user needs
  - UAT09 was very important as a readiness exercise
  - Tier 3 issues will be discussed tomorrow

Analysis Usage Growing

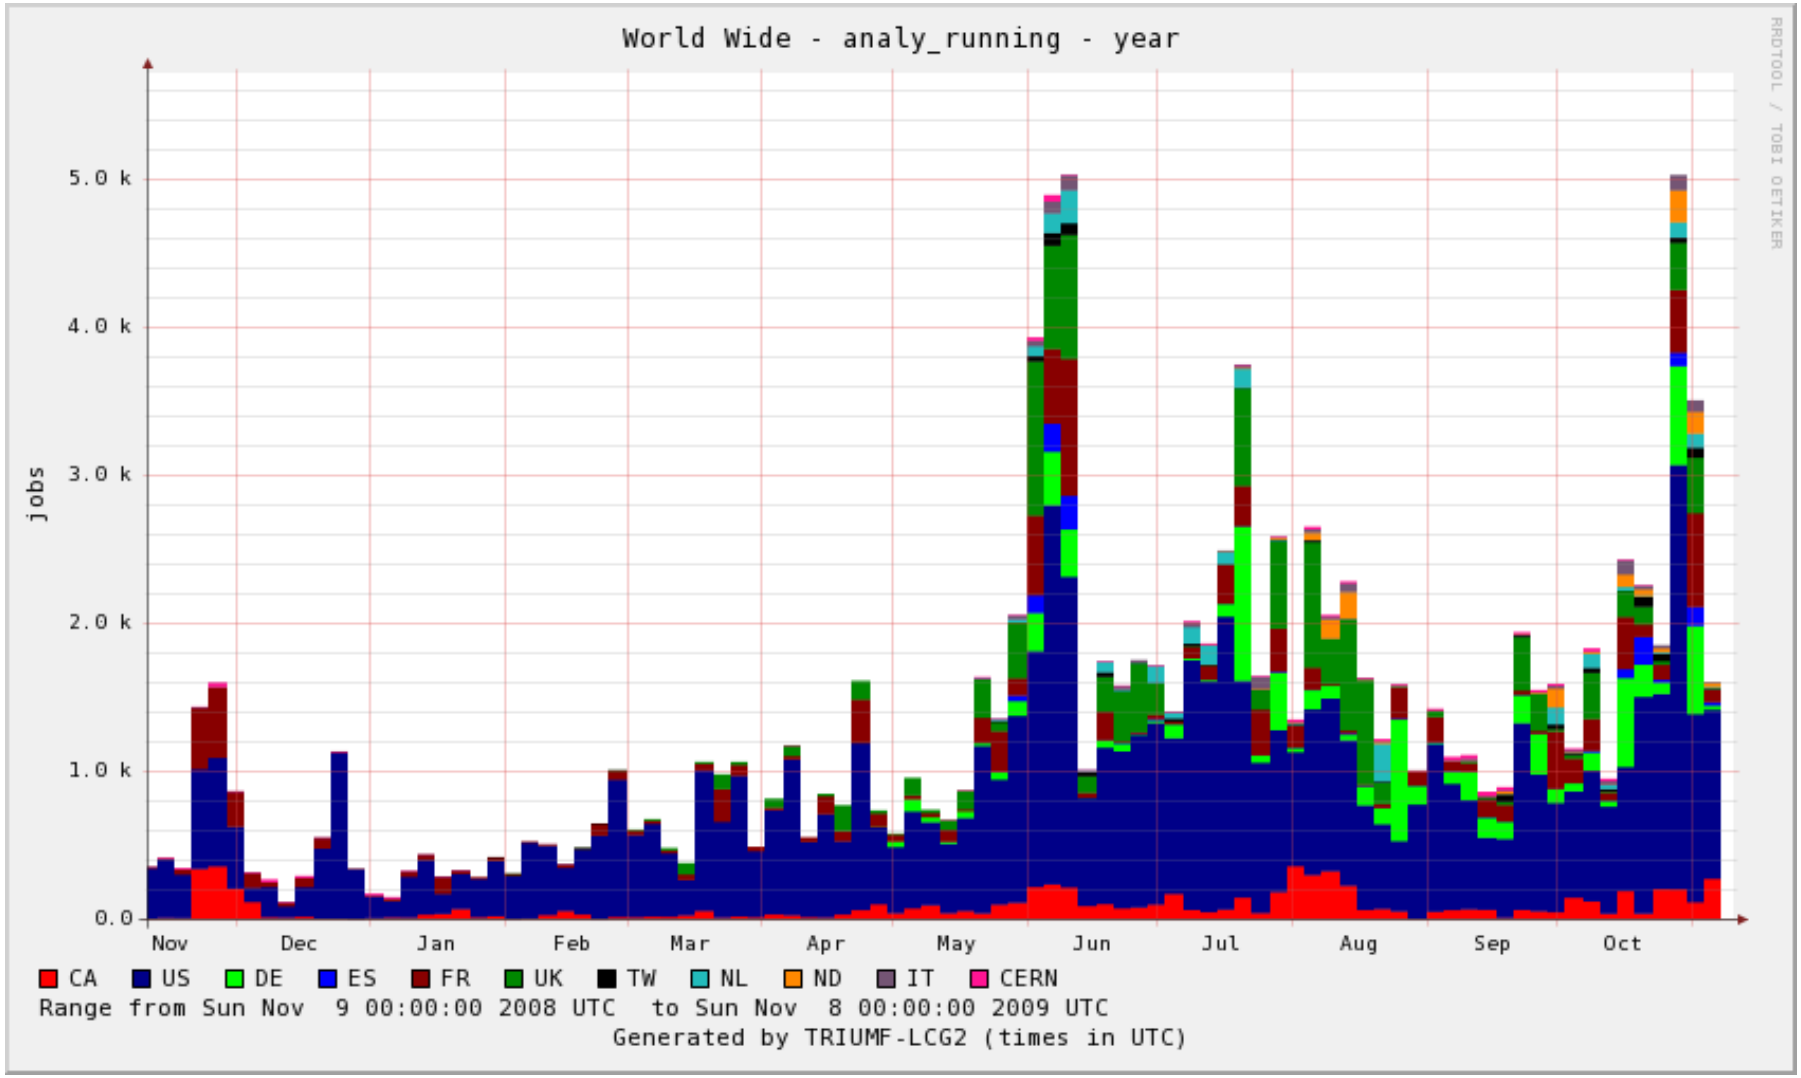

Growing Tier 2 Activity

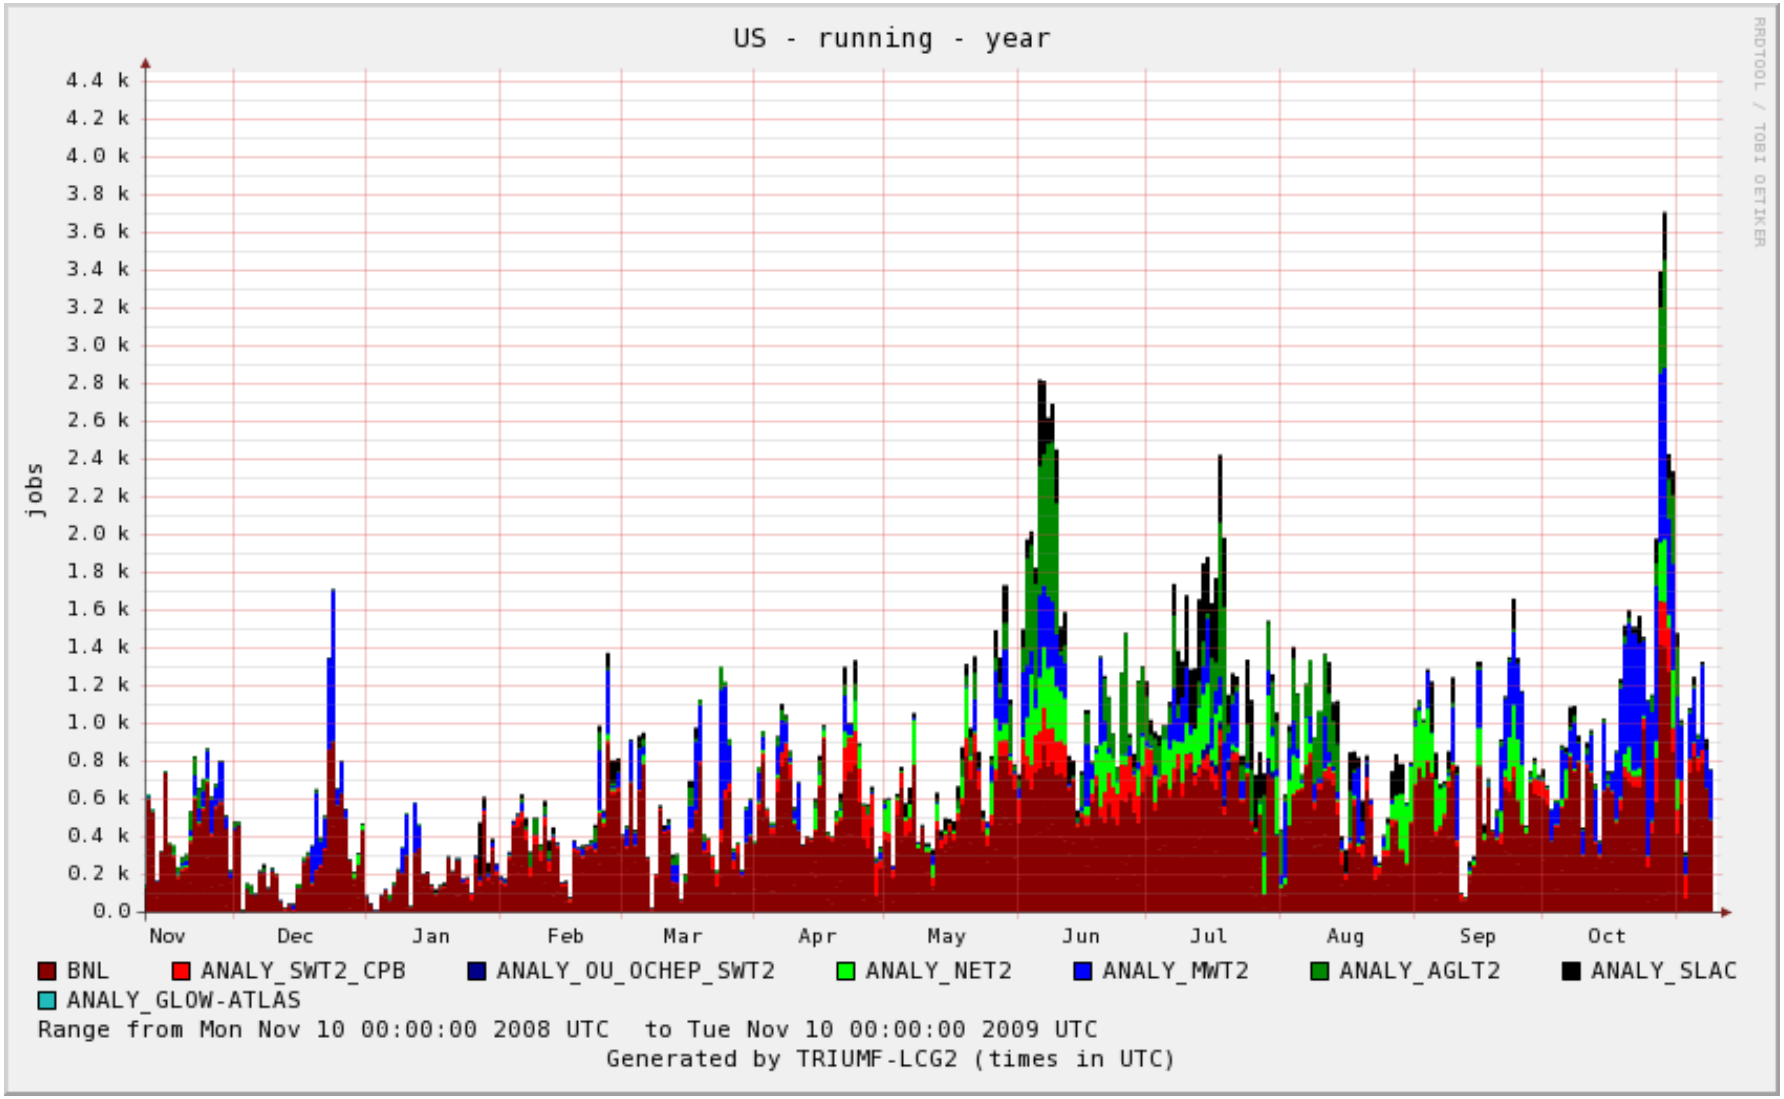

User Analysis Test - UAT09

- 520 M events, ~75 TB AOD's, generated during summer 2009 in US cloud, SM sample with jet Pt cut
- Distributed to all clouds
- Intensively analyzed by ~100 users worldwide, during a 3 day period
- More details in Nurcan's talk tomorrow

UAT09 – Pathena Jobs Worldwide

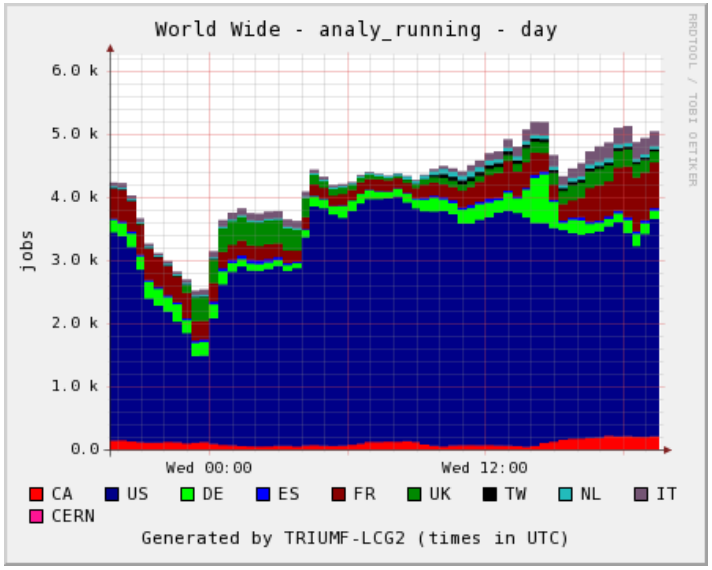

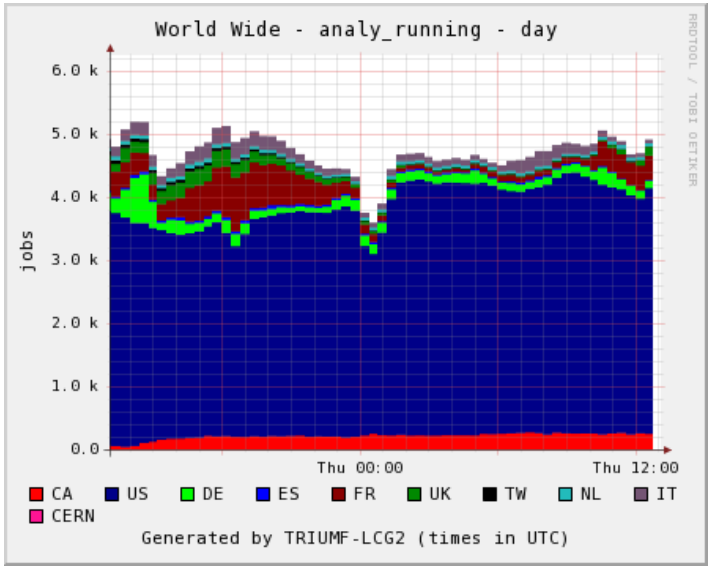

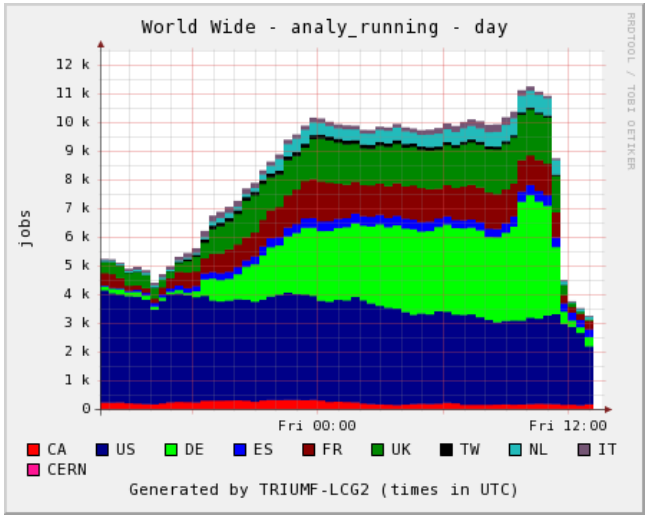

UAT09 – US Sites

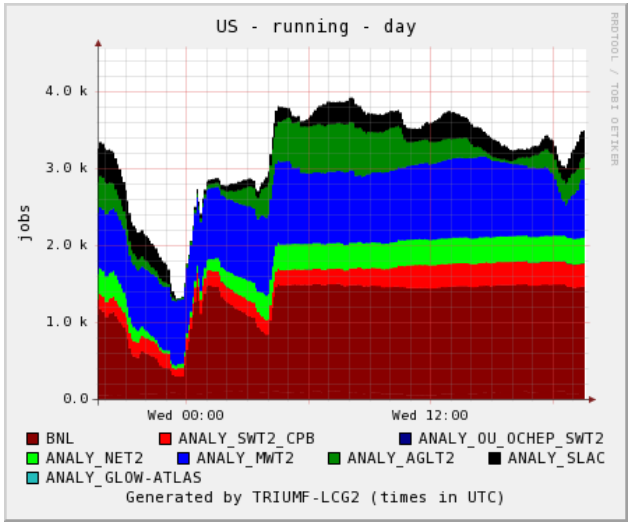

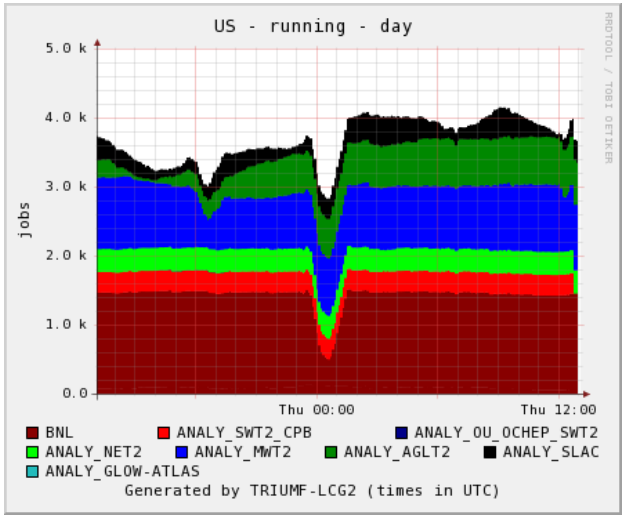

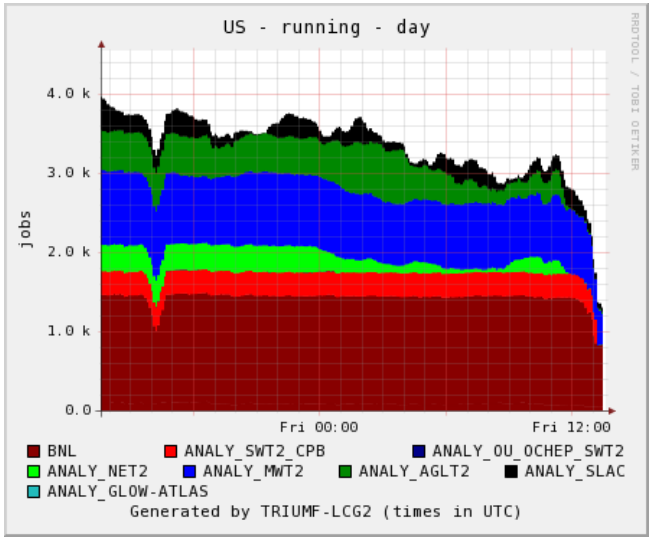

Successfully Finished UAT09 JobCount (PanDA only)

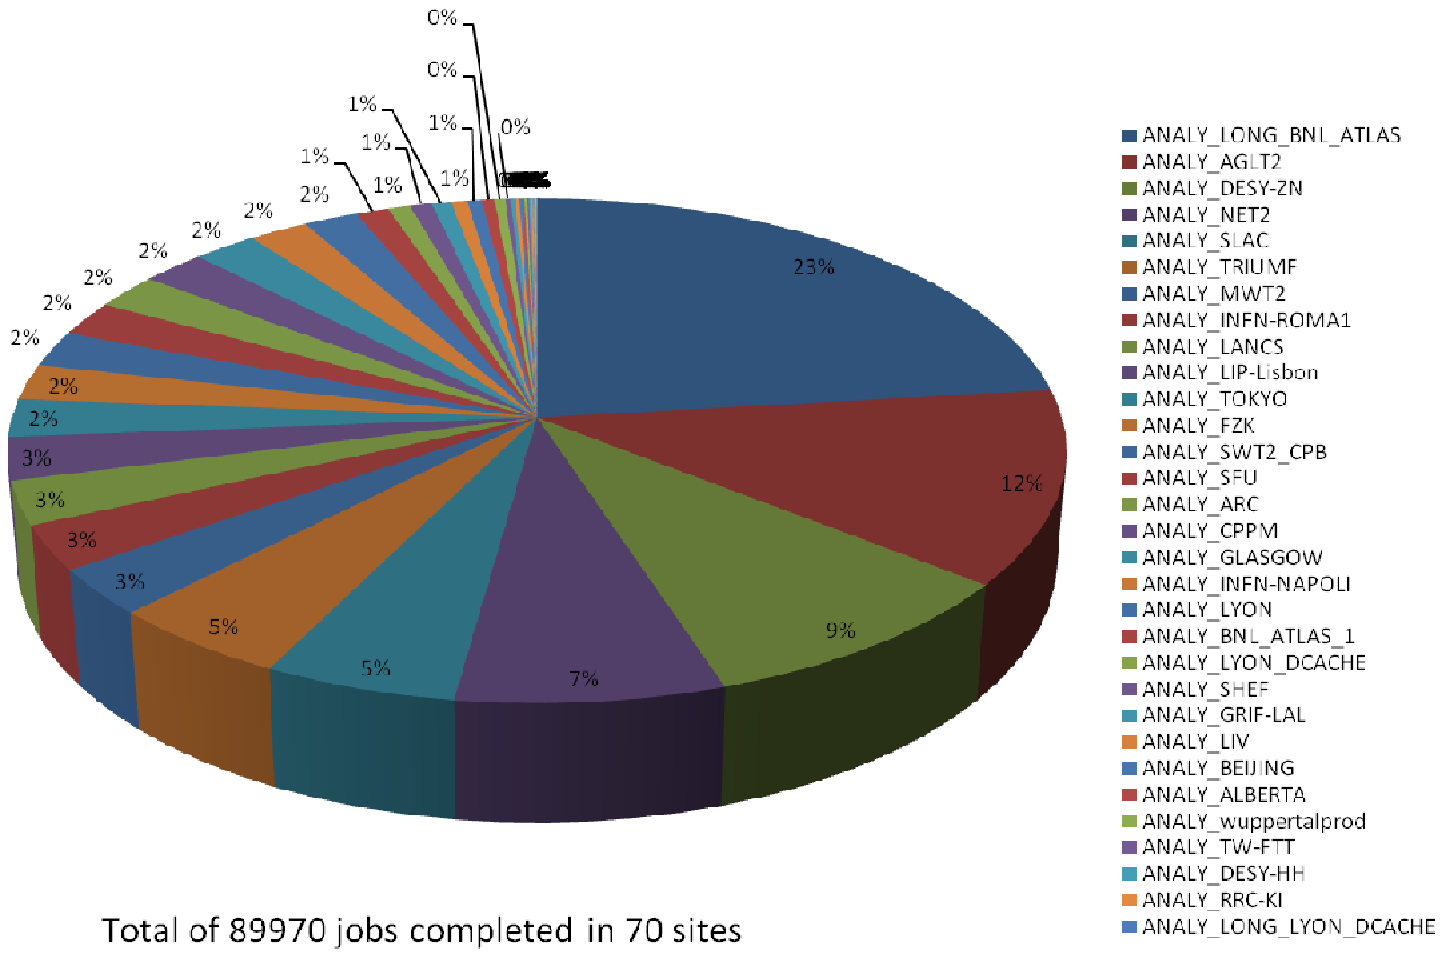

Total of 89970 jobs completed in 70 sites

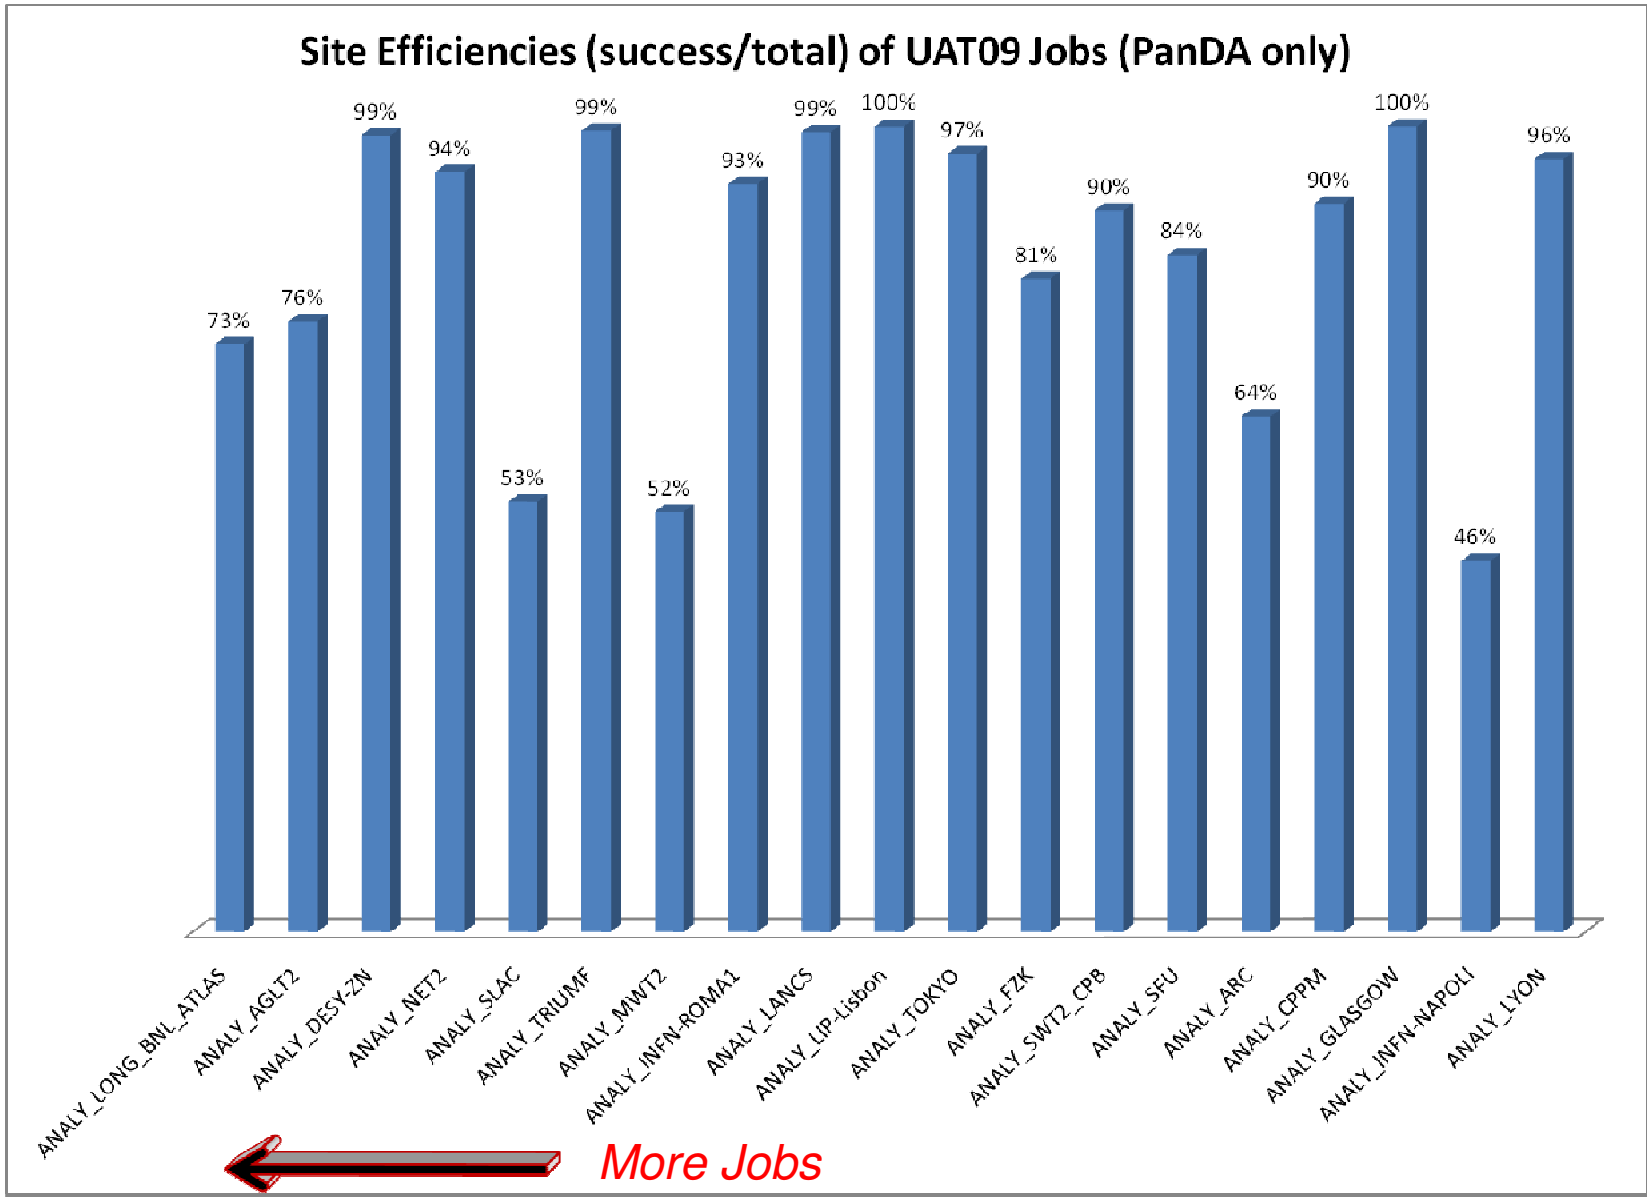
Site Efficiencies (success/total) of UAT09 Jobs (PanDA only)
73%
76%
99%
94%
53%
99%
52%
93%
99%
100%
97%
81%
90%
84%
64%
90%
100%
46%
96%
ANALY_LONG_BNL_ATLAS
ANALY_AGLT2
ANALY_DESY-ZN
ANALY_NET2
ANALY_SLAC
ANALY_TRIUMF
ANALY_MWT2
ANALY_INFN-ROMA1
ANALY_LANCS
ANALY_LIP-Lisbon
ANALY_TOKYO
ANALY_FZK
ANALY_SWT2_CPB
ANALY_SFU
ANALY_ARC
ANALY_CPPM
ANALY_GLASGOW
ANALY_INFN-NAPOLI
ANALY_LYON
More Jobs

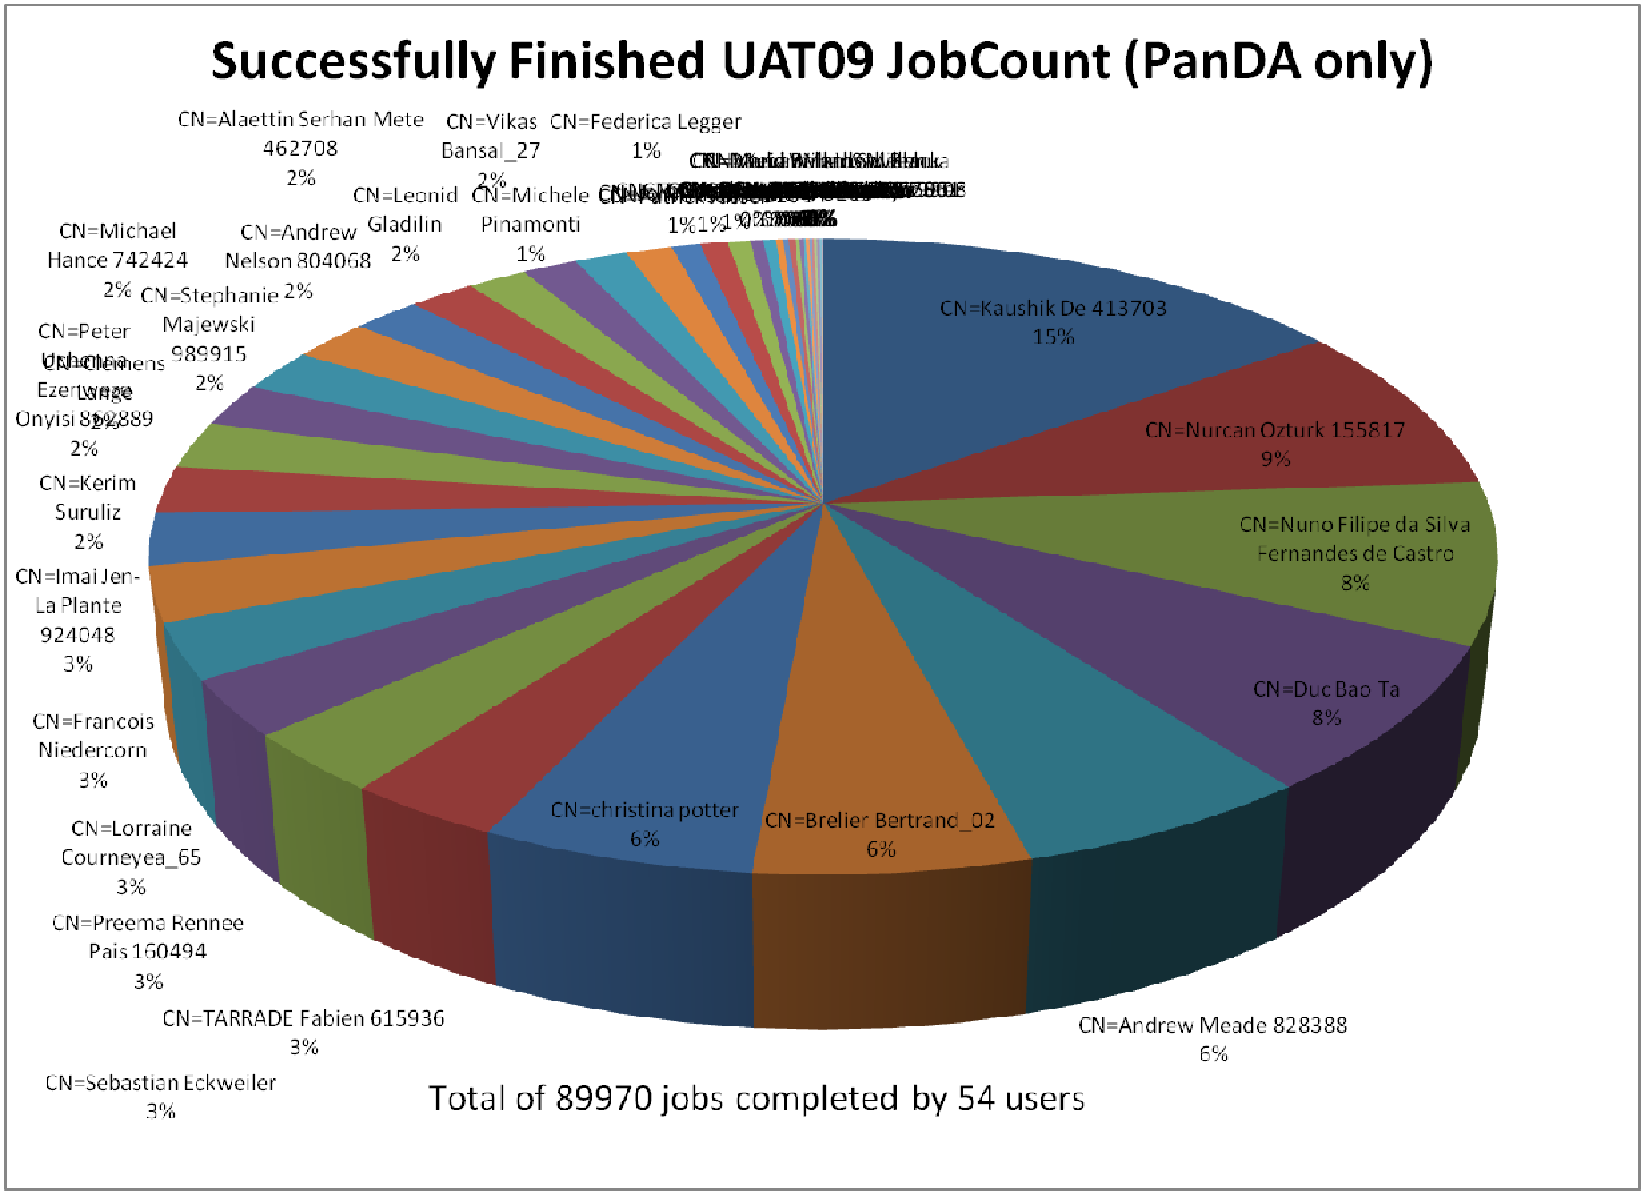
Successfully Finished UAT09 JobCount (PanDA only)
CN=Alaettin Serhan Mete 462708 2%
CN=Vikas Bansal_27 2%
CN=Federica Legger 1%
CN=Leonid Gladilin 2%
CN=Michele Pinamonti 1%
CN=Michael Hance 742424 2%
CN=Andrew Nelson 804068 2%
CN=Stephanie Majewski 989915 2%
CN=Peter
CN=Clemens
Ezeanyim
Onyisi 892889 2%
CN=Kerim Suruliz 2%
CN=Imai Jen-La Plante 924048 3%
CN=Francois Niedercorn 3%
CN=Lorraine Courneyea_65 3%
CN=Preema Rennee Pais 160494 3%
CN=TARRADE Fabien 615936 3%
CN=Sebastian Eckweiler 3%
CN=Kaushik De 413703 15%
CN=Nurcan Ozturk 155817 9%
CN=Nuno Filipe da Silva Fernandes de Castro 8%
CN=Duc Bao Ta 8%
CN=Andrew Meade 828388 6%
CN=Brelier Bertrand_02 6%
CN=christina potter 6%
Total of 89970 jobs completed by 54 users

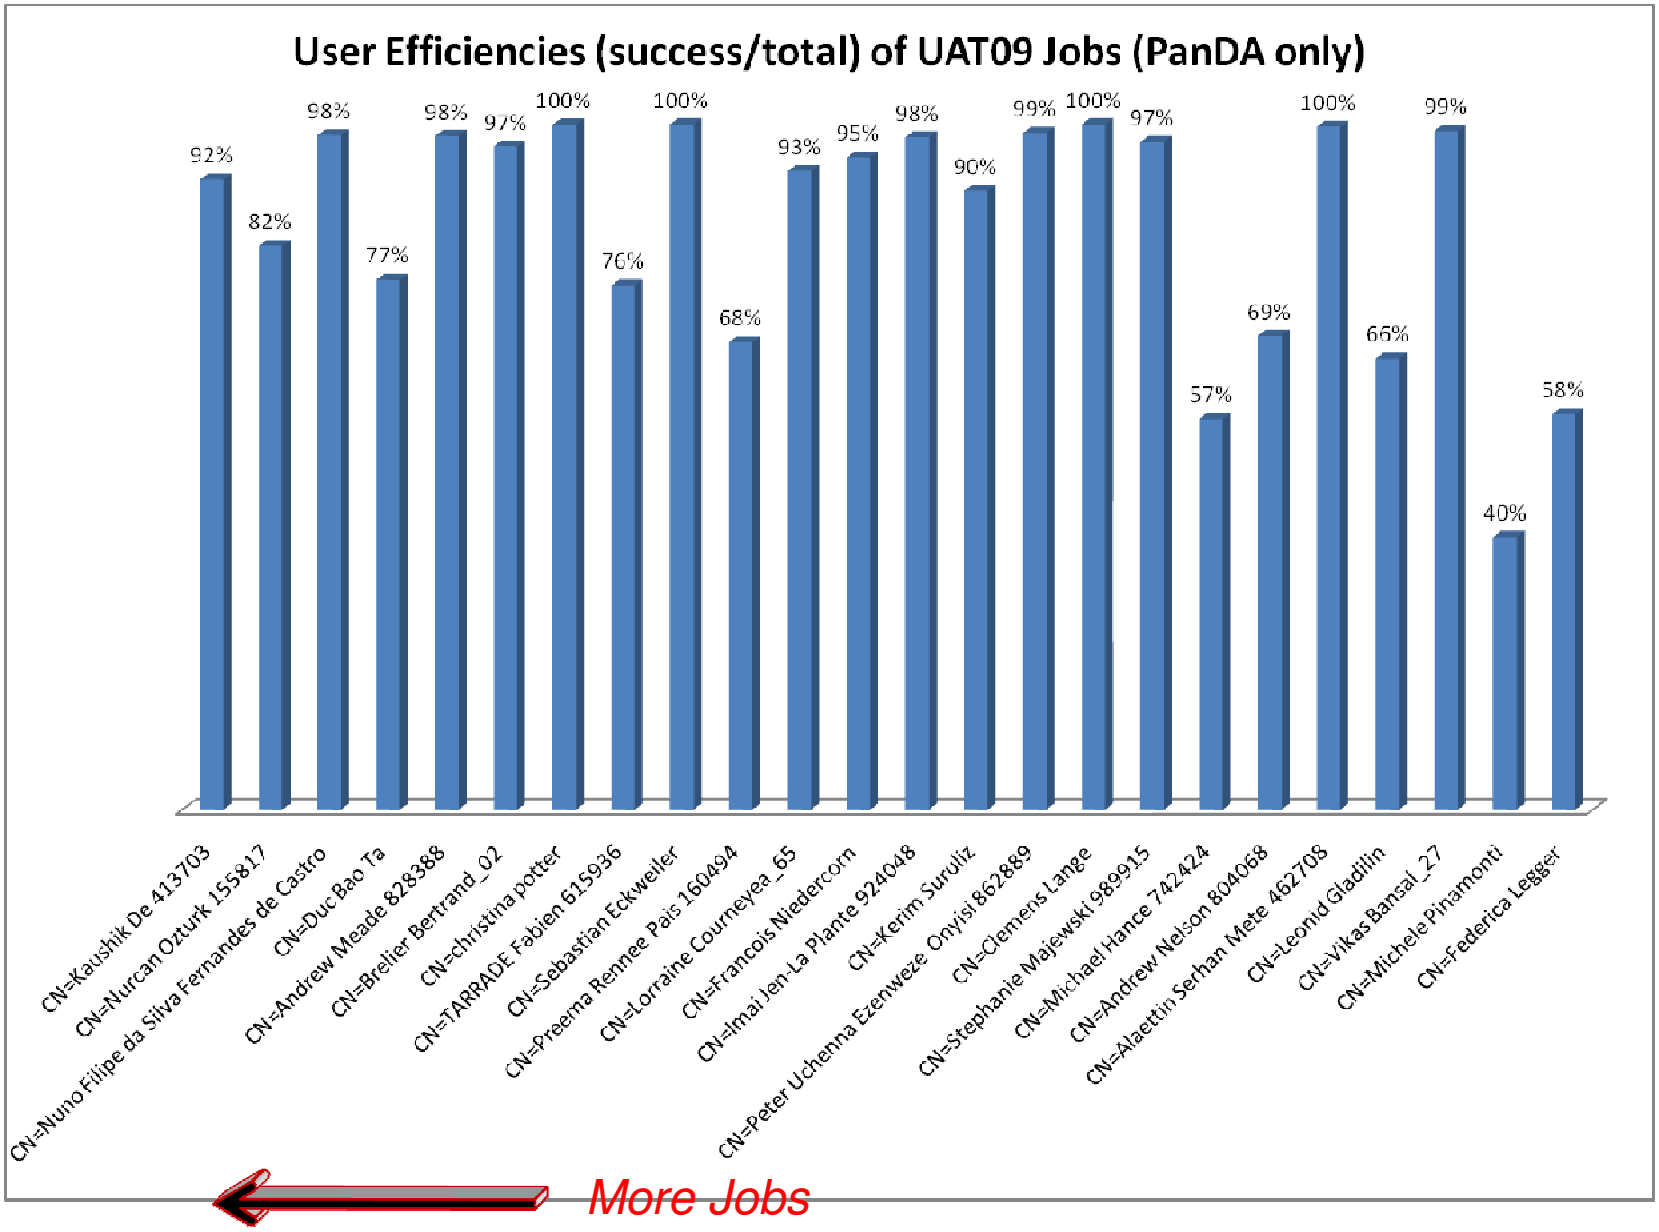
User Efficiencies (success/total) of UAT09 Jobs (PanDA only)
92%
82%
98%
77%
98%
97%
100%
76%
100%
68%
93%
95%
98%
90%
99%
100%
97%
57%
69%
100%
66%
99%
40%
58%
CN=Kaushik De 413703
CN=Nurcan Ozturk 155817
CN=Nuno Filipe da Silva Fernandes de Castro
CN=Duc Bao Ta
CN=Andrew Meade 828388
CN=Brelier Bertrand_02
CN=christina potter
CN=TARRADE Fabien 615936
CN=Sebastian Eckweiler
CN=Preema Rennee Pais 160494
CN=Lorraine Courneyea_65
CN=Francois Niedercorn
CN=Imai Jen-La Plante 924048
CN=Kerim Suruliz
CN=Peter Uchenna Ezenweeze Onyisi 862889
CN=Clemens Lange
CN=Stephanie Majewski 989915
CN=Michael Hance 742424
CN=Andrew Nelson 804068
CN=Alaettin Serhan Mete 462708
CN=Leonid Gladilin
CN=Vikas Bansal_27
CN=Michele Pinamonti
CN=Federica Legger
More Jobs

Distributed Analysis Shift Team – DAST

- User analysis support is provided by the AtlasDAST (Atlas Distributed Analysis Shift Team) since September 29, 2008. Previously, user support was on a best effort basis provided by the Panda and Ganga software developers.
- Nurcan Ozturk (UTA) and Daniel van der Ster (CERN) are coordinating this effort.
- DAST organizes shifts currently in two time zones – US and CERN. One person from each zone is on shift for 7 hours a day covering between 9am-11pm CERN time, and 5 days a week.
- Please contact Nurcan to join this effort

Conclusion

- Waiting for collisions!