

# Next Generation Workload Management System for Big Data on Heterogeneous Distributed Computing

*16th International workshop on Advanced Computing  
and Analysis Techniques in physics research (ACAT)*

A.Klimentov, P.Buncic, K.De, S.Jha, T.Maeno,  
R.P.Mount, P.Nilsson, D.Oleynik, S.Panitkin,  
A.Petrosyan, R.J.Porter, K.F.Read, A.Vaniachine,  
J.Wells and T.Wenaus

# Main topics

- **Introduction**

- Large Hadron Collider at CERN
- LHC data processing challenges
- Evolution of LHC computing model
- Workload Management System core ideas.

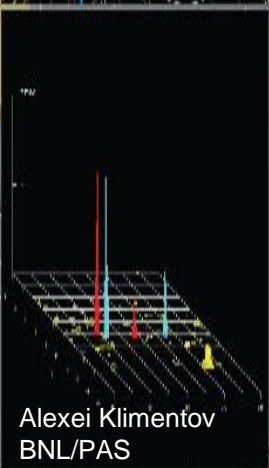
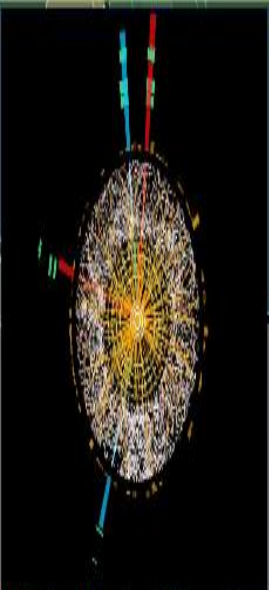
- **PanDA**

- **Next Generation Workload Management System.**

## **BigPanDA project :**

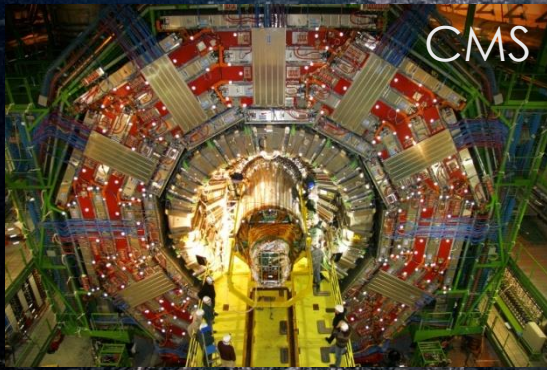
- Workload Management and networking
- Expanding PanDA beyond the Grid and High Energy Physics
- PanDA at Oak Ridge Leadership Computing Facilities

- **Summary**

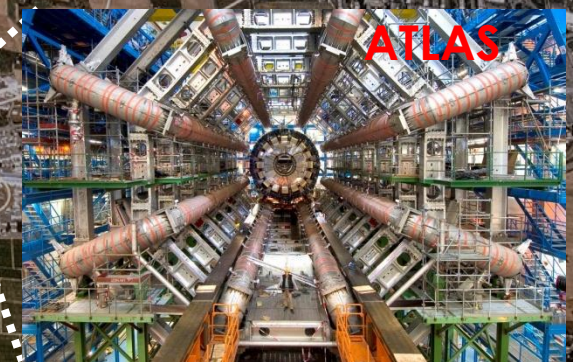


# Enter a New Era in Fundamental Science

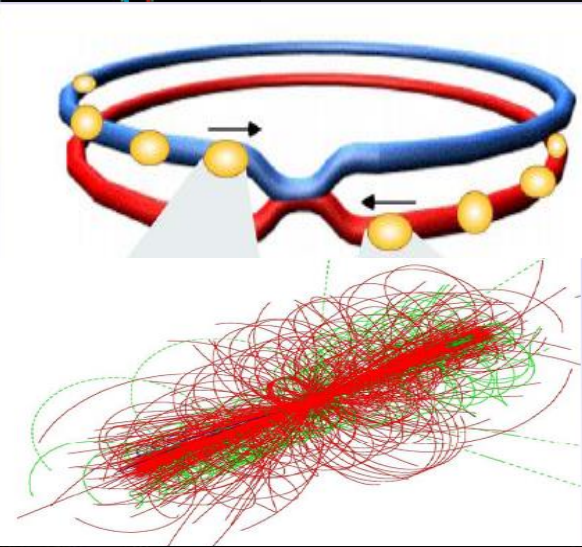
The Large Hadron Collider (**LHC**), one of the largest and truly global scientific projects ever built, is the most exciting turning point in particle physics.



Exploration of a new energy frontier  
Proton-proton and Heavy Ion collisions  
at  $E_{\text{CM}}$  up to 14 TeV



# Proton-Proton Collisions at the LHC



LHC delivered billions of collision events to the experiments from proton-proton and proton-lead collisions in the Run 1 period (2009-2013)

→ **collisions every 50 ns**

**= 20 MHz crossing rate**

▪  **$1.6 \times 10^{11}$  protons per bunch**

at  $L_{pk} \sim 0.8 \times 10^{34} / \text{cm}^2 / \text{s}$

**$\approx 35$  pp interactions per crossing – pile-up**

→  **$\approx 10^9$  pp interactions per second !!!**

▪ **in each collision**

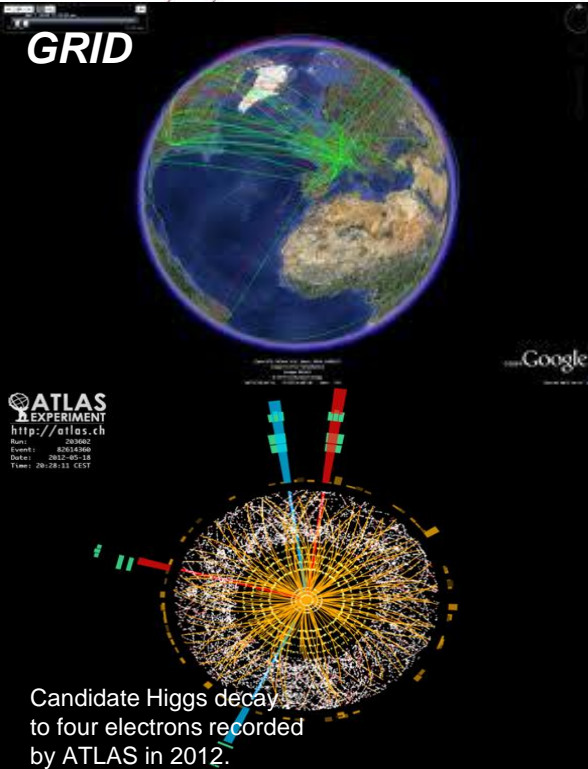
**$\approx 1600$  charged particles produced**

*enormous challenge for the detectors and for data collection/storage/analysis*

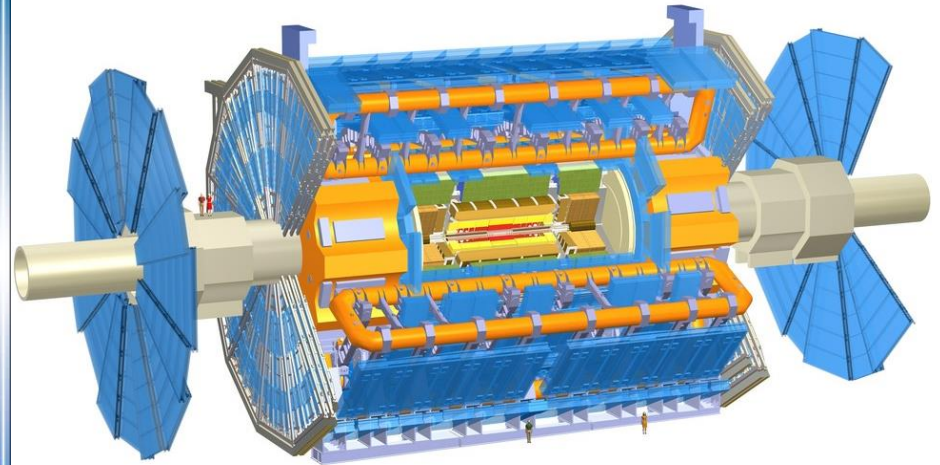
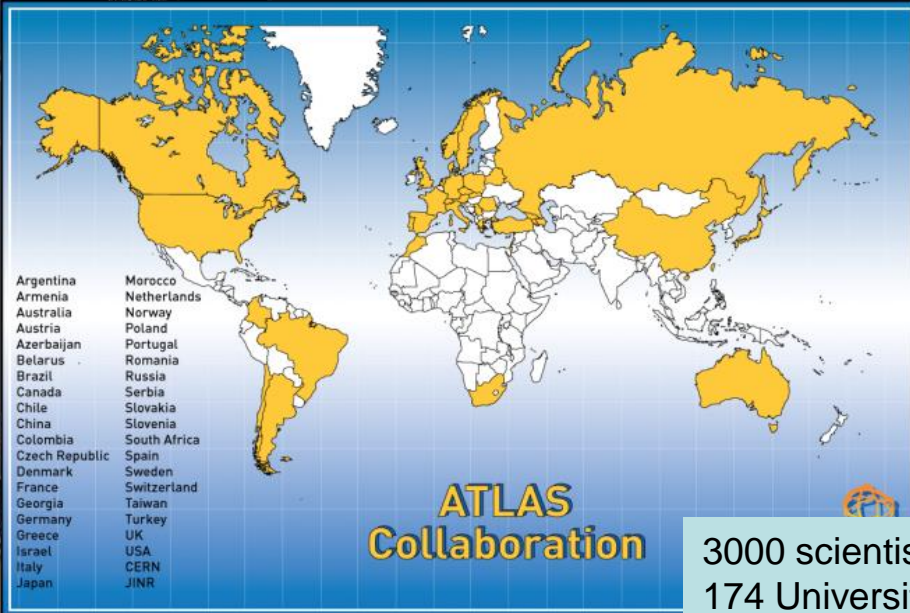
Raw data rate from LHC detector : 1PB/s

This translates to Petabytes of data recorded world-wide (Grid)

The challenge how to process and analyze the data and produce timely physics results was substantial, but at the end resulted in a great success




# The ATLAS Experiment at the LHC



3000 scientists  
174 Universities and Labs  
From 38 countries  
More than 1200 students



 The Nobel Prize in Physics 2013  
François Englert, Peter Higgs

## The Nobel Prize in Physics 2013

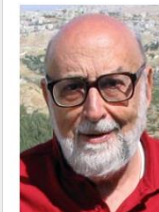


Photo: Pnicolet via Wikimedia Commons  
François Englert

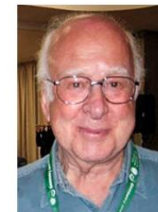
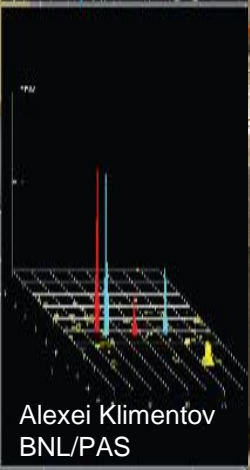
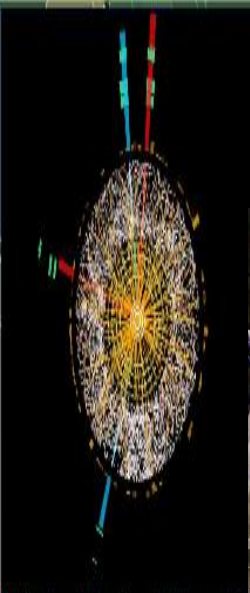


Photo: G-M Greuel via Wikimedia Commons  
Peter W. Higgs

The Nobel Prize in Physics 2013 was awarded jointly to François Englert and Peter W. Higgs "for the theoretical discovery of a mechanism that contributes to our understanding of the origin of mass of subatomic particles, and which recently was confirmed through the discovery of the predicted fundamental particle, by the ATLAS and CMS experiments at CERN's Large Hadron Collider"

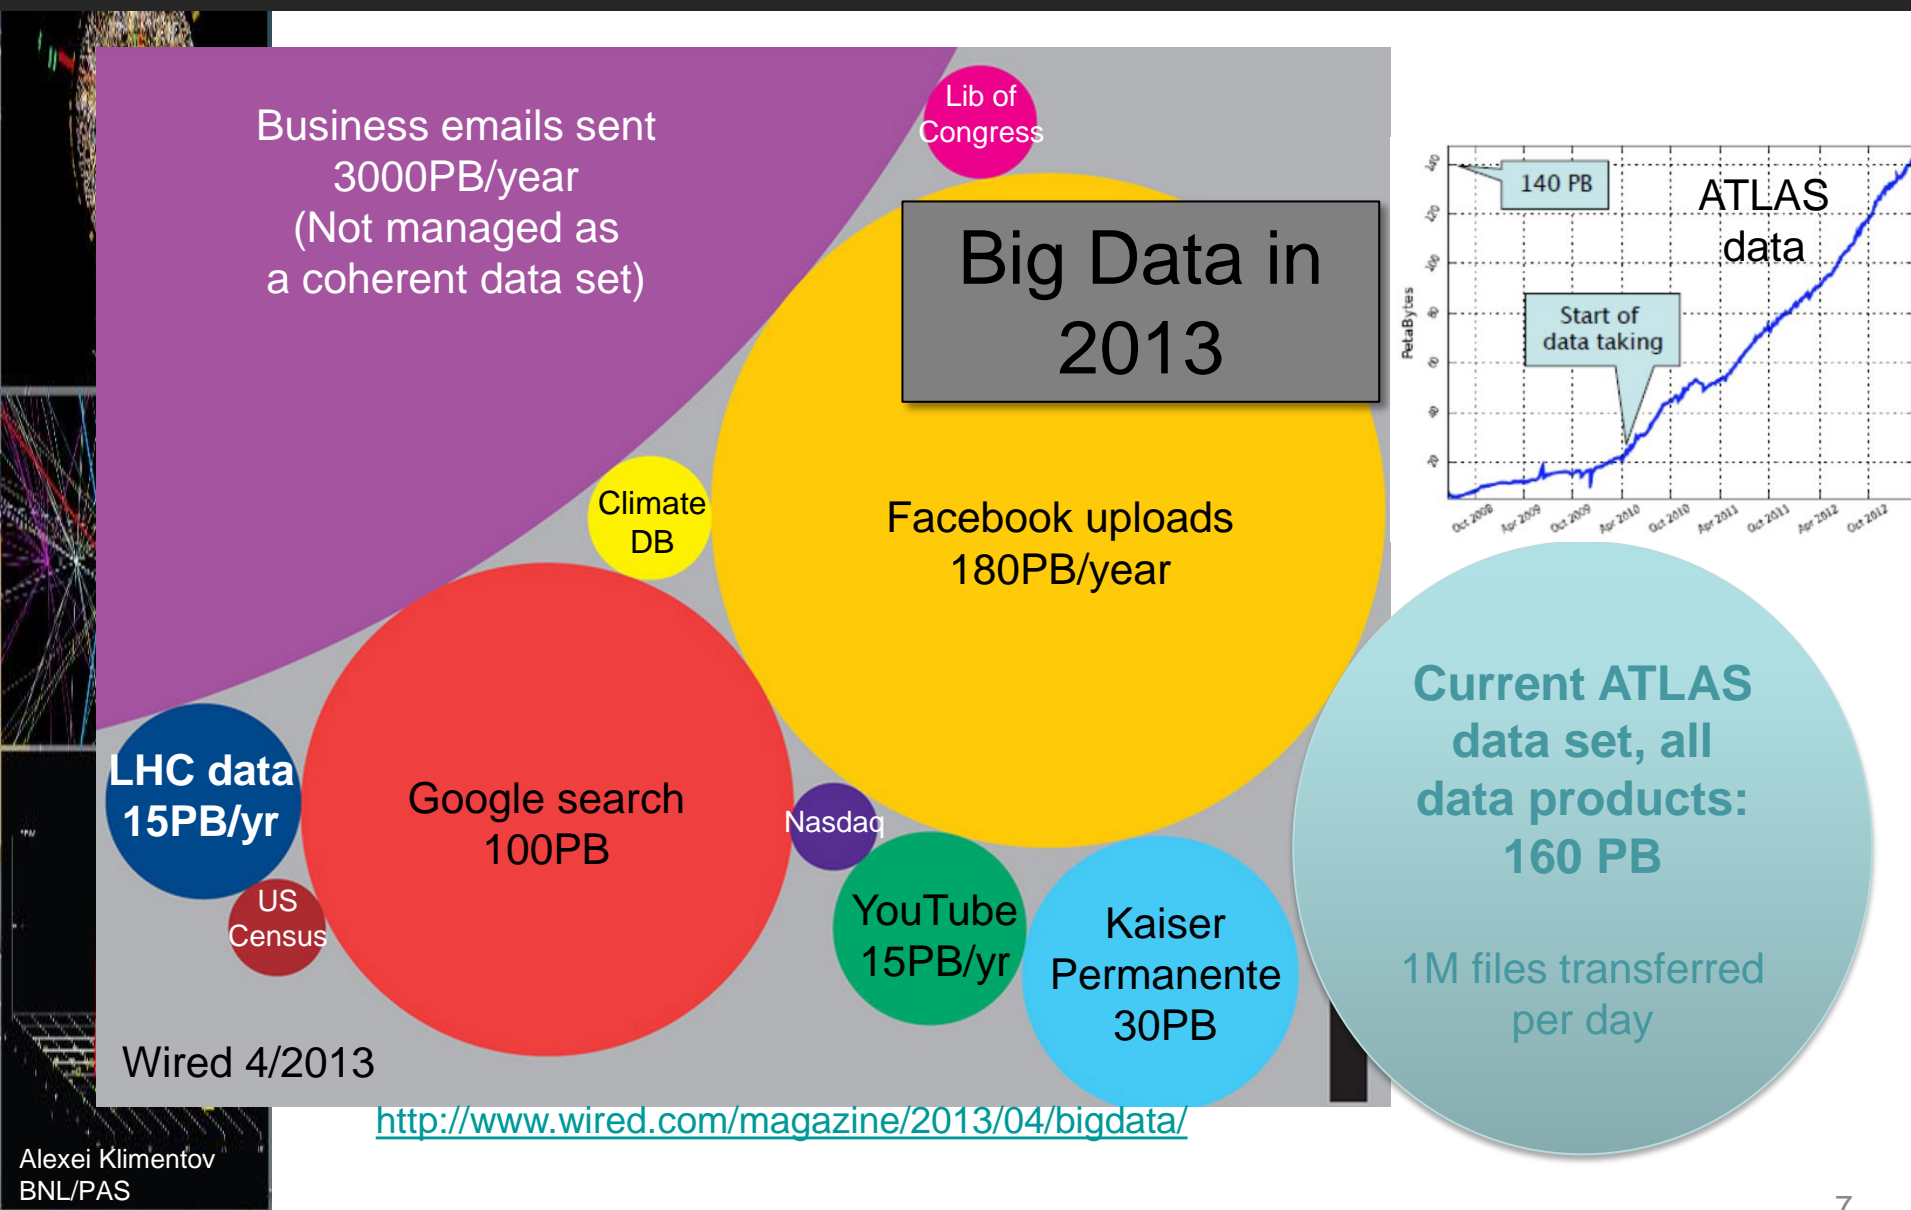
# ATLAS . Big Data Experiment

150 million sensors deliver  
data ...



Alexei Klimentov  
BNL/PAS

# Big Data: often just a buzz word, but not when it comes to ATLAS...



# LHC Computing Challenges

- **A lot of data in a highly distributed environment.**

- Petabytes of data to be treated and analyzed
  - For example ATLAS managed data volume ~160 PB, distributed world-wide to O(100) computing centers and analyzed by O(1000) physicists
  - ATLAS Detector generates about 1PB of raw data per second
  - **More than a hundred of computing centers had to work together**
- Dozens of complex applications

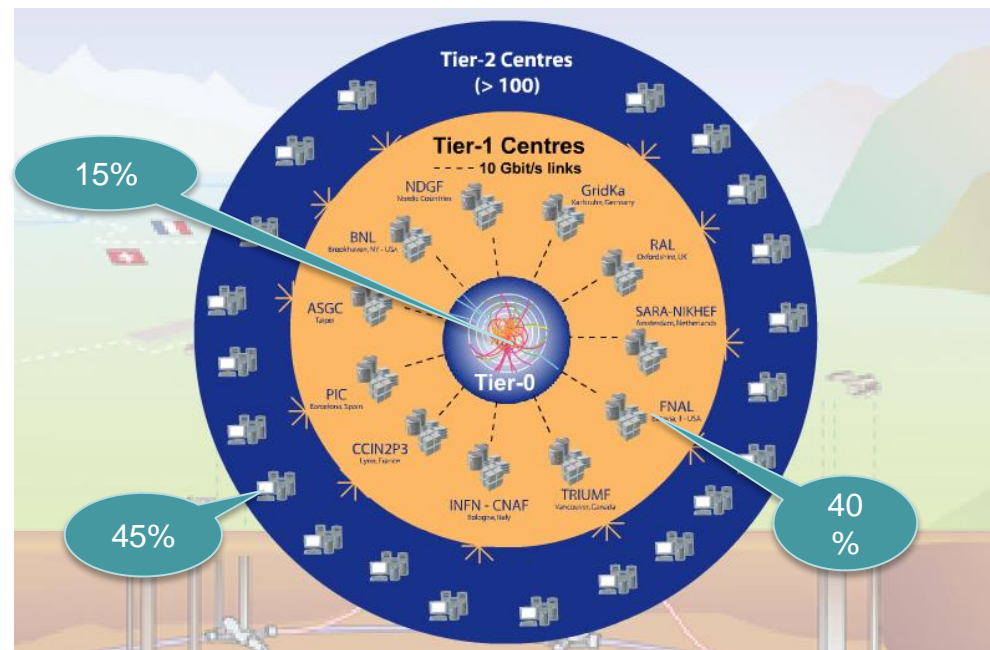
- **Very large international collaboration**

- Hundreds Institutes and Universities from many countries
- Thousands of physicists analyze the data

**LHC Experiments use grid computing paradigm to organize distributed resources;**

**A few years ago Cloud Computing RnD projects were started to explore virtualization and clouds**

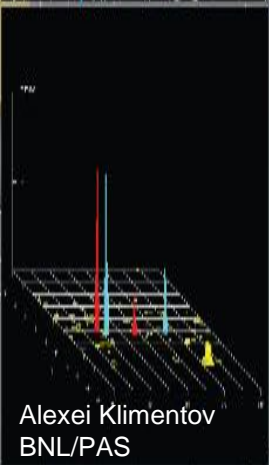
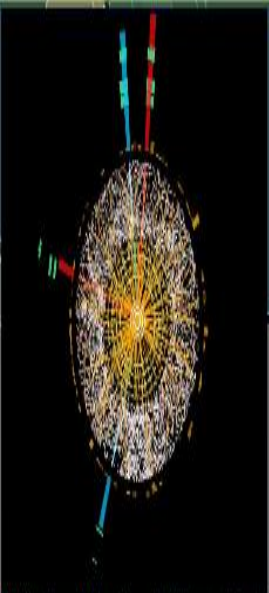
**Now we are evaluating how high-performance and super-computers can be used for data processing and analysis**



# Workload Management



- **Pilots, pilots everywhere!**
  - Approach spread from ALICE and LHCb to all LHC experiments
- **Experiment-level workload management systems successful**
  - No long term need for WMS as middleware component
- **But is WMS commonality still possible then? Yes...**
  - DIRAC (LHCb) used by other experiments, as is AliEn (ALICE)
  - PanDA (ATLAS) has new (DOE) support to generalize as an exascale computing community WMS. PanDA is under evaluation by LSST and COMPASS, and in production in AMS at ASGC (TW), also ASGC team uses PanDA for other scientific applications
- **Pilot submission and management works well on foundation of Condor, glideinWMS**
- **Multi-core and whole node support**
  - Support is pretty much there for simple usage modes – and simplicity is the aim – in particular whole node scheduling
- **Extended environmental info (eg. HS06, job lifetime)**
  - Environment variables providing homogeneous information access across batch system flavors (could have done with this years ago)
- **Virtualization and cloud computing**
  - Virtual CE: better support for “any” batch system. Essential.
  - Virtualization clearly a strong interest for sites, both for service hosting and WNs
  - Cloud computing interest/activity levels vary among experiments
    - ATLAS is well advanced in integrating and bringing cloud computing in real analysis and production workflows.
- **Super-computing (High Performance Computers and Leadership Class Facilities)**
  - ATLAS, ALICE, AMS, LSST, CERN IT collaborating on common way to submit pilots to HPC (Leadership Class Facilities - LCF)
  - ATLAS started to port PanDA and to use MIRA, Titan, Abisko, Pitz Daint, ....



Alexei Klimentov  
BNL/PAS

# Workload Management System. Core Ideas.



- **Make hundreds of distributed sites appear as local**
  - Provide central queue for users – similar to local batch systems
- **Reduce site related errors and latency**
  - Build a pilot job system – late transfer of user payloads
  - Crucial for distributed infrastructure maintained by local experts
- **Hide middleware while supporting diversity and evolution**
  - WMS interacts with middleware – users see high level workflow
  - Automation engines built into PanDA, not exposed to users
- **Hide variations in infrastructure**
  - WMS presents uniform ‘job’ slots to user (with minimal sub-types)
  - Easy to integrate grid sites, clouds, recently HPC sites
- **Use the same system for Monte-Carlo production, data processing and users analysis**
- **Similar ideas have been implemented in several independent systems developed by LHC experiments : AliEn, Dirac, PanDA**



# Production and Distributed Analysis System

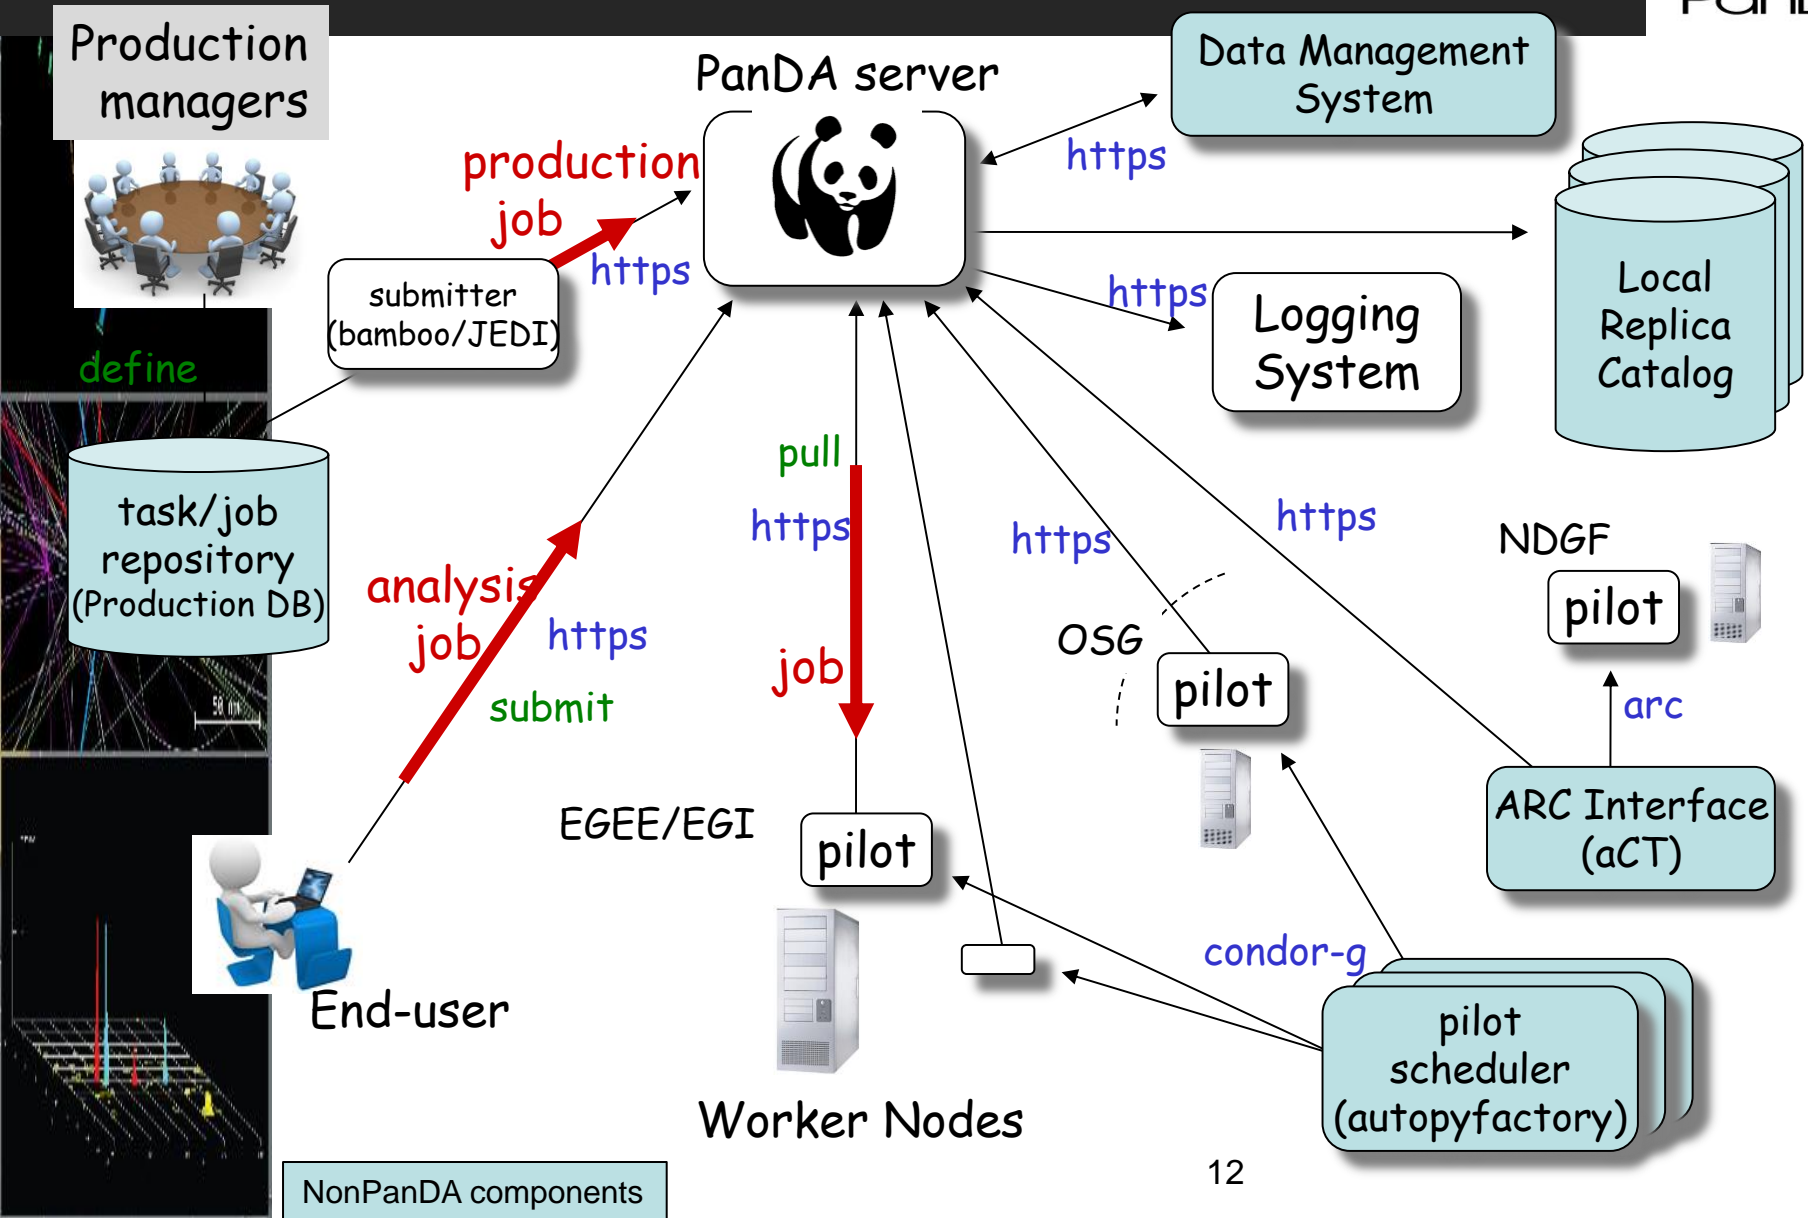


- **PanDA – Production and Distributed Analysis System**
  - Designed for the ATLAS experiment during LHC Run 1
  - Deployed on WLCG infrastructure
- **PanDA project started in 2005 by BNL and UTA groups.**
  - Production and Data Analysis system
  - An automated yet flexible workload management system which can optimally make distributed resources accessible to all users
  - Primary goal – improve user experience with distributed computing
  - Distributed computing should be as easy as local computing
  - Users should get quick results by leveraging distributed resources
  - Isolate users from heterogeneity in infrastructure and middleware
  - Fair sharing of resources among thousands of users
  - Small and large tasks, individual users and groups – same interface

*Through PanDA, physicists see a single computing facility that is used to run all data processing for the experiment, even though data centers are physically scattered all over the world.*



# PanDA Workload Management

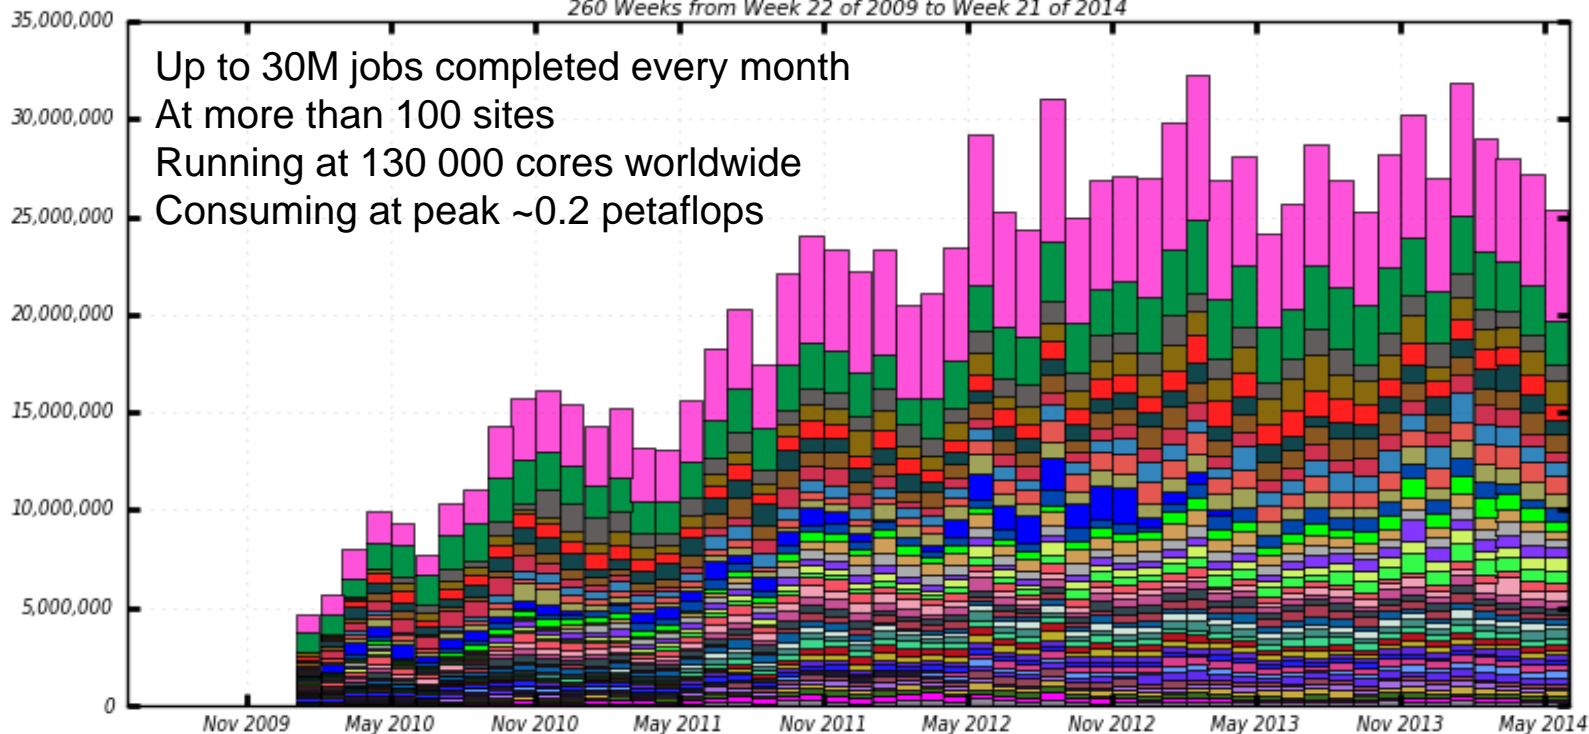


# PanDA Performance

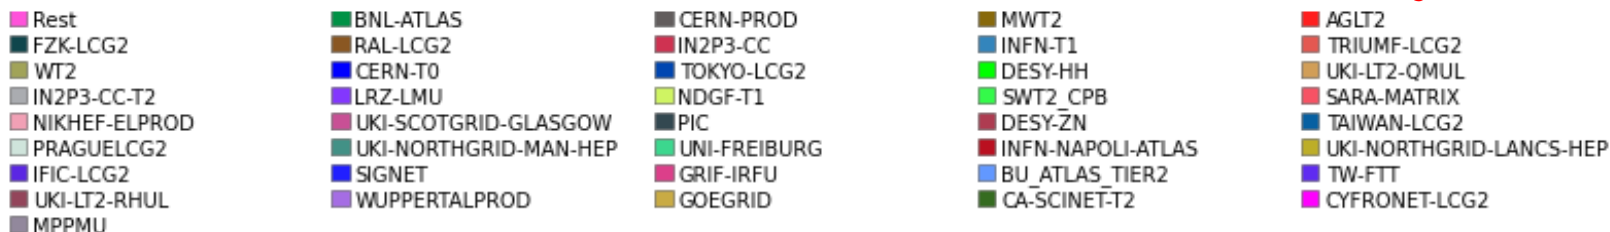


Completed jobs

260 Weeks from Week 22 of 2009 to Week 21 of 2014

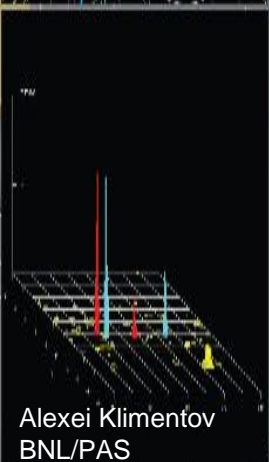
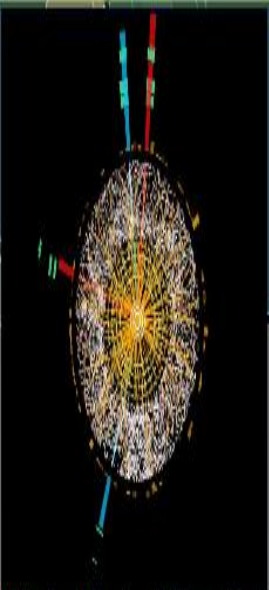
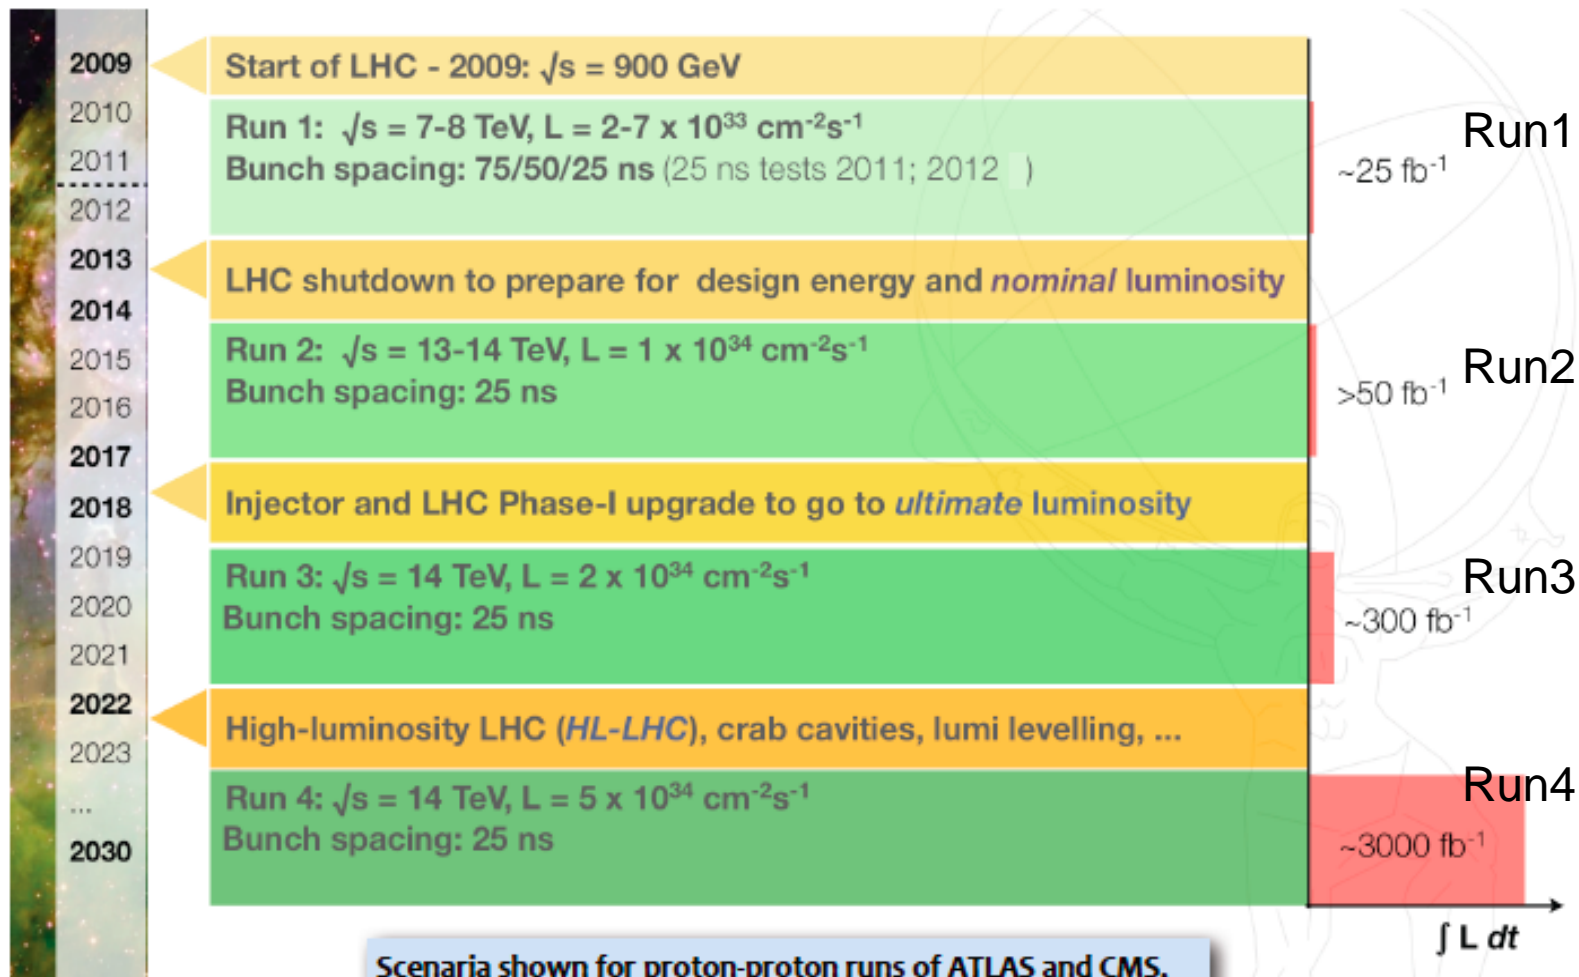


**Available resources are fully used**



**PanDA is exascale now: 1.2 Exabytes of data processed by PanDA in 2013**

# LHC Upgrade Timeline. More Data.



# LHC Computing Scale of Needs

CPU needs per event

Run1

Run2

Run3

Run4

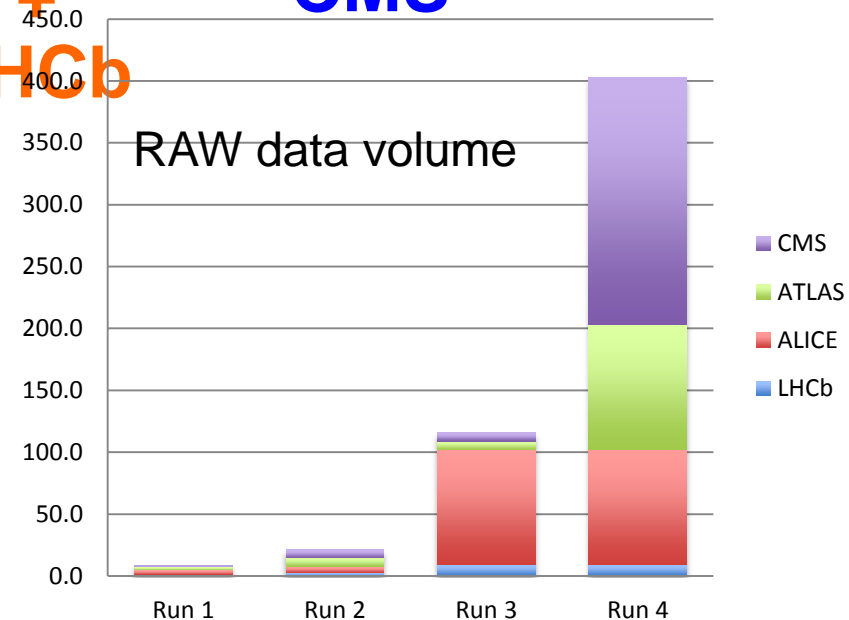
ALICE

+  
LHCb

ATLAS

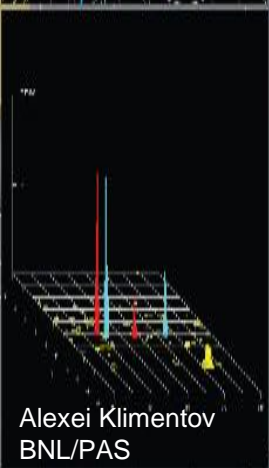
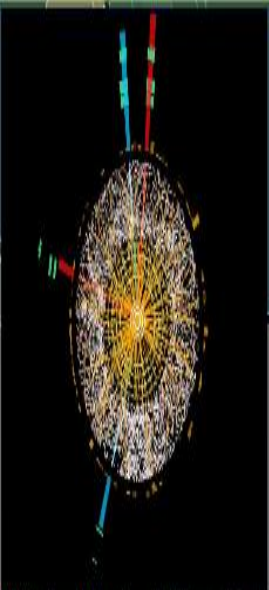
+  
CMS

- CPU needs (per event) will grow with track multiplicity (pileup) and energy
- Storage needs are proportional to accumulated luminosity
- Grid resources are limited by funding

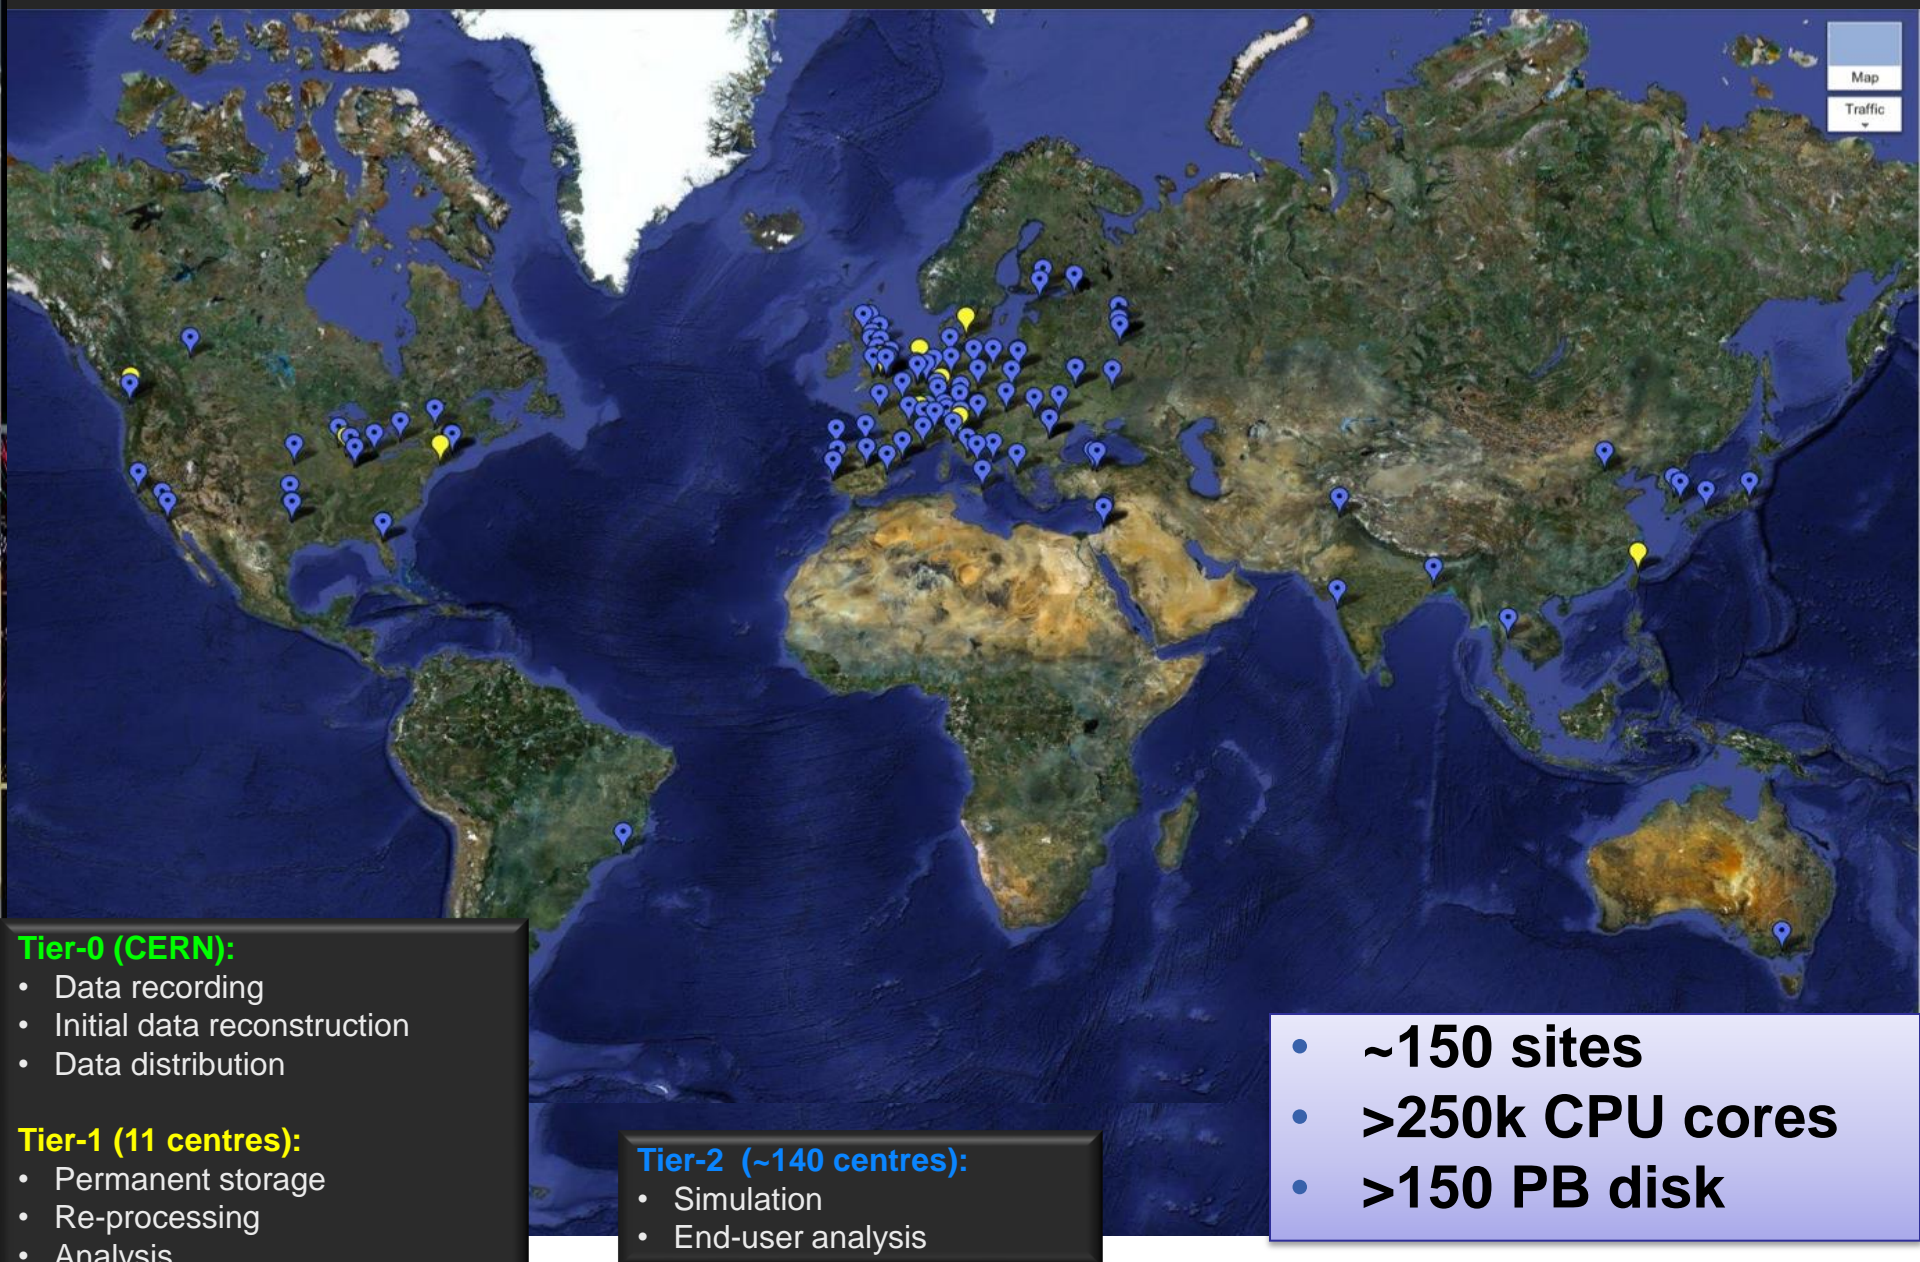


# LHC Computing Scale of Needs. Cont'd

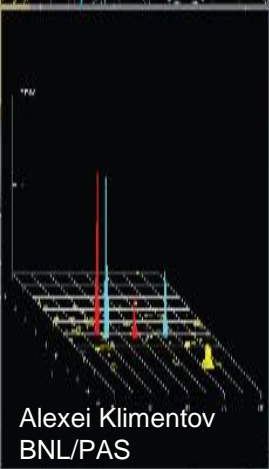
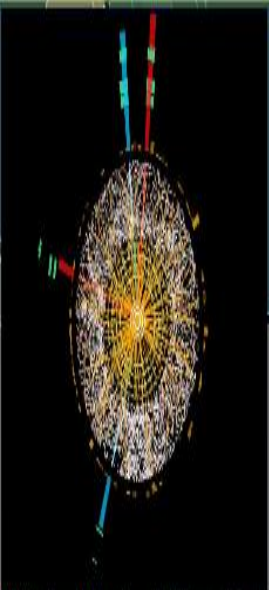
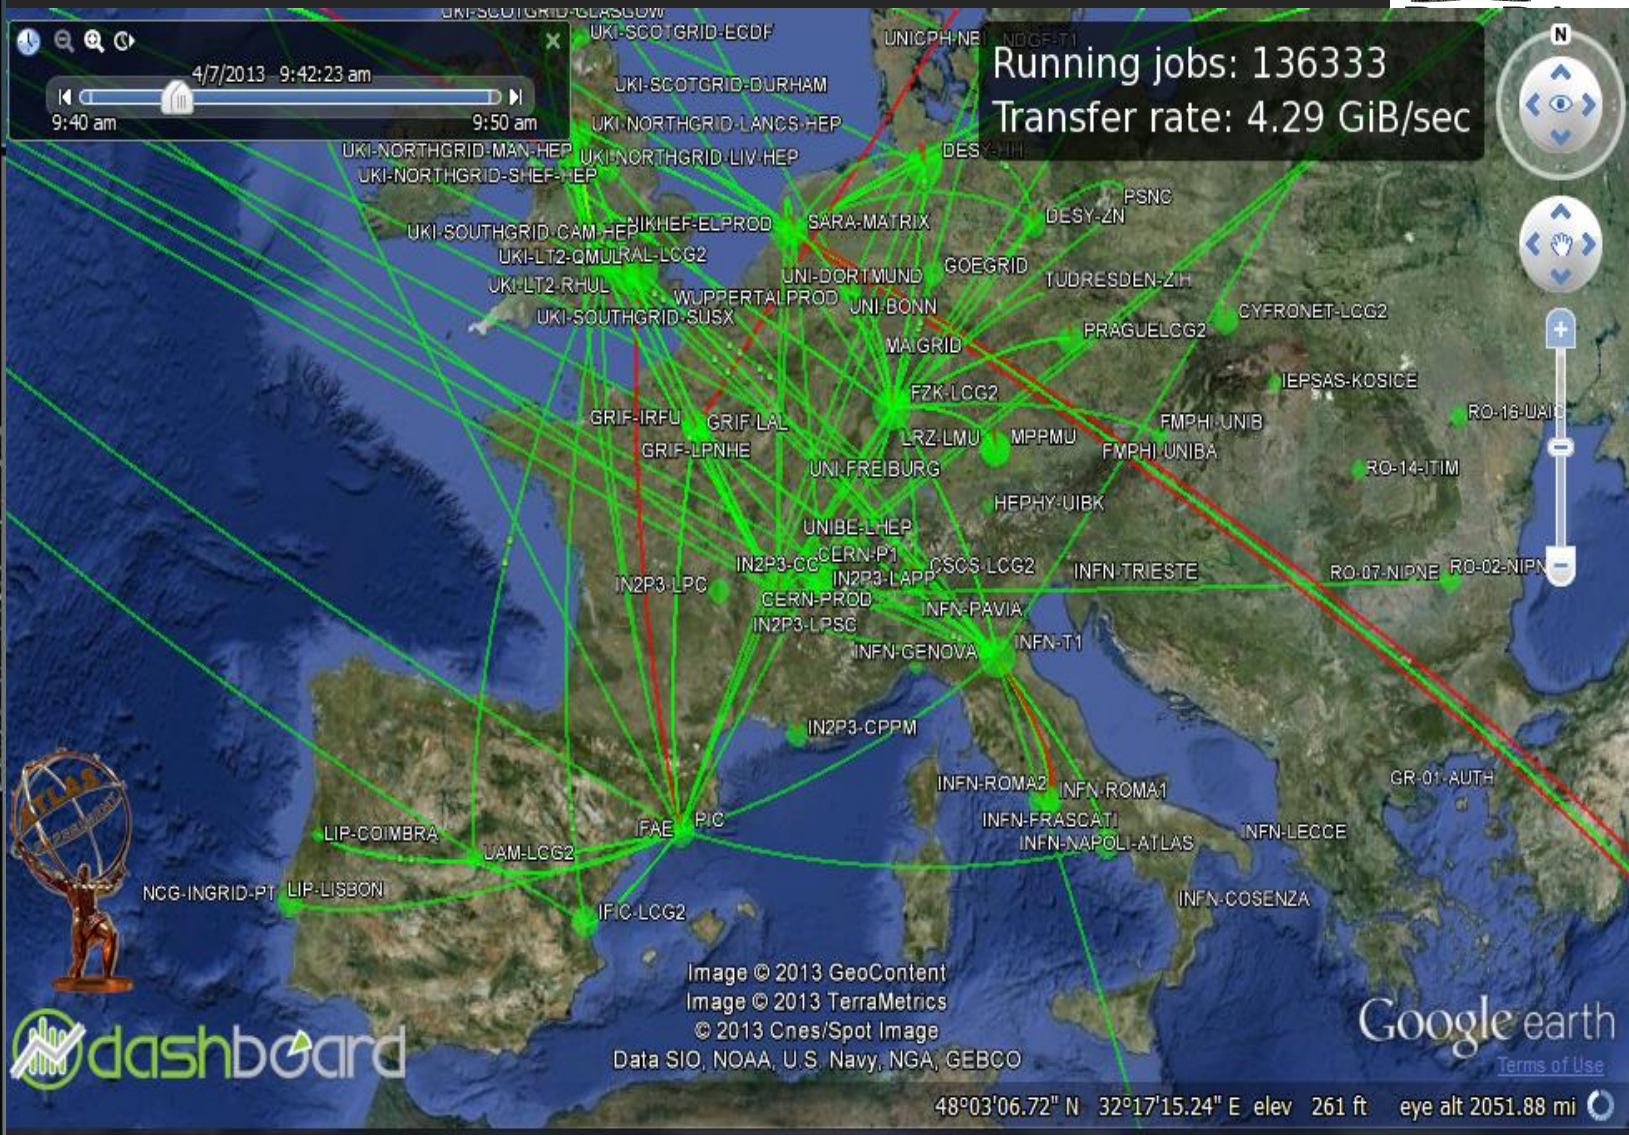
- The ATLAS experiment uses a geographically distributed grid of approximately 130,000 cores continuously, (over 1000 million core-hours per year) to process, simulate, and analyse its data.
- The need for simulation and analysis would overwhelm the expected capacity of LHC Grid computing facilities unless the range and precision of physics studies were to be curtailed.
  - Physics requires to increase rate
    - ATLAS Run1 data-taking rate 300 Hz
    - ATLAS Run2 data-taking rate 1kHz
- Leadership Class Facilities contributions of the order of 10 Million or more core hours per year become important and valuable.
- ATLAS computing can also be a close to ideal “crack-filling” application. The ATLAS Workload Management System (PanDA) is being upgraded to make it aware of dynamically changing resources, and thus able to exploit groups of processors that become available for relatively short times.
- Extending PanDA beyond the Grid will further expand the potential user community and the resources available to them.



# LHC Computing Grid: A global collaboration...



# Data fans out all over Europe



More than 25 billions bits transferred every second

# And across the Atlantic



4/7/2013 9:42:23 am  
9:40 am 9:50 am

Running jobs: 136333  
Transfer rate: 4.29 GiB/sec

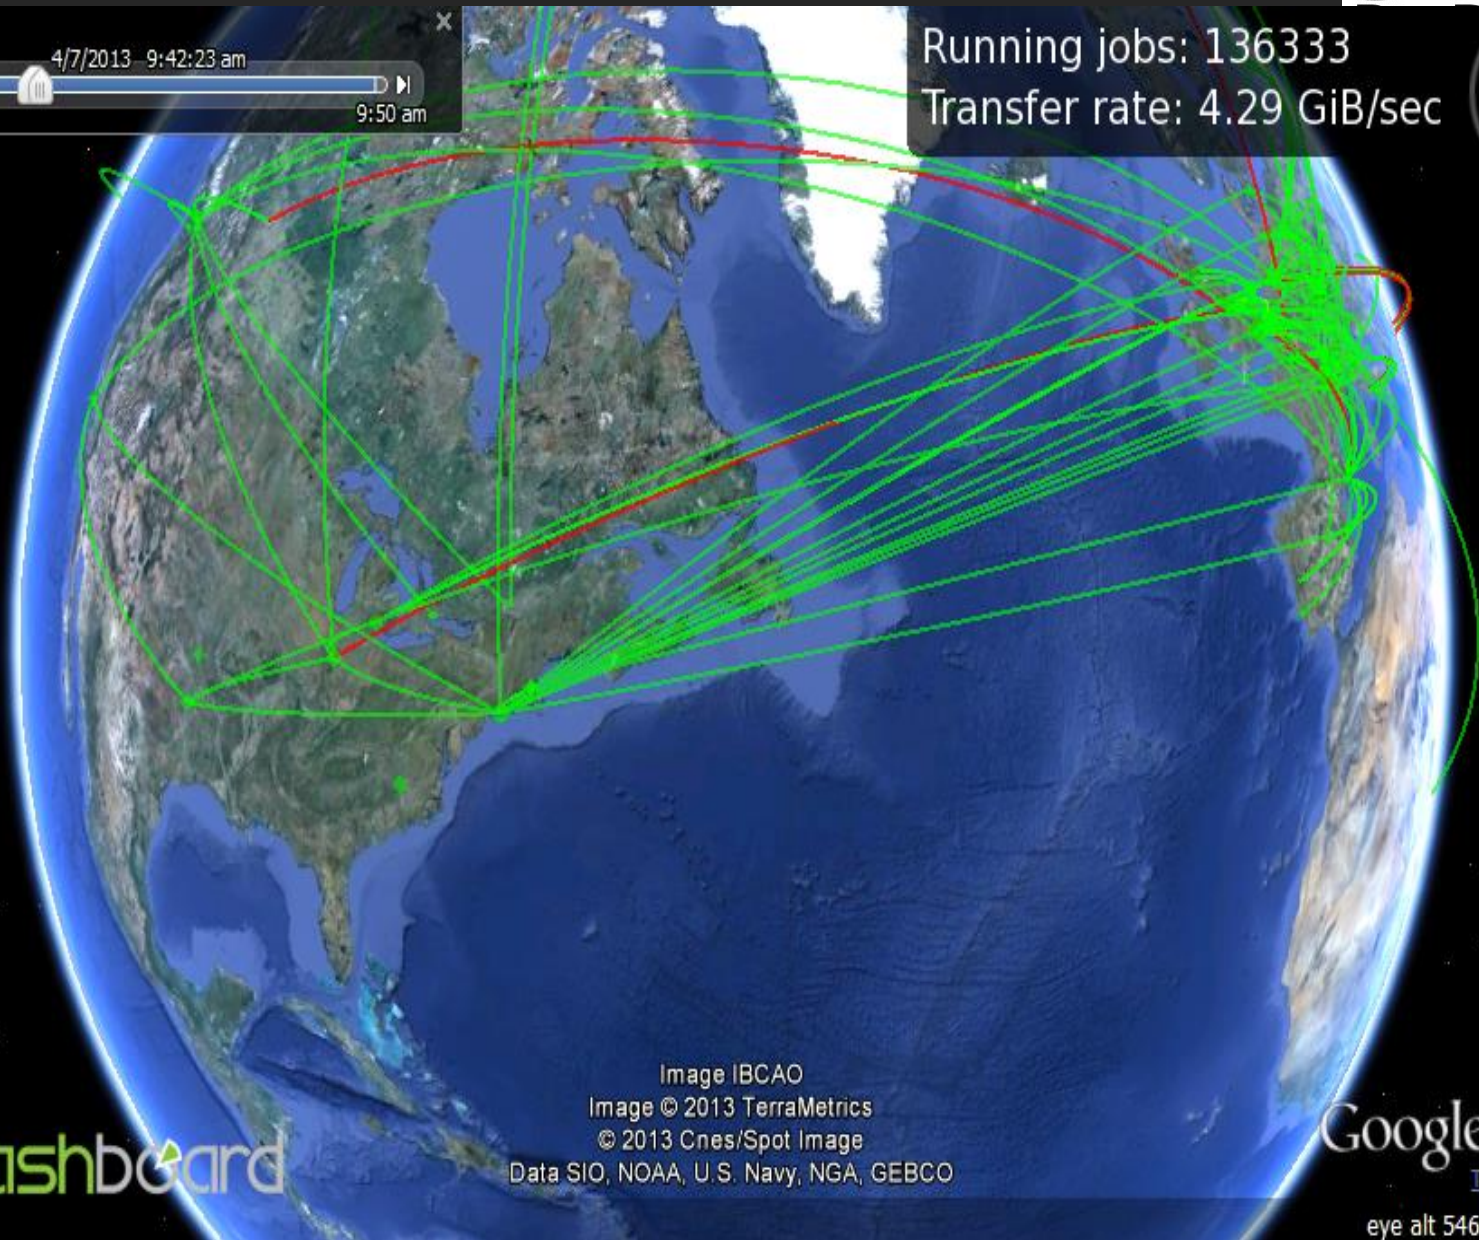


Image IBCAO  
Image © 2013 TerraMetrics  
© 2013 Cnes/Spot Image  
Data SIO, NOAA, U.S. Navy, NGA, GEBCO

Google earth  
[Terms of Use](#)

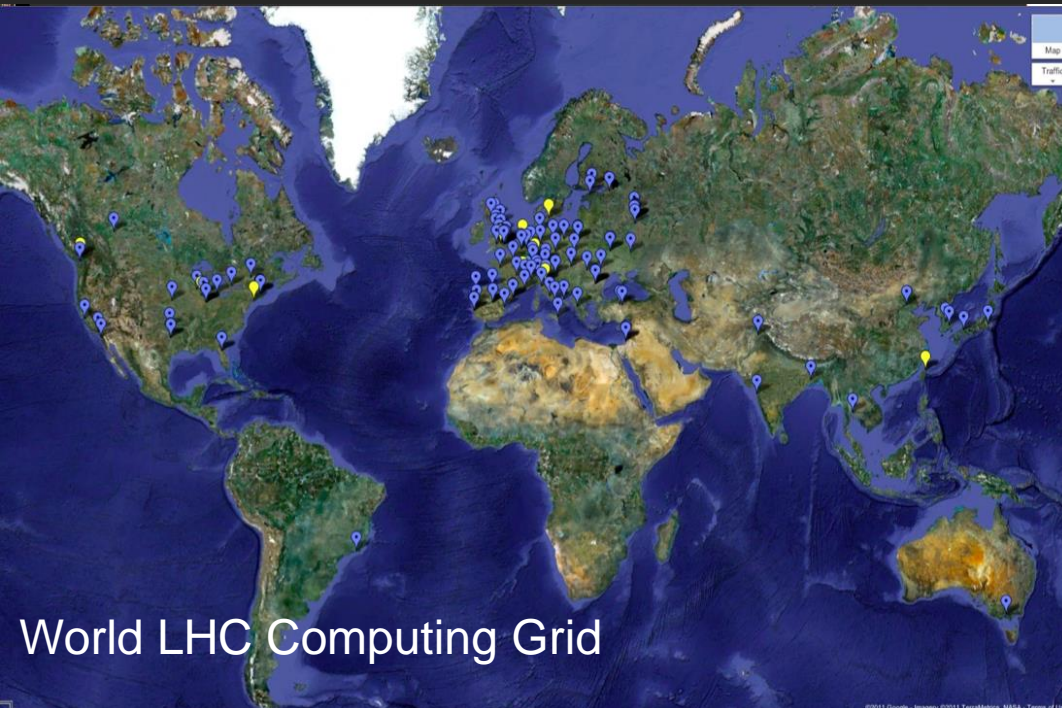
eye alt 5467.42 mi



dashboard

Alexei Klime  
BNL/PAS

# LHC, Amazon & Google Computing Centers



- One Google Data Center is estimated to cost ~\$600M
  - An order of magnitude more than the centre at CERN
- Amazon : 9 large sites/zones
  - up to ~2M CPU cores/site, ~4M total
  - 10 x more cores on 1/10 of the sites compared to our Grid
  - 500,000 users
- LHC Computing (WLCG)
  - 150 sites
  - 350k CPU cores total
  - ~5000 users





- Virtually joining together the sites based on proximity (latency) and network capacity into Regional Data Clouds
- Each cloud/region provides reliable data management and sufficient processing capability
- Dealing with handful of clouds/regions instead of the individual sites

# Virtual Storage Cloud

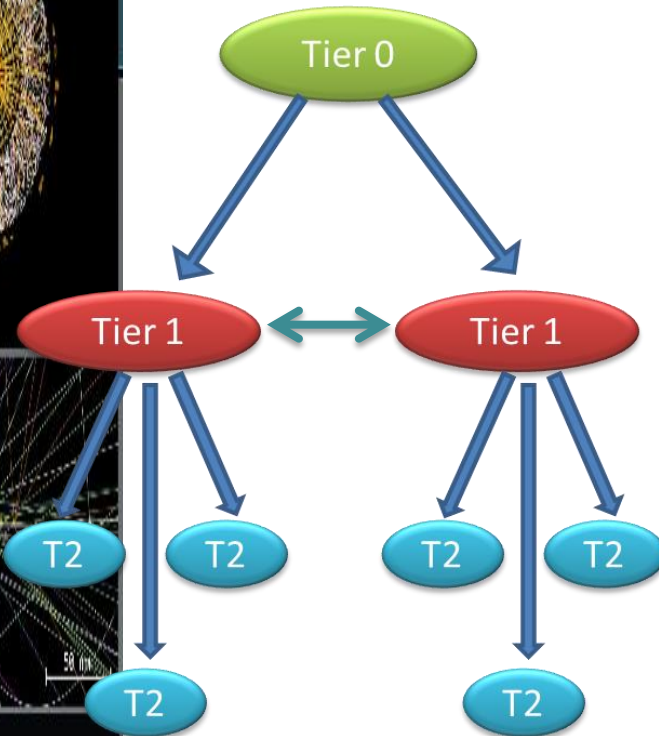
The diagram illustrates a Virtual Storage Cloud architecture. At the top, 'Clients' (represented by various OS icons) connect to 'XRootD' (a globe icon). XRootD then connects to a 'VST instance' (gear icon) labeled 'global'. This global instance is connected to a 'VST instance' (gear icon) labeled 'cloud' within 'CLOUD A'. 'CLOUD A' also contains a 'VST instance' (gear icon) labeled 'site', which is connected to a stack of server icons. 'CLOUD B' is shown to the right, containing a 'VST instance' (gear icon) and three server icons. A red arrow indicates data flow from the global instance to the cloud instance, and from the cloud instance to the site instance. A yellow box at the bottom right contains the text: 'Storage access is resolved locally or according to policies to a higher level instance.'

Storage access is resolved locally or according to policies to a higher level instance.

IT Information Technology Department

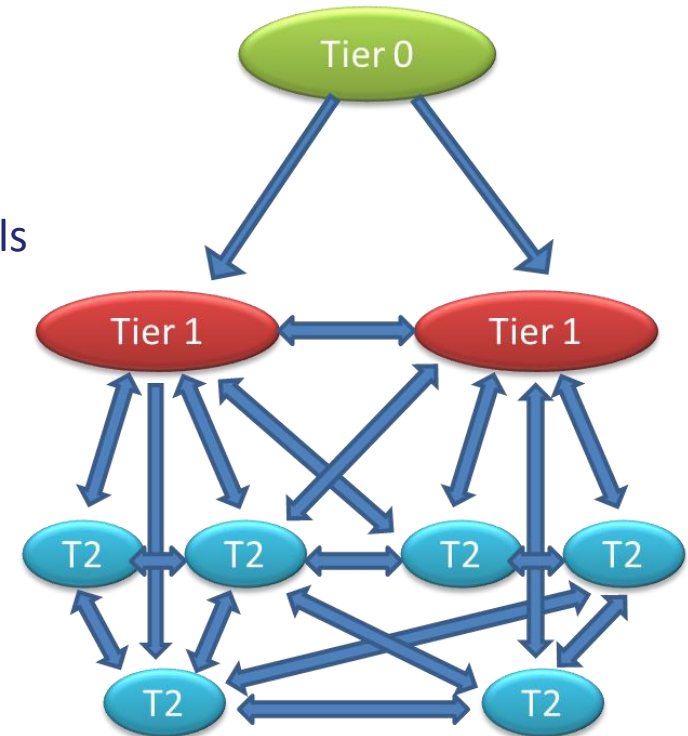


# LHC Computing Model Evolution. Tiers Hierarchy



Hierarchy

Evolution of  
computing models

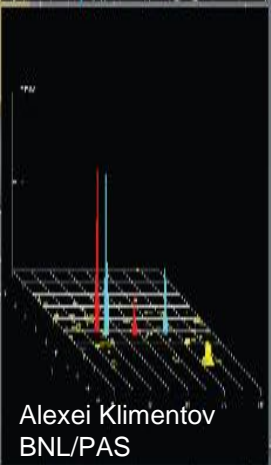
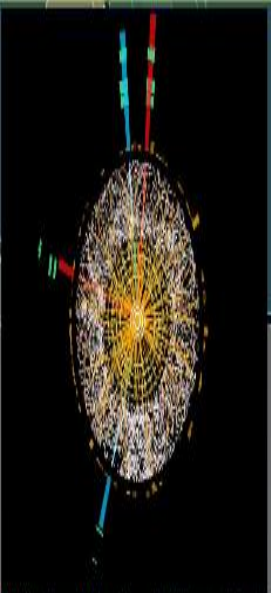


Mesh

- Network capabilities and data access technologies have significantly improved our ability to use resources independent of location
- Relaxing hierarchical model : Flat instead of Tiered Grid model

# BigPanDA. Extending PanDA beyond the Grid

- **Proposal titled “Next Generation Workload Management and Analysis System for BigData” – Big PanDA is DOE Advanced Scientific Computing Research (ASCR) and HEP funded project.**
  - Generalization of PanDA as meta application, providing location transparency of processing and data management, for HEP and other data-intensive sciences, and a wider exascale community.
    - Other efforts
      - PanDA : US ATLAS funded project
      - Networking : Advance Network Services
- **There are three dimensions to evolution of PanDA**
  - Making PanDA available beyond ATLAS and High Energy Physics
  - Extending beyond Grid (Leadership Computing Facilities, Clouds, University clusters)
  - Integration of network as a resource in workload management



# Workload Management and Networking.

- **Why WMS should care about networking ?**

- Distributed workload management systems need to transfer data (or use direct access) both for input and output
- Data transfer/access is done asynchronously : by DQ2/Rucio in ATLAS, PhEDEx in CMS, pandamover/FAX ...
- Data transfer/access systems can be optimized for network performance – PanDA will use these enhancements
- But network information can also be used directly in workflow management in PanDA at a higher level – first step to try

- **Goal for PanDA**

- Direct integration of networking with PanDA workflow – never attempted before for large scale automated WMS systems

- **Main PanDA use cases**

- Use network information for cloud selection
- Use network information for job assignment
- Use network information for site selection



# PanDA and Networking. Concept.

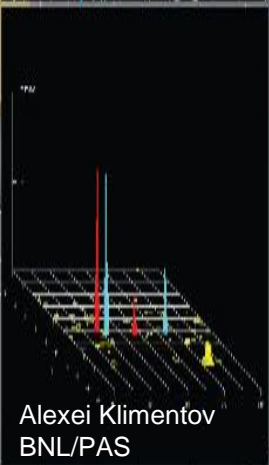
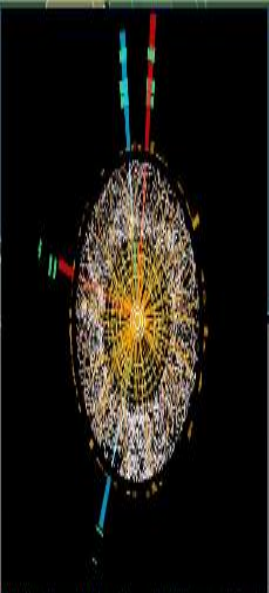


- **PanDA as workload manager**

- PanDA automatically chooses job execution site
  - Multi-level decision tree – task brokerage, job brokerage, dispatcher
  - Also manages predictive future workflows – at task definition, PD2P (PanDA Dynamic Data Placement)
- Site selection is based on processing and storage requirements
  - Can we use network information in this decision?
  - Can we go even further – network provisioning?
- Further – network knowledge used for all phases of job cycle?

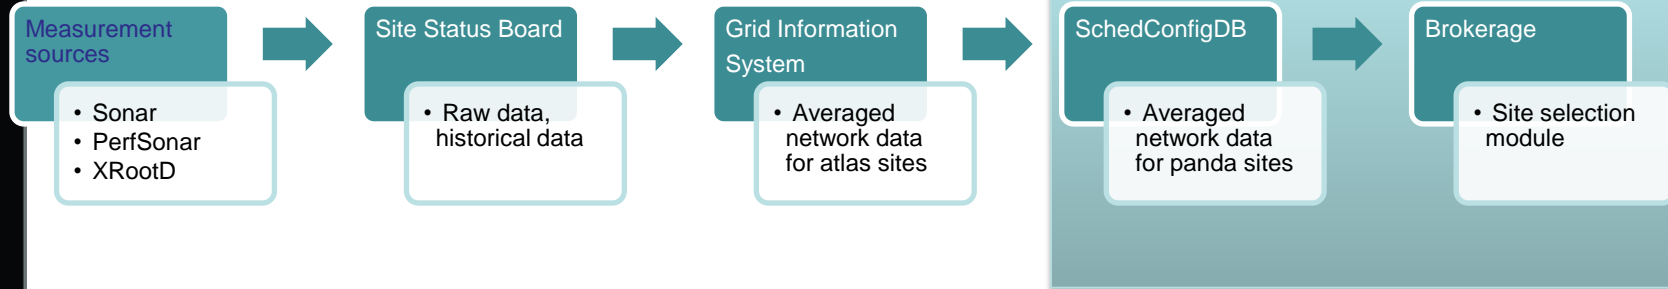
- **Network as resource**

- Optimal site selection should take network capability into account
  - We do this already – but indirectly using job completion metrics
- Network as a resource should be managed (i.e. provisioning)
  - We also do this crudely – mostly through timeouts, self throttling



# Intelligent Network Services and PanDA

- In BigPanDA we will use information on how much bandwidth is available and can be reserved before data movement will be initiated
- In Task Definition user will specify data volume to be transferred and deadline by which task should be completed. The calculations of (i) how much bandwidth to reserve, (ii) when to reserve, and (iii) along what path to reserve will be carried out by Virtual Network On Demand.



# Site Status Board. Data Injection and Monitoring



Index

Expanded Table

Show 200 entries

Copy

Print

Save

view: Sonar

Site Name	SrcSite	SrcCloud	SrcTier	DstSite	DstCloud	DstTier	Prio
AGLT2_to_OU_OCCHP_SWT2	AGLT2	US	T2D	OU_OCCHP_SWT2	US	T2	6
AGLT2_to_SWT2_CPB	AGLT2	US	T2D	SWT2_CPB	US	T2	6
AUSTRALIA-ATLAS_to_OU_OCCHP_SWT2	Australia-ATLAS	CA	T2	OU_OCCHP_SWT2	US	T2	2
AUSTRALIA-ATLAS_to_SWT2_CPB	Australia-ATLAS	CA	T2	SWT2_CPB	US	T2	2
BEIJING-LCG2_to_OU_OCCHP_SWT2	BEIJING-LCG2	FR	T2D	OU_OCCHP_SWT2	US	T2	2
BEIJING-LCG2_to_SWT2_CPB	BEIJING-LCG2	FR	T2D	SWT2_CPB	US	T2	5
BNL-OSG2_to_OU_OCCHP_SWT2	BNL-ATLAS	US	T1	OU_OCCHP_SWT2	US	T2	8
BNL-OSG2_to_SWT2_CPB	BNL-ATLAS	US	T1	SWT2_CPB	US	T2	8
CA-MCGILL-CLUMEQ-T2_to_OU_OCCHP_SWT2	CA-MCGILL-CLUMEQ-T2	CA	T2D	OU_OCCHP_SWT2	US	T2	2
CA-MCGILL-CLUMEQ-T2_to_SWT2_CPB	CA-MCGILL-CLUMEQ-T2	CA	T2D	SWT2_CPB	US	T2	5
CA-SCINET-T2_to_OU_OCCHP_SWT2	CA-SCINET-T2	CA	T2D	OU_OCCHP_SWT2	US	T2	2
CA-SCINET-T2_to_SWT2_CPB	CA-SCINET-T2	CA	T2D	SWT2_CPB	US	T2	5
CA-VICTORIA-WESTGRID-T2_to_OU_OCCHP_SWT2	CA-VICTORIA-WESTGRID-T2	CA	T2D	OU_OCCHP_SWT2	US	T2	2
CA-VICTORIA-WESTGRID-T2_to_SWT2_CPB	CA-VICTORIA-WESTGRID-T2	CA	T2D	SWT2_CPB	US	T2	5

# Network Performance Measurements



Search...

DDM Sonar						perfSONAR						FAX xrdcp rate
AvgBRS (MB/s) ↕	EvS ↕	AvgBRM (MB/s) ↕	EvM ↕	AvgBRL (MB/s) ↕	EvL ↕	MinThr (MB/s) ↕	AvgThr (MB/s) ↕	MaxThr (MB/s) ↕	MinPL ↕	AvgPL ↕	MaxPL ↕	
1.05+/-0.19	10	7.46+/-1.48	11	12.54+/-6.72	519	12.4	34.7	56.9	0.0	0.0	2.0	n/a
0.85+/-0.04	10	9.97+/-4.20	602	26.48+/-13.48	10	0.6	0.8	1.1	0.0	0.0	1.0	3.93
0.42+/-0.06	10	0.89+/-0.11	10	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.39+/-0.06	10	1.02+/-0.04	10	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.58+/-0.07	10	2.91+/-0.82	10	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.48+/-0.06	10	2.45+/-0.65	10	3.18+/-0.79	10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.12+/-0.39	465	4.13+/-1.44	1575	4.59+/-1.68	3803	164.2	172.3	180.3	0.0	0.0	0.0	n/a
2.10+/-1.88	4920	8.76+/-6.32	10075	14.05+/-23.55	4006	0.3	0.3	0.3	0.0	0.0	0.0	0.72
0.47+/-0.11	5	1.23+/-0.39	9	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.37+/-0.11	10	1.14+/-0.20	5	2.53+/-0.15	10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.67+/-0.54	10	7.53+/-3.81	10	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.56+/-0.38	10	5.95+/-2.64	10	50.52+/-9.11	10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.94+/-0.08	10	5.41+/-1.33	10	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.55+/-0.25	10	4.95+/-1.63	10	21.09+/-9.01	10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
1.13+/-0.11	10	7.17+/-1.44	510	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0.82+/-0.33	10	6.90+/-1.82	10	30.36+/-11.35	10	n/a	n/a	n/a	n/a	n/a	n/a	5.55
1.14+/-0.09	10	6.50+/-2.41	10	0.00+/-0.00	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a

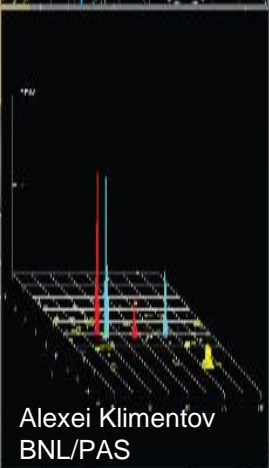
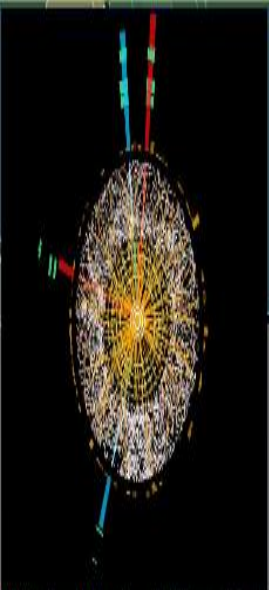
<http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview#currentView=Sonar&highlight=false>

# BigPanDA. Extending the scope. Cloud Computing.



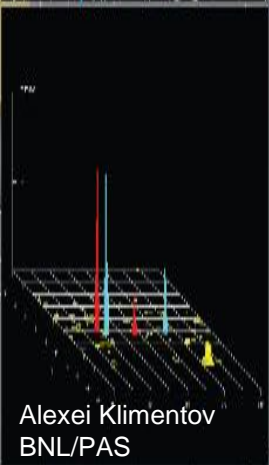
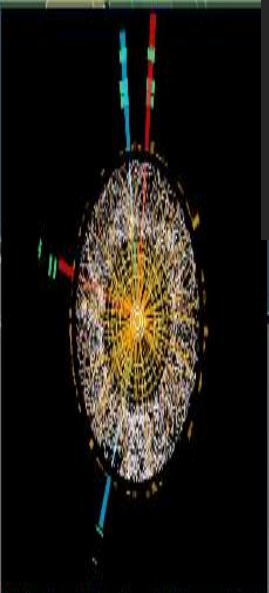
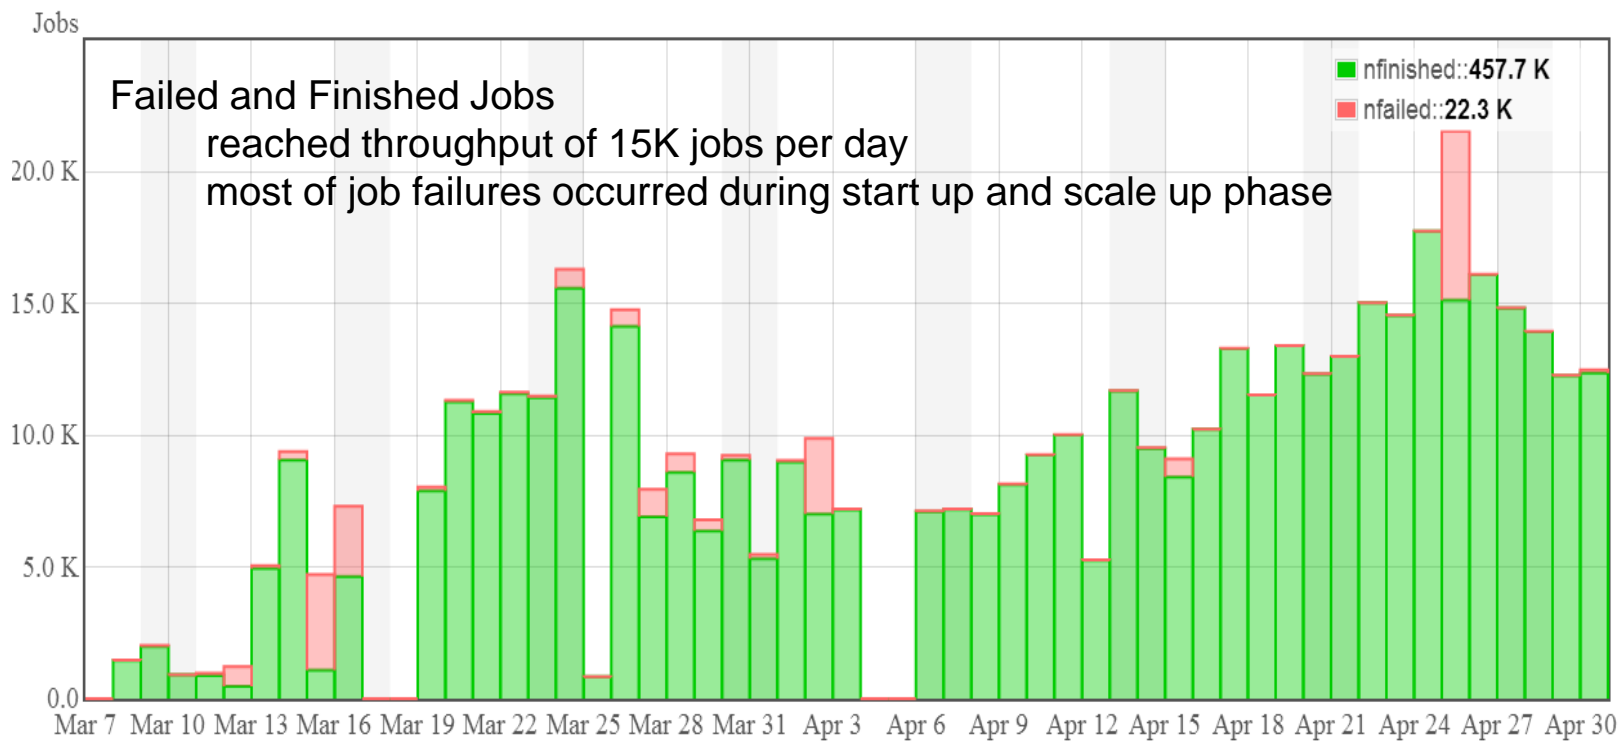
## Compute Engine (GCE) preview project

- Google allocated additional resources for ATLAS for free
  - ~5M cpu hours, 4000 cores for about 2 month, (original preview allocation 1k cores)
- Resources are organized as HTCondor based PanDA queue
  - Centos 6 based custom built images, with SL5 compatibility libraries to run ATLAS software
  - Condor head node, proxies are at BNL
  - Output exported to BNL SE
- Work on capturing the GCE setup in Puppet
- Transparent inclusion of cloud resources into ATLAS Grid
- The idea was to test long term stability while running a cloud cluster similar in size to Tier 2 site in ATLAS
- Intended for CPU intensive Monte-Carlo simulation workloads
- Planned as a production type of run. Delivered to ATLAS as a resource and not as an R&D platform.
- We also tested high performance PROOF based analysis cluster
- **PanDA based data processing and workload management**
  - Centrally managed queues in the cloud
    - Elastically expand resources transparently to users
  - Institute managed Tier 3 analysis clusters
    - Hosted locally or (more efficiently) at shared facility, T1 or T2
  - Personal analysis queues
    - User managed, low complexity (almost transparent), transient
- **Data storage**
  - Transient caching to accelerate cloud processing
  - Object storage and archiving in the cloud



# Running PanDA on Google Compute Engine

- We ran for about 8 weeks (2 weeks were planned for scaling up)
- Very stable running on the Cloud side. GCE was rock solid.
- Most problems that we had were on the ATLAS side.
- We ran computationally intensive jobs
  - Physics event generators, Fast detector simulation, Full detector simulation
- Completed 458,000 jobs, generated and processed about 214 M events



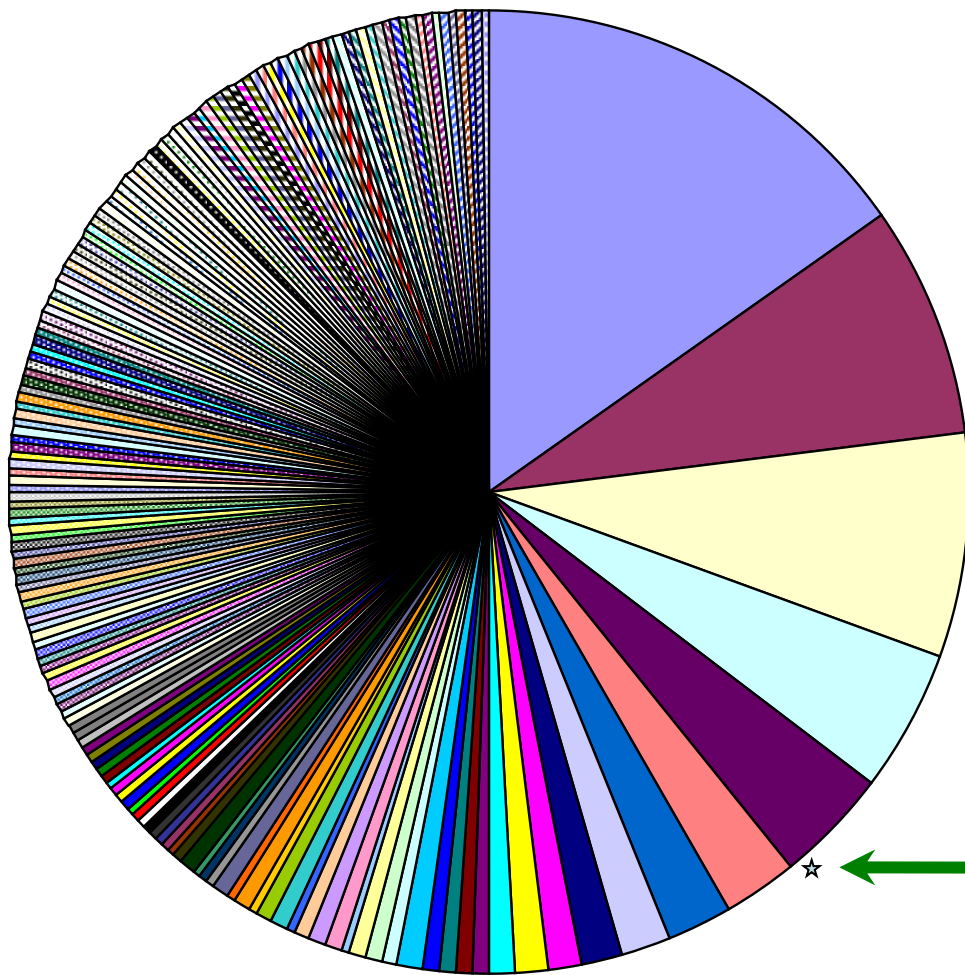
Alexei Klimentov  
BNL/PAS

# BigPanDA for Leadership Computing Facilities

- **Expanding PanDA from Grid to Leadership Class Facilities (LCF) required significant changes in our system**
- **Each LCF is unique**
  - Unique architecture and hardware
  - Specialized OS, “weak” worker nodes, limited memory per WN
  - Code cross-compilation is typically required
  - Unique job submission systems
  - Unique security environment
- **Pilot submission to a worker node is typically not feasible**
- **Pilot/agent per supercomputer or queue model**
- **Initial tests on BlueGene at BNL and ANL. Geant4 port to BG/P**
- **PanDA project at Oak-Ridge National Laboratory LCF (OLCF) - Titan**



# The Top 500 Computers



- **Most of the computational power is concentrated in a small number of machines**
  - Half the total power is in the top dozen computers
- **Equivalent ATLAS Grid use is about the size of the sector (little green star)**

# High Performance Computing (Top 10, Nov 2013)

			cores	Rmax	Rpeak	Power	
1	National Super Computer Center in Guangzhou China	<b>Tianhe-2 (MilkyWay-2)</b> - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3120000	33862.7	54902.4	17808	
2	DOE/SC/Oak Ridge National Laboratory United States	<b>Titan</b> - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560640	17590.0	27112.5	8209	★
3	DOE/NNSA/LLNL United States	<b>Sequoia</b> - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1572864	17173.2	20132.7	7890	
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	<b>K computer</b> , SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705024	10510.0	11280.4	12660	
5	DOE/SC/Argonne National Laboratory United States	<b>Mira</b> - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786432	8586.6	10066.3	3945	★
6	Swiss National Supercomputing Centre (CSCS) Switzerland	<b>Piz Daint</b> - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115984	6271.0	7788.9	2325	★
7	Texas Advanced Computing Center/Univ. of Texas United States	<b>Stampede</b> - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462462	5168.1	8520.1	4510	★
8	Forschungszentrum Juelich (FZJ) Germany	<b>JUQUEEN</b> - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	458752	5008.9	5872.0	2301	
9	DOE/NNSA/LLNL United States	<b>Vulcan</b> - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	393216	4293.3	5033.2	1972	
10	Leibniz Rechenzentrum Germany	<b>SuperMUC</b> - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM	147456	2897.0	3185.1	3423	★

★ Collaborations have members with access to these machines and to ARCHER, NERSC, ... Some HPCs are already successfully used as part of Nordu Grid (Abisko, Abel, Triolith, C2PAP, Hydra)

# HPC Scheduling

- **HEP applications (such as Geant or ROOT) can effectively use a single core**
- **HPC is full, means that the system have allocated all the cycles it is able to deliver**
  - It is probably not all cycles it has
  - Just as there is room for sand in the jar of rocks, there's room for HEP jobs on even a “full” HPC



# HPC Scheduling. Cont'd



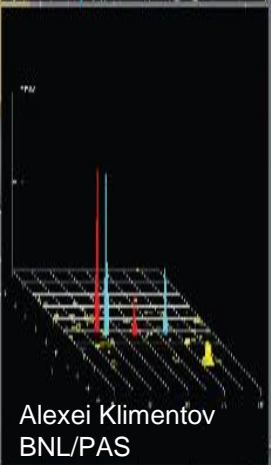
- This is not a typical workday view of HPC machine
- At this moment machine is 85% full, the largest open partition has 1024 nodes
- But the shortest job in the queue required 4096 nodes
- The scheduler will be happy to run a short job in R12
- 24hx3 backfill tests have been conducted on Titan together with OLCF team





**#2** **TOP 500<sup>®</sup>**  
SUPERCOMPUTER SITES

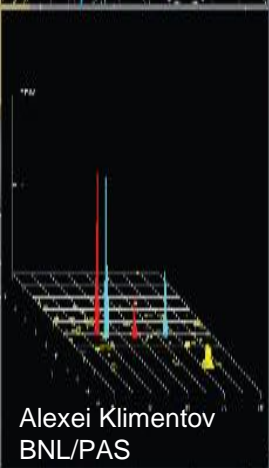
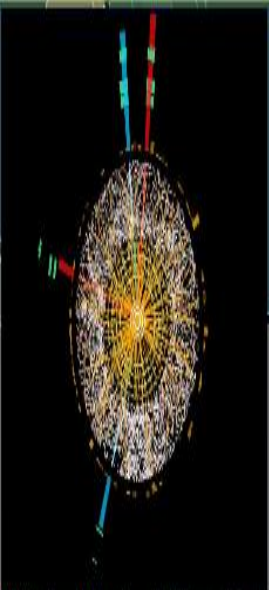
27 PFlops (Peak)  
18,688 compute nodes with GPUs  
299,008 CPU cores  
AMD Opteron 6200 @2.2 GHz (16 cores)  
32 GB Ram per node  
Nvidia TESLA K20x GPU per node  
32 PB disk storage (Luster)  
29 PB HPSS tape archive



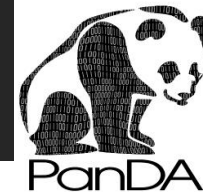
# Interfacing PanDA with Titan



- **BigPanDA project on Titan under ASCR auspices**
- **10M hours allocation for 2014-15 on Titan**
  - Access to EOS – new Cray XC30 machine at OLCF
  - Also we have access to NERSC via OSG and ATLAS allocations
- **Collaboration between ATLAS, ALICE, nEDM experiments**
- **Project members from CERN, BNL, UTA, ORNL, UTK, LBNL, ...**
  - Strong interest from OLCF, took responsibility for MPI wrapper base and docs
- **PanDA has potential to generate 300M hours per year**
- **Technology developed on Titan should be applicable for other HPC centers**
  - Functionality tests have been conducted on NERSC
  - Interest from ASGC, NRC-KI, JINR, Ostrava



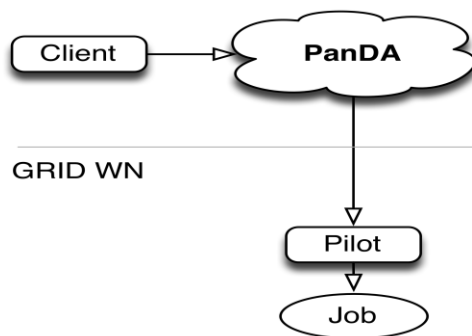
# Interfacing PanDA with Titan. Cont'd



- PanDA modular pilot augmented with HPC specific classes
- SAGA (Simple API for Grid Applications) framework as pilot's interface to HPC batch schedulers
  - <http://saga-project.github.io/saga-python/>
  - <http://www.ogf.org/documents/GFD.90.pdf>
- MPI wrapper/overlay scripts that allow to run multiple single threaded workload instances in parallel
- “Backfill” functionality in pilot

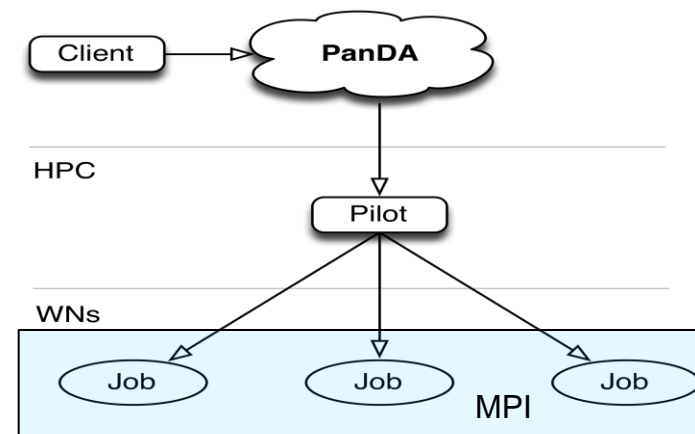
## Pilot on HPC with MPI wrapper

GRID Behavior



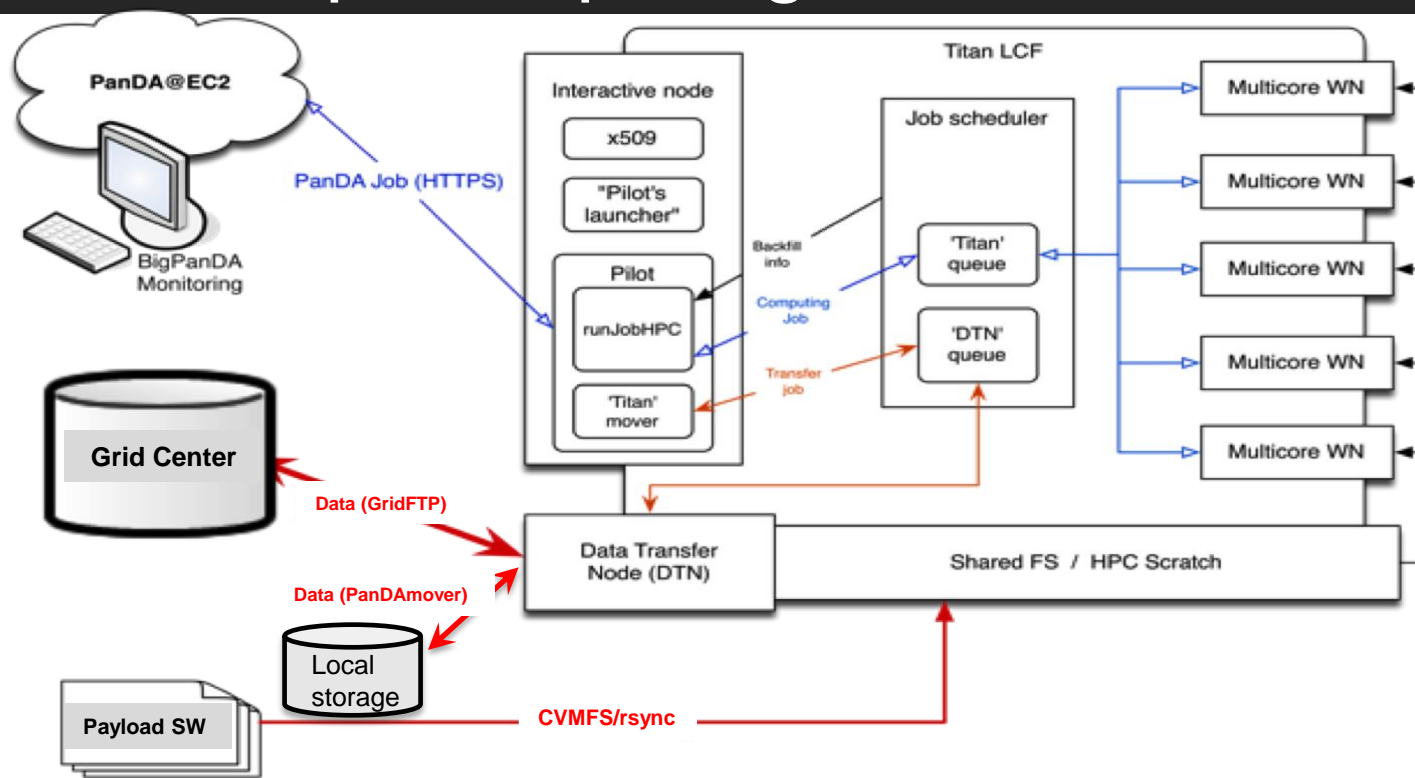
“One to One”

HPC Behavior



“One to Many”

# Extending PanDA to Oak Ridge Leadership Computing Facilities

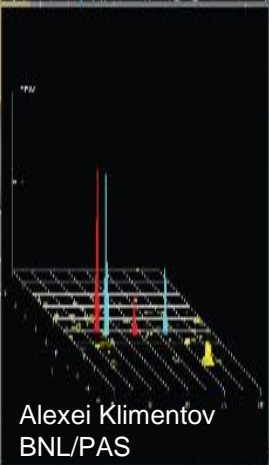
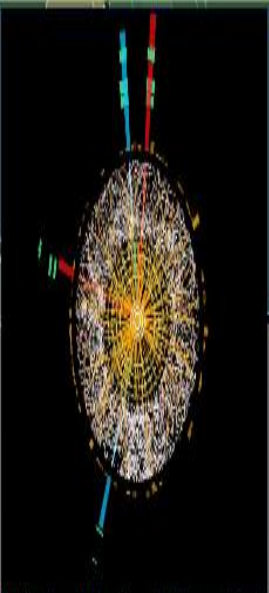


- **ATLAS (BNL, UTA), OLCF, ALICE (CERN,LBNL,UTK) :**
  - adapt PanDA for OLCF (Titan)
  - reuse existing PanDA components and workflow as much as possible.
  - PanDA connection layer runs on front-end nodes in user space. There is a predefined host to communicate with CERN from OLCF, connections are initiated from the front-end nodes
  - SAGA (a Simple API for Grid Applications) framework as a local batch interface.
  - Pilot (payload submission) is running on HPC interactive node and communicating with local batch scheduler to manage jobs on Titan.
  - Outputs are transferred to BNL T1 or to local storage

# Extending PanDA to Oak Ridge Leadership Computing Facilities



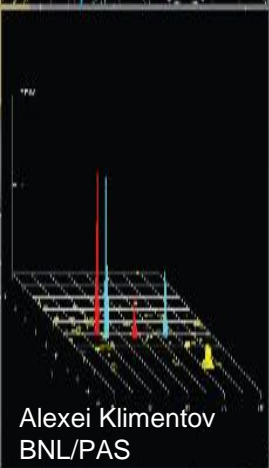
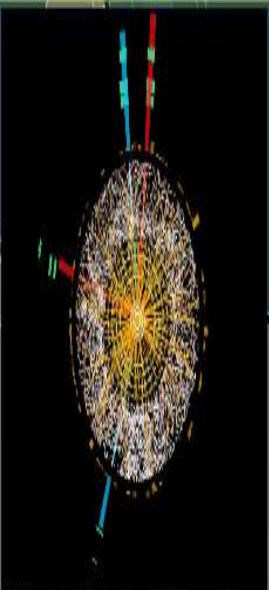
- **The initial demonstration was done for the Monte-Carlo event generators.**
  - Event generators are mostly computational and stand-alone code with small amount of stage-in/out data. In the LHC experiments 10-15% of all CPU resources are utilized to run event generators.
- **It is important to define the mode of payloads execution.**
  - One of possible scenarios to run HEP applications in backfill mode.
  - PanDA can be a vehicle to harvest Titan resources opportunistically.
- **PanDA pilot algorithm has been adopted to use backfill information and to submit jobs in backfill mode.**
  - Pilot periodically queries MOAB scheduler about available resources
  - Scheduler returns information about available (unscheduled) nodes and interval of availability
  - Pilot chooses the largest available block of nodes and submits jobs taken into account Titan's scheduling policy limitations
- **24h x 3 backfill stress test with realistic payload**
  - Main aim of tests was checking procedure for grabbing of available resources (without restriction), measure of waiting time.
  - Stable operations
  - 22k core-hours collected in 24h
  - Observed encouragingly short job wait time on Titan (4-55 min)
  - Observed IO issues on allocation of 1000 nodes or more



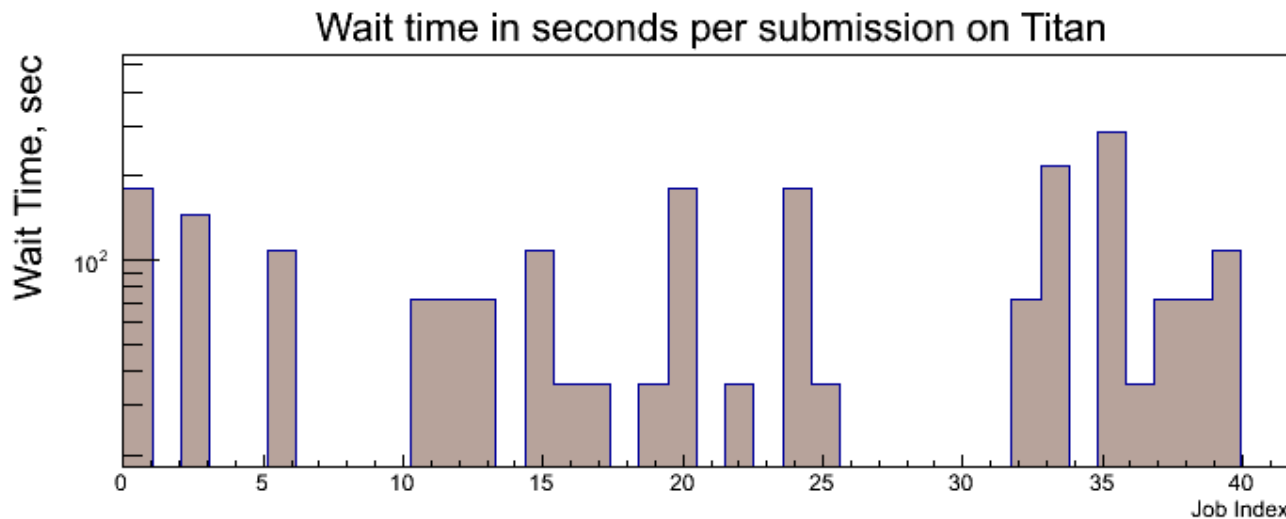
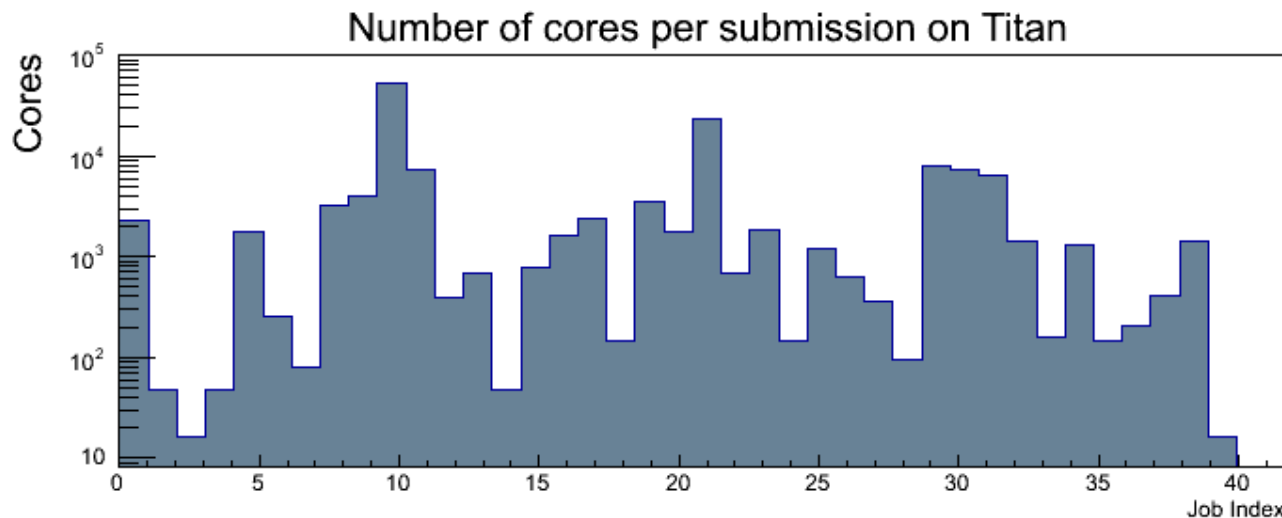
# Pilot stress tests on Titan



- **In May 2014 we ran the first 24 hour continuous job submission test via PanDA@EC2 with pilot in backfill mode, with MPI wrappers for two workloads from ATLAS and ALICE**
  - Stable operations
  - ~22k core hours collected in 24 hours
  - Observed encouragingly short job wait time on Titan ~4 minutes
- **Ran second set of tests in July 2014, with pilot modifications that were based on information obtained from the first test**
  - Limit on number of nodes removed in pilot
  - Job wait time limit introduced – 5 minutes
  - 145763 core hours collected
  - Average wait time ~70 sec
  - Observed IO related effects that need to be understood better
- **Final tests have been conducted in August**
  - Were able to collect ~ 200,000 core hours
  - Max number of nodes per job – 5835 (93360 cores)
    - Close to 75% ATLAS Grid in size!
  - Used ~2.3% of all Titan core hours or ~14.4% of free core hours



# PanDA pilot tests on Titan



Average wait time 70 seconds !

# Workloads on LCF/HPC



- Many elements of the HENP software stack (G4, ROOT, AliROOT, Alpgen, Sherpa, Madgraph, CLSHASTA, full G4 simulation for ALICE and EIC...) have been made to run on many different HPCs.
- There is a strong interest and support for the HPC activity.
  - ATLAS has been awarded something in excess of 63 million cpu-hours over the next 12 months. This is ~6% of ATLAS Grid use and half of our evgen.
- ALICE has allocation on Titan, NERSC, et al
- For the long term, jobs which use a full node (with some GPU acceleration) and scale well may be very competitive at the Leadership level, especially considering the scientific importance.



Alexei Klimentov  
BNL/PAS

# Resources Accessible via PanDA



Many  
Others



Titan System (Cray XK7)			
Peak Performance	27.1 PF 18,688 compute nodes	24.5 PF GPU	2.6 PF CPU
System memory	710 TB total memory		
Interconnect	Gemini High Speed Interconnect		3D Torus
Storage	Lustre Filesystem		32 PB
Archive	High-Performance Storage System (HPSS)		29 PB
I/O Nodes	512 Service and I/O nodes		



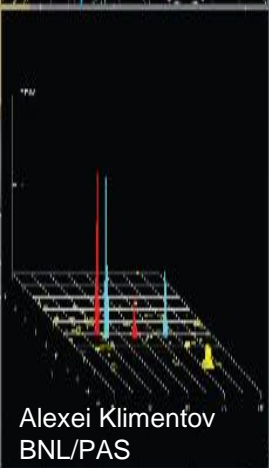
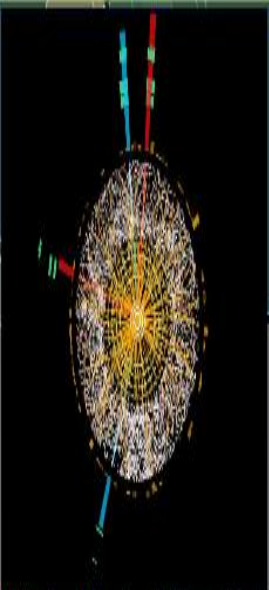
Google Cloud Platform



## ■ What is BigPanDA?

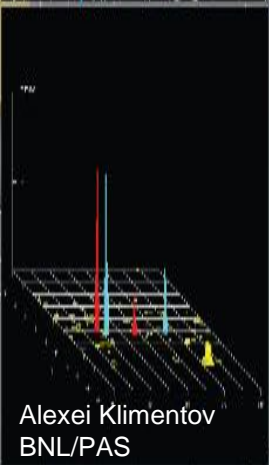
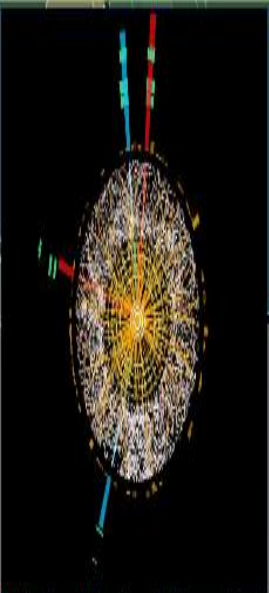
- Is it improved PanDA + ATLAS Production System after LHC Long Shutdown 1 ?
- Is it PanDA for AMS/ASGC/ALICE/CMS/LSST/NICA ?
- Is it new DOE funded BigData project to extend PanDA WMS?
- Is it PanDA integration with networking?
- Is it PanDA moving to HPC and other platforms?
- Is it PanDA for non-HEP sciences (eg. BioInfomatics) ?
- Is it Event Service?
- Is it something else?

**BigPanDA is all of the above – it is the evolution of PanDA, not just PanDA for BigData!**



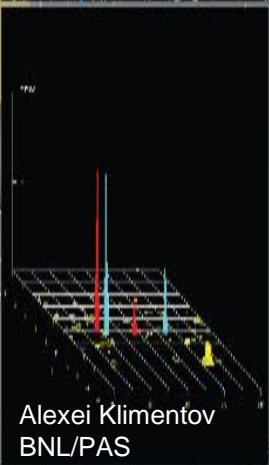
# Summary

- **2000s – The decade of the Grid**
- **First years of LHC data – Distributed Computing has helped deliver physics rapidly**
- **Entering a phase of computing evolution**
- **Challenges for computing – scale & complexity – will continue to increase dramatically**
- **The distributed computing model allows us to incorporate clouds and LCF/HPC centers and to use them efficiently for LHC Run 2 and beyond**
- **Access to the LCF coupled with collaborative help in the transformation of HEP code would be a major scientific contribution to the physics discoveries of the next ten years**
- **The work on extending PanDA to Leadership Computing Facilities has started. PanDA has been successfully ported on OLCF Titan and NERSC.**

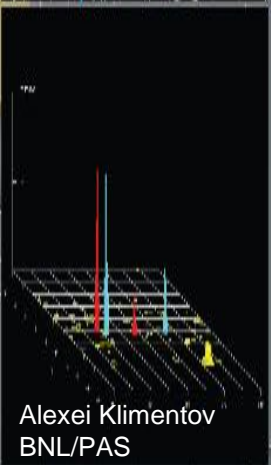
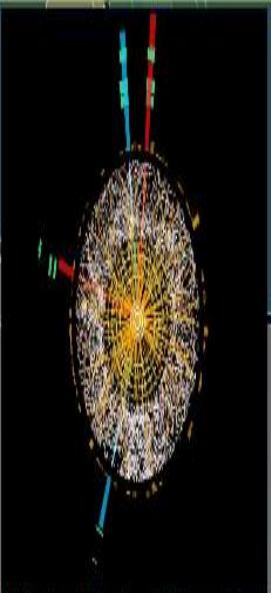


# Thanks

- This talk drew on presentations, discussions, comments, input from many
- Thanks to all, including those I've missed
- *D.Benjamin, I.Bird, D.Duellmann, R.Gardner, O.Keeble, T.LeCompte, A.Peters, M.Schulz...*



# Backup slides

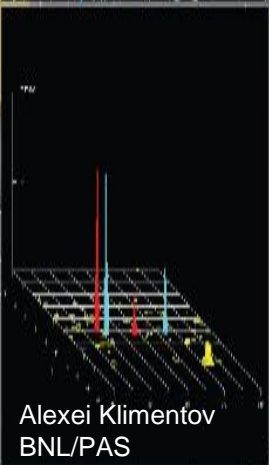
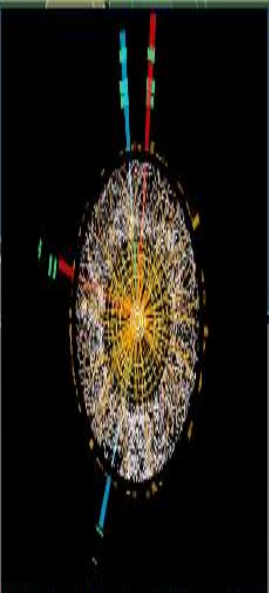


# The Growing PanDA Ecosystem



- **ATLAS PanDA**
  - US ATLAS, CERN, UK, DE, RF, ND, CA, Dubna, OSG ...
- **ASCR BigPanDA**
  - DoE funded project at BNL, UTA – 3 years
- **ANSE PanDA**
  - US NSF funded network project - CalTech, Michigan, Vanderbilt, UTA
- **High-Performance Computing and Cloud PanDA**
- **Taiwan PanDA – not only AMS**
- **LSST PanDA**
- **AliEn PanDA**
- **MegaPanDA**

*AliEn – ALICE analysis and Production framework  
AMS - Alpha-Magnetic Spectrometer experiment  
ANSE – Advanced Network Services (for LHC)  
ASCR – Advanced Scientific Computing  
LSST – Large Survey Synoptic Telescope  
OSG – Open Science Grid*



# Interfacing AliEn & PanDA



Client side

Server side

Grid resources

