Integrating Network Awareness in ATLAS Distributed Computing Using the ANSE Project

J.Batista, K.De, A.Klimentov, S.McKee, A.Petroysan

for the ATLAS Collaboration and Big PanDA team

University of Michigan

University of Texas at Arlington

Brookhaven National Laboratory

Joint Institute of Nuclear Research (Dubna)

CHEP 2015, Okinawa, Japan, April 13-17, 2015

The ATLAS Experiment at the LHC

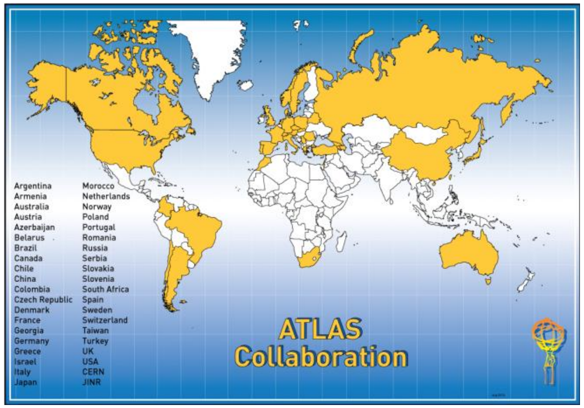

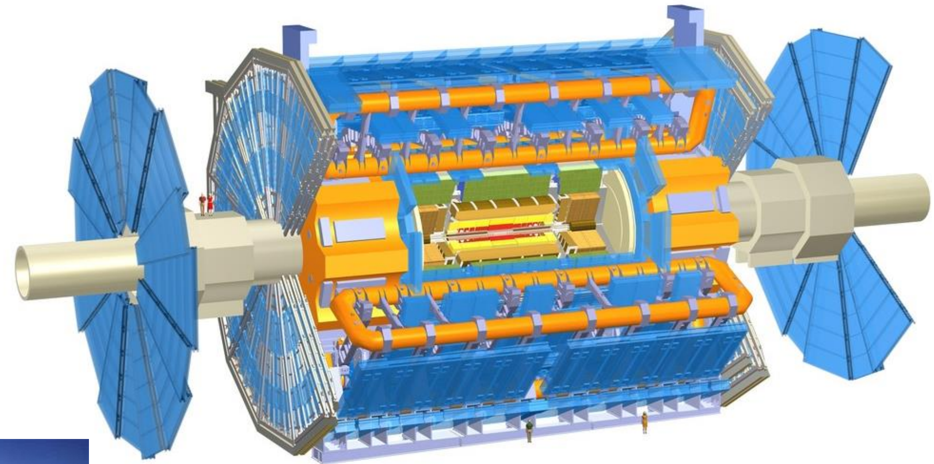

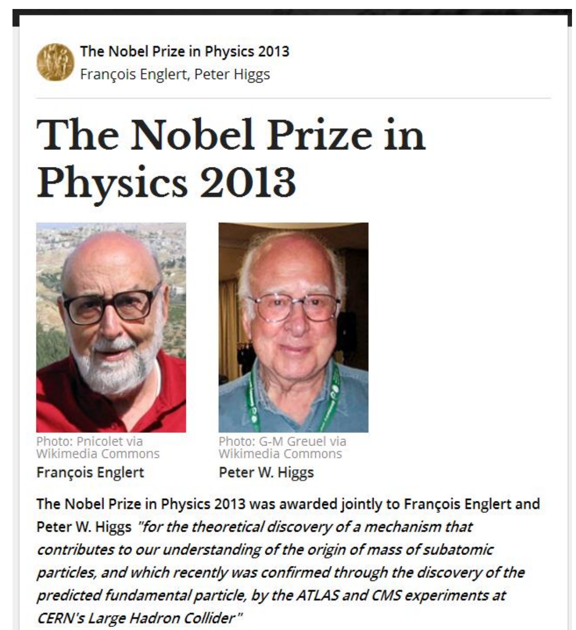
The Nobel Prize in Physics 2013
François Englert, Peter Higgs

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"*

Distributed Computing in ATLAS

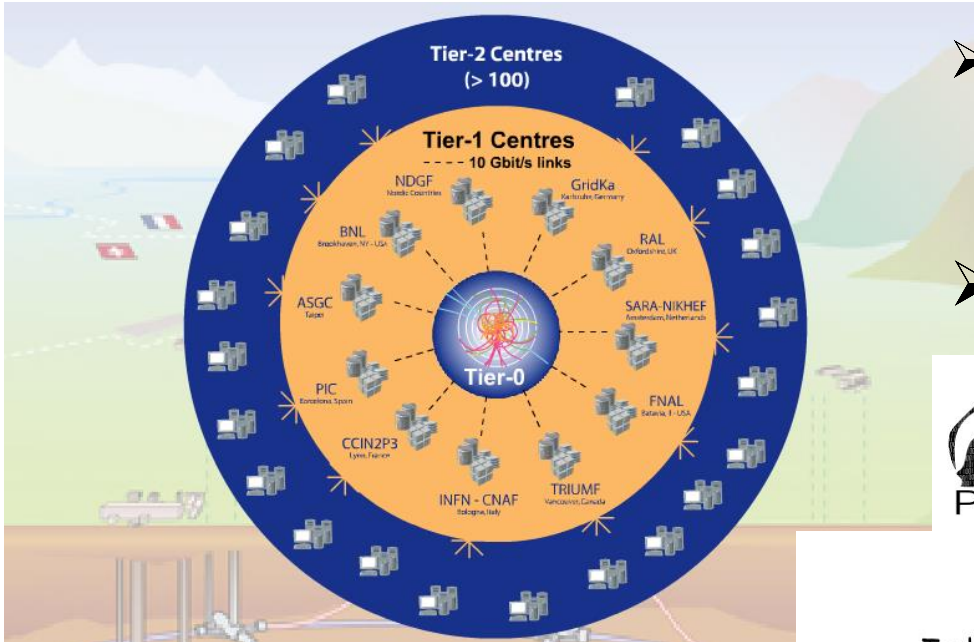

- ATLAS Computing Model :
  - 11 Clouds : 10 T1s + 1 T0 (CERN)
    Cloud = T1 + T2s + T2Ds
    T2D = multi-cloud T2 sites
- 2-16 T2s in each Cloud

Workload Management System

Task → Cloud :Task brokerage

Jobs → Sites :Job brokerage

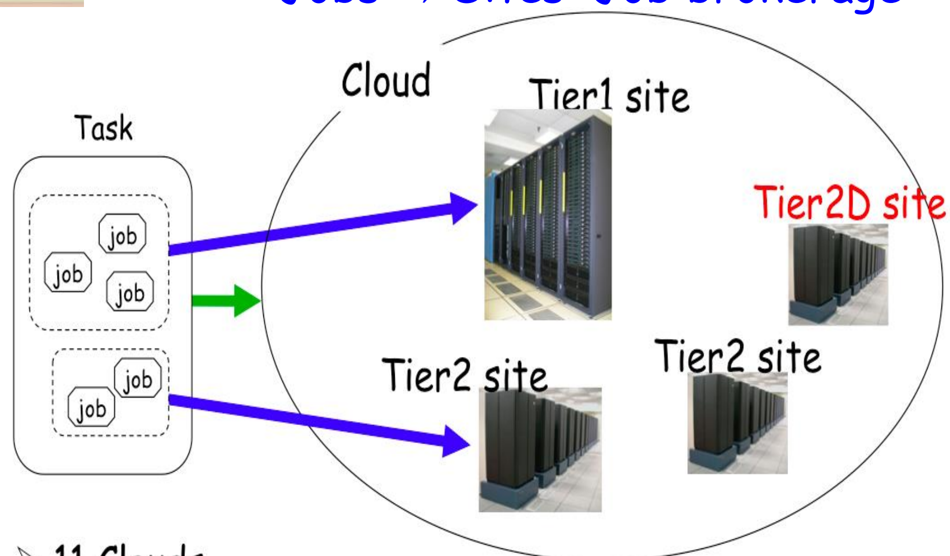

- Basic unit of work is a job:
  - Executed on a CPU resource/slot
  - May have inputs
  - Produces outputs
- JEDI – layer above PanDA to create jobs from ATLAS physics and analysis 'tasks'

Current scale – one million jobs per day

- 11 Clouds

PanDA Overview

- The PanDA workload management system was developed for the ATLAS experiment at the Large Hadron Collider
  - Hundreds of petabytes of data per year, thousands of users worldwide, many dozens of complex applications…
  - Leading to >400 scientific publications – and growing daily
  - Discovery of the Higgs boson, search for dark matter…
- A new approach to distributed computing
  - A huge hierarchy of computing centers working together
  - Main challenge – how to provide efficient automated performance
  - Auxiliary challenge – make resources easily accessible to all users
- Large Scale Networking enables systems like PanDA

- *Related CHEP Talks and Posters :*
  - *T.Maeno : The Future of PanDA in ATLAS Distributed Computing; CHEP ID 144*
  - *S.Panitkin : Integration of PanDA Workload Management System with Titan supercomputer at OLCF ; CHEP ID 152*
  - *M.Golosova : Studies of Big Data meta-data segmentation between relational and non-relational databases ; CHEP ID 115*
  - *Poster : Scaling up ATLAS production system for the LHC Run 2 and beyond: project ProdSys2*
  - *Poster : Integration of Russian Tier-1 Grid Center with High Performance Computers at NRC-KI for LHC experiments and beyond HENP; CHEP ID 100*
- *Reference :* https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA

Core Ideas in PanDA

- Make hundreds of distributed sites appear as local
  - Provide a central queue for users – similar to local batch systems
- Reduce site related errors and reduce 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
  - PanDA interacts with middleware – users see high level workflow
- Hide variations in infrastructure
  - PanDA presents uniform ‘job’ slots to user (with minimal sub-types)
  - Easy to integrate grid sites, clouds, HPC sites …
- Production and Analysis users see same PanDA system
  - Same set of distributed resources available to all users
  - Highly flexible – instantaneous control of global priorities by experiment

Resources Accessible via PanDA

LCG
WLCG
Worldwide LHC Computing Grid

egi

Many Others

NORDUGRID
Grid Solution for Wide Area Computing and Data Handling

Open Science Grid

nectar

NERSC

НАЦИОНАЛЬНЫЙ ИССЛЕДОВАТЕЛЬСКИЙ ЦЕНТР
«КУРЧАТОВСКИЙ ИНСТИТУТ»

amazon web services™ EC2

OLCF
Oak Ridge Leadership Computing Facility

FutureGrid
an XSEDE resource provider

Google Cloud Platform

About 150,000 job slots used continuously 24x7x365

PanDA Scale

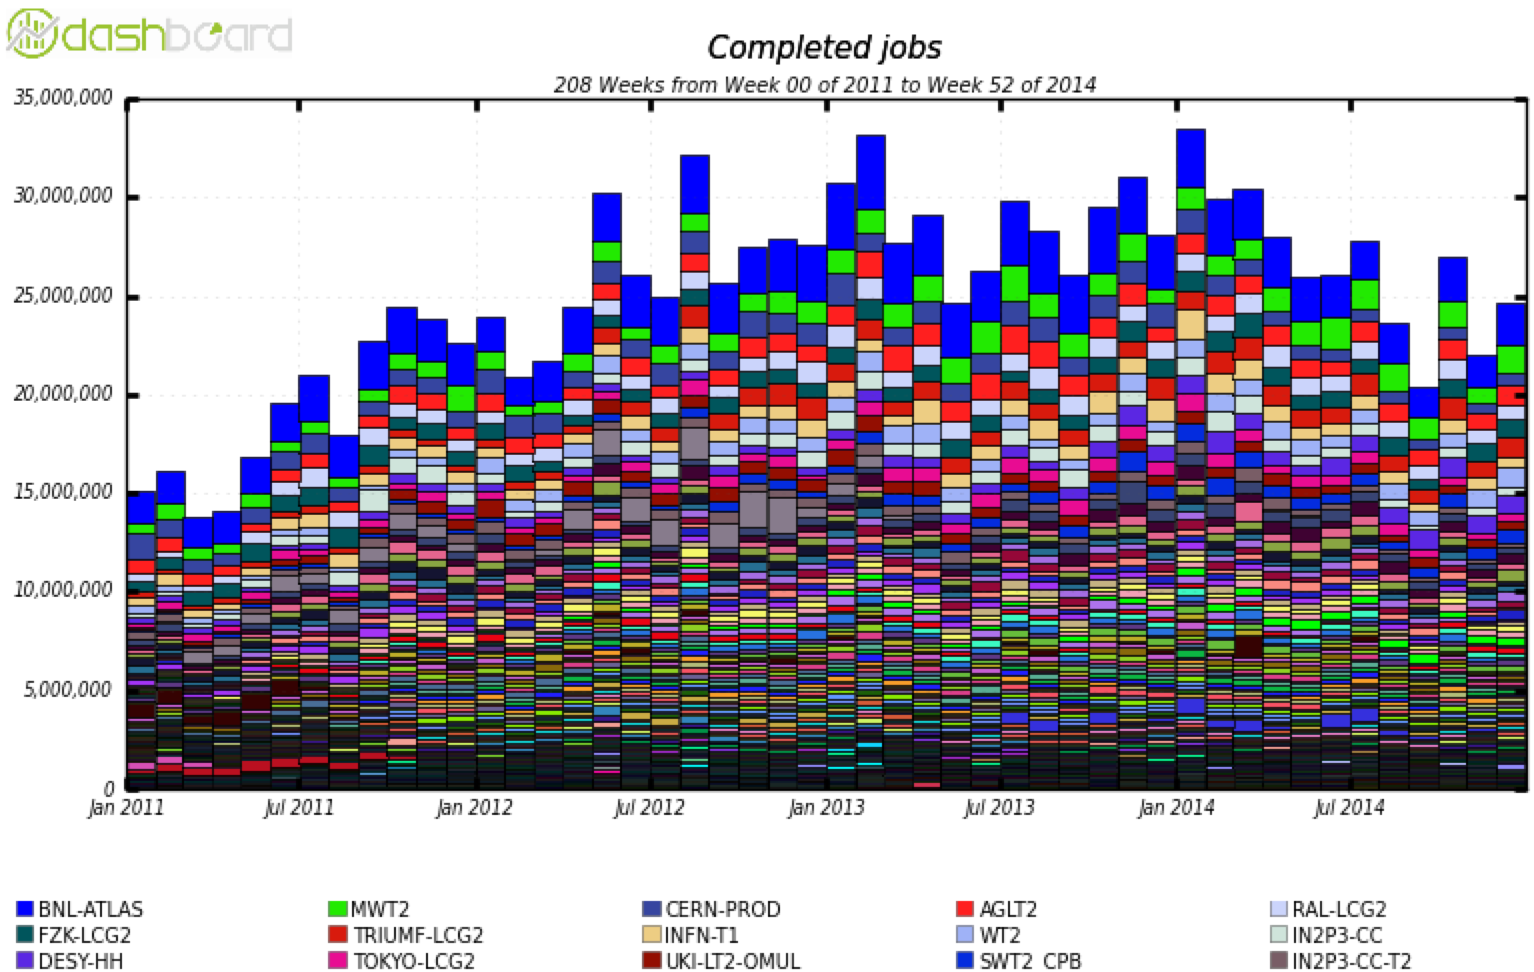

Current scale – 25M jobs completed every month at >hundred sites
First exascale system in HEP – 1.2 Exabytes processed in 2013

The Growing PanDA EcoSystem

- ATLAS PanDA core
  - US ATLAS, CERN, UK, DE, ND, CA, Russia, OSG …
- ASCR/HEP Big PanDA
  - DOE funded project at BNL, UTA – PanDA beyond HEP, at LCF
- CC-NIE ANSE PanDA
  - NSF funded network project - CalTech, Michigan, Vanderbilt, UTA
- HPC and Cloud PanDA – very active
- Taiwan PanDA – AMS and other communities
- Russian NRC KI PanDA, JINR PanDA – new communities
- AliEn PanDA, LSST PanDA, other experiments
- MegaPanDA (COMPASS, ALICE, NICA) …
  - RF Ministry of Science and Education funded project

PanDA Evolution. "Big PanDA"

- Generalization of PanDA for HEP and other data-intensive sciences
- Four main tracks :
  - **Factorizing the core :** Factorizing the core components of PanDA to enable adoption by a wide range of exascale scientific communities
  - **Extending the scope :** Evolving PanDA to support extreme scale computing clouds and Leadership Computing Facilities
  - ***Leveraging intelligent networks*** *: Integrating network services and real-time data access to the PanDA workflow*
  - **Usability and monitoring :** Real time monitoring and visualization package for PanDA

ATLAS Distributed Computing. Data Transfer Volume during Run 1.

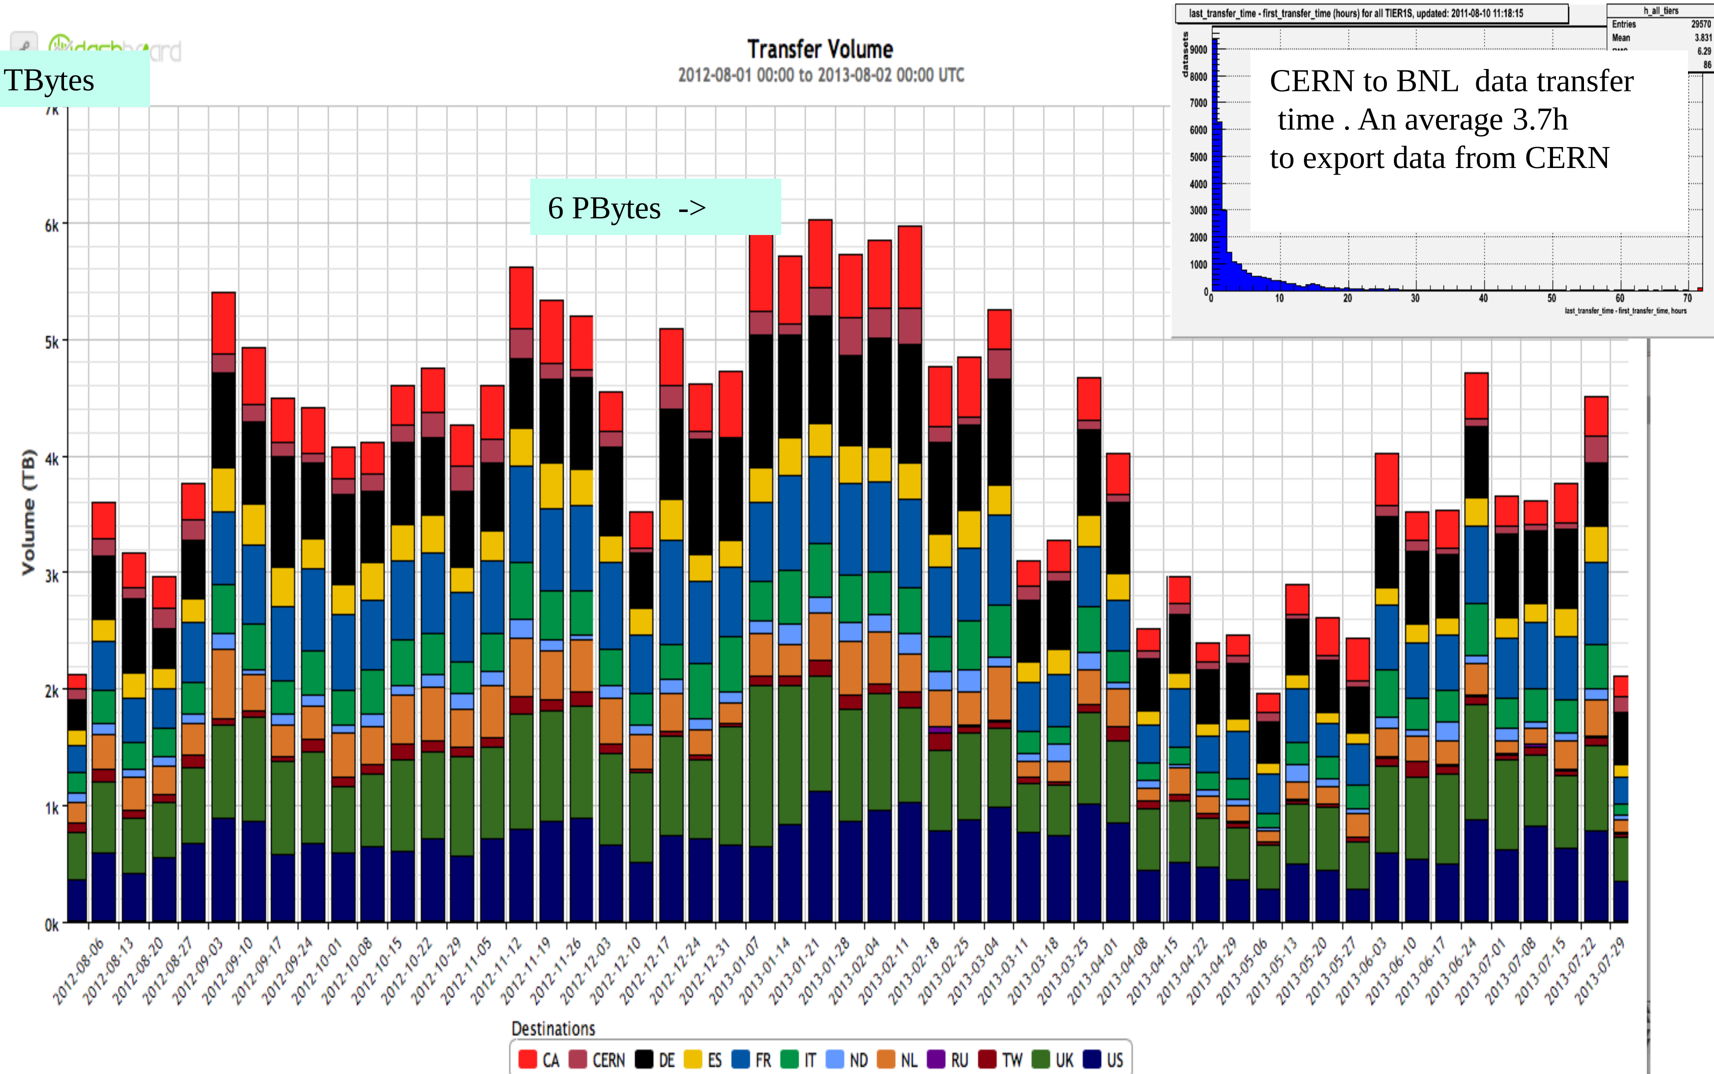

Data Transfer Volume in TB (weekly average). Aug 2012 - Aug 2013

PanDA and Networking

- PanDA as workload manager
  - PanDA automatically chooses job execution site
    - Multi-level decision tree – task brokerage, job brokerage, dispatcher
    - Also predictive workflows – like PD2P (PanDA Dynamic Data Placement)
  - Site selection is based on processing and storage requirements
    - Why not use network information in this decision?
    - Can we go even further – network provisioning?
  - Network knowledge useful 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
- Goal for PanDA
  - Direct integration of networking with PanDA workflow – never attempted before for large scale automated WMS systems

Sources of Network Information

- Distributed Data Management Sonar measurements
  - Actual transfer rates for files between all sites
  - This information is normally used for site white/blacklisting
  - Measurements available for small, medium, and large files
- perfSonar (PS) measurements
  - perfSonar provides dedicated network monitoring data
  - All WLCG sites are being instrumented with PS boxes
- Federated XRootD (FAX) measurements
  - Read-time for remote files are measured for pairs of sites
  - Using standard PanDA test jobs (HammerCloud jobs)

Network Data Monitoring. Sonar.

ATLAS dashboard | Index | Expanded Table

Show 200 entries | Copy | Print | Save | view: Sonar

Site Name	SrcSite	SrcCloud	SrcTier	DstSite	DstCloud	DstTier	Prio
AGLT2_to_OU_OCHEP_SWT2	AGLT2	US	T2D	OU_OCHEP_SWT2	US	T2	6
AGLT2_to_SWT2_CPB	AGLT2	US	T2D	SWT2_CPB	US	T2	6
AUSTRALIA-ATLAS_to_OU_OCHEP_SWT2	Australia-ATLAS	CA	T2	OU_OCHEP_SWT2	US	T2	2
AUSTRALIA-ATLAS_to_SWT2_CPB	Australia-ATLAS	CA	T2	SWT2_CPB	US	T2	2
BEIJING-LCG2_to_OU_OCHEP_SWT2	BEIJING-LCG2	FR	T2D	OU_OCHEP_SWT2	US	T2	2
BEIJING-LCG2_to_SWT2_CPB	BEIJING-LCG2	FR	T2D	SWT2_CPB	US	T2	5
BNL-OSG2_to_OU_OCHEP_SWT2	BNL-ATLAS	US	T1	OU_OCHEP_SWT2	US	T2	8
BNL-OSG2_to_SWT2_CPB	BNL-ATLAS	US	T1	SWT2_CPB	US	T2	8
CA-MCGILL-CLUMEQ-T2_to_OU_OCHEP_SWT2	CA-MCGILL-CLUMEQ-T2	CA	T2D	OU_OCHEP_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_OCHEP_SWT2	CA-SCINET-T2	CA	T2D	OU_OCHEP_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_OCHEP_SWT2	CA-VICTORIA-WESTGRID-T2	CA	T2D	OU_OCHEP_SWT2	US	T2	2
CA-VICTORIA-WESTGRID-T2_to_SWT2_CPB	CA-VICTORIA-WESTGRID-T2	CA	T2D	SWT2_CPB	US	T2	5

Network Measurements. perfSonar

Search...

DDM Sonar						perfSONAR						
AvgBRS (MB/s)	EvS	AvgBRM (MB/s)	EvM	AvgBRL (MB/s)	EvL	MinThr (MB/s)	AvgThr (MB/s)	MaxThr (MB/s)	MinPL	AvgPL	MaxPL	FAX xrdcp rate
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

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

Example: Faster User Analysis

- First use case for network integration with PanDA
- Goal - reduce waiting time for user jobs
  - User analysis jobs normally go to sites with local input data
  - This can occasionally lead to long wait times (jobs are re-brokered if possible, or PD2P data caching will make more copies eventually to reduce congestion)
  - While nearby sites with good network access may be idle
- Brokerage uses concept of ‘nearby’ sites
  - Use cost metric generated with Hammercloud tests
  - Calculate weight based on usual brokerage criteria (availability of CPU resources, data location, release…) plus new network transfer cost
  - Jobs will be sent to the site with best overall weight
- Throttling is used to manage load on network

Jobs Using FAX for Remote Access

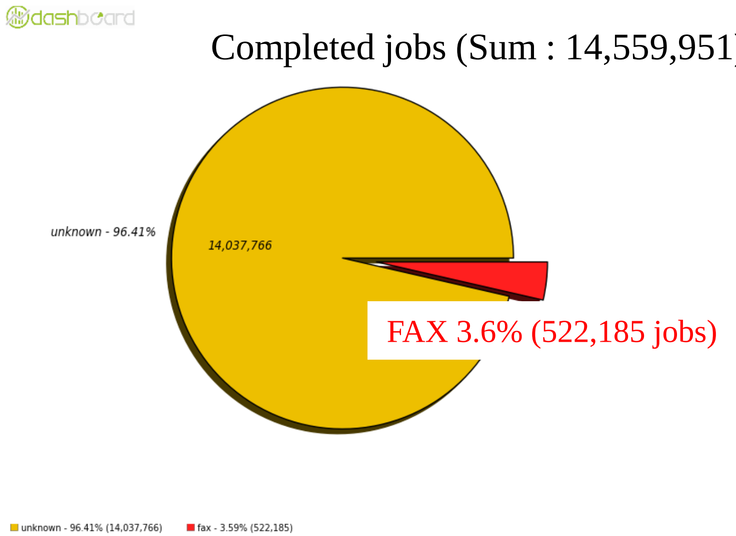

About 4% of jobs access data remotely Through federated XRootD

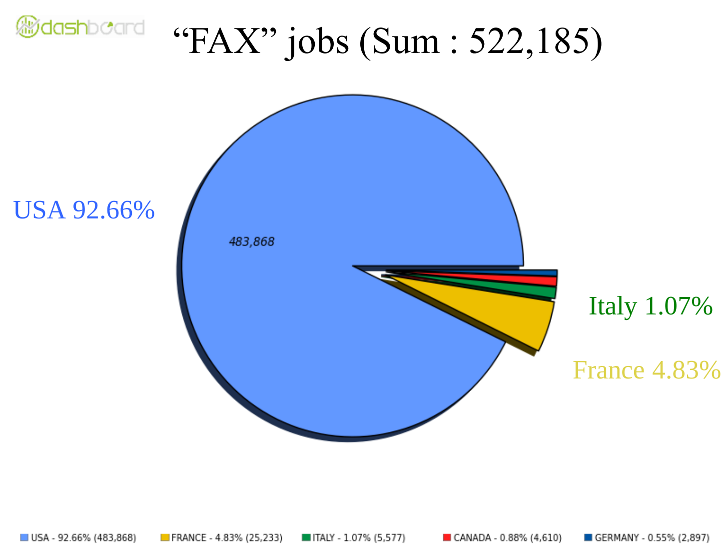

Remote site is selected based on network performance

Data for October 2014

Daily Remote Access Rates

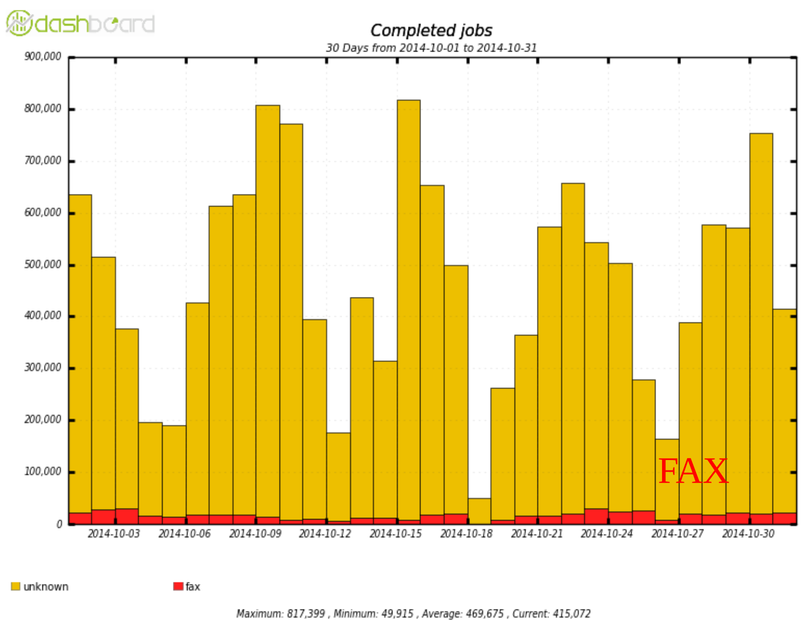

Data for October 2014

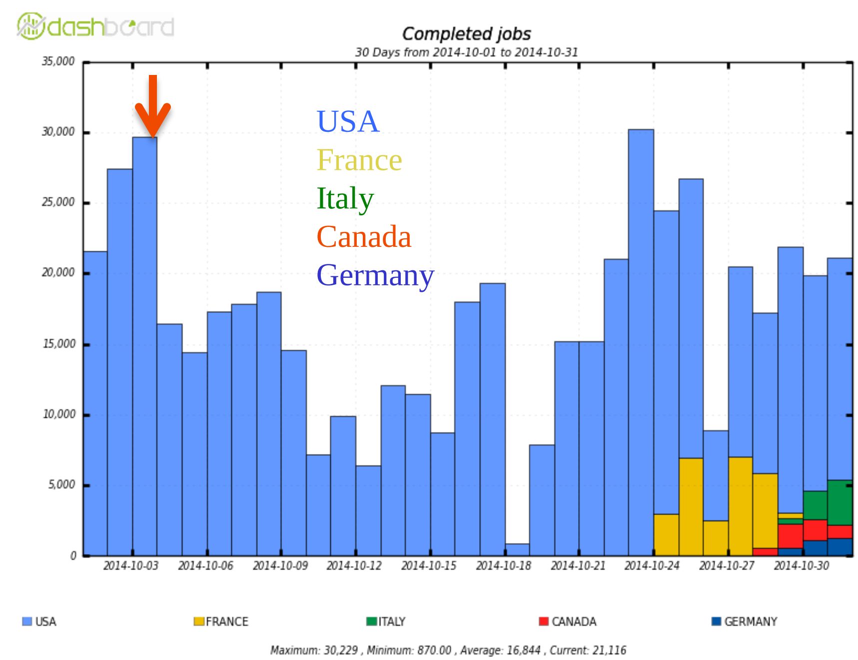

Jobs from Task 4199733

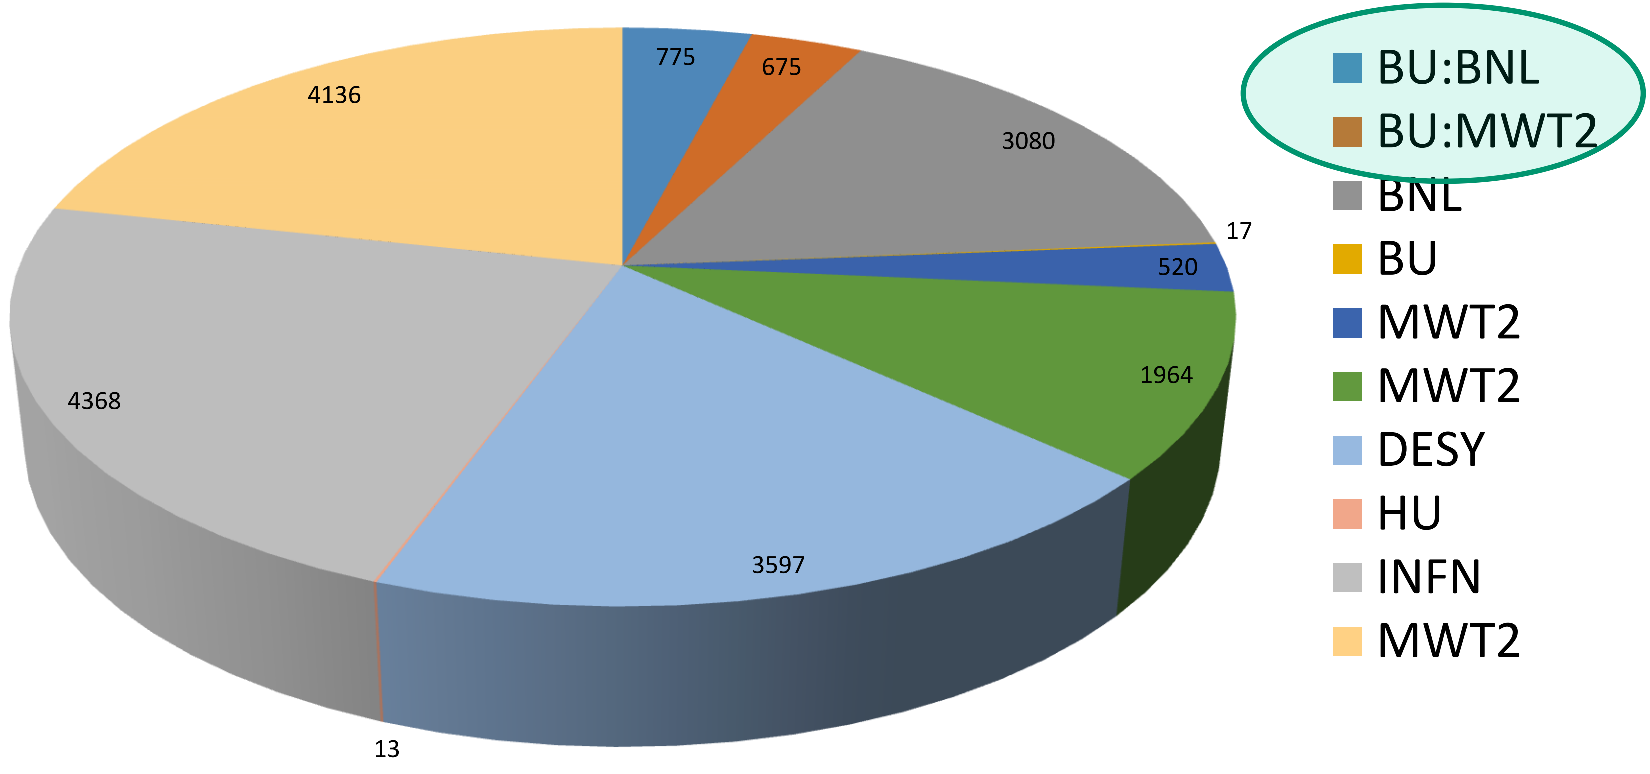

Job Wait Times for Example Task

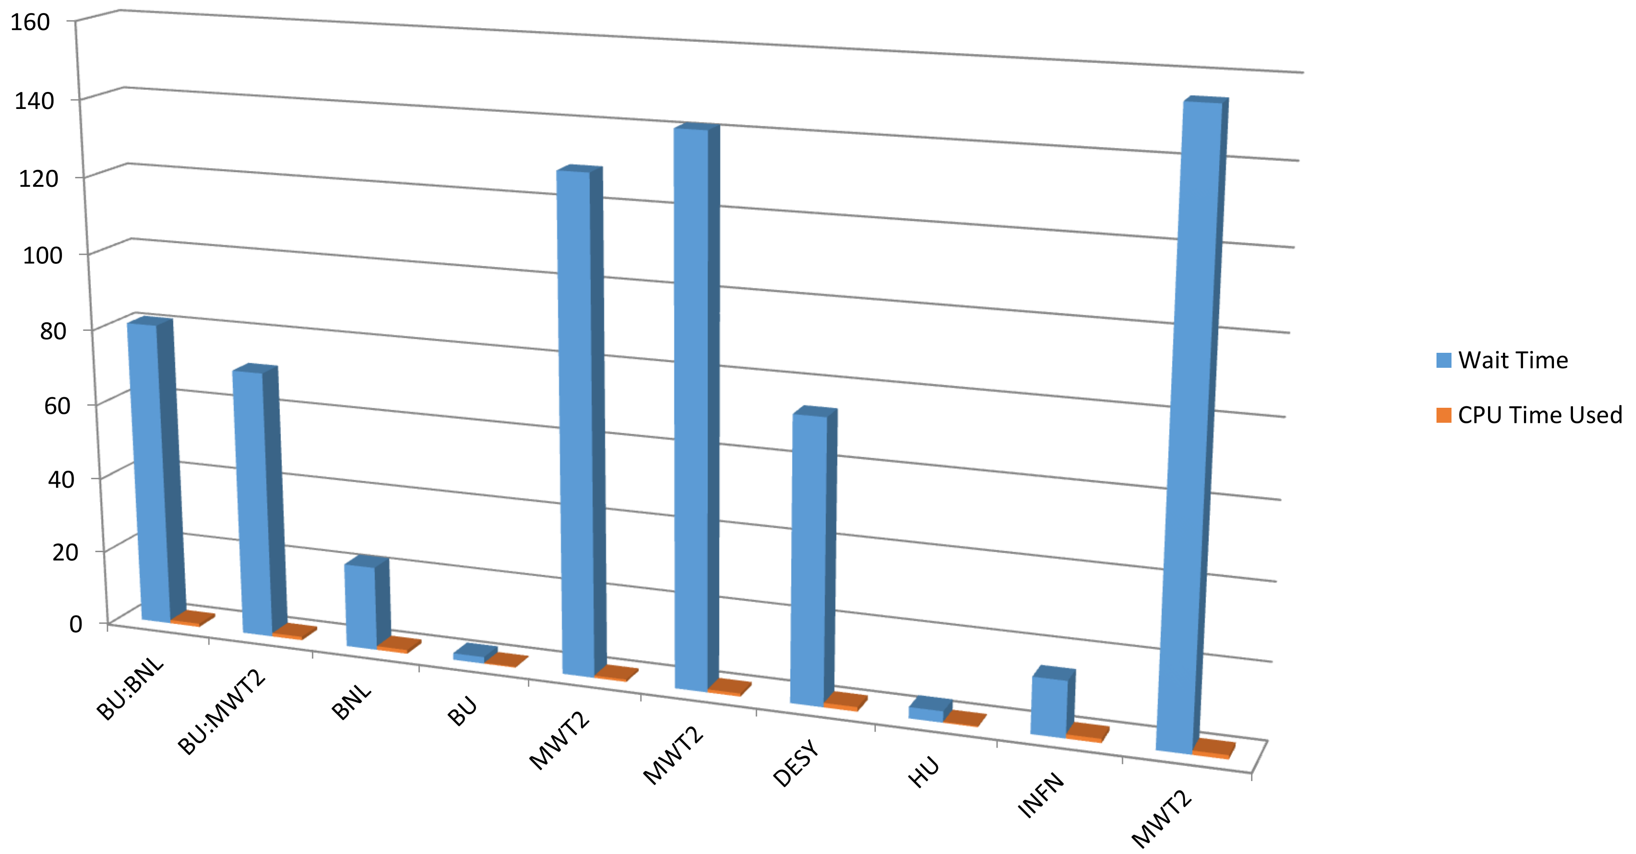

Cloud Selection

- Second use case for network integration with PanDA
- Optimize choice of T1-T2 pairings (cloud selection)
  - In ATLAS, production tasks are assigned to Tier 1’s
  - Tier 2’s are attached to a Tier 1 cloud for data processing
  - Any T2 may be attached to multiple T1’s
  - Currently, operations team makes this assignment manually
  - Automate this using network information

Example of Cloud Selection

5 best T1s for SouthWest T2 in TX, USA

	RRC-KI-T1	SARA-MATRIX	TRIUMF	ARC	INFN-T1
SWT2_CPB	40.11	37.94	25.96	25.04	18.39

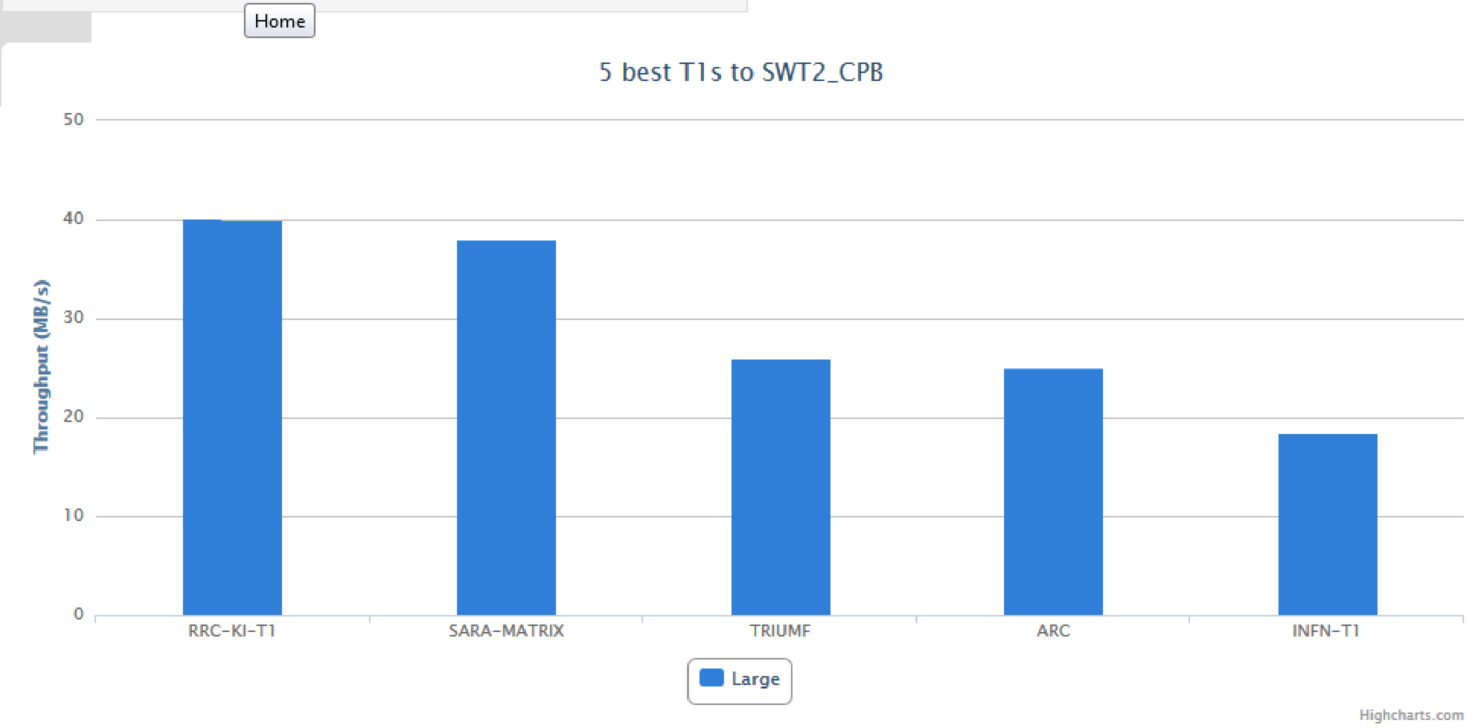

Summary

- The knowledge of network conditions, both historical and current, are used to optimize PanDA and other systems for ATLAS and allows to WMS control of end-to-end network paths can augment ATLAS's ability to effectively utilize its distributed resources.
- Optimal Workload Management System design should take network capability into account
- Network as a resource should be managed (i.e. provisioning)
- Many parts of distributed Workload Management Systems like PanDA can benefit from better integration with networking
- Initial results look promising
- Next comes SDN/provisioning
- We are on the path to use Network as a resource, on similar footing as CPU and storage

Back up slides

Dataflow

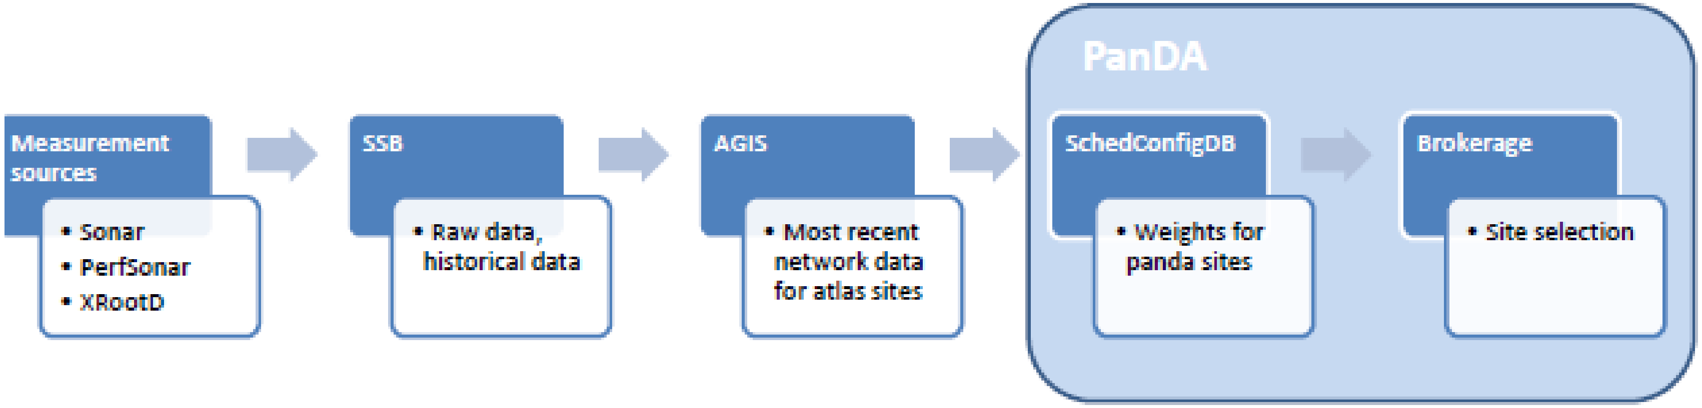

- Data is being transformed
  - Historical to most recent
  - Mb/sec to weights
  - Atlas sites to panda queues