Developments for ATLAS Supercomputing and Production system

R. Mashinistov et al

Main topics

- **Introduction**
  - WLCG
  - PanDA
  - Production System
- **Supercomputing**
  - ATLAS Pledged Resources
  - Supercomputing in HEP
  - HPC@NRC KI
- **Production System**
  - New Task interface
- **Plans**

WLCG - Worldwide LHC Computing Grid

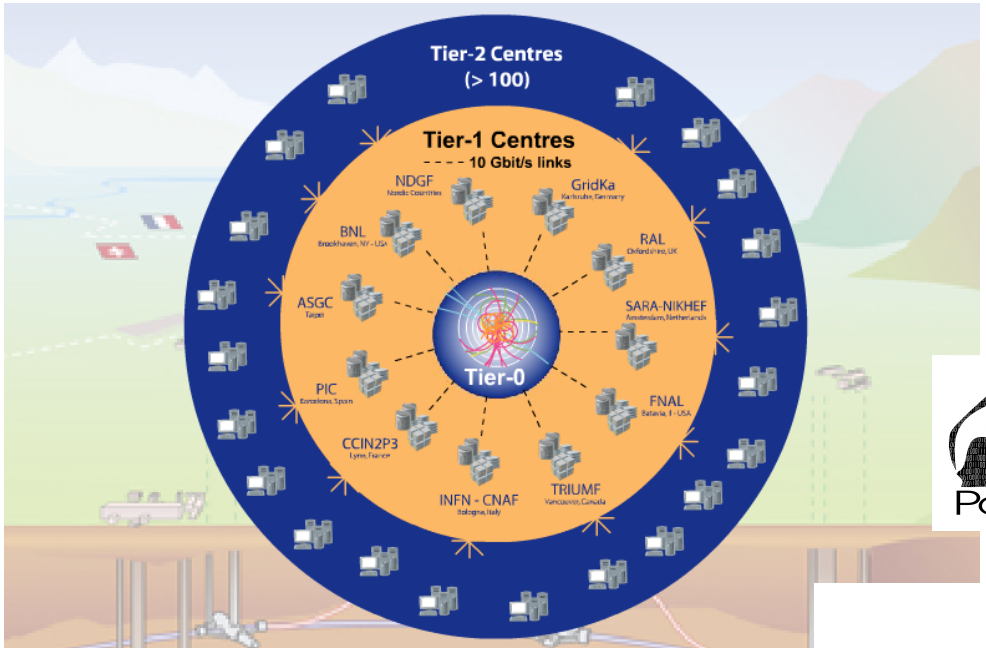

- ATLAS Computing Model :
- 11 Clouds : 10 T1s + 1 T0 (CERN)
  - Cloud = T1 + T2s + T2Ds
  - T2D = multi-cloud T2 sites
- 2-16 T2s in each Cloud

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

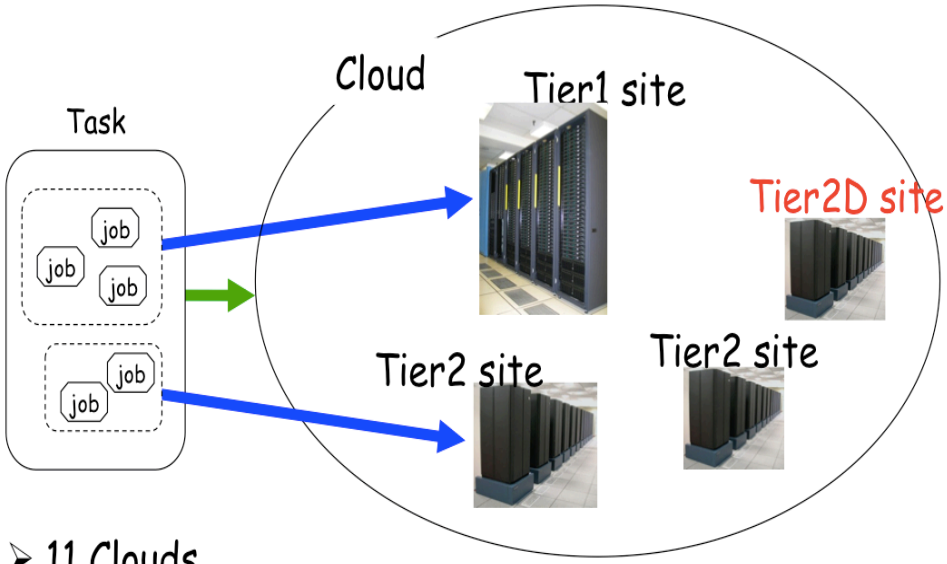

- 11 Clouds

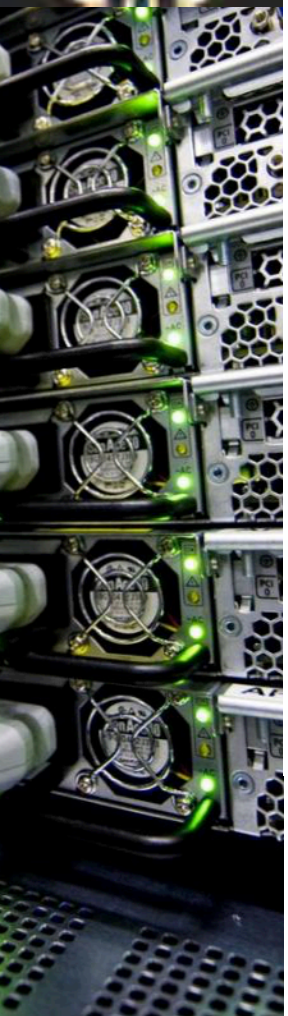

PanDA Workload Management

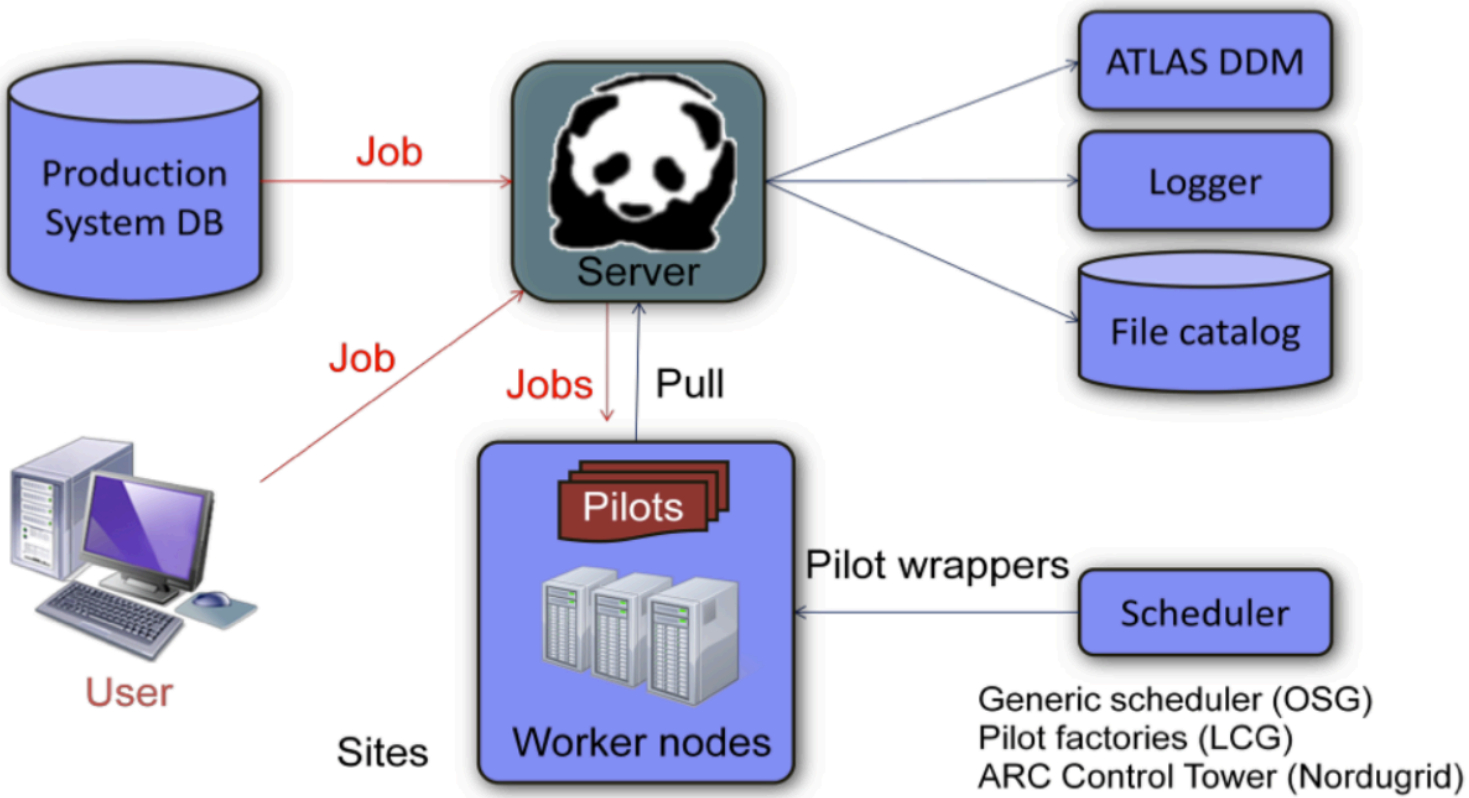

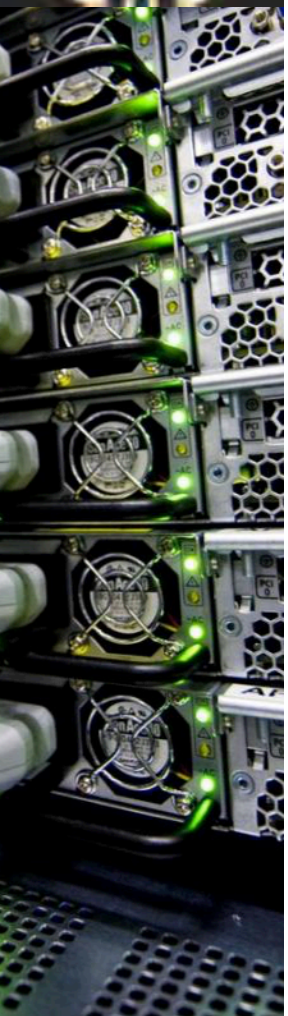

ATLAS Production System

- Prodsys2 - DEFT (Database Engine for Tasks) and JEDI (Job Execution and Definition Interface)

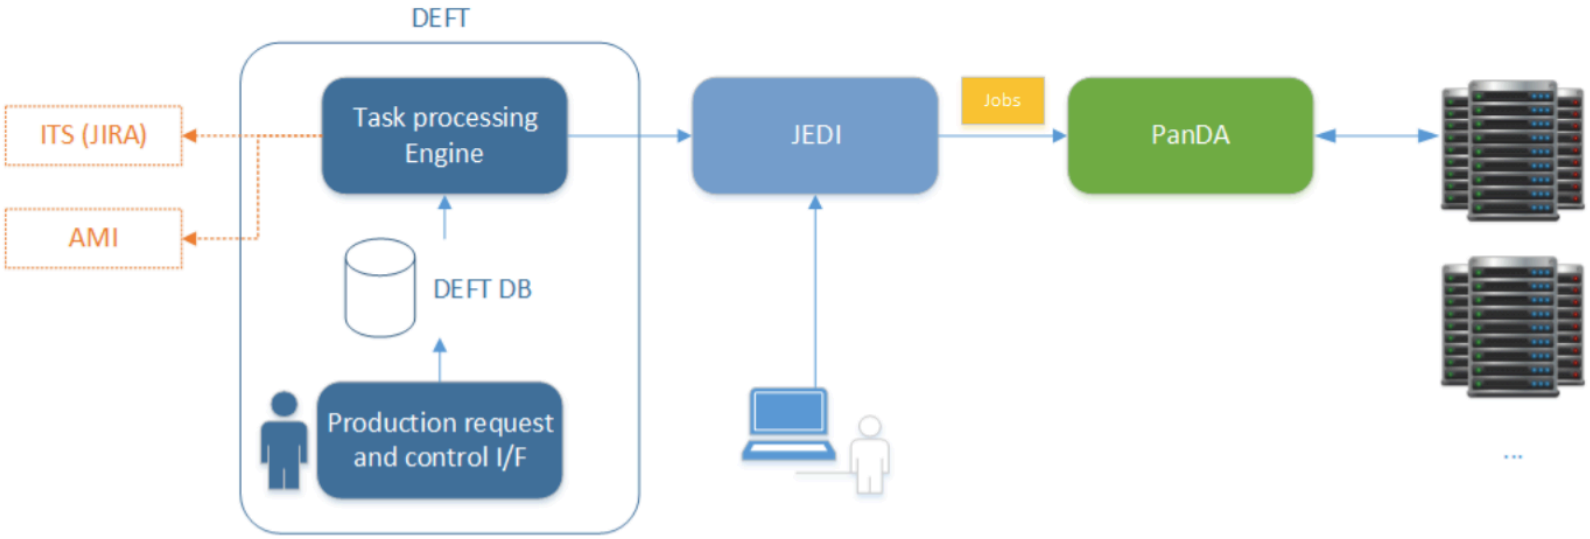

ATLAS pledged resources

Totally Tiers provide

Physical CPU	Logical CPU	HEPSPEC06	CPU Pledge
179,625	532,910	5,716,182	3,083,512
Total Online Storage,Gb	**Disk Pledge, Gb**	**Total Nearline Storage, Gb**	**Tape Pledge, Gb**
325,934,747	249,432,000	240,447,948	274,592,000

Dedicated to ATLAS:

	CPU, HEP-SPEC06	Disk, Tbytes	Tape, Tbytes
Required	1175000	103000	98000
Provided	1275226	110293	103190
% from total	22	34	43

http://wlcg-rebus.cern.ch/apps/pledges/resources/

Supercomputing in High Energy Physics. Leadership Computing Facilities

- Modern High Performance Computing involves a very large number of cores connected through a high speed network
- It is not a requirement that HPCs are able to run any possible task
- What matters is the total number of cycles that can be offloaded from the Grid
- Expanding PanDA from Grid to Leadership Class Facilities (LCF) required changes in the 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

Extending PanDA to Oak Ridge Leadership Computing Facilities

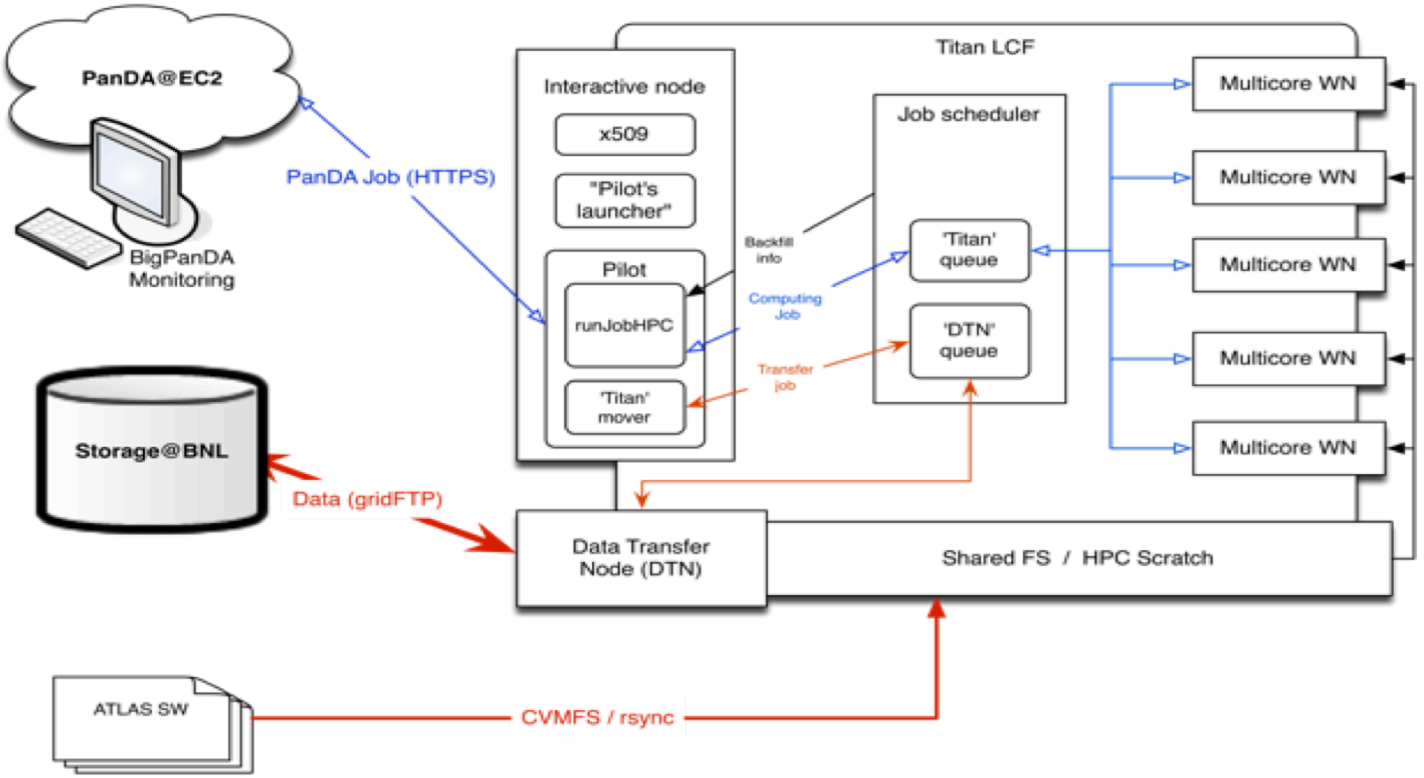

- ATLAS (BNL, UTA), OLCF, ALICE (CERN,LBNL,UTK) :
- SAGA (a Simple API for Grid Applications) framework as a local batch interface.
- Pilot (payload submission) is running on HPC interactive node and communicating with local batch scheduler to manage jobs on Titan.
- MPI wrapper/overlay scripts that allow to run multiple single threaded workload instances in parallel
- “Backfill” functionality in pilot
- Outputs are transferred to BNL T1 or to local storage

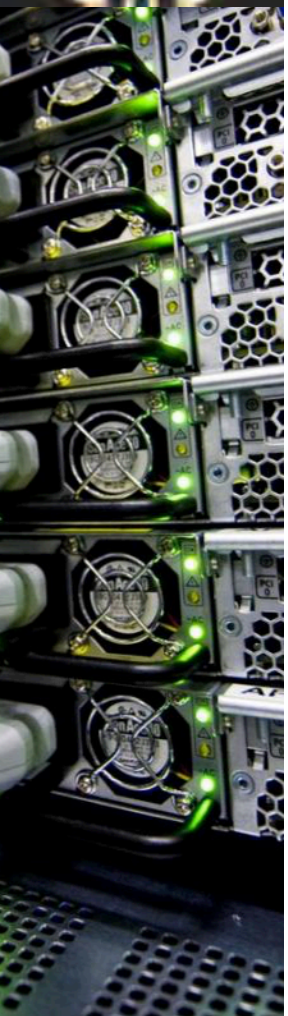

HPC2 supercomputer at NRC KI

High Performance Cluster - HPC2 second generation HPC with peak performance 122,9 TFLOPS (commissioned 2011).
#2 in 15-th issue of Russian top50 Supercomputers

- 10240 CPU cores = 1280 nodes 2x Intel Xeon E5450 3,00GHz 4 core 16 Gb RAM;
- UI node only allows to run jobs in batch system (SLURM) or to compile the code
- Shared FS Lustre for WN's and UI
- WN's has an access to WAN
- CVMFS connected to WN's
- Broadband to Tier-1 Storage Element

Integration of PanDA with HPC2 supercomputer at NRC KI

- Pilots are running on each node
- Number of Pilots = number of available cores (and number of virtual CondorLocal slots)
- Remote running via interactive node (UI)

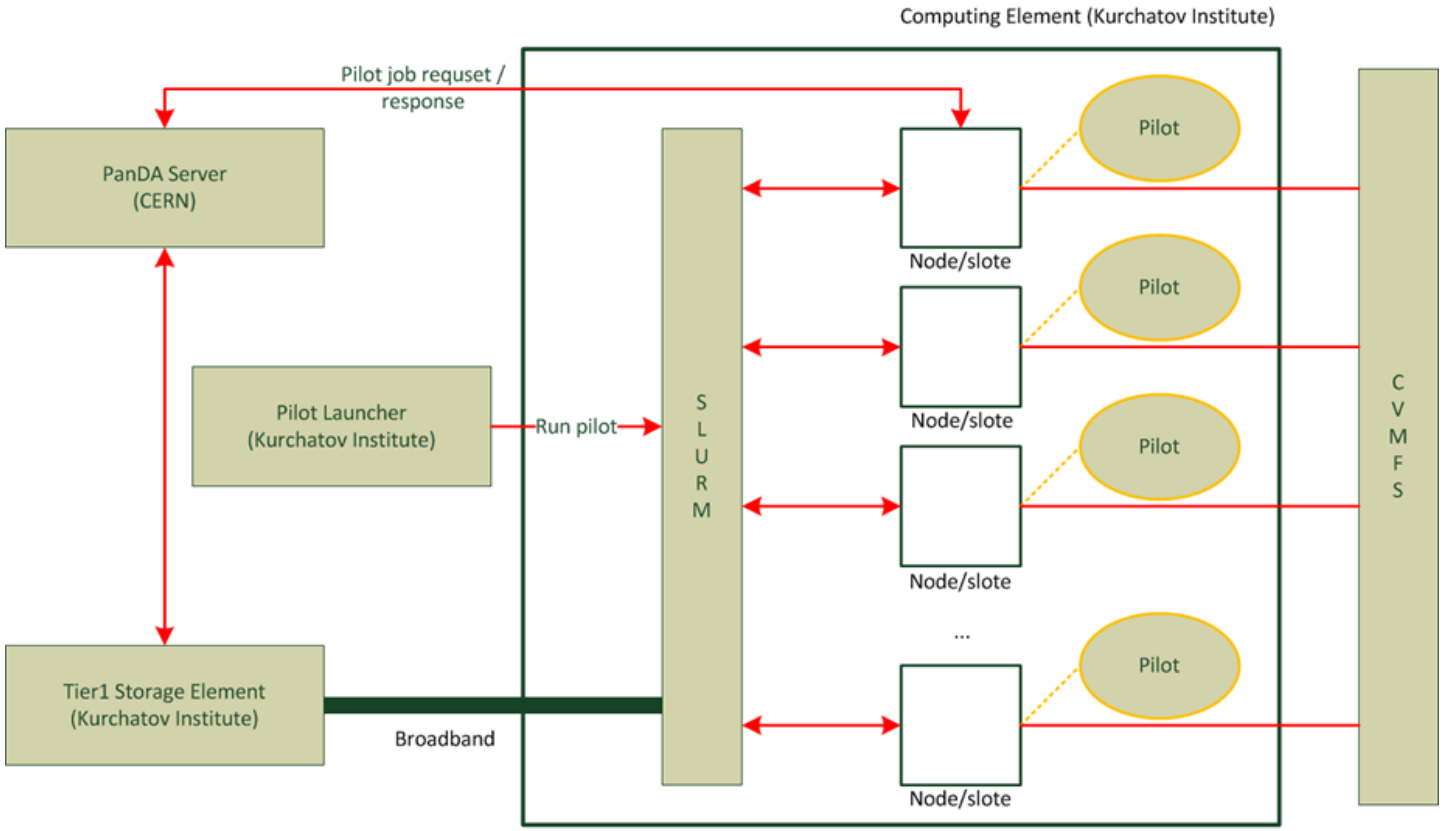

Integration of Russian Tier-1 Grid Center with High Performance Computers at NRC-KI (CHEP15)

- In 2014 the pioneering work on implementation of the portal, which will combine the Kurchatov Institute's Tier-1 site and the supercomputer HPC2 into Kurchatov Institute's Tier-1 center was done.
  - This center allows starting tasks optionally not only on Grid but on the supercomputer as well and to collect the results in the common storage.
- For ATLAS production and analysis tasks new PanDA sites within ATLAS Tier-1 site of National Research Center "Kurchatov Institute" were defined for supercomputer.
  - ANALY_RRC-KI-HPC site for user analysis – 2 nodes (16 cores, 2GB RAM per core)
    - User analysis tasks for MEPhI. RAW to D3PD reconstruction of high mu pp events (up to 70 interactions) for TRT performance at high occupancy study. (ATLAS TRT SW group)
    - Results presented at Nuclear Electronics & Computing, NEC-2015 (28 Sept- 3 Oct 2015, Montenegro)
      - https://cds.cern.ch/record/2007680
  - RRC-KI-HPC2 site for production – 32 nodes (256 cores, 2 GB RAM + 1 GB SWAP per core)

Integration scheme of Russian Tier-1 Grid Center with High Performance Computers at NRC-KI (CHEP15)

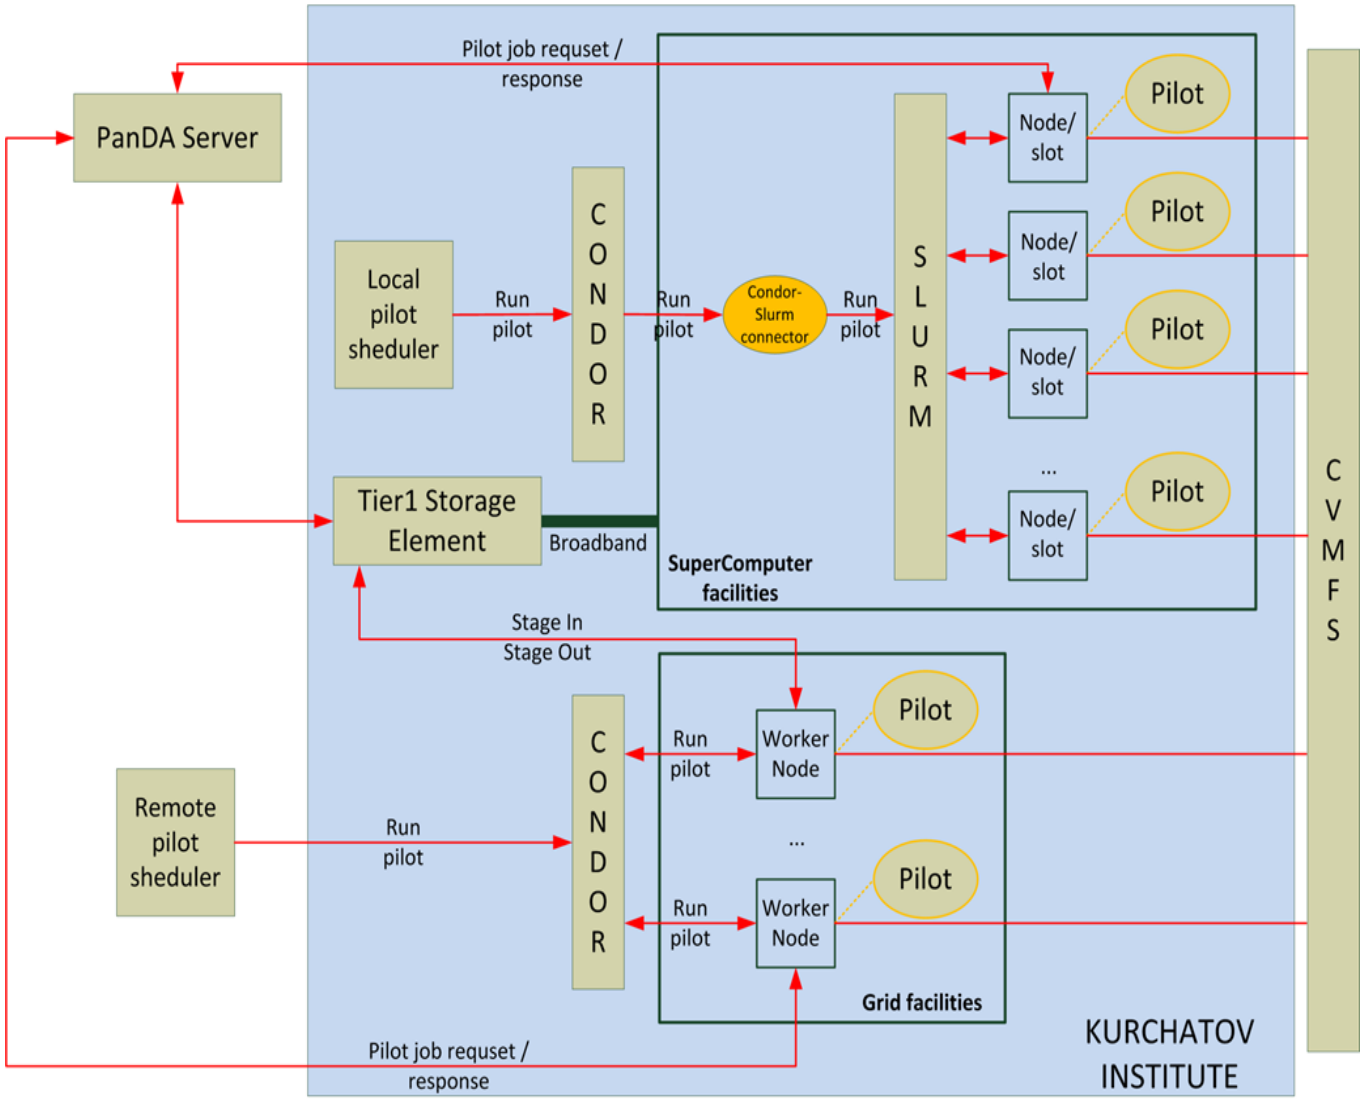

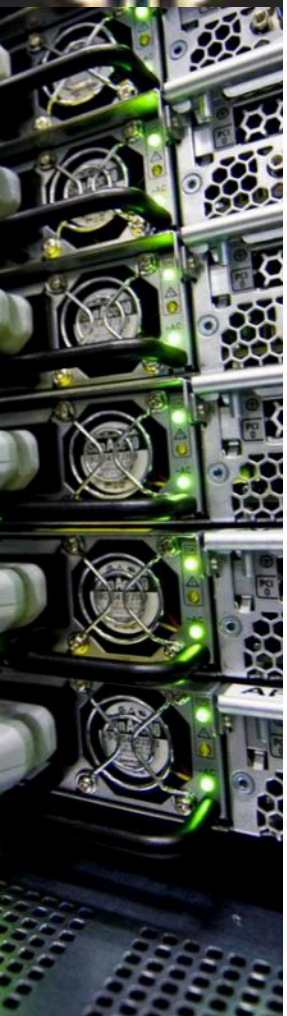

ATLAS site validation

dashboard

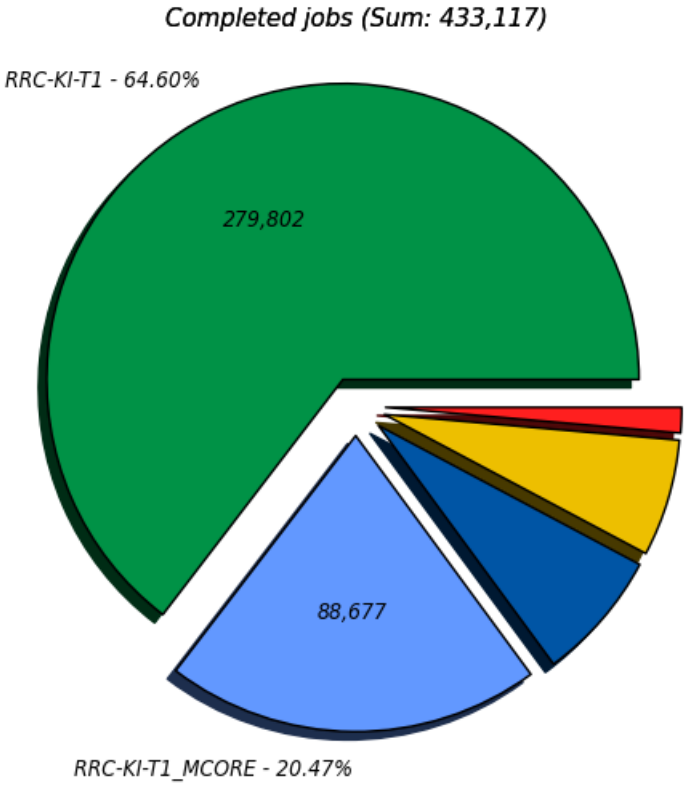

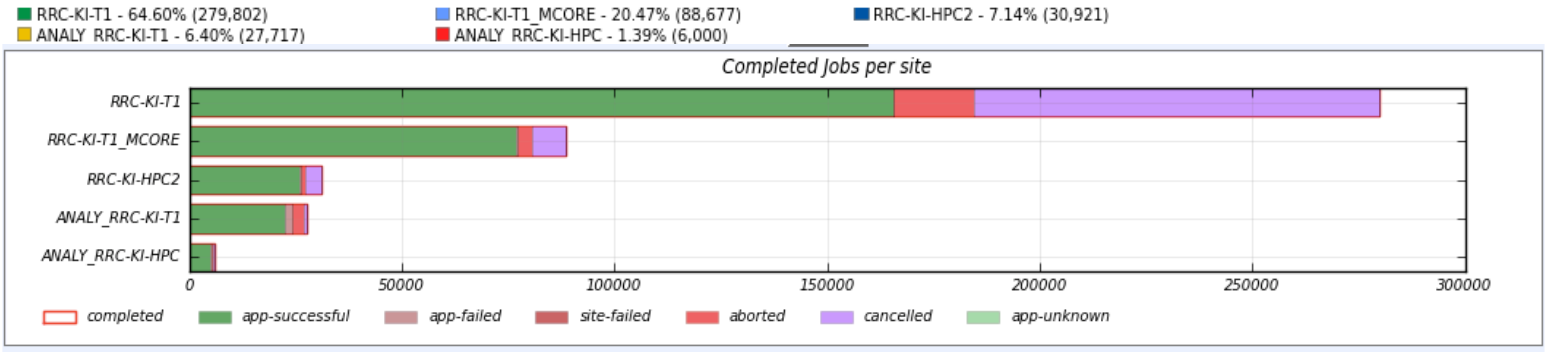

ATLAS site validation

dashboard

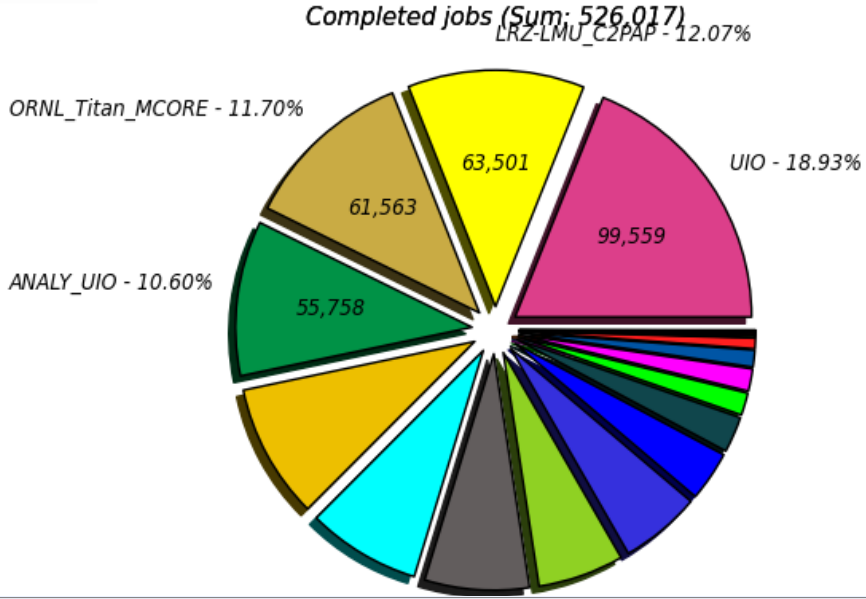

- UIO - 18.93% (99,559)
- ORNL_Titan_MCORE - 11.70% (61,563)
- ANALY_NSC - 9.22% (48,496)
- HPC2N - 7.14% (37,555)
- UIO_MCORE - 5.50% (28,939)
- NSC_MCORE - 2.45% (12,876)
- LRZ-LMU_C2PAP_MCORE - 1.51% (7,956)
- ANALY_HPC2N - 0.73% (3,818)
- BEIJING-TIANJIN-TH-1A_MCORE - 0.16% (848.00)
- CSCS-LCG2-HPC_MCORE - 0.04% (186.00)

- LRZ-LMU_C2PAP - 12.07% (63,501)
- ANALY_UIO - 10.60% (55,758)
- NSC - 7.76% (40,801)
- RRC-KI-HPC2 - 5.88% (30,921)
- HPC2N_MCORE - 3.40% (17,889)
- LRZ-LMU_MUC_MCORE - 1.53% (8,034)
- ANALY_RRC-KI-HPC - 1.14% (6,000)
- CSCS-LCG2-HPC - 0.17% (873.00)
- BEIJING-ERA_MCORE - 0.06% (294.00)

plus 2 more

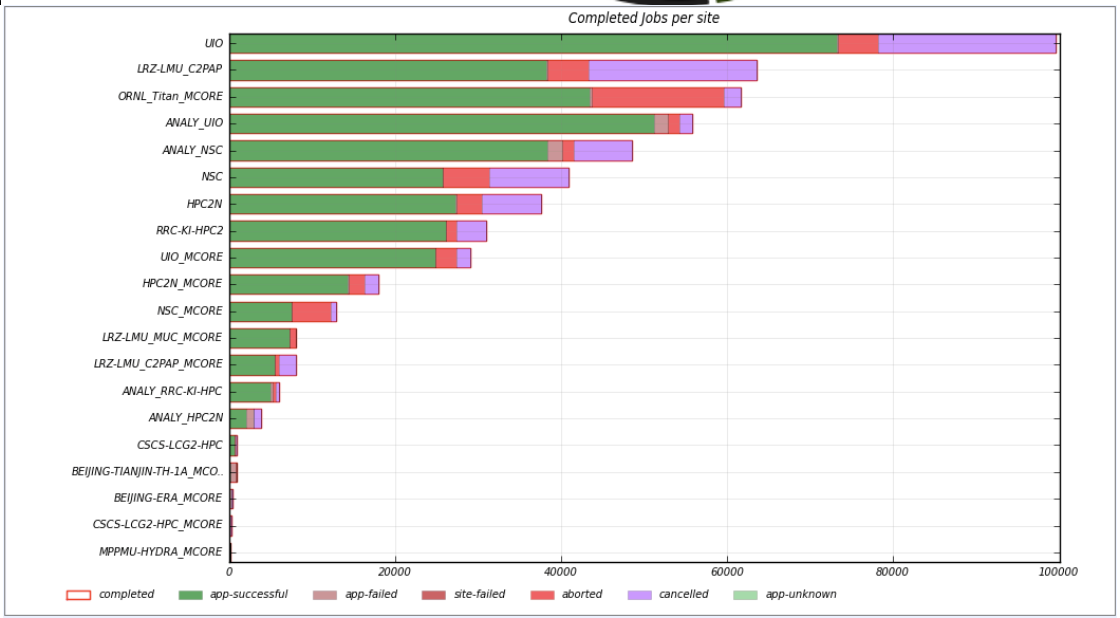

Prodsys2: Production task monitor

- Task Processing engine: translates production requests to tasks
- Task monitoring: filtering, sorting and etc.
- Task manipulation: abort, clone, finish, change priority, reassign and etc.

Tasks select parameters

Tasks type:	Production	Campaign:	
Project:		AMI tag exact:	
Owner:		Step name:	
Task name:	eg : mc.*valid.*s1773	Step output format:	
Task ID:	eg : 4003.*	Status:	All
Task ID <=		Provenance:	All
Task ID >=		Physics group:	All
Request ID:		Request type:	All
Request ID <=		Last update time period:	All time
Request ID >=		Last update time from:	01/11/2013
Chain:		Last update time to:	01/11/2016

Update table

Tasks status statistics

total	failed	assigning	finished	running	pending	broken	ready	done	submitted	obsolete	aborted	submitting	registered	exhausted	waiting
603764	10886	56	29001	929	9	12732	30	467874	6	43931	37029	267	972	9	33

Filters:

- Last update time period : all
- Task type : production

Select all | Abort | Finish | Retry | Reassign | Parameters | Obsolete | Kill jobs | View mode | 100 | Search:

Task Name	Request	Step	Task ID	Priority	Total Jobs	Done Jobs	Failure %	Events	Status	Submit time	Timestamp	Provenance	JIRA
mc15_13TeV.363075.Sherpa_CT10_llvv_renorm025.merge.e4681_p2482	5580	Reco	7409843	560	0	0	None	6000000	waiting	Jan 11 13:23	Jan 11 13:23	GP	587457
mc15_13TeV.301171.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wmintaunu_2250M2500.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408485	560	3	4	0	18770	done	Jan 11 10:03	Jan 11 13:33	GP	587366
mc15_13TeV.301157.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wplustaunu_4500M5000.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408482	560	3	4	0	18584	done	Jan 11 10:03	Jan 11 12:33	GP	587365
mc15_13TeV.301147.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wplustaunu_1250M1500.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408479	560	3	0	0	0	running	Jan 11 10:02	Jan 11 11:33	GP	587364
mc15_13TeV.301145.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wplustaunu_800M1000.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408475	560	3	4	0	15746	done	Jan 11 10:02	Jan 11 12:33	GP	587363
mc15_13TeV.301142.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wplustaunu_250M400.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408471	560	13	10	0	13805	running	Jan 11 10:02	Jan 11 12:33	GP	587362
mc15_13TeV.301163.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wmintaunu_400M600.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408468	560	5	0	0	0	running	Jan 11 10:01	Jan 11 11:33	GP	587361
mc15_13TeV.301167.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wmintaunu_1250M1500.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408464	560	3	4	0	17438	done	Jan 11 10:01	Jan 11 12:33	GP	587360
mc15_13TeV.301160.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wmintaunu_120M180.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408461	560	30	10	0	2078	running	Jan 11 10:01	Jan 11 12:33	GP	587359
mc15_13TeV.301144.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wplustaunu_600M800.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408457	560	3	4	0	14241	done	Jan 11 10:01	Jan 11 13:33	GP	587358
mc15_13TeV.301178.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wmintaunu_5000M.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408453	560	3	4	0	17631	done	Jan 11 10:00	Jan 11 13:33	GP	587357
mc15_13TeV.301168.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wmintaunu_1500M1750.merge.e3663_s2608_s2183_r7326_r6282_p2495	5577	Reco	7408449	560	3	4	0	17861	done	Jan 11 10:00	Jan 11 13:33	GP	587356

Showing 1 to 100 of 603,764 entries

Production tasks manage interface

- Current implementation of Request/Task monitor
  - Server-side
  - Initial data load takes time
  - Slow re-filtering, sorting etc.
  - Need to implement caching
- New implementation of Tasks monitor
  - Currently implemented for Tasks at given Request
  - Client-side
  - Still initial data load takes time
  - Fast filtering, sorting, search on the page
  - New features
    - Filter per column
    - Expandable fields
    - Common search including expandable fields
    - Smart row selection
  - Still under developmen

New Task monitor view

Show 10 entries

Search: Select all Select Filtered Deselect all

TaskID	Owner	RequestID	Step	Priority	Status
6563549	atlas-dpd-production	4223	Reco	880	done
Task name: data15_13TeV.00276073.physics_Main.merge.f618_m1480_p2425					
Input Dataset : data15_13TeV:data15_13TeV.00276073.physics_Main.merge.AOD.f618_m1480					
BigPanDA link : http://bigpanda.cern.ch/task/6563549					
6563551	atlas-dpd-production	4223	Reco	880	done
6563555	atlas-dpd-production	4223	Reco	880	done
6563557	atlas-dpd-production	4223	Reco	880	done
6563559	atlas-dpd-production	4223	Reco	880	done
6563560	atlas-dpd-production	4223	Reco	880	done
6563563	atlas-dpd-production	4223	Reco	880	done
6563566	atlas-dpd-production	4223	Reco	880	done
Task name: data15_13TeV.00276073.physics_Main.merge.f618_m1480_p2425					
Input Dataset : data15_13TeV:data15_13TeV.00276073.physics_Main.merge.AOD.f618_m1480					
BigPanDA link : http://bigpanda.cern.ch/task/6563566					
6563568	atlas-dpd-production	4223	Reco	880	done
6563570	atlas-dpd-production	4223	Reco	880	done
Search TaskID	Search Owner	Search RequestID	Search Step	Search Priority	Search Status

Showing 1 to 10 of 2,335 entries

Previous 1 2 3 4 5 … 234 Next

Abort Finish Retry Reassign Parameters Obsolete Kill jobs Login

Implementation for Job manage

http://bigpanda.cern.ch/jobs/?computingsite=RRC-KI-HPC2 Go

Show 10 entries Search: Select all Select Filtered Deselect all

	JobID	Attempt#	Owner	RequestID	TaskID	Transformation	Status
⊕	2729166821	1	atlas-dpd-production	5599	7418961	DAODMerge_tf.py	activated
⊕	2729166819	1	atlas-dpd-production	5599	7418961	DAODMerge_tf.py	activated
⊖	2729165856	0	gangarbt			Sim_tf.py	activated
	Job name:	ae20ec0f-90a8-422d-b4c6-ea06237c8eea					
	Input Dataset :	hc_test:mc15_13TeV.361106.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zee.evgen.EVNT.e3601_tid04972714_00					
	Output Dataset :	hc_test.gangarbt.hc20074279.RRC-KI-HPC2.346					
⊕	2729164963	0	gangarbt			AtlasG4_tf.py	running
⊕	2729163703	0	gangarbt			Sim_tf.py	running
⊕	2729160862	0	gangarbt			Sim_tf.py	running
⊕	2729156052	0	gangarbt			Sim_tf.py	running
⊕	2729154344	0	gangarbt			AtlasG4_tf.py	holding
⊕	2729152483	1	atlas-dpd-production	5612	7424565	Reco_tf.py	running
⊕	2729152482	1	atlas-dpd-production	5612	7424565	Reco_tf.py	running
	Search JobID	Search Attempt#	Search Owner	Search RequestID	Search TaskID	Search Transformatio	Search Status

Showing 1 to 10 of 593 entries

Previous 1 2 3 4 5 … 60 Next

Kill jobs Set debug Next button

Plans

- HPC: Efficiency and stability increase of the PanDA sites for ATLAS
  - Increase “maxrun” jobs number
  - New python implementation of Condor-SLURM connector with new features
  - Learning failed jobs logs
- Prodsys2: Continue to support and develop the interface

Acknowlegement

This work was funded in part by the U.S. Department of Energy, Office of Science, High Energy Physics and Advanced Scientific Computing Research under Contracts No. DE- SC0012704, No. DE-AC02-98CH10886 and Contract No. DE- AC02-06CH11357. NRC-KI team work was funded by the Russian Ministry of Science and Education under Contract No 14.Z50.31.0024. Supercomputing resources at NRC-KI are supported as a part of the center for collective usage (project RFMEFI62114X0006, funded by Russian Ministry of Science and Education).