Canadian ATLAS Tier-1

Di Qing

TRIUMF ATLAS Tier-1

TRIUMF-Helmholtz Workshop on Scientific Computing

September 16-17, 2019

DESY, Hamburg, Germany

Outline

- Overview of the Canadian ATLAS Tier-1 centre and role
- ATLAS tasks and workflows
- ATLAS resource usages at the Canadian Tier-1
- Our experience and CPU utilization improvement

Canadian ATLAS Tier-1

- Key player in large-scale distributed computing (ATLAS experiment only)
- Providing storage for the raw and secondary datasets
- Providing computing capacity for data processing, simulation, and physics group activities
- Providing 10% of worldwide Tier-1 resources
- Primary Tier-1 services and resources relocated to a new data centre at Simon Fraser University last year
- Current capacity in production:
  - 7680 cores (SFU) + 4744 cores (TRIUMF, simulation only)
  - 11 PB disk (at SFU)
  - 31 PB tape (at SFU)
- >10 years stable 24x7 operations

Tier-1 cluster at SFU

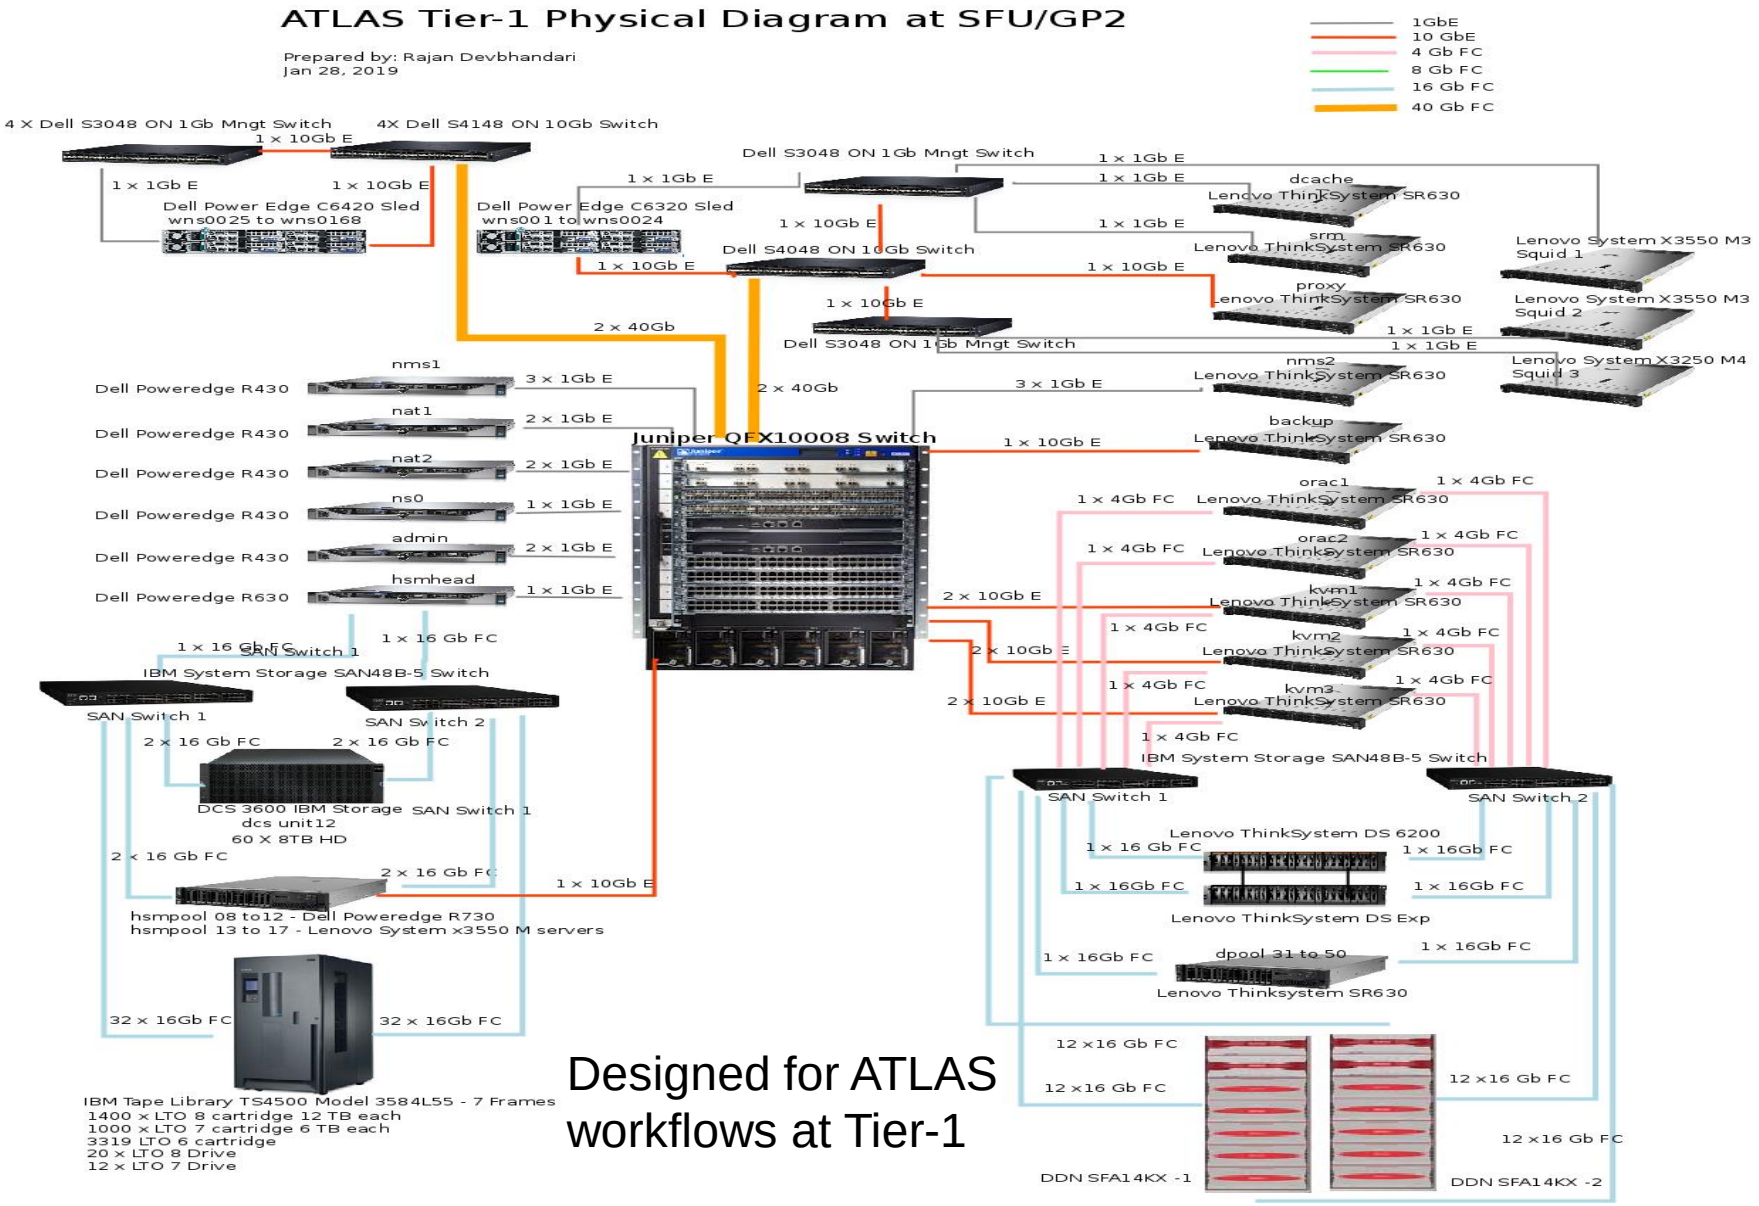

Designed for ATLAS workflows at Tier-1

ATLAS Workflow Management schematic

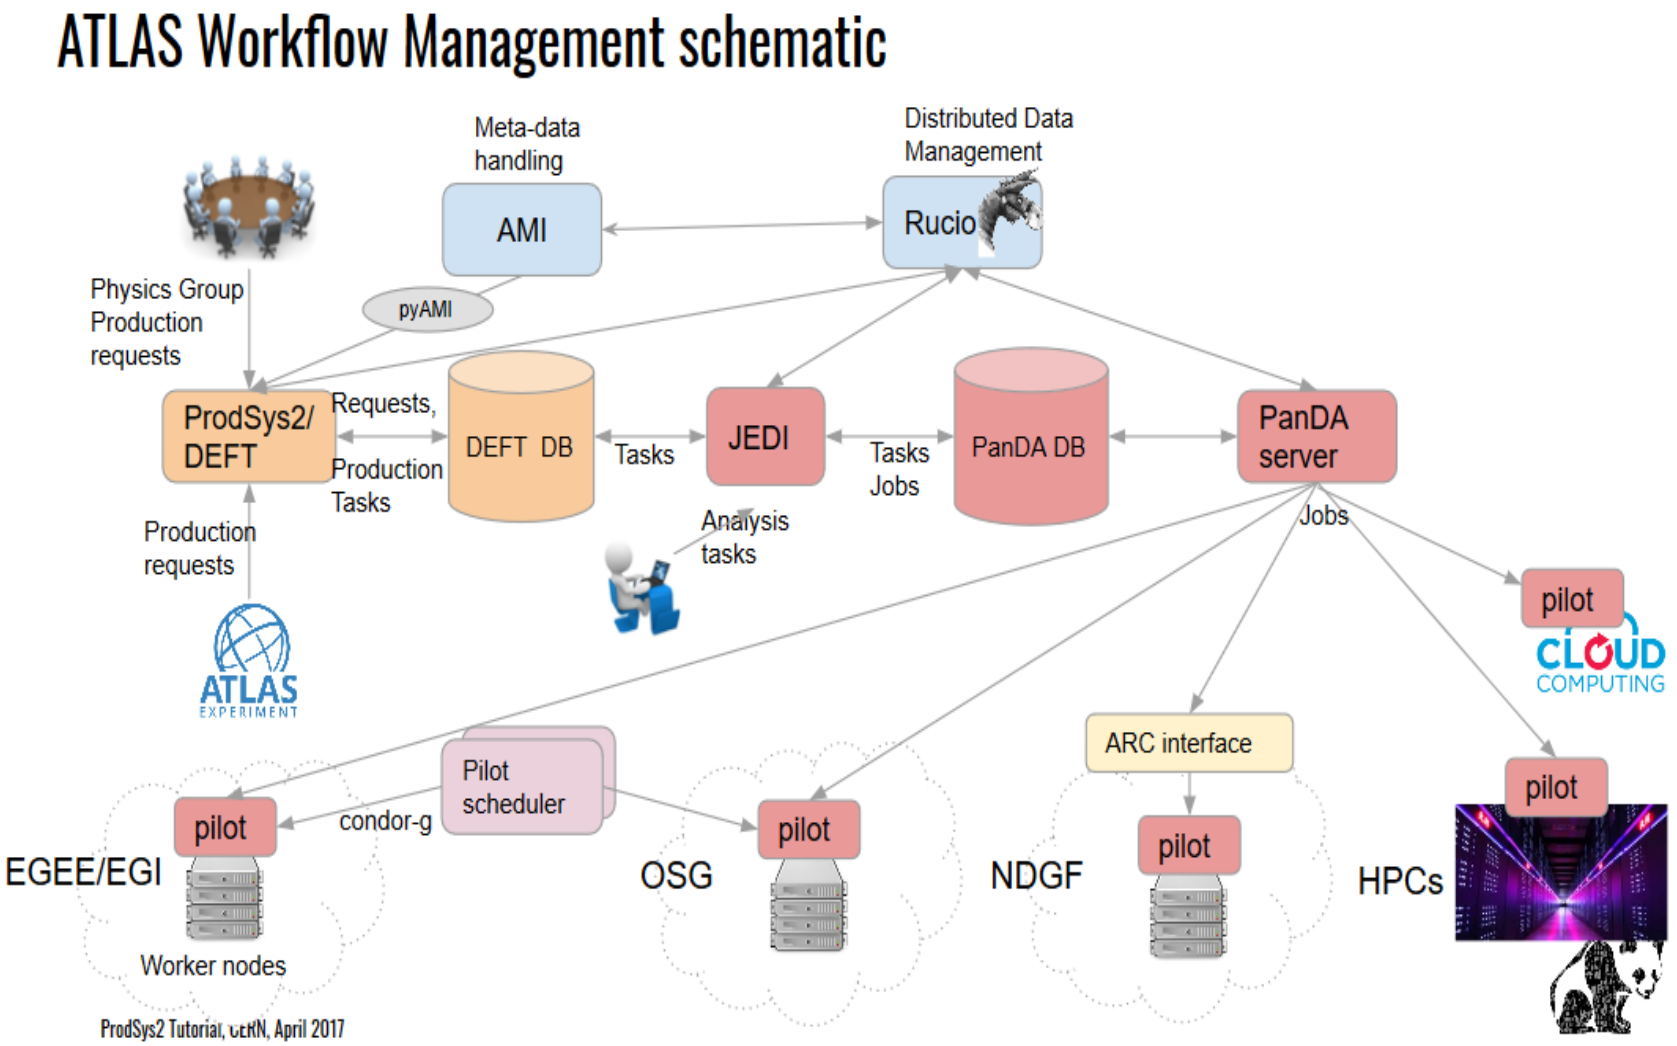

Source: ATLAS ProdSys2 tutorial

ATLAS workflows and resource requirements

Main ATLAS Tier-1 tasks

- Simulation
- Data reprocessing
- Data merging
- Derivation production (secondary dataset)
- Users analysis
- Physics groups activities

Resources required

- Single core/multiple cores
- CPU time, disk space, memory
- CPU intensive vs IO intensive
- Priority

ATLAS storage usage at Canadian Tier-1

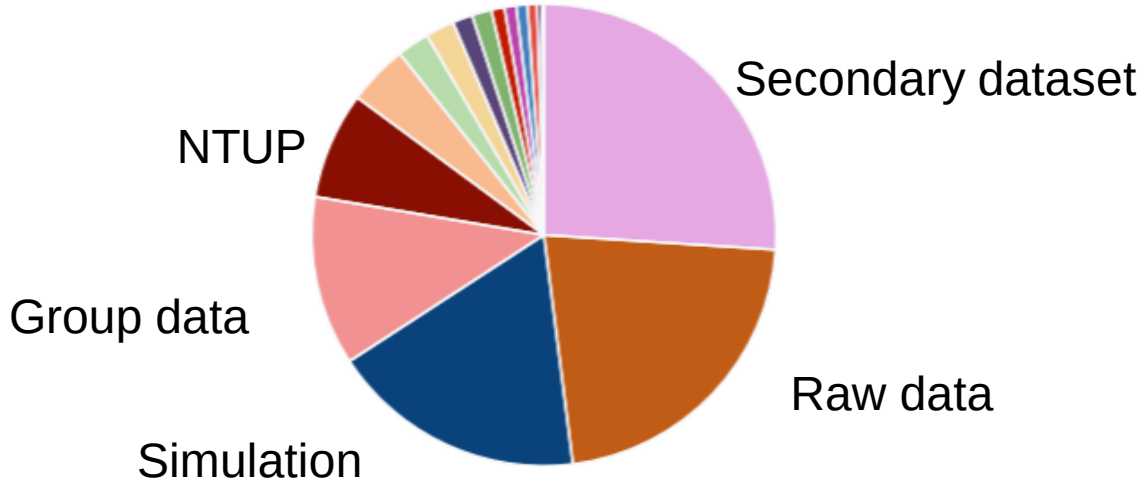

	current
AOD	6.35 PB
RAW	5.37 PB
HITS	4.34 PB
DAOD	2.897 PB
NTUP	1.828 PB
DESD	1.028 PB
ESD	548 TB
DRAW	490 TB
EVNT	340 TB
user	336 TB

- Storage usage as September 13, 2019
  - Disk usage ~8.4PB
  - Tape usage ~16.0PB
  - Total size ~24.4PB
  - 9.83% of all Tier-1s' usage

Slots of ATLAS running jobs at Canadian Tier-1

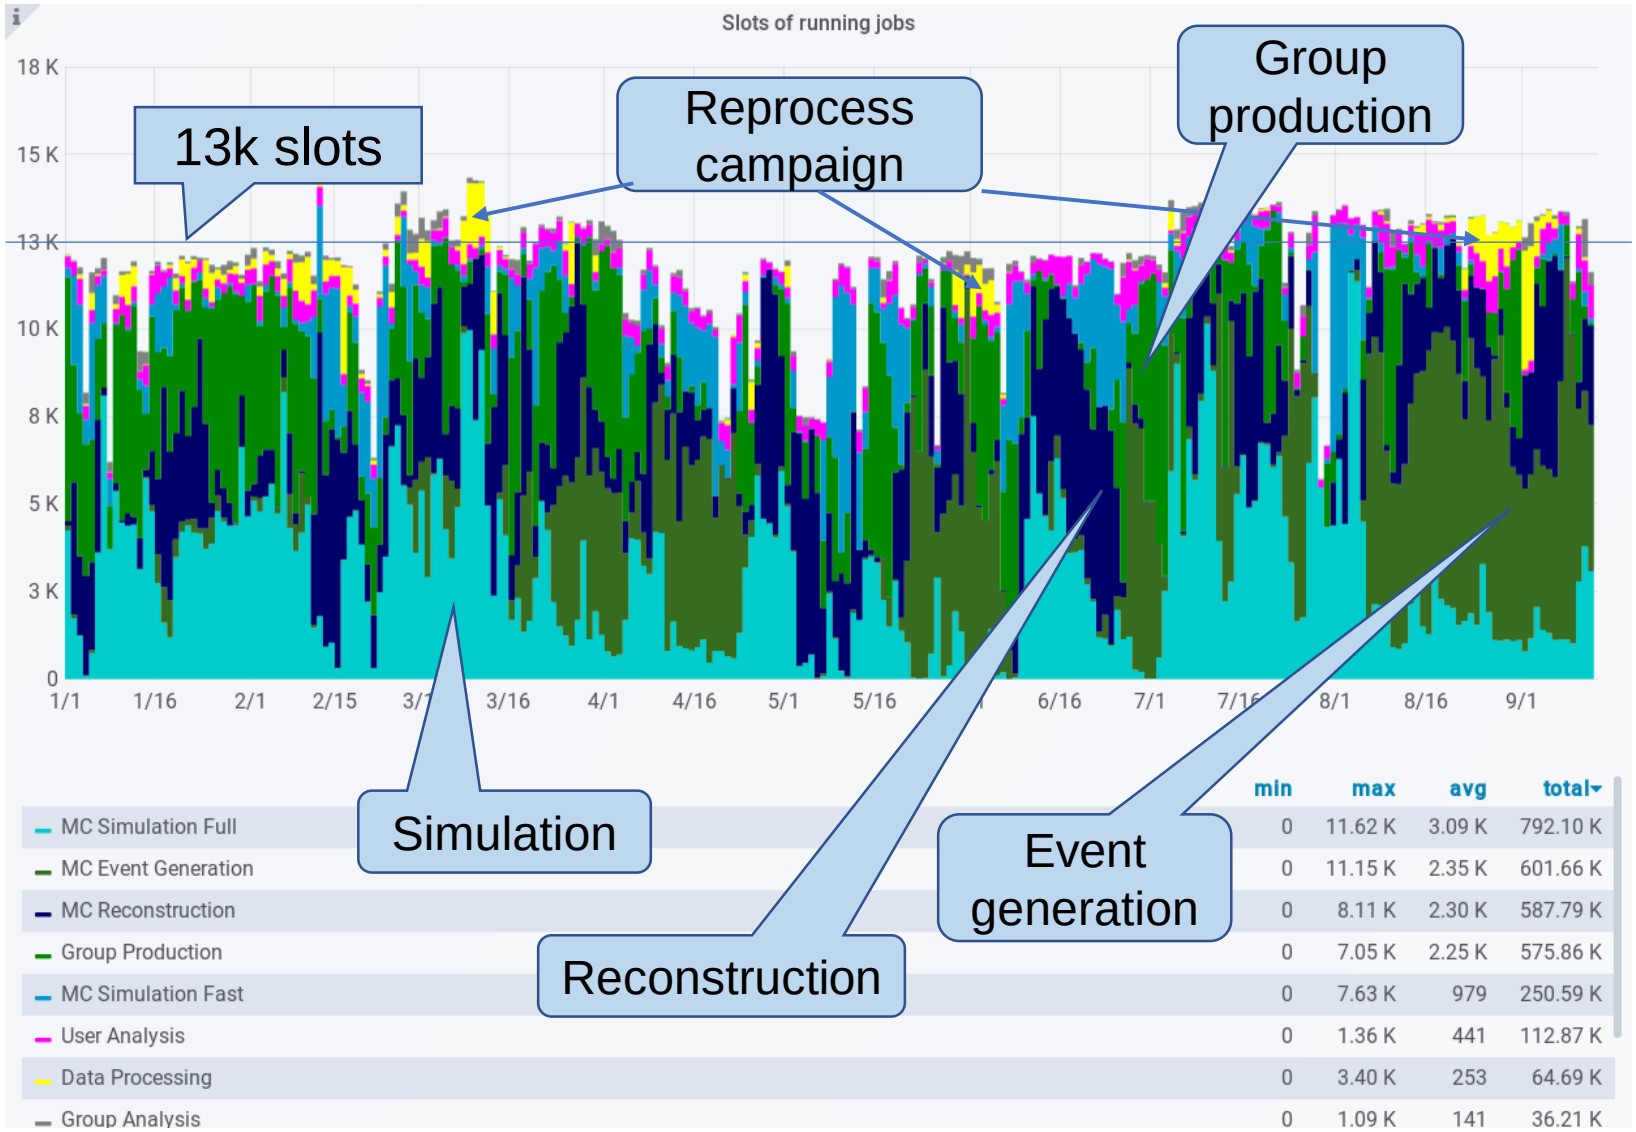

Local monitoring on CPU usages

- CPU slot usage from batch system point of view (one year)

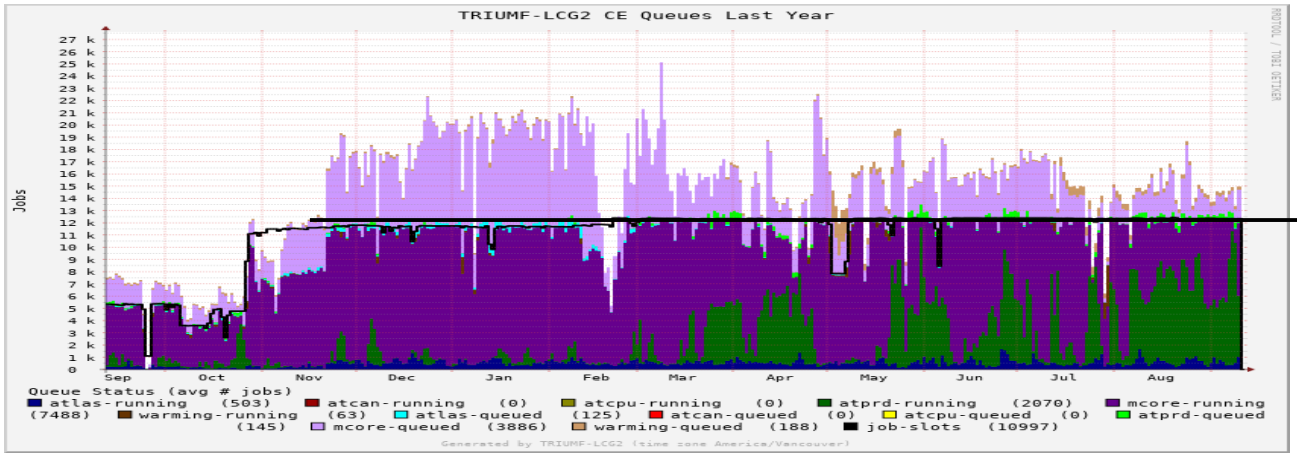

Total job slots ~12.4k

- CPU utilization of one sub cluster (1152 cores) (one year)

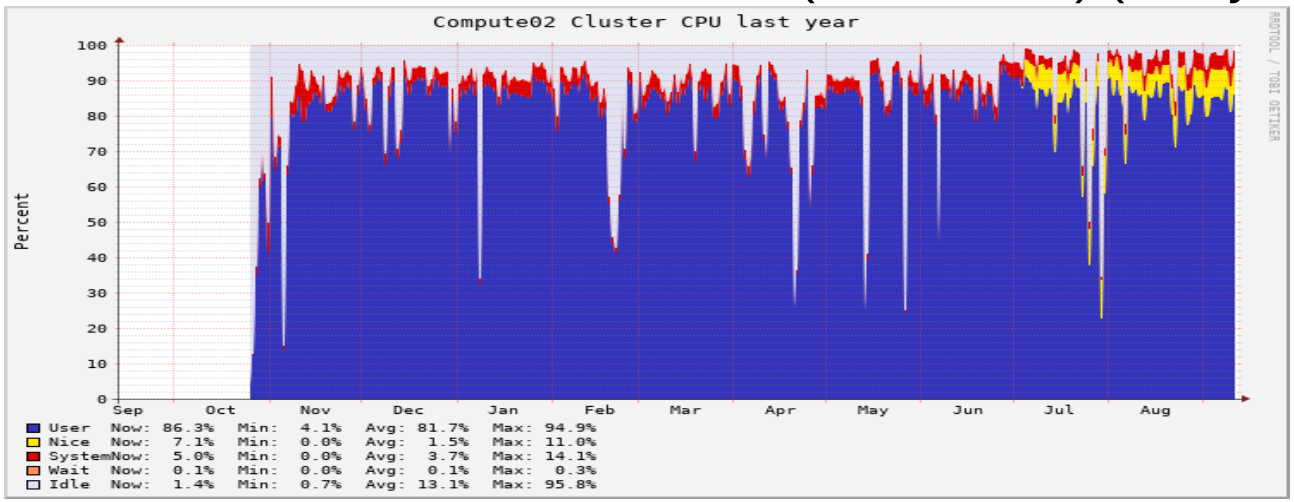

CPU resources are not fully utilized

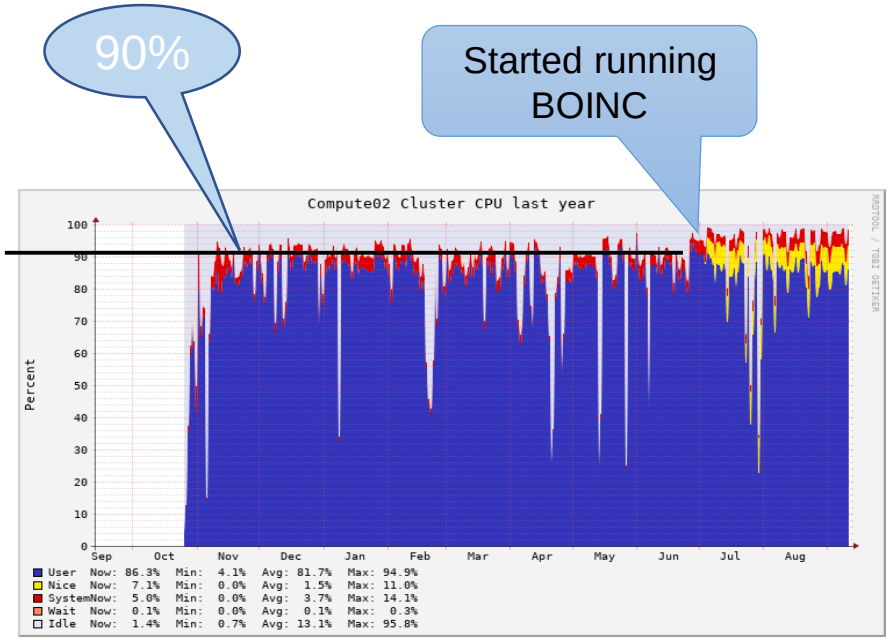

- Average CPU efficiency ~91.6% in 2019
  - CPU utilization rate even lower, ~84.6%
- Reasons
  - Staging in/out data
  - Sequential step of multi-core job
  - No payloads from ATLAS
  - Draining jobs
  - Switching between single core/multi-core productions

BOINC jobs and ATLAS@home project

- Computing with BOINC
  - A platform for distributed computing
  - Can be used in volunteer computing and in-house computing
  - Tasks run at background with low priority
  - Examples include SETI@home, Rosetta@home and Einstein@home
- ATLAS at home project
  - A project to use the internet-connected volunteer resources
  - Run simulation of ATLAS experiment - CPU intensive
  - Now also used by some grid sites and HPC to backfill nodes

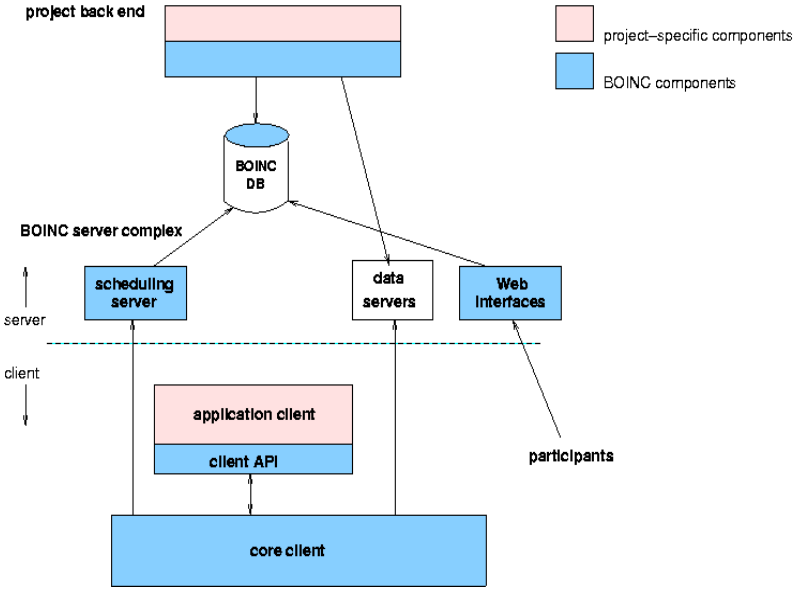

Source: https://sarwiki.informatik.hu-berlin.de/BOINC

Running ATLAS@home on TRIUMF ATLAS Tier-1 clusters

- Started in March 2018 at TRIUMF
- Started to run BOINC jobs at SFU in June 2019
- Implementations
  - Create account on LHC@home
  - Select ATLAS project
  - Install and configure BOINC client
  - Tune the configurations
    - Avoid the impacts on normal productions
    - Use CGroup to control the CPU share of BOINC jobs
    - Limit the number of cores (25%) which BOINC can use

What is gained with BOINC jobs

- Backfill nodes

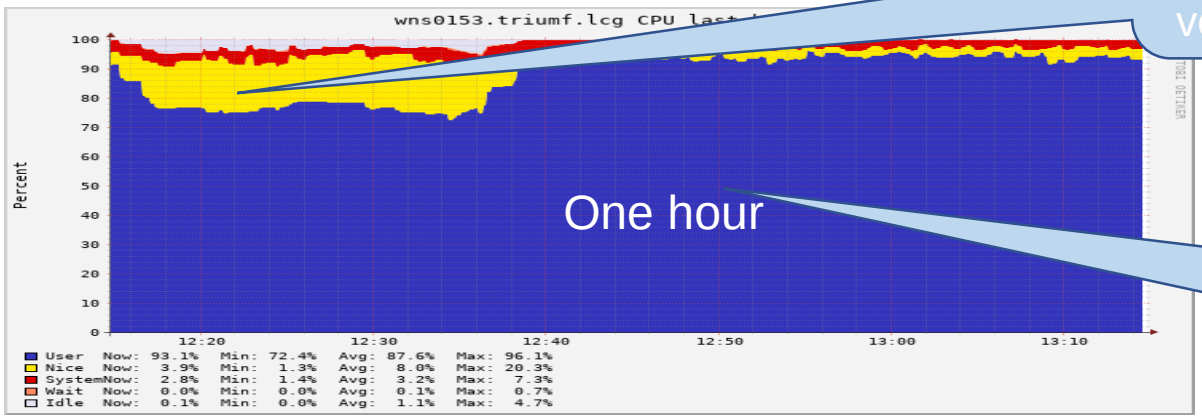

CPU usage of BOINC processes with very low priority

Processes of ATLAS jobs with normal priority

- Increase the CPU utilization of the cluster

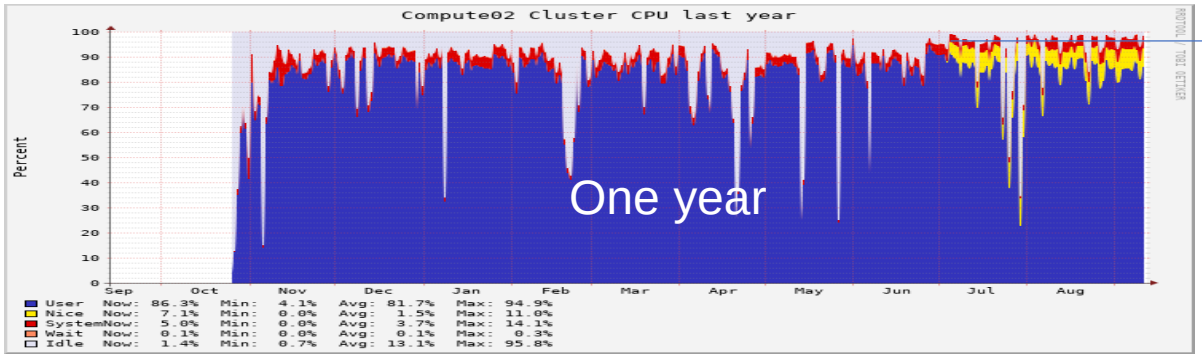

Total CPU usage

Contribution to ATLAS@home project

- Provided a lot of ‘additional’ resources to ATLAS

LHC@home Project Computing Community Jobs Site TRIUMF-LCG2 Log out

Rank	Name	Recent average credit	Total credit	Country	Participant since
1	AGLT2	6,167,351	1,505,343,681	United States	23 Jun 2014, 2:32:15 UTC
2	Agile Boincers	2,927,443	3,423,591,850	Switzerland	20 Sep 2012, 13:19:40 UTC
3	TRIUMF-LCG2	1,490,985	295,426,509	Canada	15 Mar 2018, 21:05:31 UTC
4	NDGF-T1	1,089,517	129,670,614	Norway	26 Feb 2019, 12:43:24 UTC
5	wHewitt	1,042,230	58,390,614	International	19 May 2014, 22:33:39 UTC

TRIUMF Contributions as September 12, 2019

Impact of BOINC jobs is small

- CPU efficiency of ATLAS productions slightly dropped
  - Less than 1% for our new cluster
- Potentially compete with the normal production jobs on memory
- Disk space and network bandwidth usage are trivial
- Accounting
  - The contributions of BOINC jobs are not officially accounted by ATLAS and WLCG
  - We are working on solutions

TRIUMF

Thank you!

Questions?