

# External Resources: Clouds and HPCs for the expansion of the ATLAS production system at the Tokyo regional analysis center

Michiru Kaneda<sup>1\*</sup> on behalf of ATLAS Collaboration

<sup>1</sup>International Center for Elementary Particle Physics, The University of Tokyo

**Abstract.** The Tokyo regional analysis center at the International Center for Elementary Particle Physics, the University of Tokyo, is one of the Tier 2 sites for the ATLAS experiment in the Worldwide LHC Computing Grid. The current system has 10,752 CPU cores and 16 PB disk storage. CERN plans the High-Luminosity LHC starting from 2026, which increases the peak luminosity by 5 times compared to the present value at the LHC. A shortage of computing resources for the high-luminosity LHC is expected with the current system. Therefore new ideas to expand the system are necessary. To expand the ATLAS production system at the Tokyo regional analysis center, R&D using external resources has been launched. One kind of external resources is a commercial cloud resource, such as Google Cloud Platform. Another resource is the High-Performance Computer (HPC) at the University of Tokyo. In this contribution, the current status of the R&D, the systems for these resources, and comparisons of the cost are reported.

## 1 Introduction

The Worldwide LHC Computing Grid (WLCG) [1] is a global collaboration to provide a set of computing resources to experiments at the Large Hadron Collider (LHC) [2]. WLCG consists of about 170 computing centers from 42 countries. Computing centers are classified into Tier levels from Tier 0 to Tier 3. The CERN Data Center is Tier 0 and 13 large computing centers have a Tier 1 role. Other computing centers providing resources to WLCG are Tier 2. Some computing centers which provide resources to local users are referred to as Tier 3.

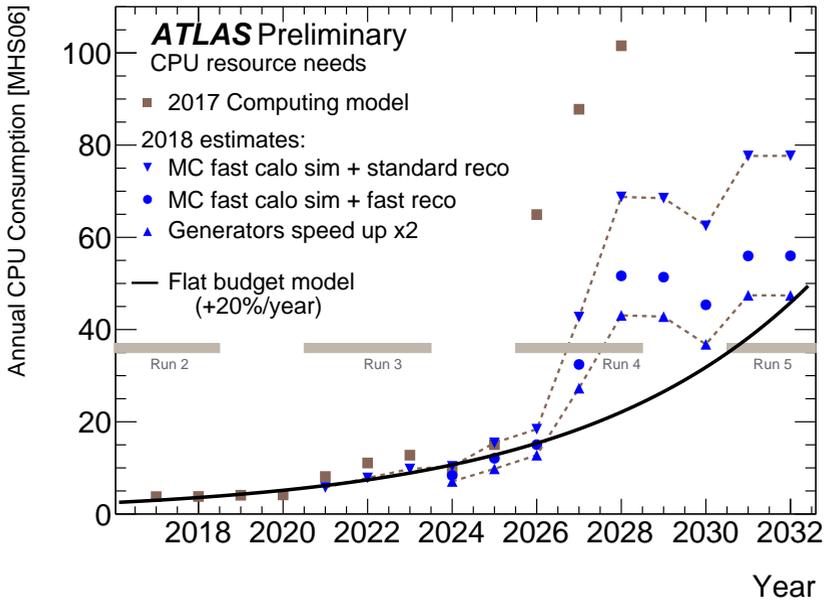
WLCG has already provided a large amount of resources successfully. But CERN plans the High-Luminosity LHC (HL-LHC) which will increase the peak luminosity to 5 times higher compared to the current LHC and therefore also the expected necessary computing resources will rise.

Expected necessary CPU resources for the next years for the ATLAS experiment [3], one of the LHC experiments, is shown in Figure 1. The solid line is the expected amount of available CPU resources and each point shows estimations of the amount of necessary resources in different scenarios. Many improvements of software have been implemented but there is still a gap between the necessary and available CPU resources. A similar problem

---

\*e-mail: [kaneda@icepp.s.u-tokyo.ac.jp](mailto:kaneda@icepp.s.u-tokyo.ac.jp)

Copyright 2020 CERN for the benefit of the ATLAS Collaboration. CC-BY-4.0 license.



**Figure 1.** Estimated CPU resources necessary for the years 2018 to 2032 for the ATLAS experiment. Each point shows the necessary computing CPU resources in different scenarios. The solid line shows the amount of resources expected to be available if a flat funding scenario is assumed [4].

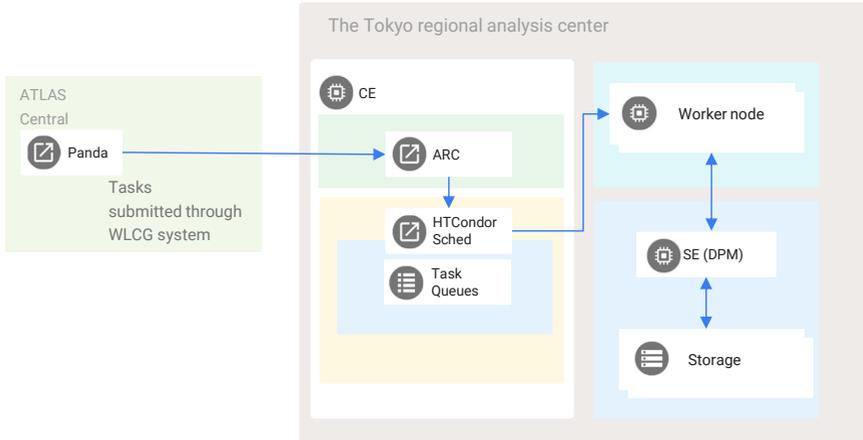
is issued for the storage resources, too. Therefore, new ideas are required to solve these problems.

Deploying external computing resources could be one of the solutions. In this contribution, R&D at the Tokyo regional analysis center to deploy external resources are reported.

## 2 The Tokyo regional analysis center

The Tokyo regional analysis center at the International Center for Elementary Particle Physics of the University of Tokyo is the Tier 2 site of WLCG and supports the ATLAS experiment. All hardware devices of the center are supplied by the three years rental and everything is renewed every three years. The current system is the fifth system and started in January 2019. The worker nodes in the current system consist of Dell Inc. PowerEdge M640 [5] which has two CPUs of Intel(R) Xeon(R) Gold 6130 CPU@2.10GHz (32 physical CPU cores). All CPU cores are provided with hyper threading (HT) off. Each node has 96 GB memory (3 GB/core). The total number of nodes is 336 and the total number of CPU cores is 10,752. For the storage system, 39 of Dell Inc. PowerEdge R640 [5] and 74 of Infotrend DS3024 [6] are used for the head nodes and disk arrays, respectively. Each disk array has 220 TB capacity and each file server has two disk arrays, which are connected by  $2 \times 16$  Gb fibre channels. The total storage size is 16 PB.

Figure 2 shows the current workflow for WLCG at the center. PanDA [7] is the central manager of the ATLAS production system. ARC CE [8] is used as the Computing Element (CE), which receives jobs as the grid front end and jobs are managed by HTCondor [9]. The Storage Element (SE) is managed by DPM [10] and transfers data from/to other WLCG sites.



**Figure 2.** The schematic view of the current production workflow at the Tokyo regional center.

### 3 Commercial clouds

Many commercial cloud services have been growing over recent years and WLCG computing sites have tried to include these resources [11][12][13][14].

For this R&D Google Cloud Platform (GCP) [15] was chosen. GCP provided two types of CPUs, Intel(R) Xeon(R) E5-2630 v4@2.20GHz and Intel(R) Xeon(R) Gold 6138 CPU@2.00GHz, at the Tokyo region. The CPUs' performances are similar to Intel(R) Xeon(R) Gold 6130, which are used at the Tokyo regional analysis center, but GCP provides the CPUs with HT switched on and the effective performance per CPU core at GCP is almost half of the one at the center. For example, the run times for 1,000 events of an ATLAS simulation job using 8 CPU cores are 5.19 hours and 9.27 hours at the center and GCP, respectively. Therefore, twice the number of CPU cores were used in order to estimate the cost at GCP compared to the center's cost.

For the cost comparison, three kind of systems, full on-premises, full cloud, and a hybrid system were considered.

For the full on-premises system, 10k CPU cores for the worker nodes and 16 PB storage attached to the same machines as the center were assumed. The total costs to buy these machines are \$4.7M<sup>2</sup> and \$1.4M for worker nodes and storage system, respectively. If the system is used for three years, the monthly cost is about \$200k. There are additional costs as power cost, hardware/system maintenance cost and other infrastructure costs but they depend strongly on the situation.

For GCP systems, 20k CPU cores were assumed, in order to make the same performance as the center's system. A preemptible instance was used which lasts for up to only 24 hours but is 80 % cheaper than a regular instance. A worker node with a 3 GB memory and 35 GB local disk per CPU core was assumed and the total cost for 20k CPU cores was estimated as \$210k/month [16]. About the storage, currently 8 PB of storage is used at the center, and 8 PB storage cost at GCP was estimated as \$184k/month [17]. There is an additional cost to extract data from GCP. The center's storage sends out 600 TB data/month and the corresponding cost

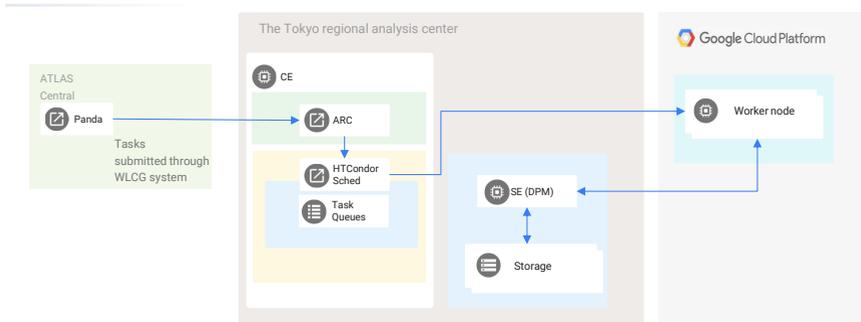
<sup>2</sup>U.S. dollar is used for the cost in this document.

of network usage at GCP is \$86k/month. As a result, the total cost of the full cloud system is \$480k/month. An additional cost is the management cost of the GCP system.

The hybrid system using GCP's worker nodes and on-premises storage can reduce the network cost. Even in this case, data extraction is needed to obtain outputs of jobs. The amount of output was estimated as 300 TB/month and the corresponding cost is \$43k/month. The total cost of the hybrid system is \$270k/month including on-premises storage. In this case, some on-premises costs are needed like the full on-premises system.

Although costs for maintenance or infrastructure depend on the situation, these three systems show costs in the same order.

In this time, the hybrid system was adopted as shown in Figure 3. The storage and job



**Figure 3.** The schematic view of the hybrid system with GCP.

management systems were located on the on-premises. As same as the on-premises system, the storage element is managed by DPM and HTCondor was used for the job management.

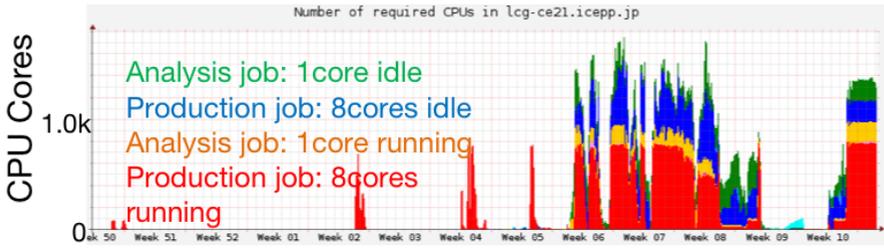
To manage worker nodes at GCP, Google Cloud Platform Condor Pool Manager (GCPM) [18] has been developed. GCPM is a Python based application and the source code is available on GitHub. GCPM checks HTCondor's task queue status and dynamically launches a worker node at GCP if a waiting job exists. After a job finishes, GCPM deletes the worker node. This management system works effectively, especially for the preemptible instances.

The R&D system consists of a maximum of 1k CPU cores and ran for a few weeks as a part of the Tokyo regional analysis center. There are two types of ATLAS jobs - the ones which use 1 CPU core and the ones which use 8 CPU cores. In the R&D system, the shares of these types were set to 20 % and 80 % for 1 CPU core and 8 CPU cores jobs, respectively. Figure 4 shows the CPU core usage of the test jobs. The far right-hand side of the graph shows the stable run with maximum 1k CPU cores. In that period, 800 (200) CPU cores are used for 8 cores (1 core) jobs and the remaining jobs are in idle state.

Real cost paid for the R&D system was \$400/day when the full 1k cores were used. It corresponds to \$240k/month for a 20k CPU core system and is consistent with the estimation.

## 4 HPC

For the HPC system, Reedbush Supercomputer System [19] at the University of Tokyo was used. Reedbush system consists of nodes with two Intel Xeon E5-2695v4 CPUs (36 physical CPU cores) and 256 GB memory. There are CPU only nodes and nodes with GPU (NVIDIA Tesla P100) and CPU only nodes were used for this R&D.

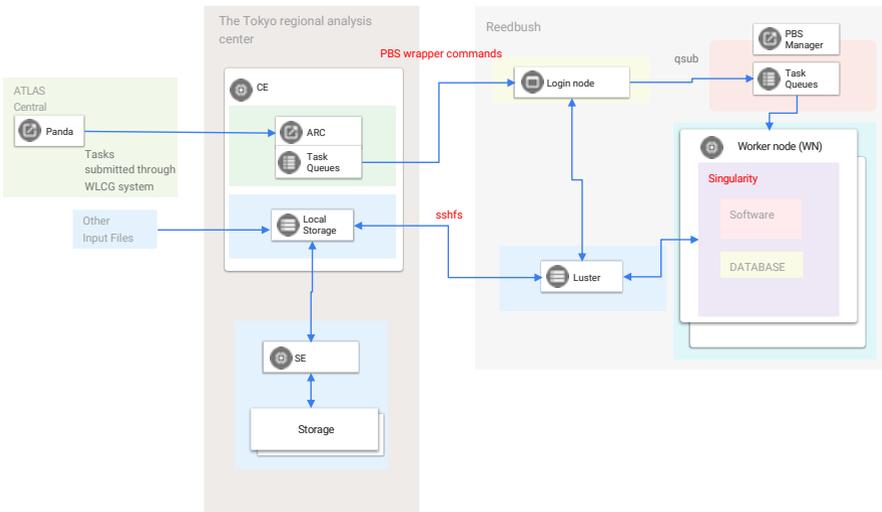


**Figure 4.** Test jobs running on the hybrid system using GCP.

PBS [20] is employed at Reedbush and the minimum unit of the job is 1 node (36 CPU cores). The PBS commands are available only on the login nodes of the Reedbush while it is not possible to run the grid front end software on the login nodes. Therefore, the system was constructed as a hybrid system with the on-premises system (Figure 5). The grid front end software (ARC CE) was placed on the on-premises system and jobs were submitted from there. To submit jobs, wrapper commands using ssh were used like<sup>3</sup>:

- `qsub: ssh <Reedbush login node> "cd $PWD && qsub $@"`

Reedbush’s disk space was mounted at the on-premises system by using sshfs and they had the same directory structure for the working space. With these wrapper commands, ARC CE can manage the PBS system of the Reedbush as a local system.



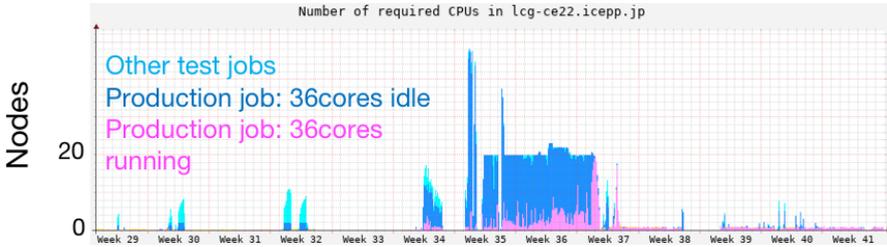
**Figure 5.** The schematic view of the system with Reedbush.

Another special situation of the Reedbush nodes is that the external network connection is not available. Normally, an ATLAS production job uses software through the CVMFS [22] and gets some data from an external database. In order to be able to run ATLAS jobs on the Reedbush system, Singularity was used with an image including the software and data. The size of the full software and database is too much to be included in one image. Therefore, the

<sup>3</sup>Original idea was taken from the Czech computing center’s work [21].

target simulation job type was chosen and the software and data only for the job type were included in the image. Input and output data were propagated through sshfs mount storage.

Special jobs using 36 CPU cores were submitted for the system. The maximum number of nodes was set to 20 and the test jobs ran successfully. Figure 6 is the monitoring graph of the number of nodes used by the system. The number of jobs were suppressed below 20 by the Reedbush system because there were other users jobs. At the end of the long time test, 20 jobs could run in parallel.



**Figure 6.** Test jobs running on the hybrid system using Reedbush.

## 5 Summary

External resources, commercial clouds and HPCs, are some of the solutions to fill the gap between necessary and available computing resources for the HL-LHC. The systems using Google Cloud Platform and Reedbush Supercomputer System at the University of Tokyo were constructed to expand the ATLAS production system at Tokyo regional analysis center. A hybrid system consisting of GCP and the on-premises systems using HTCondor as a job manager was established. GCPM has been developed in order to use GCP resources effectively. To deploy the Reedbush system, PBS wrapper command techniques, Singularity, and sshfs were used. Special jobs using 36 CPU cores were successfully executed at the system.

## References

- [1] The Worldwide LHC Computing Grid, <http://lcg.web.cern.ch/LCG>
- [2] *E. Evans, and P. Bryant*, JINST 3 S08001 (2008)
- [3] *ATLAS Collaboration*, JINST 3 S08003 (2008)
- [4] ATLAS EXPERIMENT Computing and Software Public Results, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults>
- [5] Dell EMC PowerEdge M640, <https://www.dell.com/ja-jp/work/shop/cty/pdp/spd/poweredge-m640/m640xv8301jp>
- [6] Infortrend ESDS 3024, <http://www.infortrendworks.com/DS-3024.asp>
- [7] The PanDA Production and Distributed Analysis System, <https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA>
- [8] *M. Ellert et al.*, Future Generation Computer Systems 23 (2007) 219-240
- [9] HTCondor, <https://research.cs.wisc.edu/htcondor/index.html>
- [10] DPM, <http://lcgdm.web.cern.ch/dpm>
- [11] *B. Holzman et al.*, Comput.Softw.Big Sci. 1 (2017) no.1, 1
- [12] *B. Holzman*, HEPCloud: Provisioning 160,000 Cores for Science (<http://hepcloud.fnal.gov/wp-content/uploads/2016/05/HEPCloud-DPF.pdf>)

- 
- [13] *F. Barreiro et al.*, EPJ WOC, 214, 03047 (2019)
  - [14] *C. Cristovao et al.*, Journal of Physics: Conference Series. 898. 082030 (2017)
  - [15] Google Cloud Platform, <https://cloud.google.com/>
  - [16] Google Compute Engine Pricing, <https://cloud.google.com/compute/pricing>
  - [17] Google Storage Pricing, <https://cloud.google.com/storage/pricing>
  - [18] Google Cloud Platform Condor Pool Manager, <https://github.com/mickaneda/gcpm>
  - [19] Reedbush Supercomputer System,  
<https://www.cc.u-tokyo.ac.jp/en/supercomputer/reedbush/service/>
  - [20] PBS Professional Open Source Project, <https://www.pbspro.org/>
  - [21] *S. Michal et al.*, EPJ WOC, 214, 03005 (2019)
  - [22] *A. De Salvo et al.*, J. Phys. Conf. Ser. 396, 032030 (2012)