

R&D for the expansion of the Tokyo regional analysis center using Google Cloud Platform

M. Kaneda*, J. Tanaka, T. Mashimo, R. Sawada, T. Kishimoto, N. Matsui

International Center for Elementary Particle Physics, The University of Tokyo

E-mail: kaneda@icepp.s.u-tokyo.ac.jp

The Tokyo regional analysis center is a computing center for the ATLAS experiment at the Large Hadron Collider and one of the Worldwide LHC Computing Grid Tier 2 sites. The batch system is managed by HTCondor. For the future expansion of the center, R&D project using Google Cloud Platform (GCP), which is one of commercial cloud resources, has been launched. A hybrid system is designed, where worker nodes of HTCondor are deployed on Google Cloud Platform while job manager nodes and file storages are placed on on-premises resources. To manage worker nodes on GCP, GCP Condor Pool Manager (GCPM) has been developed. The system with GCPM was integrated into the ATLAS production system and tested. In this paper, details of the R&D system are explained and experiences on Google Cloud Platform are reported.

International Symposium on Grids & Clouds 2019, ISGC2019
31st March - 5th April, 2019
Academia Sinica, Taipei, Taiwan

*Speaker.

1. Introduction

A computing system is one of the important components of high energy physics experiments. Experiments at the Large Hadron Collider (LHC) [1] produce tens of petabytes of data every year. For these experiments, the Worldwide LHC Computing Grid (WLCG) [2] was organized to provide a set of global computing resources. Currently, WLCG consists of about 170 computing centers from 42 countries. The total number of computer cores is about 1 million, and the total storage size is about 1 Eexabyte. Centers in WLCG are classified into Tier levels. Tier 0 is the CERN Data Center and 13 large computing centers have a Tier 1 role. Other universities or institutes' centers are Tier 2. Some computing centers provide local resources for the local users and these resources are sometimes referred to as Tier 3.

CERN plans the High-Luminosity LHC phase starting from 2026. It will increase the peak luminosity to 5 times higher compared to the LHC. Figure 1 shows expected computing CPU resources needed for the next years for ATLAS. Similar estimates of course exist also for other LHC experiments. Each point shows estimations of the amount of necessary resources in different




Figure 1: Estimated CPU resources needed for the years 2018 to 2032. Each point shows necessary computing CPU resources in the different scenarios. The solid line shows the amount of resources expected to be available if a flat funding scenario is assumed [3].

scenarios. The solid line is the expected available resource amount with an assumption of a flat funding scenario. Even with the most optimistic scenario, the gap to the expected amount is still large. Therefore, new ideas to fill the gap are necessary for the High Luminosity LHC.

There are developments for the grid infrastructure itself. New devices other than CPU, such as General Purpose Graphics Processing Units (GPGPU) or Field Programmable Gate Array (FPGA), have possibilities to accelerate our software. Another possibility is to deploy external computing resources, such as exploiting High Performance Computing (HPC) resources or commercial cloud resources.

Commercial cloud resources have been growing over recent years and many commercial companies have moved their on-premises systems to commercial cloud resources, some of them constructed hybrid systems of them. To expand current computing centers for LHC, some projects have tried to use commercial cloud resources [4][5][6][7]. This paper reports new site level system using Google Cloud Platform (GCP).

2. The Tokyo regional analysis center

The Tokyo regional analysis center is one of the Tier 2 sites for the ATLAS experiment [8] at LHC. All hardware devices are supplied by the three years rental and the fifth system started in January 2019. The current system has 10,752 CPU cores and 3.0 GB memory/core for worker nodes, and 15.84 PB disk storage for the file server. HEP-SPEC06 (HS06) [9] of each core is 18.97. The center provides 7,680 CPU cores and 10.56 PB disk storage for WLCG. Total HS06 base value is 145,689.6 [HS06 · cores], which corresponds to 6% of the total ATLAS Tier 2 resources. Remaining resources are provided to the ATLAS Japan group as local resources.

Figure 2 shows current work flow for WLCG at the center. PanDA [10] is the central manager of the ATLAS production system, The center receives jobs by ARC CE [11] as the grid front end and jobs are managed by HTCondor [12]. The storage system is connected to worker nodes of HTCondor and transfers data from/to other WLCG sites.

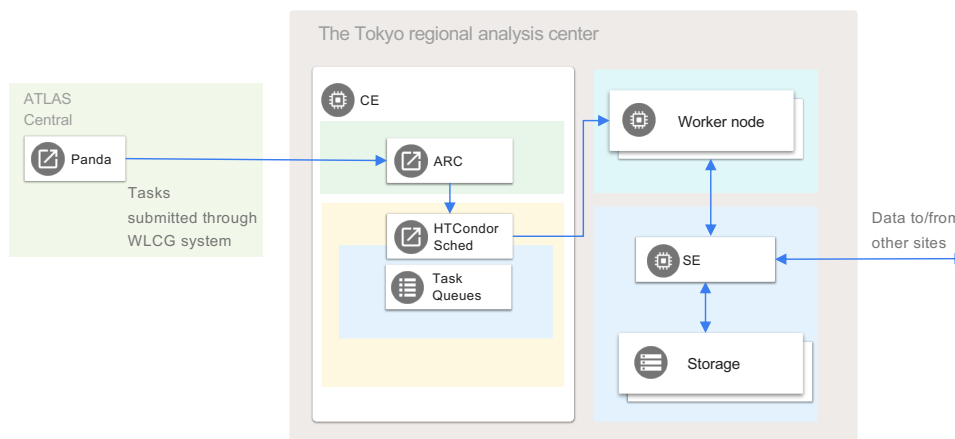


Figure 2: The schematic view of current production work flow at the Tokyo regional center.

3. Commercial clouds

There are a lot of commercial cloud providers nowadays, such as Google Cloud Platform (GCP) [13], Amazon Web Services (AWS) [14] and Microsoft Azure [15]. GCP provides instances for which number of cores and memory are fully customizable, while CPU is almost uniform¹. On the other hand, AWS and Azure provide various sets of CPU and memory. Most instances of them are operated with hyper threading (HT) on. Only Azure provides HT off instances.

One of the interesting points is that HTCondor supports AWS resources as worker nodes from version 8.8.0 released in January 2019. But it does not support other cloud resources. In this R&D, GCP was used and a new tool to manage HTCondor worker nodes has been developed.

4. Google Computing Element

To estimate cost performance, firstly, Google Computing Element (GCE) in GCP was compared to resources of the center. The center provides the system with HT off while GCP provides instances with HT on. For the comparison, a HT off machine was prepared at the center. For GCE, two different CPUs were tested and multi core and single core instances were prepared. SPEC CINT2006 Rates [16], HS06, and a processing time for the same ATLAS simulation 1000 events were checked. SPEC CINT2006 values were extracted from the data sheets. The center uses Dell Inc. PowerEdge M640. For GCP, the system is unknown and values of Dell Inc PowerEdge R630 and Dell Inc. PowerEdge M640 for Broadwell and Skylake, respectively, are used. The ATLAS simulation uses 8 cores in 1 job. To measure the processing time at the center. 4 and 8 parallel jobs were processed at HT off and HT on systems respectively, and their mean time is shown. Table 1 and 2 show CPU types and performances for each system, respectively.

Table 1: CPU types of the center and GCP

System	CPU
The Tcenter's system	Intel(R) Xeon(R) Gold 6130 CPU@2.10GHz
GCE (Broadwell)	Intel(R) Xeon(R) E5-2630 v4@2.20GHz
GCE (Skylake)	Intel(R) Xeon(R) Gold 6138 CPU@2.00GHz

At the center, HT on system shows better performances than half of HT off system as a result of benefit of HT. Two different GCE CPUs show almost same spec although Skylake is a little better. GCE instances shows almost half spec compared to the center. In the following estimations, GCE's spec is assumed to be the half of the center. 1 core instance of GCE shows much better value in HS06. This may be because multi core instances use same physical cores for each 2 virtual cores and these cores are always fully used in the simulation job or HS06 measurement. But 1 core instance shares a physical core with other instance. If other 1 core instance is not fully used the test instance can use almost fully its physical core without any interruption. the test instance can use almost full physical core.

¹At Tokyo region, only Intel Broadwell (2.20 GHz) and Skylake (2.00GHz) are available and their performances are almost the same.

Table 2: Computing Performances of the center and GCP. SPEC CINT2006 Rates were taken from [16]. At the center, a system is Dell Inc. PowerEdge M640 which has two CPUs of Intel(R) Xeon(R) Gold 6130 and total 32 physical cores. For GCP instances, the systems are unknown and values of Dell Inc PowerEdge R630 and Dell Inc. PowerEdge M640 for Broadwell and Skylake, respectively, are shown as reference values. HS06 values were evaluated by running HS06 test jobs for each machine. The ATLAS simulation uses 8 cores in 1 job and processing times for 1000 events are shown. To evaluate processing times at the center, 4 and 8 same jobs were processed in parallel at HT off and HT on systems, respectively, to fill CPU cores and the mean times of these parallel jobs are shown. The larger value means better performance at SPEC CINT2006 and HS06 and the shorter time means better performance at the ATLAS simulation processing time.

System	Cores	SPEC CINT2006 Rates/core	HS06/core	ATLAS simulation 1000 events processing time (hours)
The center HT off	32	46.25	18.97	5.19
The center HT on	64	N/A	11.58	8.64
GCE (Broadwell)	8	(39.75)	12.31	9.32
GCE (Broadwell)	1	(39.75)	22.73	N/A
GCE (Skylake)	8	(43.25)	12.62	9.27

To suppress the cost, GCP provides a low cost preemptible instance. The preemptible instance lasts for up to 24 hours and it is sometimes shut down before 24 hours depending on GCP situation. But the cost is up to 80% cheaper than regular instances. Jobs of the ATLAS experiment finish within 24 hours in most cases. Therefore, preemptible instances were used for the R&D.

5. Cost estimation

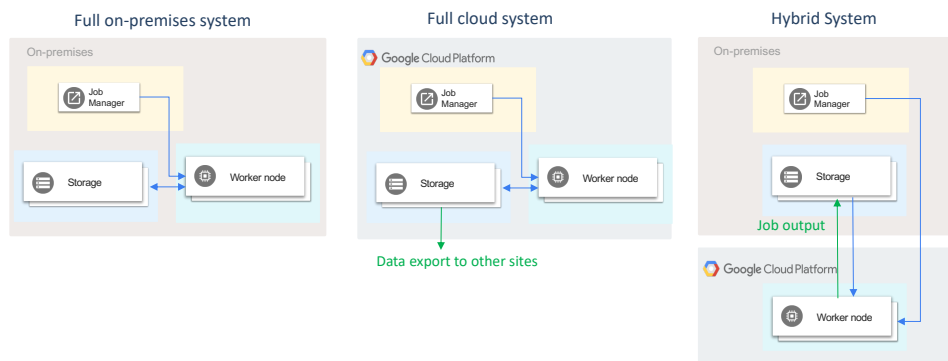


Figure 3: Schematic views of the full on-premises, the full cloud, and the hybrid systems.

Costs were estimated for three cases: the full on-premises system, the full cloud system and the hybrid system shown in Figure 3. For the full on-premises system, Dell machines and Infortrend machines were used to estimate hardware costs for computing nodes and storages, These machines are same machines used in the center's system. GCP costs were officially taken from [19] and [20]. For CPU cores, twice the number of CPU cores were used for the estimation to compare to the center because the spec of the CPU core at GCP is the almost half of the center's one. The storage in GCP is assumed as the half of the on-premises system because currently half of the center storage is used. For GCP, an additional cost of the network to export data from GCP will be needed to external places. At the full cloud system, there will be a large amount of transfers of data from the site to other WLCG sites every day and it will be about 600 TB/month. At the hybrid system, outputs of jobs will be transferred when every job finishes. The total amount is expected to 280 TB/month. This value highly depends on job types. These transfer costs were estimated from the center's history logs in 2018. Monthly costs for the hardware were estimated by assuming that the system is used for 3 years for on-premises system. Cost estimations were summarized in Table 3, 4 and 5. Although costs for facilities, infrastructures and hardware maintenances depend strongly on the situation, these costs will be a few hundreds dollars per month. Therefore, these three system seem on the same order. In other word, GCP provides enough reasonable prices compared to on-premises system.

Table 3: The full on-premises system cost estimation. For computing nodes, a server of PowerEdge M640 with two Intel(R) Xeon(R) Gold 6130 CPUs (32 CPU cores) was considered for the estimation. The price of one server was \$1.5k [17]. For 10k CPU cores, 315 servers are needed and total cost of computing nodes was estimated as \$4.7M. For the storage system, a disk array of Infortrend DS3024, which has 24 hard disk slots, was assumed for the estimation. Assuming that each disk array is set up with 24×10 TB hard disk with RAID6 (200 TB capacity), 16 TB is available with 72 disk arrays. The price of Infortrend DS3024 was \$1.22k [18] and 10 TB hard disk's price was assumed as \$300. The total cost of disk arrays was estimated as \$1.4M. Monthly costs for the computing nodes and the storage system were calculated by assuming that these machines are used for 3 years. The power cost was estimated with the assumption that each server use 200 W and \$0.3/kWh. For 500 machines, the cost was estimated as \$20k/month.

Resource	Cost/month
Computing node: 10k CPU Core, Memory 3 GB/core, Local disk 35 GB/core	\$130k
16 PB disk storage	\$40k
Power cost	\$20k
Facility/infrastructure cost Hardware maintenance cost etc...	N/A
Total	\$190 + α

For the R&D, the hybrid system was chosen. Storage price and data export cost are very high so that the hybrid system might be a good choice in term of cost.

Table 4: The full cloud system cost estimation. The number of CPU Core was set to twice as much as the full on-premises system to obtain same performance with machines of the hyper threading on. The storage capacity was set to the half of the on-premises system because the cost is needed only for the used space and the current on-premises system uses the half of the full capacity. Network usage corresponds to exports of data in the storage to the other sites. The amount of exports was estimated from the center’s log. For the computing node resources (CPU, memory, and local disk), prices were taken from **Preemptible price of Custom vCPUs and memory** and **Standard provisioned space of Persistent disk pricing** for Tokyo (asia-northeast1) region in Google Compute Engine Pricing [19]. For the disk storage and the network usage, prices were taken from **Regional Storage of Data storage** and **Egress to Asia Destinations (10+TB) of General network usage** in Cloud Storage pricing [citestorage](#).

Resource	Unit price	Cost/month
20k CPU Core	\$6.45/core month	\$130k
20k × Memory 3 GB	\$0.86/GB month	\$52k
20k × Local disk 35 GB	\$0.052/GB month	\$36k
8 PB disk storage	\$0.023/GB month	\$184k
Network 600 TB	\$0.12/GB	\$72k
Cloud maintenance cost		N/A
Total		\$474k + α

6. GCPM

To manage worker nodes of HTCondor in GCP, Google Cloud Platform Condor Pool Manager (GCPM) [21] has been developed. GCPM is a Python based application and the source code is available on GitHub².

Figure 4 is the schematic view of the system with GCPM. GCPM first prepares necessary machines such as cache servers before starting any worker nodes. GCPM then checks HTCondor’s task queues and creates a new instance in GCP if there are idle jobs and no worker nodes are ready, HTCondor has the list of worker nodes which can join in the HTCondor pool. The list is dynamically updated by GCPM with a new instance. Then HTCondor can submit a job to the worker node in GCP. To use preemptible instances effectively, an instance is better to be deleted when the first job has finished to avoid to run new job in the machine with a limited operating time until the preemption. This function is implemented by sending `condor_off -peaceful -startd` command 10 min after the instance was launched. The second job is submitted at the instance only if the first job finishes within 10 min otherwise the instance is deleted and a new instance is created for the new job.

7. R&D system and real cost example

Figure 5 shows full overview of this R&D system using GCP. For the R&D, the max number of total cores was set to 1000. 2.5 GB memory is assigned to each core. Disk sizes are 50 GB and 150 GB for 1 core and 8 core instances, respectively.

A real cost example of one day with the R&D system is summarized in Table 6. To compare

²The application can be installed by pip command: `$ pip install gcpm`.

Table 5: The hybrid system cost estimation. The system was assumed that computing node resources (CPU Core, Memory and Local disk) are on GCP and the storage is on the on-premises system. The number of CPU Core was set to twice as much as the full on-premises system to obtain same performance with machines of the hyper threading on. Network usage corresponds to transfers of outputs of jobs to the on-premises storage. The amount of transfers was estimated from the center’s log. For the computing node resources (CPU, memory, and local disk), prices were taken from **Preemptible price of Custom vCPUs and memory** and **Standard provisioned space of Persistent disk pricing** for Tokyo (asia-northeast1) region in Google Compute Engine Pricing [19]. For the network usage, prices were taken from **Egress to Asia Destinations (10+TB)** of **General network usage** in Cloud Storage pricing citestorage. For the storage system, a disk array of Infortrend DS3024, which has 24 hard disk slots, was assumed for the estimation. Assuming that each disk array is set up with 24×10 TB hard disk with RAID6 (200 TB capacity), 16 TB is available with 72 disk arrays.

Resource	Unit price	Cost/month
20k CPU Core	\$6.45/core month	\$130k
20k \times Memory 3 GB	\$0.86/GB month	\$52k
20k \times Local disk 35 GB	\$0.052/GB month	\$36k
Network 280 TB	\$0.12/GB month	\$34k
On-premises 16 PB disk storage		\$40k
Other on-premises costs		N/A
Cloud maintenance cost		N/A
Total		\$292k + α

Table 6: Hybrid system one day real cost example

Resource	Usage/day	Cost/day	$\times 30 \times 20$
Core	20057 Core·Hours	\$177	\$106k
Memory	47581 GB·Hours	\$56	\$34k
Disk	644898 GB·Hours	\$50	\$30k
Network	559 GB	\$78	\$47k
Other services	N/A	\$30	\$18k
Total	N/A	\$391	\$236k

with the estimation of Table 5, scaled values ($\times 30$ (days) $\times 20$ (core ratio)) are also listed in the table. The real cost shows reasonable agreement with the estimation.

8. Instance preemption impact on jobs

Figure 6 and 7 show job status of data production and user analysis jobs, respectively. Data production jobs take typically 10 hours for the processing, while user analysis jobs are shorter in the most cases. For the production jobs, GCP jobs with preemptible instances showed 35% failure rate and the center showed 11% failure rate in the same time. There was a short non-preemptible instance test at GCP and it showed 6% failure rate. Most of failures in the center be considered as software problems. Therefore, preemptible instances deployment results in a few tens % increase of failure rates as site problems. For analysis jobs, GCP jobs with preemptible instances showed

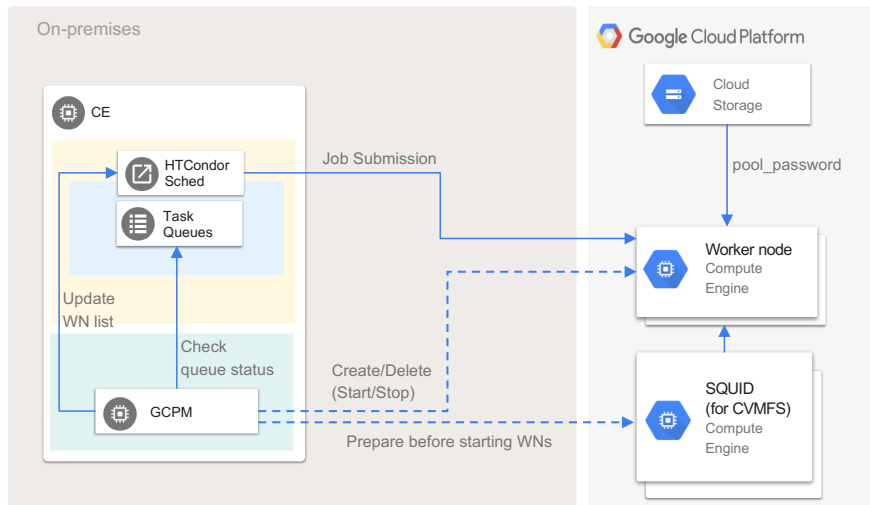


Figure 4: Schematic view of the system with GCPM.

19% failure rate and the center showed 8% failure rate. Non-preemptible test at GCP showed 14% failure rate. Analysis jobs are less affected by preemption because these jobs are shorter and can be finished before the preemption occurs. A few tens % increase of failure rate is still acceptable as the cost is up to 80% cheaper than regular instances.

9. Summary and discussion

R&D for the expansion of the Tokyo regional analysis center using GCP has been started. The system uses HTCondor for the job management and GCPM has been developed to manage worker nodes on GCP. The system was implemented in the ATLAS production system and a performance was checked. Estimated costs for the full on-premises, the full cloud and the hybrid system were same order. One day real cost example for the hybrid system is reasonably agreed with the estimation. Preemptible instances make failure jobs by the preemption and the failure rate for the ATLAS production jobs was increased by a few tens % in the test. This is still acceptable as the cost is up to 80% cheaper compared to regular instances. Job selections might reduce such an effect as shorter jobs are less affected. Another optimization can be done for the network cost. It is better to have jobs which produces smaller output files with the same CPU usage. Jobs of such as merging files or archiving files should be avoided to be submitted to GCP.

Although this R&D was tested with a large computing system, small computing systems may gain more benefit by using commercial cloud resources. GCPM is pure HTCondor tool and it can be used for any systems which use HTCondor. WLCG resources are used fully all the time. On the other hand, computing systems at the laboratory level are usually not used fully. By using

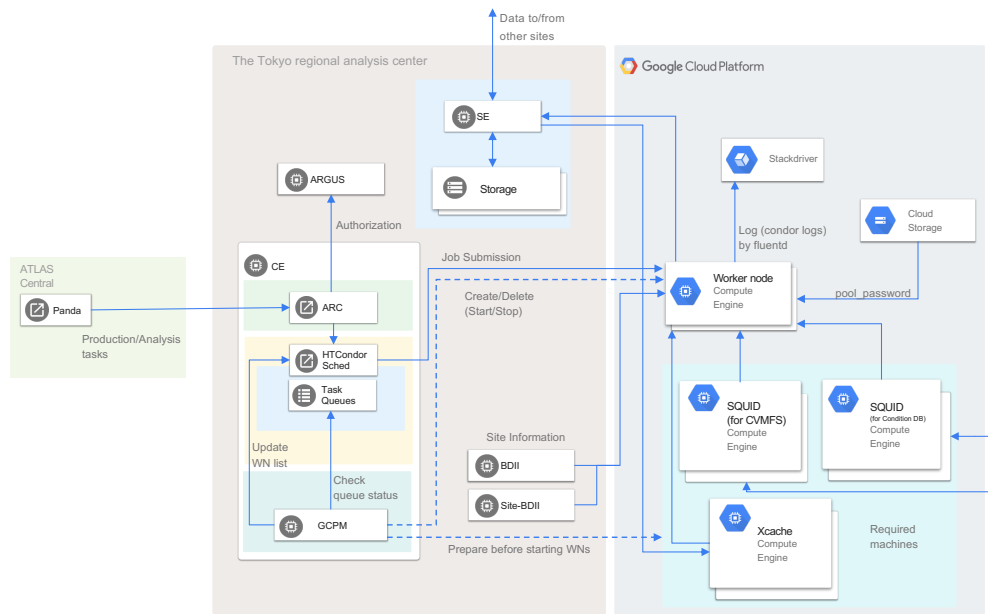


Figure 5: Full schematic view of R&D system.

commercial cloud resources, these systems can save costs when the systems are not used and increase worker nodes as many as needed at the peak time.

References

- [1] *E. Evans, and P. Bryant*, JINST 3 S08001 (2008)
- [2] The Worldwide LHC Computing Grid, <http://lcg.web.cern.ch/LCG>
- [3] ATLAS EXPERIMENT Computing and Software Public Results, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults>
- [4] *B. Holzman et al.*, Comput Softw Big Sci (2017) 1: 1.
- [5] *B. Holzman*, HEPCloud: Provisioning 160,000 Cores for Science (<http://hepccloud.fnal.gov/wp-content/uploads/2016/05/HEPCloud-DPF.pdf>)
- [6] *F. Barreiro et al.*, EPJ Web of Conferences 214, 03047 (2019)
- [7] *C. Cristovao et al.*, Journal of Physics: Conference Series. 898. 082030 (2017)
- [8] *ATLAS Collaboration*, JINST 3 S08003 (2008)
- [9] HEP-SPEC06 (HS06), <https://w3.hepivx.org/benchmarking.html>
- [10] The PanDA Production and Distributed Analysis System, <https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA>
- [11] *M. Ellert et al.*, Future Generation Computer Systems 23 (2007) 219-240
- [12] HTCondor, <https://research.cs.wisc.edu/htcondor/index.html>
- [13] Google Cloud Platform, <https://cloud.google.com/>




Figure 6: Job statuses of production jobs. Blue and green regions represent succeeded and failed jobs, respectively. Left and right figures show jobs in GCP and the center worker nodes, respectively.

[14] AWS, <https://aws.amazon.com/>

[15] Microsoft Azure, <https://azure.microsoft.com/>

[16] All SPEC CINT2006 Rates Results Published by SPEC, <https://www.spec.org/cpu2006/results/rint2006.html>

[17] Dell EMC PowerEdge M640, <https://www.dell.com/ja-jp/work/shop/cty/pdp/spd/poweredge-m640/m640xv8301jp>

[18] Infortrend ES DS 3024, <http://www.infortrendworks.com/DS-3024.asp>

[19] Google Compute Engine Pricing, <https://cloud.google.com/compute/pricing>

[20] Google Storage Pricing, <https://cloud.google.com/storage/pricing>

[21] Google Cloud Platform Condor Pool Manager, <https://github.com/mickaneda/gcpm>




Figure 7: Job statuses of analysis jobs. Blue and green regions represent succeeded and failed jobs, respectively. Left and right figures show jobs in GCP and the center worker nodes, respectively.

Pos (ISGC2019) 017