

Improvements in utilisation of the Czech national HPC center

Michal Svatoš^{1,*}, Jiří Chudoba^{1,**}, and Petr Vokáč^{2,***}

¹Institute of Physics of the Czech Academy of Sciences, Na Slovance 1999/2, Prague, 18221, Czech Republic

²Czech Technical University in Prague, Faculty of Nuclear Sciences and Physical Engineering, Břehova 7, Prague, 115 19, Czech Republic

Abstract. The distributed computing system of the ATLAS experiment at LHC is allowed to opportunistically use resources at the Czech national HPC center IT4Innovations in Ostrava. The jobs are submitted via an ARC Compute Element (ARC-CE) installed at the grid site in Prague. Scripts and input files are shared between the ARC-CE and a shared file system located at the HPC centre via sshfs. This basic submission system has worked there since the end of 2017. Several improvements were made to increase the amount of resource that ATLAS can use. The most significant change was the migration of the submission system to enable pre-emptable jobs, to adapt to the HPC management's decision to start pre-empting opportunistic jobs. Another improvement of the submission system was related to the sshfs connection which seemed to be a limiting factor of the system. Now, the submission system consists of several ARC-CE machines. Also, various parameters of sshfs were tested in an attempt to increase throughput. As a result of the improvements, the utilisation of the Czech national HPC center by the ATLAS distributed computing increased.

1 Introduction

The distributed computing of the ATLAS experiment [1] at the Large Hadron Collider (LHC) opportunistically uses computing resources of the Salomon HPC cluster located at the Czech National HPC Center IT4Innovations (IT4I) in Ostrava. When Salomon was commissioned, it was ranked 39th in Top500 [2] (in June 2015). In the list published in June 2019, it was ranked 282nd [3]. Worker nodes of the HPC available to the ATLAS Distributed Computing (ADC) have the following hardware specifications:

- 24 cores of Intel Xeon E5 CPUs
- 128 GB of RAM
- Infiniband (56 Gbps)

The batch system is PBS Professional [4].

*e-mail: Michal.Svatos@cern.ch

**e-mail: Jiri.Chudoba@cern.ch

***e-mail: Petr.Vokac@cern.ch




Figure 1. A scheme of the job submission system

2 Settings

The infrastructure providing interface between ATLAS and Salomon HPC is shown on Figure 1. The process starts with the ARC Control Tower (aCT) obtaining job description from the ATLAS workflow management system and submitting it to one of the ARC-CE machines installed at Czech Tier2 site (praguelcg2) [5]. The ARC-CE translates the job description into a PBS script. It puts job’s script and its other dependencies into the session directory. The session directory is located in file system of the ARC-CE where a directory from the HPC’s storage is permanently mounted via sshfs. This way they are accessible for jobs running in the PBSpro. The ARC-CE also provides the Data Delivery Service which obtains the input files and stores them in the session directory. The service can obtain the input files from either the local DPM grid storage or from storage of other grid sites. The input files are then stored in the ARC-CE cache (with 60 days retention) located on Lustre storage, as one file can be reused by many jobs. The file is then linked from the cache to the session directory. Then the ARC-CE submits the job to the batch system via an ssh connection to Salomon’s login node. Software used by running jobs is also located on the Lustre storage. Once the job is finished, the output (located in the session directory) is transferred via sshfs back to the ARC-CE. For standard jobs, both the job output and log tarball are moved to the local DPM grid storage. For Event Service jobs (see section 3.1), both the job output and log tarball are moved to the CERN S3 Object Store and a copy of the log tarball is stored at the local DPM storage.

This submission system is described in more detail in [6].

3 Improvements

While the basic system worked well, it sometimes did not manage to keep the available resources full. Several things were tried to improve the situation.

3.1 Pre-emption

In October 2018, the management of the IT4Innovations decided to change the conditions of opportunistic usage. Under the new policy, jobs using Salomon opportunistically could be pre-empted.

The ADC has a system which handles such jobs, called Event Service [7]. The submission system ran in that mode for several months while receiving gradual updates and tweaks. During that time, the submission system demonstrated its preparedness for pre-emption.

The new pre-emption policy of the IT4Innovations management was not enforced at that

time. Therefore, the submission system was switched back to standard jobs, as those are more CPU efficient. The reason is that Event Service jobs process far fewer events and therefore are shorter.

3.2 Sshfs

The sshfs protocol seems to be the bottleneck of the submission system. Observing the traffic, it is clear that the speed reaches a plateau around 60 Mbps (Figure 2). But when a throughput test (transfer of few big files via scp) was performed, the plateau appeared at level of 500 Mbps. The cause of this is probably huge number of small files in shared area (1k+ files per one Event Service job).




Figure 2. Network traffic at one of the ARC-CE machines (each bin is 5 min) reaching saturated state at ~60 Mbps. The majority of the traffic is caused by the sshfs.

Several parameters of the sshfs were tested in an attempt to improve throughput:

- no compression
- faster encryption (aes128-ctr)
- caching (tested parameters: kernel_cache, noauto_cache, cache_timeout=300, cache=no, no_readahead)

The tests were performed on two identical ARC-CE machines running on the same hardware with the same configuration. One parameter of the sshfs setting was changed on one machine and then the network usage was compared between both machines for a couple of weeks.

Results of the testing are as follows: Changes in compression and encryption showed no noticeable difference. When cache=no and no_readahead were used, the volume of data being transferred decreased but the plateau remained. These two parameters are currently used in production on all ARC-CE machines submitting jobs to Salomon.

3.3 Number of machines

Currently, there are 4 ARC-CE machines submitting jobs into Salomon. The reason being that two user accounts are used to submit jobs with a limit of 100 single-node jobs in the batch system per user. Thus, each machine has limit of 50 jobs and 2 sshfs connections (session directory and cache of input files). At this level, the sshfs transferring capacity is sufficient for the submission system to fill the slots at reasonable speed and capable of keeping them full (example of filling is shown in Figure 3).

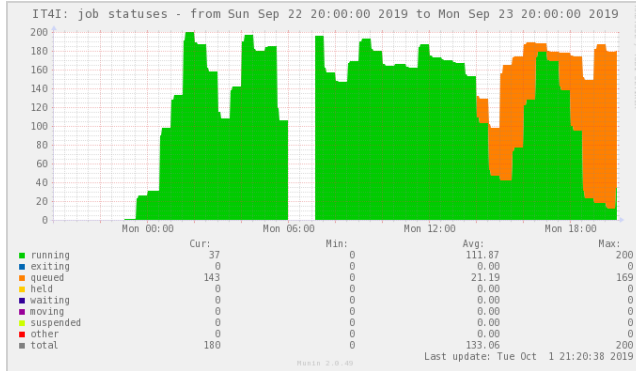


Figure 3. Statuses of jobs in the Salomon batch system (each bin is 5 minutes but values are updated every 30 minutes). The submission system is capable of filling all available resources (the limit is 200 jobs) within few hours.

3.4 Number of PBS requests

ARC-CE periodically gathers information from the batch system. This introduces a significant number of interactions between each ARC-CE and Salomon above basic job submission and deletion. The Salomon PBS configuration limits number of requests per user to approximately 6 requests per minute. Exceeding this limit will result in job failures because the job submission command will not be accepted.




Figure 4. Histograms of the number of requests per minute before (left) and after (right) modification of ARC-CE commands. Before the change, the ARC-CE usually exceeds 6 PBS requests per minute limit. After the change, it is rare.

At the default frequency, the ARC-CE exceeds this limit by doing, most often, 8 requests to the PBS server per minute (left plot on Figure 4). This behaviour cannot be tuned in version 5 of the ARC-CE which is used by the machines (this has been fixed in the ARC-CE 6.1). Moreover, there are two ARC-CE machines submitting jobs to each user account (as explained in section 3.3).

Therefore, a work-around was applied. ARC-CE commands which are submitting `qstat`, `pbsnodes`, and `qmgr` commands to Salomon’s login node were modified to slow down the

interaction with the HPC. The modification was a simple 60 second sleep before execution of the PBS pro command. After the modification, the number of PBS requests rarely went above the limit (right plot on Figure 4).

4 Performance

In January 2019, the submission quota increased from 100 jobs (i.e. 2.4k cores) to 200 (i.e. 4.8k cores). The number of running jobs, as well as the number of processed input files, increased (left plot on Figure 5) while the number of processed events decreased (right plot on Figure 5). The cause of this unexpected behaviour comes from another configuration update made at about the same time. Between the end of November 2018 and the beginning of September 2019, Salomon ran the Event Service jobs that process less events per job.




Figure 5. Number of processed files and events. Number of processed events decreased since November 2018 because the system was submitting Event Service jobs. Since January 2019 the number of processed files increased because of increase in job quota.

Figure 6 shows that the number of failed jobs is reasonable (13%). In comparison, the whole ADC had 12% failed jobs in the same time period, i.e. the failure rate of jobs at Salomon is in line with the rest of the ADC. Also, the error distribution is similar to normal grid queue (software errors, storage problems, problems of central services, etc.).




Figure 6. Number of finished (green) and failed (red) jobs. The failure rate is stable and relatively low.

Figures 7 and 8 show that Salomon provided significant amount of unplugged resources to ATLAS (about 13% of CPU time of the prague2 site).




Figure 7. CPU consumption of jobs running under pragueclg2. The orange color represents contribution of the Salomon. It opportunisticly provided 13% of CPU time which pragueclg2 delivered to ATLAS.




Figure 8. Slots of running jobs of the pragueclg2. The orange color represents contribution of the Salomon. It is fluctuating due to its opportunistic nature but the contribution is significant.

5 Summary and Conclusion

Several ways of improving the opportunistic usage of the Salomon HPC cluster by the distributed computing of the ATLAS experiment at LHC have been investigated. The system is ready for job pre-emption (by switching to Event Service jobs). The tuning of the sshfs parameters caused a decrease in network usage. A sufficient number of the ARC-CE machines were deployed to ensure the slots are filled at reasonable speed and kept full. After a work-around was applied to deal with the limit on the number of PBS request, jobs have been running smoothly.

With those improvements, the submission system can fill the opportunistic resources and contribute a significant amount of unpledged resources to the ADC.

Computing resources at FZU are co-financed by projects Research infrastructure CERN (CERN-CZ) and OP RDE CERN Computing (CZ.02.1.01/0.0/0.0/16013/0001404) from EU funds and MŠMT. This presentation is supported by project LTT17018.

We gratefully acknowledge National Supercomputing Center IT4Innovations (IT4I) for computing time provided by the Salomon supercomputer.

Copyright

Copyright 2020 CERN for the benefit of the ATLAS Collaboration. Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license.

References

- [1] ATLAS Collaboration, *JINST* **3**, S08003 (2008)
- [2] *Top500 list - june 2015*, <https://www.top500.org/list/2015/06/>
- [3] *Top500 list - june 2019*, <https://www.top500.org/list/2019/06/>
- [4] *Pbs professional open source project*, <https://www.pbspro.org>
- [5] J. Chudoba, M. Adam, D. Adamová, T. Kouba, A. Mikula, V. Říkal, J. Švec, J. Uhlířová, P. Vokáč, M. Svatoš, *J. Phys. Conf. Ser.* **898**, 082034 (2017)
- [6] M. Svatos, J. Chudoba, P. Vokac, *EPJ Web Conf.* **214**, 03005 (2019)
- [7] D. Cameron, A. Filipčič, W. Guan, V. Tsulaia, R. Walker, T. Wenaus (ATLAS), *J. Phys. Conf. Ser.* **898**, 052010 (2017)