

Effective Dynamic Integration and Utilization of Heterogeneous Compute Resources

Max Fischer¹, Manuel Giffels^{1,*}, Andreas Heiss¹, Eileen Kuehn¹, Matthias Schnepf¹, Ralf Florian von Cube¹, Andreas Petzold¹, and Günter Quast¹

¹Karlsruher Institute of Technology (KIT), Germany

Abstract. Increased operational effectiveness and the dynamic integration of only temporarily available compute resources (opportunistic resources) becomes more and more important in the next decade, due to the scarcity of resources for future high energy physics experiments as well as the desired integration of cloud and high performance computing resources. This results in a more heterogeneous compute environment, which gives rise to huge challenges for the computing operation teams of the experiments.

At the Karlsruhe Institute of Technology (KIT) we design solutions to tackle these challenges. In order to ensure an efficient utilization of opportunistic resources and unified access to the entire infrastructure, we developed the Transparent Adaptive Resource Dynamic Integration System (TARDIS). A scalable multi-agent resource manager providing interfaces to provision as well as dynamically and transparently integrate resources of various providers into one common overlay batch system. Operational effectiveness is guaranteed by relying on COBaLD – the Opportunistic Balancing Daemon and its simple approach of taking into account the utilization and allocation of the different resource types, in order to run the individual workflows on the best-suited resource respectively.

In this contribution we will present the current status of integrating various HPC centers and cloud providers into the compute infrastructure at the Karlsruhe Institute of Technology as well as our experiences gained in a production environment.

1 Introduction

Nowadays computing in high energy physics (HEP) is predominantly relying on homogeneous resources provided by the World LHC Computing Grid (WLCG) [1] based on a flat-budget funding model. However, recent studies of the HEP Software Foundation [2], legitimately assuming a continuity of the flat-budget funding model, show that the expected technology advance will not be sufficient to meet the computing requirements of future HEP experiments. One promising approach to narrow the gap is to supplement the WLCG with resources not permanently dedicated, but temporarily available for HEP computing tasks. So-called opportunistic resources are mainly provided by High Performance Computing (HPC) Centres as well as commercial and public cloud providers. In addition, the exploration of utilising HPC

*e-mail: Manuel.Giffels@kit.edu

and Cloud resources in HEP computing is increasingly desired and also supported by funding agencies.

2 Opportunistic Resources and their Challenges

In contrast to the homogenous resources provided by the WLCG, utilising opportunistic resources results in a more heterogenous computing environment not fully-controlled by WLCG policies and thus imposing huge challenges to the computing operation teams of the experiments. The amenities provided by the WLCG like the dedicated operating system, the available HEP software environment, the Grid Computing Elements (CE) as unified entry points to resources and the dedicated high-bandwidth wide area network (WAN) connections are not necessarily available anymore. In addition, opportunistic resources differ a lot among each other. Starting with the provided hardware environment (CPUs, memory, available accelerators, varying WAN bandwidth, etc.), which have an impact on their suitability for a given task (e.g. MC simulation, reconstruction, analysis, etc.), through to the customised setups and special configurations needed for their integration into the HEP workflow management systems. Furthermore, the situation is even aggravated by their temporary and not well-predictable availability. Altogether leading to enormous resource management challenges the typically small computing operation teams cannot comply with, in particular if we assume that each of the roughly 150 WLCG sites will have at least one or even more co-located opportunistic resources to be integrated.

While the provisioning of the required operating system and the deployment of the needed HEP software can be solved by using existing modern container technologies like Singularity [3] and the CERN Virtual Machine File System (CVMFS) [4], tackling the resource management challenges require the development of novel concepts and approaches described in the following sections.

3 Towards Collaborative Integration and Unified Access to Opportunistic Resources

At KIT, we developed a concept on how the resource management challenges described above can be overcome by moving away from a centralised approach of resource integration by the computing operations teams. The basic principle here is to associate opportunistic resources to existing WLCG sites close by and to move towards a national or even regional approach of resource integration. This brings along several benefits, the effort of resource integration is distributed among multiple sites. The proximity to the actual resource providers simplifies the communication between the involved parties, because people potentially know and trust each other. In addition, local or regional negotiations with providers of opportunistic resources are in our experience more efficient and promising than experiment wide negotiations for the same reason. On the other hand that means additional work for the sites, that needs to be minimised by providing common guidance and tools for the resource integration. Furthermore, the acquirement and integration of additional resources needs to be worthwhile for the sites, for example by acknowledging those efforts as service work for the experiments.

For the actual integration of opportunistic resources into the WLCG computing, our work is focusing on solutions working for all experiments supported by WLCG, in order to avoid additional integration efforts for sites supporting multiple experiments. Furthermore, well established technologies like the Grid CEs as entry points should be re-used in order to minimise the development efforts and to simplify the integration into the existing submission infrastructure of the experiments. The Grid CEs at a given site act as a single point of entry

and first abstraction layer to all opportunistic resources associated with that site, respectively. This allows for hiding the complexity of submission to multiple opportunistic resources by providing one well-known job submission interface to the experiments. A so-called Overlay Batch System (OBS) supporting dynamic as well as transparent allocation, integration and management of resources e.g. HTCondor [5] is following as a second abstraction layer behind the Grid CE. This allows for dealing with the temporary availability and different suitability of the integrated opportunistic resources. The arriving workloads of the experiments can be dynamically matched by the OBS to the currently available and most suitable opportunistic resources. Whether a resource is suitable or not for a given workload can be determined by evaluating the HTCondor ClassAds [6] provided by the arriving workloads. Admittedly, the pilot mechanisms used by the experiments makes this challenging and an important prerequisite for an optimal matchmaking would be the possibility to know or even better to select, which workflows are executed in a given pilot based on workflow requirements like its type, estimated I/O rates, etc.. This is currently a matter of ongoing work together with the CMS experiment. The interaction between the OBS and the HTCondor based pilot system will be realized using the CMS site local configuration mechanism, a set of configuration files stored on the worker node itself. HTCondor ClassAds describing suitable tasks are read from the configuration and evaluated as HTCondor START expression by the CMS pilot, so that only selected payloads are allowed to run on the resource.

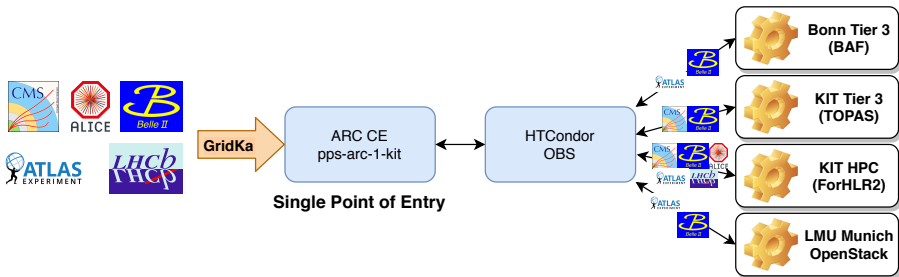


Figure 1. Multi-experiment setup to allow for a dynamic and transparent integration of opportunistic resources into WLCG computing at the German Tier-1 centre GridKa at KIT.

In figure 1 the currently deployed setup at the German Tier-1 centre GridKa at KIT is depicted. For reasons of development, the operation of opportunistic resources is currently decoupled for the Tier-1 production system. That means all supported experiments are submitting their workloads to a dedicated ARC-CE [7] at GridKa with a dedicated OBS running behind it doing the matchmaking between jobs and resources. However, a seamless integration of opportunistic resources into the Tier-1 batch system is intended in the future.

One additional aspect to consider is data access and storage for the opportunistic resources. In the setup described in figure 1, the associated Grid site is providing storage services that can be remotely accessed by the opportunistic resources via the XRootD protocol [8], while stage-out is usually done via SRM [9]. Remote data access can also be assisted by deploying XRootD based caches on storage capacity potentially available at a given opportunistic resource as described in [10].

At the moment of writing four different providers of opportunistic resources are integrated into the WLCG via this setup. The ForHLR II HPC cluster [11] located at KIT, TOPAS a high-throughput optimised analysis cluster located at KIT [12], the Bonn Analysis Facility (BAF) Tier-3 cluster at the University of Bonn as well as OpenStack [13] based cloud re-

sources located at the Leibniz Supercomputing Centre in Munich. As shown in figure 1, the resource providers usually do not support all experiments, since they are member of particular ones and hence only allow those to use their compute resources. The selection of eligible experiments is also done by evaluation of the HTCondor ClassAds described above.

4 The TARDIS Resource Manager

In order to allow for dynamic provisioning as well as a transparent integration of opportunistic resources of diverse providers into one common OBS, the Transparent Adaptive Resource Dynamic Integration System (TARDIS) [14] has been developed at KIT. TARDIS has a modular design depicted in figure 2 and is thus easily modifiable and extendable. TARDIS itself is composed of site adapters, batch system adapters and plugins. The site adapters provide interfaces to various cloud providers and batch systems in order to allocate and provision resources utilising pre-built virtual machines or containers. At the moment of writing, CloudStack [15] and the OpenStack APIs are supported as well as Moab [16], SLURM [17] and HTCondor batch systems. The batch system adapter provides an interface to the OBS for the integration and management of resources. At the moment of writing only HTCondor is supported as OBS. An implementation of a batch system adapter to support SLURM as OBS is ongoing work at the University of Freiburg. In addition, two plugins are currently available in TARDIS. The SQLite [18] plugin offers a persistent storage of the TARDIS resources states to allow for restarting TARDIS without losing track of managed resources. The Telegraf [19] plugin can be used for monitoring purposes.

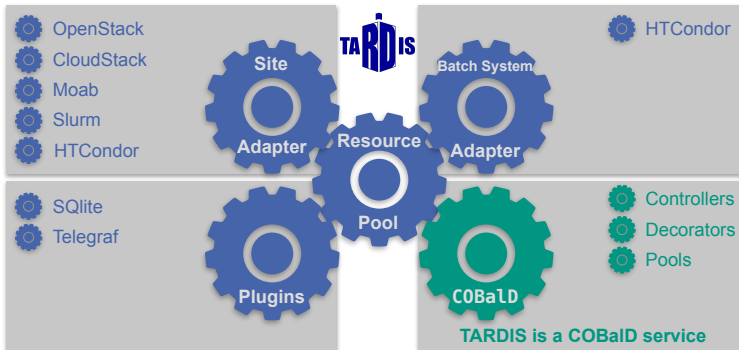


Figure 2. Modular and easily extendable design of the TARDIS resource manager.

TARDIS relies on COBaID – the opportunistic balancing daemon [20] which is also developed at KIT. COBaID is a light-weight opportunistic balancing daemon taking into account the actual utilisation and allocation of resources to decide whether they fit to the current job mix or not. Therefore, COBaID does not depend on a usually unreliable prediction of which resources might be suitable for the current workloads. Instead, it is just following a simple logic to increase the number of well-used resources and to release them otherwise. TARDIS then takes care of the resource life cycle management following the decisions made by COBaID. The actual matchmaking of jobs to resources is up to the OBS, which introduces only a loose coupling between the OBS and COBaID allowing for late binding of resources. COBaID achieves a native composition and therefore reduces management overheads by organising resources of the same type in pools managed as one. In addition, multiple COBaID agents can manage pools for the same OBS. This multi-agent approach together with the asynchronous

implementation of COBaLD and TARDIS ensures a very good horizontal scalability. Further details on COBaLD can be found in [21].

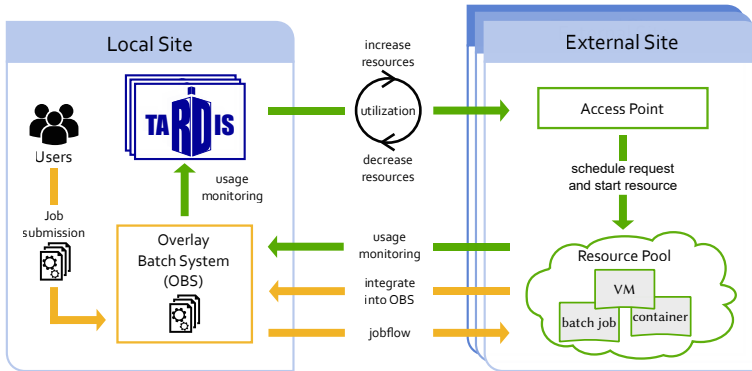


Figure 3. Simplified schematic diagram of the interplay between job submission, overlay batch system, TARDIS resource manager and opportunistic resource providers.

In figure 3 a simplified schematic diagram of the interplay between job submission, overlay batch system, TARDIS resource manager and opportunistic resource providers is shown.

5 Modus Operandi for Different Resource Providers

5.1 Dedicated Share at High Performance Computing Centres

The traditional way to get compute time on a HPC cluster is to write a dedicated project proposal that is reviewed by a HPC steering group. Once the project is granted a dedicated share is assigned to this project and resources can be requested on-demand. Within the scope of the bwForCluster initiative of the state of Baden-Wuerttemberg, a shared HPC centre for Neuroscience, Elementary Particle Physics, Microsystems Engineering and Material Science (NEMO) was procured at the University of Freiburg. Up to 33% (~6600 cores) of the resources can be granted to HEP projects. The KIT Institute of Particle Physics was granted a share of 8% (~1600 cores) at the NEMO HPC centre. In addition, implemented fair-share policies allow to exceed your share in case of otherwise idle resources are available. However, is not allowed to run arbitrary WLCG jobs on those resources due to political policies, since only employees at Baden-Wuerttemberg's universities are entitled to use them. Therefore, the NEMO HPC cluster is not depicted in figure 1, since it is only available for KIT user jobs. As shown in figure 4 a significant amount of opportunistic resources can be made available to the users at the KIT Tier-3 depending on the overall utilisation of the NEMO HPC cluster.

Our experience with the NEMO HPC cluster is very positive in all aspects. NEMO supports the utilisation of modern virtualisation and container technologies, which is an indispensable requirement for HEP computing on HPC clusters. The HEP group at KIT is strongly involved in project since the very beginning, so that the needs of HEP computing have been taken into account in the design decisions made for the NEMO HPC cluster.

5.2 Back-filling at High Performance Computing Centres

Resource schedulers of HPC clusters are usually optimised to schedule large-scale scientific simulation and calculation workflows distributed among many cores and nodes. This leads

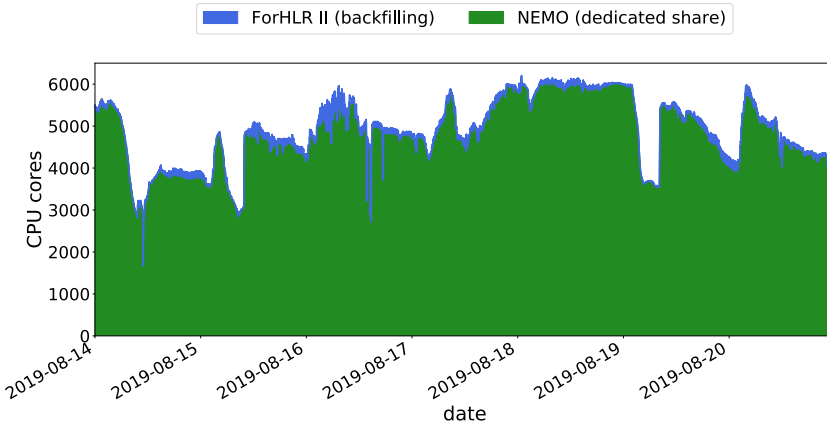


Figure 4. Opportunistic resources allocated at NEMO HPC cluster in Freiburg and the ForHLR II HPC cluster at KIT.

inevitably to unused resources due to fragmentation and draining of machines in order free nodes for the next large-scale workflow. One possibility to increase the overall utilisation is to enable back-filling with short running HEP jobs. However, those slots are usually subject to pre-emption in order to release resources once sufficient slots for the next large-scale HPC task are available. Unfortunately, the minority of WLCG experiments are supporting pre-emption in their experiment software yet, which would be useful for an efficient utilisation of those resources as well.

Our experience with back-filling of the KIT HPC cluster ForHLR II is positive as well. Singularity support has been enabled on request and the resource utilisation could be significantly increased, which is appreciated by both sides.

5.3 Back-filling of Tier-3 Resources

In Germany and potentially also in other countries, universities usually operate a significant amount of local compute resources for their analysis users, so-called Tier-3 centres. Most often there no Grid services deployed at these sites, due to a lack of person power. Hence, those resources are not available for WLCG computing. However, exhausting user job requirements and fluctuations in utilisation lead to unused resources, which can be easily back-filled with WLCG jobs. At the moment of writing, the Tier-3 centres at the University of Bonn and KIT have been made available to the WLCG via the approach described in section 3. As for back-filling of HPC resources, the support of pre-emption in the experiment software is inevitable for an efficient utilisation of those resources.

Our experience with back-filling of Tier-3 resources is very positive in all aspects. The physics group at the University of Bonn eagerly provides Tier-3 resources opportunistically to the ATLAS and Belle 2 experiments. They actively contribute to the development of COBa1D/TARDIS and give essential feedback that helps to further improve our approach.

5.4 Integration of Cloud Resources

The idea to develop COBa1D and TARDIS as well as the concept presented in section 3 originates from the participation of KIT in the Helix Nebula Science Cloud [22] in order to

dynamically integrate commercial cloud resources into the WLCG. Currently, only Open-Stack based cloud resources at the Leibniz Supercomputing Centre in Munich are integrated. In our experience the integration of commercial as well as public cloud resources is usually straightforward and hassle-free.

6 Conclusion and Outlook

We have presented a general multi-experiment solution on how to integrate opportunistic resources into the WLCG computing by associating them to existing WLCG sites close by and utilising well established Grid computing elements as single point of entry for the experiments. We showed that the dynamic resource provisioning and transparent integration of such resources into one common overlay batch system is efficiently handled by TARDIS and COBaLD taking into account the actual utilisation of the resources. For the future we have already concrete plans to integrate further resources utilising our tools into the WLCG computing. Further topics in terms of developments are to offer tools to account for providing opportunistic resources in APEL, the official accounting portal for sites participating in EGI and WLCG infrastructure. Furthermore, we plan to develop an extension of TARDIS to dynamically manage the deployment of Squid caches on opportunistic resources.

Acknowledgements

The authors acknowledge support by the state of Baden-Wuerttemberg through bwHPC and the German Research Foundation (DFG) through grant no INST 39/963-1 FUGG (bwFor-Cluster NEMO). Parts of this work were performed using the supercomputer ForHLR funded by the Ministry of Science, Research and the Arts Baden-Wuerttemberg and by the Federal Ministry of Education and Research. In addition, the authors acknowledge the fruitful collaboration with Peter Wienemann and Oliver Freyermuth from the University of Bonn.

References

- [1] Computing for the Large Hadron Collider, Ian Bird, Annual Review of Nuclear and Particle Science, 61, 99-118 (2011), doi:10.1146/annurev-nucl-102010-130059
- [2] A Roadmap for HEP Software and Computing R&D for the 2020s. Albrecht, J., Alves, A.A., Amadio, G. et al., Comput Softw Big Sci 3, 7 (2019). <https://doi.org/10.1007/s41781-018-0018-8>
- [3] Singularity: Scientific containers for mobility of compute. Kurtzer GM, Sochat V, Bauer MW (2017), PLoS ONE 12(5): e0177459, <https://doi.org/10.1371/journal.pone.0177459>
- [4] Distributing LHC application software and conditions databases using the CernVM file system, J Blomer et al., 2011 J. Phys.: Conf. Ser. 331 042003, <https://doi.org/10.1088/1742-6596/331/4/042003>
- [5] Condor – A Distributed Job Scheduler, Todd Tannenbaum and Derek Wright and Karen Miller and Miron Livny, Beowulf Cluster Computing with Linux , MIT Press, 2001, pp.307-350
- [6] Matchmaking Frameworks for Distributed Resource Management, Rajesh Raman, PhD thesis, University of Wisconsin, October 2000.
- [7] The NorduGrid project homepage, URL <http://www.nordugrid.org> [accessed 2020-03-09]

- [8] XROOTD/TXNetFile: A Highly Scalable Architecture for Data Access in the ROOT Environment, Dorigo A, Elmer P, Furano F and Hanushevsky A, Proceedings of TELE-INFO'0546:1–46:6, (2005)
- [9] Storage resource managers: Middleware components for grid storage, Shoshani, Arie, Alex Sim and Junmin Gu., NASA Conference Publication. NASA; 1998, 2002.
- [10] Dynamic Resource Extension for Data Intensive Computing with Specialized Software Environments on HPC Systems, C. Heidecker, et al., Proceedings of the 5th bwHPC Symposium, 2019, <http://dx.doi.org/10.15496/publikation-29051>
- [11] ForHLR: a New Tier-2 High-Performance Computing System for Research, Barthel, R. and S. Raffener (2017), Proceedings of the 3rd bwHPC-Symposium. Universitätsbibliothek Heidelberg, pp. 73–75, doi:10.11588/heibooks.308.418.
- [12] Setup and commissioning of a high-throughput analysis cluster, R. Caspart, et al., Available in this proceedings.
- [13] OpenStack Open Source Cloud Computing Software, URL <https://www.openstack.org> [accessed 2020-03-09].
- [14] TARDIS – Transparent Adaptive Resource Dynamic Integration System. Manuel Giffels, Matthias Schnepf, et al.. <https://doi.org/10.5281/zenodo.2240605>
- [15] Apache CloudStack – Open Source Cloudcomputing, URL <https://cloudstack.apache.org/> [accessed 2020-03-10].
- [16] Moab HPC Suite, URL <https://adaptivecomputing.com/cherry-services/moab-hpc/> [accessed 2020-03-10].
- [17] SLURM: Simple Linux Utility for Resource Management, Andy b. Yoo, et al., Job Scheduling Strategies for Parallel Processing pp. 44 – 60, 2002, Springer Berlin Heidelberg, https://doi.org/10.1007/10968987_3.
- [18] SQLite Database Engine Hipp, R, et. al., URL: <https://www.sqlite.org> [accessed 2020-03-10]
- [19] Telegraf documentation URL: [tps://docs.influxdata.com/telegraf](https://docs.influxdata.com/telegraf) [accessed 2020-03-10]
- [20] COBa1D – the Opportunistic Balancing Daemon. Max Fischer, et al.. <http://doi.org/10.5281/zenodo.1887872>
- [21] Lightweight dynamic integration of opportunistic resources, M. Fischer, et al., Available in this proceedings.
- [22] Europe's Leading Public-Private Partnership for Cloud, URL: <https://www.helix-nebula.eu> [accessed 2020-03-12]