T 66: Grid-Computing

Zeit: Dienstag 16:45–18:50 Raum: JUR 372

Gruppenbericht T 66.1 Di 16:45 JUR 372

Computing strategy to cope with the upcoming massive HEP and HI data collection — •Thomas Kress[1] and Kilian Schwarz[2] — [1]RWTH Aachen University, Physics Institute III B — [2]GSI, Helmholtzzentrum für Schwerionenforschung, Darmstadt

The LHC scientific program has led to numerous important physics results. This would not have been possible without an efficient processing of PetaBytes of data using the Worldwide LHC Computing Grid (WLCG). In the periods following the accelerator and detector upgrades, a huge increase in the data rate is expected. In addition, other big experiments like BELLE-2 and the FAIR collaborations will also take large amounts of data during the next years. So far the LHC computing strategy, based on Grid computing as a distribution of data and CPUs over a few hundred of dedicated sites, has met the challenges. However, to cope with substantially increased data volumes and correspondingly higher CPU requirements, new techniques like cloud computing and the usage of opportunistic resources are necessary. In parallel a reorganisation of the interplay of the computing sites is presently addressed by the evolving computing models of the affected experiments. Recently the Technical Advisory Board of the WLCG German Tier-1 site GridKa in Karlsruhe organised a meeting aimed to identify the guidelines for keeping German HEP and Heavy Ion computing excellent for future requirements. In a follow-up meeting working groups were launched in order to effectively organise the work on the above topics. The presentation will address the challenges, the German strategy, and the current status of the work packages.

T 66.2 Di 17:05 JUR 372

ATLAS simulation on ARM-64 bit servers; porting, validation and benchmarks — Gen Kawamura, Arnulf Quadt, and •Joshua Wyatt Smith — II. Physikalisches Institut, Georg-August-Universität Göttingen

The ATLAS experiment explores new hardware and software platforms that, in the future, may be more suited to its intensive workloads. An example is simulation; a CPU intensive workload that would profit drastically if it was more “portable” and therefore usable on a wider variety of platforms. We will present the latest results of the port of the ATLAS software stack onto new prototype ARM 64-bit servers. Patches were needed to introduce this new architecture into the build as well as correct for platform specific code that caused failures on non-x86 architectures. We will show selected results from the validation of the physics outputs on these ARM 64-bit servers. CPU, memory and IO intensive benchmarks using ATLAS specific environment and infrastructure have been performed, with a particular emphasis on the performance vs. energy consumption. From our results it is clear that the prototype ARM 64-bit server outperforms the standard Intel server in terms of Events/KWh.

T 66.3 Di 17:20 JUR 372

Data intensive workflows in the Cloud — Gen Kawamura[1], Oliver Keeble[2], Arnulf Quadt[1], and •Gerhard Rzehorz[1,2] — [1]II. Physikalisches Institut, Georg-August Universität Göttingen — [2]IT Department, CERN

Cloud computing in the Infrastructure as a Service case means to rent computing resources from commercial providers. For Monte-Carlo simulations, which record low data in- and output, running on the Cloud is mostly understood. This stems from the fact that no permanent storage is required and the network is not stressed. On the other hand, for data intensive workflows network and storage I/O can be crucial bottlenecks. Therefore, they are usually only computed on Grid sites that either have the required input data in their storage (fast interconnect) - or in a small number of cases the input is downloaded from other sites. This talk will provide answers on how to set up a Cloud to get the best workflow performance, addressing whether to implement a Cloud site with traditional Grid storage, whether to use object storage or even go storage-less altogether. Going further into detail, questions that will be answered are: How big is the influence of the computing to storage distance on the overall event throughput of the site? Can latency effects be mitigated by optimisation techniques? In order to assess this, the Workflow and Infrastructure Model with the output metric ETC = Events/Time/Cost is formed and applied. This output quantifies the different workflow and infrastructure configurations that are tested against each other in an understandable and comparable fashion.

T 66.4 Di 17:35 JUR 372

On-demand provisioning of HEP compute resources on shared HPC centers — •Frank Fischer, Günther Erli, Manuel Giffels, Thomas Hauth, and Günter Quast — Karlsruher Institut für Technologie

Rather than solely relying on dedicated HEP computing centers, it is nowadays more reasonable and flexible to utilize remote computing capacity via virtualization techniques.

Since last year the remote HPC center (NEMO Cluster, Freiburg University) was scaled by a factor of 20. This contribution reports on challenges, experiences and recent developments with scaling dynamic deployment from a small prototype system to a TOP500 cluster.

The remote and local systems are tied together with the ROCED scheduler [1] such that, from the user perspective, local and remote resources form a uniform, virtual computing cluster with a single point-of-entry. This environment serves the needs of researchers participating in the CMS, AMS and Belle II experiments at KIT. A job volume of over 2 million CPU hours per month (in average) is managed with this system.

[1] O. Oberst et al. Dynamic Extension of a Virtualized Cluster by using Cloud Resources, J. Phys.: Conference Ser. 396(3)032081, 2012

T 66.5 Di 17:50 JUR 372

Experience with Docker Container at Tier-3 Operations — •Matthias J. Schnepf, Christoph Heidecker, Frank Fischer, Max Fischer, Manuel Giffels, and Günter Quast — Karlsruhe Institue of Technology

Nowadays High Energy Physics (HEP) computing and analysis can profit a lot from available opportunistic resources and new technologies. The biggest challenges in taking advantages of these resources and technologies are the dedicated operating systems and software requirements by HEP.

A promising way to meet these challenges is the usage of Docker container technology, which allows to run HEP software on operating systems that are not officially support by HEP experiments.

The Institut für Experimentelle Kernphysik (IEKP) uses Docker container in combination with the batch system HTCondor to run physics analyses on desktop PCs and on dedicated worker nodes independent from their operating system.

This talk gives an overview about the experience with Docker containers in our infrastructure and possible upcoming projects.

T 66.6 Di 18:05 JUR 372

Advantages of coordinated caching for run 2 analyses — •Christoph Heidecker, Max Fischer, Manuel Giffels, and Günter Quast — Karlsruher Institut für Technologie

The current run 2 of the LHC delivers an enormous amount of data waiting for further processing. This includes not only big challenges to the storing capabilities, but also to the distribution of data to the end-user for further analyses. At the KIT, a coordinated caching on SSDs was designed to process large datasets on short turn around cycles. The typical workflow of an analysis at high energy physics requires an iterated execution of the analysis tools. Usually, this is done via a batch-system which allows a parallel execution on different worker nodes. The High Throughput Data Acquisition system caches the necessary datasets directly after the first execution on SSDs located at the worker nodes. This allows an faster processing of data from the second execution of the analysis. Hence, the processing time is reduced dramatically which allows a prompt deployment of new results. The iterated deployment of the jet energy corrections allowed an detailed testing of the caching system with an increasing amount of data. In this context, the developments and experiences of the Karlsruhe workgroup using this coordinated caching will be presented.

T 66.7 Di 18:20 JUR 372

Virtualization of the ATLAS Software Environment on the bwForCluster NEMO — •Ulrike Schnoor[1], Felix Buehrer[1], Anton Gamel[1], Konrad Meier[1,2], and Markus Schumacher[1] — [1]Albert-Ludwigs-Universitaet Freiburg, Physikalisches Institut —

[2]Albert-Ludwigs-Universitaet Freiburg, Rechenzentrum

High-Performance Computing (HPC) and other research cluster computing resources provided by universities can be useful as a supplement to the collaboration's own Grid computing resources for data analysis and production of simulated event samples. In Freiburg, the HPC cluster NEMO has been established as a part of the Baden-Württemberg (BW) HPC infrastructure and can be used by the BW particle physics community. The talk describes the concept and implementation of virtualizing the ATLAS software environment to run both data analysis and production on the NEMO host system. Main challenges include the integration into the NEMO and Tier-3 schedulers in a dynamic, on-demand way, providing access to the local user environment, as well as the automatic generation of a fully functional virtual machine image.

T 66.8 Di 18:35 JUR 372

GridKa Tier-1 Zentrum - Status und Zukunft — •Max Fischer und Andreas Petzold — Karlsruher Institut für Technologie, Karlsruhe, Deutschland

Als eines der weltweit 13 Tier-1 Zentren ist das GridKa am KIT ein zentraler Bestandteil des LHC Computing Grid. Mit den vier großen LHC-Experimenten, ALICE, ATLAS, CMS und LHCb, sowie anderen Kollaborationen als Nutzergruppen erfüllt das GridKa ein breites Spektrum an Anforderungen. Neben technischen Herausforderungen bildet es als Multi-VO Zentrum auch eine Basis für die Zusammenarbeit mit und zwischen Kollaborationen. Die Nähe zu Forschungsgruppen des KIT stellt dabei eine einzigartige Kombination aus Produktiv- und Forschungsbetrieb dar.

Um aktuellen und zukünftigen Anforderungen gerecht zu werden, wird die Infrastruktur am GridKa stetig weiterentwickelt. In diesem Rahmen werden beispielsweise die Rechenkapazitäten schrittweise auf das HTCondor Batch-System umgestellt. Als Tier-1 für mehrere VOs stellt das GridKa besonders auch Speicherkapazität für unterschiedlichste Ansprüche bereit. Hohe Anforderungen an langfristige Verfügbarkeit, Volumen und Durchsatz spiegeln sich sowohl in Technologie als auch Infrastruktur der verwendeten Speicher wieder.

Dieser Beitrag befasst sich mit dem aktuellen Status des GridKa als Multi-VO Tier-1. Fokus liegt dabei auf den Herausforderungen, um zukünftigen Anforderungen gerecht zu werden. Dies umfasst generell anwendbare Erfahrungen mit aktueller Technologie im WLCG-Umfeld.