Письма в ЭЧАЯ. 2016. Т. 13, № 5(203). С. 1041–1045

GRIDS AND CLOUDS IN THE CZECH NGI

J. Kundrát [a, b, 1], *M. Adam* [c], *D. Adamová* [c], *J. Chudoba* [a], *T. Kouba* [a], *M. Lokajíček* [a], *A. Mikula* [a], *V. Říkal* [a], *J. Švec* [a], *R. Vohnout* [b]

[a] Institute of Physics of the CAS, Prague

[b] CESNET, a.l.e., Prague

[c] Nuclear Physics Institute of the CAS, Řež, Czech Republic

There are several infrastructure operators within the Czech Republic NGI (National Grid Initiative), which provide users with the access to high-performance computing facilities over the grid and cloud interfaces. This article focuses on those where the primary author has personal first-hand experience. We cover some operational issues as well as the history of these facilities.

PACS: 07.05.-t; 07.05.Bx

INTRODUCTION

There are several high-performance computing (HPC) facilities in the Czech Republic and each of them operates with a specific focus. Some of them try to accommodate all scientific users through offering a generic set of services, while others are limited in scope to providing raw power to specific projects. There are various laboratories providing experimental, cutting-edge services with latest releases of the control plane software and state-of-the-art management, as well as more conservative infrastructures tasked with providing a reliable, production-grade substrate.

The infrastructures described in this work were selected based on personal first-hand experience of the primary author. The aim of this paper is to shed light on the day-to-day operational challenges rather than enumerating all possible services.

This work does not pretend to be an exhaustive reference. The MetaCentrum [1] and CERIT-SC [2] projects are just a few examples of prominent infrastructures which are not covered here.

PRAGUE TIER-2 LHC SITE

The Prague Tier-2 LHC computing center was founded in 2002 at the Institute of Physics of the CAS [3]. The first worker nodes (WNs) were thirty-four HP LP1000r 1U servers with dual-socket Intel Pentium III CPUs, 1 GB of RAM, and 18 GB SCSI hard drives. Over the years, the infrastructure has been regularly extended. In 2015, the entire cluster comprises 258 WNs (or 4756 CPU cores), 4.25 PB of disk storage, and a tape library. These resources

[1] E-mail: jan.kundrat@cesnet.cz

are hosted in a dedicated data center (62 m^2 floor space) with one 200 $kV \cdot A$ and two 100 $kV \cdot A$ UPSes and a 350 $kV \cdot A$ diesel generator, a matching number of computer room air conditioning (CRAC) units, and water-chilled racks. The network connectivity is provided by three separate 10 Gbps links providing IPv4 and IPv6 access to the public Internet and LHCONE [4] networks.

The Prague LHC site is a distributed facility. Additional hardware is hosted in a different data center provided by CESNET [5], the Czech National Research and Education Network (NREN). The CESNET site hosts 31 WNs providing a total of 920 CPU cores, 2416 GB RAM, and 180 TB of disk storage. Further circa 500 TB of disk space is hosted at the Nuclear Physics Institute of the CAS.

The local batch system control has been historically implemented with the PBSPro [6] software which got changed into the Torque [7]. Grid access from the WLCG computing grid is provided through a standard set of middleware software stack.

The system administrators have used various tools and helpers for managing the compute servers. In 2008, ad-hoc shell scripts were replaced by cfengine (version 2.x) [8]. In order to better participate in the collaboration with other computing centers in the HPC particle physics community, a migration to Puppet was started in 2014. By 2015, the majority of the grid services had been already migrated to Puppet [9]. There were experiments [10] with developing in-house applications for improving the inventory management, but they have been deprecated since.

Monitoring is implemented using a mixture of standard open-source tools as well as in-house scripts. Performance metrics are tracked using Munin [11] and Ganglia [12], while alerting is provided through Nagios [13]. The network traffic is monitored through a combination of MRTG [14] and Netflow [15] as well as by CESNET's FTAS [16] systems. Batch system statistics are provided through an in-house tool.

OPERATIONAL ISSUES

Similar to any other infrastructure of a comparable size, we had to solve our share of unexpected issues during the day-to-day operation of the Prague site.

One of these issues were problems related to the electricity distribution. The data center in which the Prague site is hosted has evolved from its original design (a liquid gas processing workshop) into a high-power–density facility. However, the electrical infrastructure was not that easy to upgrade, and as a result of that some of the distribution cables and fuses were being run at more than 90% of their rated load for a period of time. One evening, a 100 A fuse blew up, cutting power to one phase within one of the two isles. No important service was affected thanks to the built-in redundancy in power distribution and dual PSUs (Power Supply Units) in the non-WN servers. However, one of the most power hungry appliances was a blade chassis with three-phase PSUs. With one phase offline, all energy was being delivered through the remaining two phases which led to a cascading failure resulting in a meltdown of the second 100 A fuse. Thankfully, at that time, the blade chassis realized that there is no point in trying to run any further, and the third fuse remained operational. Needless to say, the next day was entertaining for all system administrators.

A long-standing operational issue were Ethernet ports on a central router. With no obvious pattern, but typically following a power cycle of the remote host, the Ethernet port in a switch's line card shuts down and refuses to come back up. Over the years, the vendor replaced all parts of the router, including the line cards, the PSU modules, the supervisor and firewall modules, and even the router chassis, but we still see occasional troubles and the root clause remains undiagnosed as in early 2016.

The Prague site has been one of the first sites running IPv6 in production [17, 18]. Some of the issues we faced were related to missing software support in various layers of the stack, as well as to operational issues in other parts of the network. For example, two of the nodes within a round-robin pool of machines at another computing center were rejecting IPv6 connections even though they were installed in an identical manner to the rest of the pool, etc. That said, our experience with IPv6 is positive in general, and we are happy that it relieves the pressure on our NAT (Network Address Translation) router.

FIWARE OPENSTACK CLUSTER AT CESNET

In contrast to the grid computing which usually provides access to a preinstalled machine with a well-known and centrally managed software stack, cloud computing or the Infrastructure-as-a-Service (IaaS) paradigm provides access to raw virtual machines (VMs). The end user is then tasked with the responsibility to install, operate and maintain an operating system (OS) and a complete application stack on the resources they rent. One of the public clouds operated at CESNET is offered through the FIWARE Lab [19], a part of the FI-PPP effort [20].

The hardware consists of eight servers (the SuperMicro FatTwin) offering the total capacity of 192 AMD Bulldozer 6344 cores and eight Apple Mac Pro sporting a total of 48 E5-1650v2 cores. The total RAM capacity is 1 TB, and the total storage consists of 1 TB of PCI-e SSDs, 3.84 TB of SATA SSDs, and 8.25 TB of 10,000 RPM HDDs. The topology is admittedly unusual, with the Apple hardware commissioned due to licensing issues with the Mac OS X software which is only legally allowed to run on Apple-manufactured hardware. This brought operational challenges to the table, with the Mac Pros having just a single PSU and no out-of-band management, for example. There is also no native support for 10 Gbps networking, so we are currently using a 4×1 Gbps trunking (with half of these ports implemented as Thunderbolt NICs).

Software-wise, the cluster is running on a mixture of Scientific Linux 6 and CentOS 7. The only supported access is through the OpenStack [21] APIs which offer native cloud access over HTTP.

One of the operational challenges was working with a centrally hosted Keystone, a requirement imposed by the federated architecture of the FIWARE cloud. The support for IPv6 also leaves something to be desired, especially in the Grizzly release. We had to backport some patches to make nova-network operate efficiently with our centrally provided Router-Advertisement infrastructure and flat L2 network.

MONITORING

At the core of the monitoring there is the collectd [22], a high-performance metric-gathering application. Thanks to the efficient use of the rrdcached, we are able to plot more than 10,000 metrics, most of them every ten seconds, with only negligible load on the central monitoring VM. This is in stark contrast to the Munin which was mentioned earlier where we had to resort to, e.g., operating on top of tmpfs, an in-memory filesystem, and a plotting interval of five minutes in order to bring the monitoring performance into a reasonable resource envelope.

Alerting is handled by Nagios, similar to the Prague WLCG site. There are also some internal monitoring applications, the majority of which use collectd's plugin features. As a rule of thumb, we are gathering all performance-related data in collectd's databases and are alerting on top of these rather than utilizing several competing systems.

ADD-ON SERVICES

Starting in 2015, we have deployed a new service built on top of this public cloud, the Continuous Integration (CI) for software developers [23]. The overall architecture is heavily inspired by the OpenStack project, with Gerrit [24] and Zuul [25] being at the core. As a proof-of-concept, this infrastructure is offering CI coverage for some projects within the KDE umbrella [26]. We have received positive comments about the CI functionality, with developers praising early feedback on proposed patches before they hit the official repository.

CONCLUSIONS

This article provided a short overview of some of the computing resources which are available through a collaboration under the umbrella of the Czech NGI. Together, these facilities empower the users to accessing high-performance computing resources as well as to self-service cloud-computing infrastructure.

REFERENCES

1. MetaCentrum NGI. http://www.metacentrum.cz/en/.
2. CERIT Scientific Cloud. http://www.cerit-sc.cz/en/.
3. Institute of Physics of the AS CR. http://www.fzu.cz/.
4. LHC Open Network Environment. http://lhcone.net/.
5. CESNET. https://www.cesnet.cz/?lang=en.
6. PBS Works: http://www.pbsworks.com/.
7. Terascale Open-Source Resource and QUEue Manager (TORQUE). http://www.adaptivecomputing.com/products/open-source/torque/.
8. *Kundrát J.* Using Cfengine for GRID Site Management. http://indico.twgrid.org/getFile.py/access?contribId=109&sessionId=43&resId=0&materialId=paper&confId=471.
9. Puppet. https://puppetlabs.com/.
10. *Kundrát J. et al.* Deska: Tool for Central Administration of a GRID Site. 2010.
11. Munin. http://munin-monitoring.org/.
12. Ganglia Monitoring System. http://ganglia.sourceforge.net/.
13. Nagios. https://www.nagios.org/.
14. MRTG — Multi Router Traffic Grapher. http://oss.oetiker.ch/mrtg/doc/mrtg.en.html.

15. FlowViewer. https://sourceforge.net/p/flowviewer/wiki/Home/.
16. Large Infrastructure CESNET: Institute of Physics Academy of Sciences of the Czech Republic — Network Connections. http://netreport.cesnet.cz/netreport/hep-cesnet-experimental-facility2/.
17. *Eliáš M. et al.* IPv6 Testing and Deployment at Prague Tier 2 // J. Phys. Conf. Ser. 2012. V. 396. P. 042034.
18. *Kouba T., Chudoba J., Eliáš M.* Enabling IPv6 at FZU — WLCG Tier2 in Prague // J. Phys. Conf. Ser. 2014. V. 513. P. 062027.
19. FIWARE Lab, the Open Innovation Lab. https://www.fiware.org/lab/.
20. Future Internet Public Private Partnership. https://www.fi-ppp.eu/.
21. OpenStack Open Source Cloud Computing Software. https://www.openstack.org/.
22. Collectd — The System Statistics Collection Daemon. https://collectd.org/.
23. *Kundrát J.* One Year with Gerrit. https://conf.kde.org/en/akademy2015/public/events/269, 2015.
24. Gerrit Code Review. https://www.gerritcodereview.com/.
25. Zuul — A Project Gating System. http://docs.openstack.org/infra/zuul/.
26. The KDE Community. https://www.kde.org/.