

The Czech e-Infrastructure and the European Grid Infrastructure Perspective

Ludek Matyska

CERIT-SC, Masaryk University, Brno, Czech Republic and CESNET, Prague, Czech Republic
ludek@ics.muni.cz

Abstract. National e-Infrastructures are playing an increasingly important role in the support of complex computational and data requirements from all scientific disciplines, environmental informatics not excepting. Since 1996, such an e-Infrastructure is developed and operated in the Czech Republic, with its emphasis shifting from a shared uniform distributed infrastructure to a more user-tailored environment. Its development relates (and in some cases precedes) the evolution of European Grid Infrastructure (EGI), with its current vision of an Open Science Commons concept. While the current e-Infrastructure concept and its implementation is of a very generic nature, it can be tailored to specifically cover different needs of environmental applications. This paper gives an overview of the Czech national e-Infrastructure, its connection to EGI and a number of applications from environmental science domains.

Keywords: e-Infrastructures, high throughput computing, compute and storage resources, weather forecast and modelling.

1 Introduction

Nowadays, high quality research and education is practically impossible without an extensive information and communication technology support. As there is an ever increasing request for more resources, it is necessary to combine compute and storage resources from different sites leading to the development of a concept of an e-Infrastructure (known also as Cyberinfrastructure in the USA). The e-Infrastructure is a distributed system that could span an institution, country or even a continent, combining resources sufficient for even the most demanding needs.

2 The Czech e-Infrastructure

2.1 Brief History

The history of a modern complex e-Infrastructure in the Czech Republic dates back to 1996, when Masaryk University through its Supercomputing Centre Brno coordinated the first project to build a distributed shared computing infrastructure known as “MetaCentrum” [1]. In 1998, the MetaCentrum coordination was transferred to CESNET,

the Czech National Research and Educational Network provider (NREN). The Meta-Centrum evolved gradually into the Czech National Grid Initiative.

The primary goal of the MetaCentrum was to build a distributed computing environment that exposed a uniform interface to its users; the emphasis was on uniformity, hiding the complexity of the used technology, but in the same time forcing users to adapt to the environment which MetaCentrum provided. MetaCentrum created a true grid – a combination of resources provided by different resource providers (public universities and institutes of Academy of Science) that together created a uniform compute and storage infrastructure. Recently, CESNET complemented the computing resources (clusters) with a distributed infrastructure of large scale hierarchical storage systems to cover also needs for big data storage and processing (the dawn of big data analytics).

MetaCentrum connects some 10 000 cores at 6 major and several more minor locations. The largest is at CERIT-SC. All the cores use x86_64 architecture, some nodes are also equipped with NVidia graphics. The total storage capacity are near 30 PB at four major installations throughout the country.

The SCB at Masaryk University, while offloading the responsibility for the management of the large scale distributed e-infrastructure to CESNET, continued to develop new concepts for the e-Infrastructure. In 2011, it started a transformation into the center CERIT Scientific Cloud, a highly flexible computing and data center focusing on user-tailored e-Infrastructure development and operation. CERIT-SC provides some 4500 cores and additional 4 PB of storage capacity; among the resources is also a SGI Ultraviolet 2 system with 6 TB of RAM shared among its 284 cores (a slightly larger system is expected to be delivered in the mid of November 2014).

Approximately at the same time, Technical University of Ostrava started its “IT4Innovations” project, to build a full scale national supercomputing center. In October 2014, SGI won a tender on the national supercomputer with its ICE X system that will include almost 600 basic nodes with nearly 900 Xeon Phi multicore processors—this will be the largest Xeon Phi installation in Europe.

Currently, the national e-Infrastructure landscape of the Czech Republic is composed of these three centers: CESNET with its MetaCentrum and extensive storage capacities, IT4Innovations supercomputing center and CERIT-SC as the grid and cloud center for collaborative research. All these centers are connected through the high speed network backbone provided by CESNET. The network provides capacity of ten of gigabits per second (with the main trunks recently upgraded to 100 Gbps lines). The network connects all the university cities and has sufficient capacity to support even the most demanding data transfers within the country.

2.2 Technology and Access

The basic distributed e-Infrastructure provided currently by MetaCentrum and CERIT-SC is based on a lightweight model of shared resources offered through a scheduling system (i.e., targeting a job submission mechanism). The system, built around the Torque schedulers [2], offers several queues that differ in their priority, the maximum time each job is allowed to run, and the queue capacity (how many jobs an

individual user can submit simultaneously). When submitting a job, users must select the appropriate queue and could annotate their job with additional information like number of cores needed, amount of needed disk space or extent of RAM needed for successful job completion. To support complex requirements, a specific web service is provided that can be used to check whether the used combination could be actually fulfilled by the e-Infrastructure, which combination of resources could serve the job with the given annotation (and how many such combinations exists). This service helps users to optimize their job descriptions to best meet their needs while taking into account the limits of the e-Infrastructure and its capacity.

Practically the whole computing infrastructure is virtualized, with part accessible through standard cloud interface as provided by OpenNebula [3]. This gives the system higher flexibility, as users are not restricted to the “standard” environment (operating system—in this case Debian—set of defined libraries plus an extensible set of applications), but they could submit whole virtual machines (either from a pre-defined pool or their own) with their own specific environment (including a completely different operating system like MS Windows).

With the virtualized nodes and the cloud interface, users are able to ask also for interactive access to the individual nodes or their sets. MetaCentrum provides even full virtual clusters [4,5], where a set of nodes could be connected through a virtualized private network and used as a whole. Use of network virtualization guarantees separation from other users, creating thus a secured (encapsulated) environment that is accessible through specified interfaces only (or not at all).

2.3 Scheduling and Fairshare

The whole e-Infrastructure is available free of charge to all bona-fide scientists and academicians in the Czech Republic (this includes also all university students). Users can just register with MetaCentrum and are immediately given access to the resources. With its near 900 users the demands of the e-Infrastructure exceed its capacity; therefore some scheduling is needed. The Czech national e-Infrastructure uses a specific fairshare mechanism [6] that dynamically adjust individual user’s priorities with respect to the extent of past computations.

When a user joins MetaCentrum, he is given some basic priority. The priority is automatically decreased when he uses the e-Infrastructure above average use; and it is increased when the usage is low. This means all users get a “fair” access to the e-Infrastructure. To support excellent research, user’s priority can be (permanently) increased if a result (usually a publication) is registered with MetaCentrum. This priority increase is immediately taken into account by the fairshare scheduler, giving such users better access to the resources. Users with high (and high quality) research output are thus visibly prioritized without any unnecessary bureaucracy (i.e. this system does not need to evaluate any a-priori user’s proposals for the use of e-Infrastructure while guaranteeing to highly productive users that they will have access regardless of the number of average users).

2.4 Flexibility and CERIT-SC

The center CERIT-SC has a specific position in the national landscape. It is not only the largest individual resource provider in the Czech Republic, but the center focuses on collaborative research with other user communities and also on further development of the e-Infrastructure itself. The center promotes the wide adoption of cloud technologies to provide highly flexible and adaptable compute and storage environment. The center works closely with its partners from other research fields to analyze the problem and to find out the most efficient ways of the use of the e-infrastructure. At the same time, it supports extensive modifications and adaptation to the e-Infrastructure itself to support its most innovative use.

The basic mode of work is the establishment of joint research teams together with users. Computer Science experts contributed by the center complement domain scientists from the application area; such interdisciplinary teams are best prepared to maximize the potential offered by the e-Infrastructure. These teams usually involve also bachelor, master and especially PhD students (both from the Computer Science and the particular research discipline) that are thus given an opportunity to directly participate in high quality research. See [7] for examples of successful collaborations.

3 European Grid I

The Czech e-Infrastructure is a part of the European Grid Infrastructure (EGI) [8]. The current EGI infrastructure is a result of 15 years of development. It currently connects 347 data centers in 57 participating countries (not all of them in Europe, one is even in Australia) and it provides 487,600 CPU cores, 286 PB disk and 118 PB tape storage. Its measured reliability is above 99.6 % and it runs around 1.5 million jobs each day (this is equivalent of almost 5 million core-hours per day). Since 2011, more than 2000 publications with EGI acknowledgment have been produced in almost all scientific disciplines in more than 200 projects.

Since the initial EU activities in the DataGrid project in late nineties, CESNET has been a strong partner in all major activities related to the European grid infrastructure. CESNET was a coordinator of the EGI Design Study project that led to the proposal EGI and its first consolidated project, EGI InSPIRE.

More recently, CESNET together with CERIT-SC are major players in the Federated Cloud activity that allows multiple IaaS cloud providers to become visible to the user as a single system that scales to user needs, provides resilience, prevents provider/vendor lock-in and could be targeted towards the research community. The Federated Cloud infrastructure currently includes resources from 12 countries; the Czech partners provide not only necessary resources but also play a strong role in the development of individual software components—the major contribution lies in the development of the rOCCI implementation to guarantee standards compliance of the federated Cloud approaches [9].

CESNET and CERIT-SC are also working together in the area of AAI and Identity management, especially through a joint development of the Perun system. Perun

serves as an extensive identity management and consolidation system supporting also authorization through VO and group concept, both at the national and also at international levels [10].

3.1 EGI Vision

As the primary grid (i.e., batch processing oriented) the EGI infrastructure is considered too static and inflexible by an increasing number of research communities: Therefore a new EGI vision has been developed and presented recently.

The vision states:

Researchers from all disciplines have easy and open access to the digital services, data, knowledge and expertise they need to collaborate and perform excellent research.

This vision found its expression through **Open Science Commons**, a concept based on three complementary pillars:

e-Infrastructure Commons: a flexible and dynamic ecosystem providing integrated services through interoperable infrastructures;

Open Data Commons: Observations, results, applications of scientific activities available for anyone to use and reuse;

Open Knowledge Commons: Easy access to knowledge and expertise to address challenges in education and research.

EGI positions itself to be the key provider of open access to compute, storage, data, knowledge and expertise complementing the community specific capabilities (the EGI Mission).

This mission is currently being implemented through a portfolio of platforms based on a middleware agnostic infrastructure. It covers and integrates all previous approaches like desktop grids, high throughput and high performance computing systems and is primary based on the cloud Infrastructure as a Service (IaaS) concept.

4 Environmental Applications

While the original grid-based e-Infrastructure provided a rather uniform environment that forced its users to adapt their applications to the environment, the recent trends are much more user-friendly. Using techniques of virtualization of all e-Infrastructure component, smart scheduling and cloud access policies, it is the e-Infrastructure that is adapted and tailored to best fit the applications, data manipulation, and workflows. Such user orientation is naturally beneficial for the environmental sciences, with their complex workflows, whose elements have extremely different requirements on the e-infrastructure.

An example is the SDI4Apps project (*Uptake of open geographic information through innovative services based on linked data*; <http://sdi4apps.eu/>, co-funded by the EU). Its main objectives are to integrate a new generation of spatial data infrastructure (SDI) based on user participation and social validation, support easy

discovery and accessibility of spatial data for everybody, and link spatial and non-spatial data using the Linked Open Data principles. To actually implement these objectives, the SDI4Apps rely extensively on a scalable cloud infrastructure, whose architecture is developed and continuously updated with the input from six pilot applications. These applications focus on easy access to data, tourism, sensor network, land use mapping, education and ecosystem services evaluation. The projects is trying to bridge the 1) top-down managed world of INSPIRE, Copernicus and GEOSS initiatives with 2) the bottom-up mobile world of voluntary initiatives and thousands of micro SMEs and individuals developing ad ho applications based on geospatial information. CERIT-SC is responsible for the architecture, development, and operation of the cloud infrastructure for SDI4Apps, making it a natural part of the national e-Infrastructure landscape. The experience built through the development for Federated Cloud infrastructure is directly used by the SDI4Apps, while the know-how achieved through this project implementation is transferred to other areas and disciplines, esp. in the area of storage and processing of geospatial data in a federated cloud environment.

The *Platform for the provision of specialized meteo-predictions for power plants*, a project funded by the Technology Agency of the Czech Republic, targets the development and implementation of a modular software system for prediction of electricity production from solar and wind power plants. It combines numerical models of weather forecast, whose outputs are among other intensity fields of global microwave radiation, cloudiness, air humidity, wind and temperature. The individual models are cross-correlated and the forecast values are used to remove errors and improve precision of the inputs into the model of power prediction [11]. The reliability and universality of the power production prediction is increased through “real-time” verification, using both the real-time weather and power production data to modify the forecast. The resulting highly reliable, modular and at the same time flexible system can be used by wide range of power generating plants. The development of this system extensively relies on the national e-Infrastructure that supports fast combination of different models, collection of data from real-time sensors and on demand processing in a virtualized environment.

The national e-Infrastructure is used to support reconstruction of 3D models of forest covers from full-waveform LiDAR and multispectral scans. This ultimate goal of such a reconstruction is to increase prediction precision for ecological models that simulate processes in natural and also human controlled ecosystems, to analyze the photosynthetic activity in the large vegetation covers etc. The precise 3D models of studied ecosystems play essential role in the prediction. These 3D models could be created from a combination of data provided through different techniques of surface and air scans (full-waveform LiDAR, hyperspectral scans, thermal scans etc.), eventually combined with “in-situ” measurements. Again, the target is to combine all the available data and improve thus analysis of independent individual data sets. The collaboration between experts in both environmental prediction and computer science already lead to the development of highly optimized 3D reconstruction algorithms of individual trees that are two orders of magnitude faster than other currently available methods [12].

The Atmospheric processes and modelling group, part of the Research Center for Toxic Compounds in the Environment (Recetox), studies processes which influence concentration of persistent organic pollutants in the atmosphere and other parts of the environment. The group aims to understand the processes important for the geographical distribution of pollutant and study time trends of their concentration. The studies use complex atmospheric models, e.g. to study air quality in Europe between 2009 and 2010 with hourly temporal resolution, space resolution 12x12 km and 28 vertical layers. The study uses the chemical transport model CMAQ (Community Multiscale Air Quality Model), the weather is forecast using the WRF (Weather Research and Forecasting Model). The WRF produces meteorological input fields for the CMAQ model which then performs modelling of advection, diffusion and chemistry of gas phase, aerosols and clouds [7:2011-2012].

All these examples of successful collaboration between domain science (in this case environment sciences) and e-Infrastructure experts demonstrate the validity and strength of the most recent evolution in the e-Infrastructures, namely their increased ability to support large scale data collections combined with extensive simulations on a tailored e-Infrastructure.

5 Conclusion

e-Infrastructures are becoming an indispensable tool for environmental sciences. The previous phase of the e-Infrastructure development required massive adaptation of applications and workflows to fit the provided e-Infrastructure; this led to rather disappointing experience that resulted in rather reluctant uptake of the large scale e-Infrastructures. Recently, with the shift to more user-tailored and user-friendly e-Infrastructures, using the most advanced techniques of virtualization and clouds to provide highly flexible and adaptable environments, the environmental sciences widely benefit from the existence of large scale e-Infrastructures. The Open Science Commons concept defines a novel way for the joint development of the e-Infrastructure and the ways it is used. At the Czech national level, CERIT-SC is the forerunner of such activities, through extensive collaboration with research communities and focusing on mutual co-development of both the e-infrastructure and the applications that are expected to run on it.

References

1. Matyska, L., Vocu, M., Krenek, A.: MetaCenter—Building a Virtual Supercomputer. ERCIM News, vol. 45 (2001)
2. Toth, S., Ruda, M.: Practical Experiences with Torque Meta-Scheduling in the Czech National Grid. In: Cracow Grid Workshop 2011, Cracow, Poland, pp. 33–45 (2012)
3. Sotomayor, B., Montero, R.S., Llorente, L.M., Foster, I.: Virtual Infrastructure management in private and hybrid clouds. *IEEE Internet Computing* 13(5), 14–22 (2009)
4. Antos, D., Sitera, J., Matyska, L., Holub, P.: Metacenter Virtual Networks. In: Cracow Grid Workshop 2008, Cracow, Poland, pp. 86–93 (2009)

5. Antos, D., Matyska, L., Holub, P., Sitera, J.: VirtCloud: Virtualizing Network for Grid Environments—First Experience. In: The 23rd IEEE International Conference on Advanced Information Networking and Applications (AINA) 2009, pp. 876–883. IEEE Comp. Soc., Bradford (2009)
6. Isard, M., Prabhakaran, V., Currey, J., Wieder, U., Talvar, K., Goldberg, A.: Quincy: Fair scheduling for distributed computing clusters. In: ACM SIGOPS 22nd Symposium on Operating Systems Principles, pp. 261–276 (2009)
7. The MetaCentrum Yearbooks,
<http://www.metacentrum.cz/en/about/results/yearbooks/>
8. Kranzmueller, D., Lucas de, J.M., Oster, P.: The European Grid Initiative (EGI). In: Remote Instrumentation and Virtual Laboratories, pp. 61–66. Springer (2010)
9. Rochwerger, B., et al.: The reservoir model and architecture for open federated cloud computing. *IBM Journal of Research and Development* 53(4), 1–11 (2009)
10. Prochazka, M., Licehammer, S., Matyska, L.: Perun—Modern Approach for user and Service Management. In: IST4Africa Conference Proceedings, Mauritius, 11p. (2014)
11. Stepanek, P., Zahradnicek, P., Farda, A.: Experiences with data quality control and homogenization of daily records of various meteorological elements in the Czech Republic in the period 1961–2010. *Időjárás. Quarterly journal of the Hungarian meteorological society* 117(1), 123–141 (2013)
12. Sloup, P., Rebok, T.: Automatic reconstruction of Norway spruce tree models from LiDAR data. In: ISPRS2013-SSG (2013)