

Optimisation

Introduction defining context and scope

System-level environmental science involves large quantities of data, often diverse and dispersed insofar as there are many different kinds of environmental data commonly held in small datasets. In addition, the velocity of data gathered from detectors and other instruments can be very large. Data-driven experiments require not only access to distributed data sources, but also parallelisation of computing tasks for the processing of data. The performance of these applications determines the productivity of scientific research and some degree of optimisation of system-level performance is urgently needed by the RI projects in ENVRI+ as they enter production.

This topic focuses on how to improve many of the common services needed to perform data analysis and experiments on research infrastructure, with an emphasis on how data is delivered and processed by the underlying e-infrastructure. There needs to be consideration the service levels offered by e-infrastructure, and of the available mechanisms for controlling the system-level quality of service (QoS) offered to researchers. This topic should therefore focus on the mechanisms available for making decisions on resources, services, data sources and potential execution platforms, and on scheduling the execution of tasks. The semantic linking framework developed in Task 5.3 on linking data, infrastructure, and the underlying network can be used to embed the necessary intelligence to guide these decision procedures (semi-)autonomously.

Ultimately, based on the relevant task (7.2) of the ENVRI+ project, we will need to:

1. Provide an effective mapping between research-level quality attributes (ease-of-use, responsiveness, workflow support) to infrastructure-level quality attributes on computing, storage and network services provided by underlying e-infrastructures.
2. Define test-bed requirements for software and services, and identify conditions for operating final software and services inside each domain, and between multiple domains.
3. Extend and customise existing optimisation mechanisms for computing and storage resources, and provide an effective control model between processes of data analysis and the underlying e-infrastructure resources, making the application performance as easy as possible to control at runtime.

Thus the focus of the technology review in ENVRI+ from the optimisation perspective is to determine two things:

1. What the RI projects already have at their disposal for effective data access, delivery and processing.
2. What mechanisms can be used to meet RI projects' [processing](#) and [optimisation requirements](#).

The optimisation section of the ENVRI+ technology review focuses on the second point above; the first point should be addressed in other sections, particularly [data processing](#).

Change history and amendment procedure

The review of this topic will be organised by [Paul Martin](#). He will partition the exploration and gathering of information and collaborate on the analysis and formulation of the initial report. Record details of the major steps in the change history table below. For further details of the complete procedure see item 4 on the [Getting Started](#) page.

Note: Do not record editorial / typographical changes. Only record significant changes of content.

| Date | Name | Institution | Nature of the information added / changed |
|-----------|-----------------------------|-------------|--|
| 3/1/2016 | Paul Martin | UvA | Provided introduction, context and scope for optimisation topic. |
| 21/3/2016 | Paul Martin | UvA | Initial draft for technology review report. |

Sources of information used

Analysis of state of the art and trends

In principle, optimisation can be conducted at every level of interaction—at the social level between investigators, at the interface level between researchers and their tools, at the service level, at the functional level, at the infrastructure level, and so forth. Any number of optimisations can be applied at each of these levels based on an understanding of the technologies and engineering extant at that particular level—a thousand different bespoke manipulations in order to ensure perfect operation.

In reality, while there will always be scope for hand-crafted solutions to every problem where the payoff is sufficient to offset the effort required to understand and produce those solutions, what is increasingly necessary is the ability to produce *generically* optimisable systems. As described in the [optimisation requirements analysis](#), there exist different ways for human experts to embed their insight into the operation of a system:

- The investigator engaging in an interaction can directly configure the system based on their own experience and knowledge of the infrastructure—this is the bespoke optimisation already alluded to.
- The creator of a service or process can embed their own understanding in how the infrastructure operates—this is key to producing high quality software/middleware/infrastructure, but is not always applicable to broader contexts.
- Experts encode their expertise as knowledge stored within the system, which can then be accessed and applied by autonomous systems embedded within the infrastructure—this is the approach that is being adopted in ENVRIplus, in its formal modelling, semantic linking, interoperable architecture design, and provenance support.

To embed knowledge into the system, it is necessary to do so at multiple levels, and it is necessary to link those different levels—from the abstract requirements of researchers to the fundamental characteristics of the infrastructure. This has been the focus of the technology review for optimisation in this instance.

Optimisation is conducted according to certain metrics measured at various levels from different perspectives. From the high-level user perspective, these metrics concern *quality of service (QoS)*.

Most experimental or analytical tasks, especially when distributed, are subject to degraded performance when limited by the underlying infrastructure, especially when that infrastructure is shared with other applications. Thus most QoS research is focused on telephony and the Internet: the International Telecommunication Union defined a standard for telephony QoS in 1994, to be revised in 2008 (ISO 2008); the ITU later defined a standard for information technology QoS in 1997 (ISO 1997). Regardless of context, QoS requirements are generally the same; the application requires certain levels of performance in terms of speed, stability, smoothness, response, etc. Advancements in distributed computing drive research into service-based infrastructures that provide assets on-demand, reacting to changes in the system in real-time (Menychtas et al. 2009). Thus the notion of QoS, wherein an application requires a certain level of performance (speed, stability, smoothness, etc.) from components, has been subjected to greater scrutiny of late as the demand to move more and more quality-critical applications onto the Internet raises reliability issues that may not be resolvable by blanket over-provisioning of computational and network resources. Li et al. (2012) proposes a taxonomy for cloud performance which can be generalised to Grid and other virtual infrastructure contexts, constructed across dimensions of performance features and experiments. Aceto et al. (2013) stress the importance of monitoring of virtualised environments.

If a system provides the ability to prioritise different applications, processes, users, or data-flows as opposed to simply making a best-effort attempt to do everything, then technical factors that influence the ability to fulfil QoS requirements include the reliability, scalability, effectiveness, sustainability, etc. of the underlying infrastructure and technology stack. Other factors however include the information models used to describe applications and infrastructure that then can be used to infer how to manage QoS requirements; for example (Kyriazis 2008) demonstrates how QoS might be specified and verified when mapping workflows onto Grid environments.

On the platform level, the QoS of the application and QoE of users are ensured by dynamically allocating resources with the fluctuations of workload. There are only limited resources and the computing and networking infrastructures also have a maximum capacity. Therefore all the resources have to be shared in a virtualized manner. So the challenge is to determine the resource requirements of each application and allocate resources most efficiently. The state of the art of this problem can be classified into resource *provisioning*, resource *allocation*, resource *adaptation* and resource *mapping* (Manvi and Shyam 2014).

Workflows provide a means for researchers and engineers to configure multi-stage computational tasks, whether as part of the generic operation of a research infrastructure or as part of a specific experiment. Workflows are typically expressed as directed (a)cyclic graphs. A key property is that workflows provide a means to manage dataflow. There are a number of different workflow management systems that could be enlisted by research infrastructure for framing workflows (Deelman et al. 2009)—e.g. Taverna, Pegasus and Kepler. The specification of workflows for complex experiments provides structural information to the operating environment about how different processes interrelate, and thus provides guidance as to how data and processes need to be staged in order to better support research activities. Given information about all the different workflows concurrent in a system, it is also then possible to regulate the scheduling of resources to best optimise *overall* system performance.

Conscripting elastic virtualised infrastructure services permits more ambitious data analysis and processing workflows, especially with regard to 'campaigns' where resources are enlisted only for a specific time period. Resources can be acquired, components installed, and processes executed with relatively little configuration time provided that the necessary tools and specifications are in place. These resources can then be released upon the completion of the immediate task. However in the research context, it is necessary to minimise the oversight and 'hands-on' requirement for researchers, and to automate as much as possible. This requires specialised software and intelligent support systems; such software either does not currently exist, or operates still at too low a level to significantly reduce the technical burden imposed on researchers, who would presumably rather concentrate on research rather than programming.

Finally, the adoption and collection of precise provenance information permits deep analysis of historical data/resource use, which can be used to refine decision procedures and so enhance the overall performance of the system.

The longer-term horizon

In the longer term, the increasing complexity and use of virtualised infrastructure will widen the gulf between researchers and the hands-on engineering necessary to manually configure the acquisition, curation, processing and publication of datasets, models and methods. Thus context-aware services will be required at all levels of computational infrastructure to manage and control the staging of data and the provisioning of resources for researchers autonomously, and these services will have to be aware of the state of the entire systems, catering not to the whims of individual researchers, but taking into account the wider use of the system by entire communities. The establishment of such topics will be wholly dependent on integrative thinking—taking heed not just of developments in individual areas of (for example) workflow management, provenance and cataloguing, but also the development of techniques to promote interoperation between all parts of research infrastructure.

Relationships with requirements and use cases

The optimisation topic is strongly related to the compute, storage and networking topic, the processing topic and the provenance topic in particular:

- The focus of optimisation is on more efficient use of underlying e-infrastructure, especially of the kind provided by initiatives such as EGI.
- The target of optimisation is on better data retrieval and processing.
- Autonomous optimisation relies on knowledge embedded in the datasets, services and resources involved in data retrieval and processing tasks—a significant portion of which is generated as part of provenance services.

There are a number of ENVRI+ use-cases for which the optimisation task is a potential contributor (see <https://envriplus.manageprojects.com/projects/wp9-service-validation-and-deployment-1/notebooks/625/pages/324>):

- The **data subscription service**, for the transport and staging of data onto cloud resources.
- Implementing a prototype **cross-RI provenance model** using workflow management systems and EUDAT services requires intelligent data movement and resource management.
- **Re-processing of data by users using their own algorithms** requires smart resource control.

Summary of analysis highlighting implications and issues

It is possible to automate large portions of research activity—however this is contingent on there being good formal descriptions of data and processes, and on there being good tool support for initiating and informing the automated procedures with regard specific experiments and applications.

The optimisation of resources is dependent on the requirements of researchers. The quality of service offered is based on certain taxonomies used to frame constraints that are then translated into requirements for the configuration of networks and infrastructure. Three branches can be distinguished in a classical performance taxonomy (Barbacci et al. 1995):

- *Concerns* list quality of service attributes that may be of concern to researchers.

- *Factors* lists properties of the environment that may impact concerns.
- *Methods* lists the mechanisms at the disposal of the system that can be used to monitor concerns.

It is necessary to identify the concerns of researchers in specific use-cases investigated within ENVRI+, and to analyse the factors dictating performance in current research infrastructures. The role of Task 7.2 in ENVRI+ is to provide methods for monitoring and responding to selected concerns.

The broader implications of generic optimisation of infrastructure and resources extends to the increasing prevalence of and reliance upon virtualised infrastructure and networks. Being able to generate a deeper understanding of how different kinds of task impose different requirements on different underlying infrastructure by being able to reason from the level of user-level quality constraints down to physical resource specifications is invaluable if we wish to be able to handle ever more extensive computational research. This is particularly true if we want to keep the accessibility of research assets as open to the broader research community as possible, rather than within the hands of a few well-resourced experts—in this light, we need to consider infrastructure as a utility, one that is intelligent and self-organising.

Bibliography and references to sources

Aceto, Giuseppe, et al. "Cloud monitoring: A survey." *Computer Networks* 57.9 (2013): 2093-2115.

Brooks, Peter, and Bjørn Hestnes. "User measures of quality of experience: why being objective and quantitative is important." *Network*, IEEE 24.2 (2010): 8-13.

Barbacci, Mario, et al. *Quality Attributes*. No. CMU/SEI-95-TR-021. CARNEGIE- MELLON UNIV PITTSBURGH PA SOFTWARE ENGINEERING INST, 1995.

Deelman, Ewa, Dennis Gannon, Matthew Shields, and Ian Taylor. "Workflows and e-Science: An overview of workflow system features and capabilities." *Future Generation Computer Systems* 25, no. 5 (2009): 528-540.

International Telecommunications Union. 1997. ITU-T X.641, information technology—quality of service: framework.

International Telecommunications Union. 2008. ITU-T E.800, definitions of terms related to quality of service.

Kyriazis, Dimosthenis, Konstantinos Tserpes, Andreas Menychtas, Antonis Litke, and Theodora Varvarigou. 2008. An innovative workflow mapping mechanism for grids in the frame of quality of service. *Future Generation Computer Systems* 24, no. 6, 498-511.

Li, Zheng, et al. "Towards a taxonomy of performance evaluation of commercial Cloud services." *Cloud Computing (CLOUD)*, 2012 IEEE 5th International Conference on. IEEE, 2012.

Manvi S S, Shyam G K. Resource management for Infrastructure as a Service (IaaS) in cloud computing: A survey. *Journal of Network and Computer Applications*, 2014, 41: 424-440.

Menychtas, Andreas, Dimosthenis Kyriazis, and Konstantinos Tserpes. 2009. Real-time reconfiguration for guaranteeing QoS provisioning levels in Grid environments. *Future Generation Computer Systems*, 25(7), 779-784.