COMMIT

WORKPLAN

WORKPACKAGES

DELIVERABLES

BUDGET

E-INFRASTRUCTURE VIRTUALIZATION FOR E-SCIENCE APPLICATIONS (P20)

Projectleader prof.dr. Henri Bal, Vrije Universiteit Amsterdam

1. Background

Increasingly, experimenters resort to the e-Science paradigm to design and implement their research; the e-Science paradigm builds a problem-solving environment, also known as a *virtual laboratory* that enables and stimulates multidisciplinary collaborations. Biobanking, bioscience, food, ecology, high-energy physics, and astronomy applications all require distributed collaborations and are examples of current applications using such laboratories (see projects P24 and P26). In the virtual laboratory heterogeneous resources must be accessed by the experimenter. Example resources are Grid and Cloud computing facilities, visualization displays, massive storage and knowledge databases, sensor networks, instrumentation and observatories, connected by optical wide-area networks. However, using such heterogeneous resources in a virtual laboratory is currently a tedious human process taking considerable effort in programming, search, allocation and manual configuration. In addition, the power consumption and carbon footprint of large-scale computer experiments is becoming a major issue in science and society that is usually ignored by application scientists. This problem is highly urgent.

Project P20 will address these problems and make it much easier for application scientists to find and access resources, program them, and use them into applications. The optimal e-Science architecture is one where scientists can specify what they want to achieve and where applications can take automated decisions while composing the virtual lab for an experiment, taking performance and energy consumption into account.

This project continues work that has been done in the BSIK project VL-e, which finished at the end of 2009. VL-e tried to bridge the gap between application scientists and the complicated underlying hardware/software infrastructure. In P20, we study ways to simplify (virtualize) the infrastructure, ease programming, and ease integration. P20 builds on web- and grid-technology developed by international research groups, so it is near the ``bottom'' of the virtual lab, closest to the infrastructure. Other COMMIT projects build on P20 and are closer to the application (the ``top'' of the virtual lab), in particular project P23 (Knowledge Management). These projects will make it possible to let application scientists use the e-Science infrastructure in an efficient, productive and even energy-preserving way, thus making the large investments in this national infrastructure much more effective.

The participants (UvA, VU, TU Delft, TNO, and Logica,) have many years of experience in large-scale distributed computing, heterogeneous grid resources, advanced networking, workflow management systems and generic middleware models. They have produced multiple-award winning software and they have collaborated for over 12 years in the Distributed ASCI Supercomputer (DAS) projects of ASCI (Advanced School for Computing and Imaging). The partners have complementary points of interest and thus can take advantage of each other’s expertise to better achieve the overall goals of this project

Our work will build on the state-of-the-art in the fields of Grid computing and Web services and on an important new development called Cloud computing. The Grid community builds much of the lower-level software needed for e-Science, in particular middleware that allows sharing of resources. Much of this software is standardized by the Open Grid Forum (OGF). The XtreemOS and Contrail EU projects (in which we participate) have developed a grid operating system (based on Linux) that exposes distributed resources using the operating system, and are applying this operating system to cloud computing platforms. The Semantic Web "provides a common framework that allows data to be shared and reused across application, enterprise, and community boundaries". In an e-Science infrastructure we can leverage this idea to build ontologies that describe the various resources and allow e-Science applications to consume this meaning. Cloud computing aims at delivering long-running and migration capable services and at virtualizing resources in a utility-like way, which also is useful for scientific applications.

2. Problem description

More and more research institutes and industries depend on a state-of-the-art hardware- and software-infrastructure for e-Science applications, the so-called e-Infrastructure. These infrastructures, however, are becoming increasingly difficult to build, scale, manage and use:

- The data explosion causes vastly increasing demands for storing, transporting, and processing distributed data.
- The infrastructures are becoming very heterogeneous, due to innovations in both hardware and software. Hardware examples are photonic networks, sensor networks, new processor architectures and accelerators (e.g., graphical processing units, multiprocessor systems-on-a-chip), each with widely different programming models and performance characteristics. Software examples are new Grid/Cloud operating systems, middleware, protocols, workflow systems, etc.
- There is a large gap between the applications and the infrastructure that steadily increases as the infrastructures become more complicated. Application scientists should not need the expertise to deal with the low-level details of the infrastructure. A top-down approach is required for applications.
- The power consumption and carbon footprint of the ICT infrastructures is a major concern.

Since many industries are faced with these problems, it is essential to better understand the underlying issues and to obtain new knowledge and solutions. The energy problem is of direct concern to society, especially as the ICT industry is overtaking the airline industry in energy consumption and pollution. There are many other projects in COMMIT that face the problems described below (see Section 5, Collaborations).

To address these problems, we will exploit several directions:

- We will use our expertise in photonic and programmable networks to develop application programming models for deterministic transport and processing of data in sensor, storage and processing resources. We will investigate distributed data-intensive applications, exploiting new technologies like light paths, network processors, next generation Ethernet and accelerators (GPUs, FPGAs, Cell).
- We will use the concept of virtualization to address heterogeneity and in general reduce complexity. Virtualization comes in many forms. Virtual machines (e.g., the JVM, Xen) can abstract from the underlying processor/operating system platform. Cloud computing builds on this technology to virtualize services and resources. We intend to take this work one step further and use an ontology-based approach to describe all components in an e-Science infrastructure and let application scientists construct exactly the virtual lab they need from such virtualized resources.
- We use workflow systems to bridge the gap between applications and infrastructure, to orchestrate the processing of data over the distributed e-Infrastructure using advanced programming systems.
- We will contribute to solving the carbon footprint problem by studying techniques that map e-Science applications onto heterogeneous resources, such as scheduling based on power-consumption monitoring and adding power-related attributes to resource information. Recent studies have shown that GPU accelerators have the potential of a disruptive performance gain at less energy consumption. Also, the systems in the supercomputer top500 (www.top500.org) have widely different energy/performance ratios.

We perform this research in collaboration with several major application areas including astronomy, multimedia image processing, the IJkdijk sensor network, and many others. Our results are applied in the design of the next generation of GigaPort and its international collaborations through the Lambda Exchange in Amsterdam. We use the NWO-funded DAS-4 system to test and validate our ideas. DAS-4 is a distributed system with 6 clusters connected by a photonic wide-area network. The clusters have many different types of accelerators. DAS-4 gives us a controlled environment for reproducible experiments. Throughout the project, we will use our expertise from VL-e, GigaPort, earlier NWO projects DAS-3, StarPlane, SCARIe, GUARD-G, and international collaborations on networks, grids, Clouds, and Green-IT.

The research questions we will study are the following:

1. How to describe the different types of resources using concepts of Semantic Web and Resource Description Framework (RDF), such that the search and allocation of compatible and connected resources can be automated?
2. How to program and tune instrument and network resources to realize a functional and sustainable virtual laboratory, with respect to functionality, robustness, predictable and knowable performance and collective behavior?

3. How to schedule and co-allocate the resources (computers, storage, and network links) in large-scale heterogeneous distributed systems for a variety of applications?
4. How can we map data-intensive e-Science applications onto large-scale hybrid systems, taking performance and energy-consumption into account?
5. How to develop programming systems that allow elastic scalability on clouds?
6. How to find and combine data and services and to automatically compose workflows out of such data and services? How to add provenance to workflow models and make scientific results reproducible?
7. How can middleware-independent APIs for applications as well as for programming environments and workflow systems be integrated into a consistent set that allow flexible functionality for both applications and environments?

3. Objectives

Project's goal

The modern distributed computer infrastructures used by scientists have become extremely complicated. The infrastructures use a diversity of processors (normal CPUs, but also accelerators like GPUs) and networks (e.g., optical networks, sensor networks). Also, these infrastructures contain many different computing resources from different organizations or providers, integrated by complicated grid or cloud software. The infrastructures have to be used for ever more challenging tasks, such as dealing with the data explosion, but also for privacy-sensitive applications. The P20 project will allow application scientists to effectively use such infrastructures. It will create a flexible e-Science architecture that can describe and optimize the programmable infrastructure components. It will also make tools that empower data-intensive distributed applications to exploit accelerators (like GPU cards) or optical networks to optimize their efficiency and reduce their energy consumption. The project will also build a workflow system that supports sharing of components, reproducibility, and interoperability. The workflow system will also enable efficient collaboration in development, running and sharing complex e- Science applications; this will reduce the effort in building new applications and requirements for computing and storage resources.

Planning of all dimensions

The main results of our project will be papers in top venues (ACM/IEEE journals and top conferences), software releases, demonstrators, and awards at competitions. An important result is an easy-to-use method to find, access, program, and use the resources of a distributed e-Science architecture. Our impact & valorization strategy is largely based on having well designed open-source software distributions with adequate support, encouraging researchers elsewhere to use our software. Dissemination takes place through popular and widely read scientific magazines (like IEEE Computer) and courses at all academic levels. International embedding takes place through many FP7 projects and through demonstrators (e.g., at the yearly SC conference) and competitions. We expect the results of P20 to be used at least by P6, P23, P24, and P26.

10 end goals of the project:

- Models for infrastructure descriptions (including virtualization and abstraction layers) with a set of intelligent dynamic selection algorithms for resource selection aimed at content-aware applications and data-centric applications. Ways to describe or allocate (virtual) resources using system descriptions, which may be shared and searched, such that attacks on the infrastructure can be avoided.
- Elastically scalable applications for at least 3 application areas that run in production on a grid or cloud.
- At least 3 mobile applications available on an app market that use cloud offloading and have a large number of users (downloads).
- An open source programming system for developing compute- and data-intensive e-Science applications +on large-scale hybrid distributed systems, including a variety of processor accelerators; demonstrators +implemented using the programming system - for (at least) the domains of multimedia computing (with P6, University of Amsterdam), astronomy (with Leiden University), and remote sensing (with University +of Extremadura, Spain).
- Set of intelligent mechanisms (integrated in the programming system) for static and dynamic selection of equi-kernels/hardware for increasing application performance and/or reducing energy consumption.
- A decentralized scheduler that is cloud-aware, accelerator-aware, cost-aware, and power-aware deployed on the DAS-4 and being able to access public clouds.
- A set of new scheduling policies and mechanisms for a set of (escience) application types and areas.
- A new comprehensive model of workflow applications and workflow management systems.
- This model will be innovative; it will cover different phases of workflow life cycle: building complex workflow applications, scheduling, monitoring of execution on distributed infrastructures, collecting and analyzing data. This model will take into account the collaborative nature of e-Science.
- A comprehensive set of interfaces and a system of components for implementing, orchestrating, and deploying scientific, many-task computing applications like workflows or parameter sweep applications

Results

P20 will research and create the methods and tools that allow application scientists to more easily find and access state-of-the-art and emerging ICT resources, +program them, and use them for applications. With the participation in many EU projects and many international publications, P20 supports the COMMIT objective of cutting-edge research. Several P20 results will have a societal impact; in particular, our development of open source software toolboxes to master e-Infrastructures and to process workflows on data in a highly distributed infrastructure will benefit also industrial communities and +ICT service providers. Also, we make important contributions to dissemination through demonstrators, which we will also use at international conferences and competitions.

Deliverable Impact and Valorization

An important valorization result will be that we will have built a rapid prototyping environment based on DAS-4 that can be used by a large number of computer scientists and application scientists in the Netherlands to do research on modern hybrid distributed computing systems. At the end of the project, we expect to have the right methods, tools, and software that will allow a wide range of researchers to effectively use such advanced but complex infrastructures in an effective way. In contrast, hybrid systems today are extremely complicated to use: essentially all parts of the pipeline are too difficult for many users, ranging from finding and obtaining resources, allocating, scheduling, programming, and deploying them. This environment will also serve as an example for other environments that aim to offer production facilities for end-user application scientists. With the forthcoming e-Science research center, there will be many of such scientists. TNO will use the P20 results in its applied research activities and transfer part of the technology to companies, either directly (e.g., the monitoring of energy grids, windmills, human health care) or through the UrbanFlood FP7 project on dike safety, or through its participation in FloodControl 2015. TNO will use the knowledge for its consultancy activities in telecommunications and data centres. The aim of our collaborations with the many application domains not only is a further understanding of our own methods and tools, but also to provide these +domains with the software capabilities to make a big leap forward in their research direction, research efficiency, and generation +of innovative (even groundbreaking) results. The new model of the workflow process and management system should result in improvement of the WS-VLAM system. This should result in broader usage of this system in other countries and to build a research community around this system.

Deliverable Dissemination

We plan to disseminate our results primarily to two broad groups:

(1) Computer scientists who are interested in setting up or programming modern e-Science infrastructures, and;

(2) Application scientists who want to use such infrastructures for solving their problems. Besides regular publications, a major dissemination mechanism is demonstrators. We will create working demonstrators for several domains (e.g., multimedia, astronomy, medicine, flood risks, and remote sensing), which will be presented at top venues. In particular, we plan to consistently provide demo's for the yearly SC conference, as this is by far the largest conference in our area (over 10,000 participants). In addition, we will participate in competitions (like SCALE) and provide demo's at various other venues (e.g., HPDC, CCGrid, e-Science, ICCS, CineGrid workshops). We also aim to publish innovative results in high-impact popular journals, e.g. IEEE

Computer. We will organize several workshops, including a DAS-4 opening workshop. We will also create software and demo's of distributed smart phone applications. Although most of our software was primarily designed for large-scale high-performance systems (like grids), some of it also is very suitable for mobile platforms or combinations of mobiles and clouds. Our earlier success with the Ibis-based PhotoShoot application +(44,000 downloads from the Android market) illustrates the advantage of targeting this huge market. Finally, the results of our work in P20 will also be used in courses of the UvA, VU, and TUD and of the ASCI research school.

International Imbedding

There is hardly any other country that provides a common distributed infrastructure for computer science like DAS. In virtually all other countries (except perhaps for France), computer scientists can only use production systems, which are shared with (many) application scientists, making it difficult to do controlled computer science experiments. The new DAS-4 system is designed for "diversity'' and gives project P20 an outstanding platform for more systematic research on virtualization, programmable networks, clouds, hybrid computing, scheduling, and workflow. This gives Dutch Computer Scientists a strong position in international collaborations. For example, we participate in several large FP7 projects that started recently and that will thus run in parallel with COMMIT. The CONTRAIL project studies cloud computing platforms. The UrbanFlood project investigates the use of sensors within flood embankments to support an online early warning system. The results of P20 will be used for the UrbanFlood Early Warning system for Dike Failures. The MAPPER project develops a system for building and running multi-scale applications from physiology, fusion, engineering, nano material science, and computational biology; this application are of workflow type. The NOVI and the Geysers FP7 projects create the formal models and the software for integration and federation of virtualized networking and computing infrastructures. The gSLM project aims to improve Service Level Management (SLM) in grid systems. VPH-Share will provide the organizational fabric, realized as a series of services, to expose and to manage data, information and tools, to enable the composition and operation of new Virtual Physiological Human workflows. We will also extend our collaboration with researchers and e-Science domains part of the EU COST ComplexHPC project. In particular, within this project we aim to expand the user-base of our Ibis software framework (e.g. by way of Spring School organization and tutorials). Also, members of P20 are very active in international program committees and steering committees of major conferences like HPDC, CCGrid, SC, and ICCS. +In 2012, we will organize the 20-th ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), which is a top conference in our field. In 2013, we will organize the 13-th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid). Both conferences (at TU Delft) also are ideally suited to show COMMIT results. Annually, we co-organize the Workshop on Large-scale System and Application Performance (LSAP) in conjunction with HPDC. We will also participate in international competitions. At least once every two years we aim to participate in the IEEE International

Scalable Computing challenge (SCALE), held in conjunction with the CCGrid conference. Other possible competitions are those for mobile applications and semantic web challenges.

Deliverable Synergy

- Collaboration P20-P1

  We will collaborate with (WP6 of) P1 (Information Retrieval for Information Services) on how to provide efficient application-type specific scheduling support for the map-reduce paradigm, and specifically for the query plan rewriting techniques of PigLatin as well as the S4 distributed stream computing platform. This support will be tested and analyzed with applications provided by P1.
- Collaboration P20-P6

  We will collaborate with P6 (in particular Dr. Cees Snoek) to enable scalable solutions for multimedia problems, in particular on hybrid distributed architectures. The collaboration is aimed at a combined deliverable with P6, in particular in the form of a software system (and demonstrator) in the context of the yearly international TRECVID benchmark evaluation.
- Collaboration P20-P9

  We intend to collaborate with P9 (Very Large Wireless Sensor Networks for Well-Being) on efficient resource management and scheduling for the offline analysis of the data obtained from monitoring sensor-based social networks in multi-cluster grids and clouds.
- Collaboration P20-P23

  We will collaborate intensively with P23 (Knowledge management) on efficient distributed execution of reasoning engines. P20 will produce rule based distributed reasoners (in different years) that will be used by P23. Also, in Y2 we will obtain from P23 (WP2) a proof-of-concept search and query application that will validate the ontologies developed in P20 against the tools developed in P23.
- Collaboration P20-P24

  We intend to collaborate with P24 (e-Biobanking with Imaging) on efficient resource management and scheduling for medical imaging applications (e.g., proteomics) in multi-cluster grids and in clouds.
- Collaboration P20-P24/26

  An extended survey of workflow design and validation methods will be available to other COMMIT projects, especially to the complex applications being developed in P24 (e-Biobanking with Imaging) and P26 (e-Science in Agrifood). The results of the developments in workflow will be implemented to develop and run workflow applications in these two COMMIT projects. The collaboration will start from the beginning of the project and will concern collections of requirements, using applications of P24 and P26 for testing and validation of our workflow system model, and finally we expect that they will be efficiently run with the system developed in WP8. P24 (in particular dr. Silvia Olabarriaga) will use our software to enable workflow provenance for medical imaging applications.

4. Economic and social relevance

The importance of a well-organized e-Science infrastructure is clearly explained in the ICTregie report (Dec. 2008) ``Towards a competitive ICT infrastructure for the scientific research in the Netherlands'' ([1]) that resulted in very positive reactions from the Dutch government and from analyst media in the US; e.g. the ``Cook report'' ([2]).

The NWO BigGrid proposal ([3]) argued that "it is very likely that a grid service provisioning industry will emerge and mature in the coming decade, which offers excellent opportunities for new economic activities." P20 will generate many innovations for such an industry as further explained by the application proposers in BIGGRID:
"It enhances the excellent position of Dutch academic hospitals in patient data collections using the grid for biobanking. It enables major advances in drug discovery through combining data and through availability of massive compute resources for modelling. It allows industrial research labs, such as Philips, to both contribute to and profit from the available resources for engineering sciences."

The e-Infrastructure enables academic and industrial research communities to communicate, collaborate and otherwise profit from the installed computing and storage resources. P20 enhances the excellent position of the Netherlands in distributed processing, networking and middleware to empower virtual laboratories.

A 'knowledge economy' such as the Netherlands is becoming increasingly dependent on automated techniques to produce, transfer, analyse, store, and manage vast amounts of digital data. The e-Science infrastructure is orders of magnitude more complicated than using normal PCs. The goal of our project is to engineer the e-Science infrastructure in such a way that it becomes much more accessible to non-experts. Our generic (domain-independent) approach ensures that solutions will be beneficial for industrial and scientific problems in a wide variety of application areas. Therefore, results obtained in this project will have broad potential socio-economic impact. The ability to transparently access and use large-scale computing infrastructures, whilst at the same time being able to pose requirements on the speed and cost (in terms of direct finances and/or power consumption) of the computations, will give scientific and industrial users clearly a decisive edge over competitors.

Some concrete socio-economical results from this project are:

- open source software toolboxes for the e-Science and industrial communities and ICT service providers to master e-Infrastructure and process workflows on data in a highly distributed infrastructure

[1](see http://www.ictregie.nl/publicaties/nl_08-NROI-258_Advies_ICT_infrastructuur_vdef.pdf
[2] search for february 2009 issue on http://cookreport.com/
[3]http://www.biggrid.nl/fileadmin/documents/proposal/BIG-GRID-executive-summary.pdf and http://www.nikhef.nl/grid/BIG/BIG-GRID.pdf

- demonstrations that show better scaling of scientific and industrial applications on the e-Infrastructure using the e-Science methodology and our software
- prototyping hybrid computing infrastructures to demonstrate how an intelligent choice of accelerators can reduce energy consumption for applications.

5. Consortium

The consortium consists of:

- The High Performance Distributed Computing group of Prof. Bal at the VU University in Amsterdam, which has over 15 years of experience with distributed programming environments.
- The Parallel and Distributed Computing group of Prof. Sips and Dr.Epema of TU Delft, which is one of the world-leading groups in scheduling and resource management systems for grids.
- The System and Network Engineering science group of Prof. de Laat at the University of Amsterdam, which has pioneered RDF ontologies to describe complex data networks and is leading research on Networks in the GigaPort project. The group has extensive expertise in Networking, security and privacy on the Grid, and Authorization architectures as is shown by their track record, contributions in standardization at IETF and OGF, visibility in the Global Lambda Integrated Facility (GLIF) community, and published papers in the respective fields.
- The Section Computational Science of Prof. Sloot at the University of Amsterdam, including Prof. Dr.Bubak, who has over 15 years of experience in parallel, distributed and grid computing and had major roles in the CrossGrid, K-WfGrid, CoreGRID, Virolab and GREDIA EU IST projects.
- Prof. Meijer of TNO has pioneered programmable networks to obtain deterministic behavior of sensor networks in the IJkdijk project.
- Logica has experience in automating the detection of anomalies in large scale computer infrastructures.

The groups have different expertise, on networking, distributed computing, grids, workflow, scheduling, and sensor networks. The groups have participated extensively in the VL-e project and in the various DAS projects. The Ibis and KOALA software tools they have developed are also integrated with each other. The groups have a substantial number of staff members and postdocs available. At the VU, Dr. Kielmann is a well-known grid expert with an excellent international reputation. Dr. Seinstra runs a large (1MEuro VU-ERC) project that applies grid computing to multimedia-content analysis. At the UvA, dr. Paola Grosso is a leading researcher (together with Prof. de Laat) in the Research on Networks program in GigaPort, and dr. Belloum has many years of research experience with workflow systems (e.g., VLAM-G). The VU and UvA have international (top-)master programs on Parallel and Distributed Computing Systems and System and Network Engineering and Computer Science , which generate many high-quality PhD candidates. Dr A. Iosup of the PDS group at TU Delft has ample experience in research in scheduling and resource management in grids.

Logica has developed software agents that analyze log files and reduce analysis efforts normally done by system administrators. Currently they are investigating how to extend this work in the direction of the detection of suspected patterns (intrusion, malicious usage), and case specific pro-active network management (black hole worker nodes). In parallel, using a social and human-machine perspective, they are interested in the interaction, integration and collaboration between software agents and system administrators. The experiences obtained in a previous project has also geared their motivations for future work in the domain of *cloud computing*. The aspect of elasticity and flexibility in price and number of resources (WP4) make cloud-infrastructures extremely volatile. We expect performance management challenges at organizational- as well as technical levels. Stability, reliability, and robustness in cloud infrastructures are key to their further growth. Logica focuses on research in the direction of governance and *pro-active cloud management*.

We will collaborate through the usual means (scientific meetings, workshops, joint papers, etc.), but also by setting up a common infrastructure (around DAS-4) and by exchanging and using each other's software. We have many years of experience with this approach. The software tools that we produce are an important result of our work which still is useful after the project finishes. For example, many of the tools we produced in VL-e have been distributed to other groups (partially in the form of the VL-e Proof-of-Concept distribution), so others can use them, even in production-like environments like Big Grid. In addition, we participate in standardization committees (e.g., we serve on core positions in OGF), which also has long-lasting effects.

Collaboration with other COMMIT projects:

- We collaborate with the e-Science application programs P24 and P26. We have six years of experience in VL-e with such collaborations, which were vital for the success of our research and for the software tools (e.g., Ibis, JavaGAT, KOALA) we produced. In particular our work with P24 focuses on efficient resource management and scheduling for medical imaging applications (e.g., proteomics) in multicluster grids and in clouds.
- We collaborate intensively with P23 (Knowledge management) on efficient distributed execution of reasoning engines. A new joint Ph.D. student between Prof. Bal and Prof. van Harmelen already obtained impressive results (including the CCGrid 2010 SCALE award and several publications) about a web-scale forward reasoning engine for RDFS and OWL. We will strengthen this new collaboration within COMMIT.
- We will collaborate with P6 (in particular Dr. CeesSnoek) to enable scalable solutions for multimedia problems, in particular on hybrid distributed architectures.
- We will collaborate with P1 (Information Retrieval for Information Services) on how to provide efficient application-type specific scheduling support for the map-reduce paradigm, and

- We intend to collaborate with P9 (Very Large Wireless Sensor Networks for Well-Being) on efficient resource management and scheduling for the offline analysis of the data obtained from monitoring sensor-based social networks in multicluster grids and clouds.

Collaboration with external partners:

- We collaborate with NIKHEF (through a joint Ph.D. student) on large-scale data management in high-energy physics.
- We collaborate with ASTRON (through a joint Postdoc) on the usage of GPUs for LOFAR and SKA applications
- We will present our results to the CineGrid collaboration during the annual workshops and to SURFnet via contacts in the GigaPort-RON project.
- We collaborate with LifeWatch
- We collaborate with many other institutes, such as Leiden Observatory (prof. S. Portegies-Zwart) and HR Wallingford, UK (dr. B Gouldby)

Since virtually all these partners have a strong international embedding, this also increases the international orientation of P20.

6. Workplan

The project is structured into the following Work Packages:

P20-WP1	Semantic Description of Infrastructure (SemDIF)
P20-WP2	Programmable Infrastructure (PIF)
P20-WP3	Security of (Virtual) e-Science Infrastructure (SeSI)
P20-WP4	Programming Systems for elastic Cloud applications (ProSysClouds)
P20-WP5	e-Science applications on large-scale hybrid distributed systems (SALSY)
P20-WP6	Generic scheduling and co-allocation (SCHOALA)
P20-WP7	Application type-specific Scheduling Support (ATYSUP)
P20-WP8	Workflow Process Modeling & Management (WOPMOM)
P20-WP9	Workflow Sharing and Reproducability (WSAR)
P20-WP10	Workflow and Application Component Integration (WACI)

WP1-WP3 enable resources and make them available to applications. WP4-WP7 support programming distributed programming and scheduling. WP8-WP-10 studythe integration with applications.

Measurable results per Workpackage:

- SemDIF: Toolsets and modules that generate, publish and retrieve information in RDF and algorithms that reason about different aspects of the e-Infrastructure to fulfil resource requests by applications.
- PIF: A programmable network architecture and libraries/tools that allow to program deterministic behavior in networks that can be integrated in the workflows of applications.
- SeSI: Methods for creating, finding or allocating (virtual) resources using system descriptions based on an application's security and privacy requirements.
- ProSysClouds: A set of programming systems supporting elastic scaling.
- SALSY: A working prototype system that automatically and dynamically maps applications onto a diversity of hardware architectures, taking speed and energy consumption into account.
- SCHOALA: working prototype of a reliable energy-aware co-allocating version of the KOALA scheduler for federations of heterogeneous inter-operating grids and clouds.
- ATYSUP: The operational support for several application types (e.g., data-intensive, workflows) in the KOALA scheduler.
- WOPMOM: A prototype of the workflow management environment that supports reusability across multiple workflow management systems
- WSAR: Services for capturing workflow provenance
- WACI: A working prototype of components for integrated application and workflow deployment across heterogeneous platforms.

Measureable results for the project as a whole:

- Software
- Publications in high-level conferences and journals
- An RDF framework that is able to describe the infrastructure components
- Algorithms that optimize the ICT infrastructure for typical usage scenarios
- Programming interfaces to configure the ICT services
- A state-of-the-art distributed computing system supporting modern hybrid processor architectures on an advanced optical wide-area interconnect for cloud and green-it experiments.

Organizational risks:

- Integration of the solutions developed for each of the issues at hand may be difficult.
- Going from application-specific to generic may be difficult – or otherwise ad-hoc – because our available set of example applications may not cover all common data access and execution patterns.
- Too few external applications and early adopters; - Frequent changes in the specification of an implemented standard;
- Integration problems among different virtual laboratory frameworks;

DELIVERABLES

Project number P20	
WP title & acronym	Semantic Description of Infrastructure (SemDIF)
WP leader	Dr. Paola Grosso, UvA–SNE

Objectives

WP1 addresses the research question R1: *How do we describe the different types of resources in RDF, such that the search and allocation of compatible and connected resources can be automated*? The current challenge is to capture the virtualization of devices at all levels in the infrastructure, and to handle multiple levels of abstractions when composing sets of resources in multi-technology multi-domain environments. This works build on the experience gained with the development of NDL in the GigaPort project; it also creates research synergies with the ongoing efforts on models for e-Infrastructures in European projects such as NOVI and Geysers. The goal of the workpackage is to produce ontologies for the different types of resources present in e-Infrastructure. The possibility to reason about resources opens up ways to more tightly coupled application demands and infrastructure capabilities. The PhD student will develop the schemas and the ontologies and will validate them using existing e-Science applications and simulation models. One promising area where we can apply our result is in the development of content-aware and data-centric network applications. The student will also develop a search mechanism that can logically reason about the resources based on their descriptions and automatically compose a virtual lab consisting of compatible components.

Project number P20	
WP title & acronym	Programmable Infrastructure (PIF)
WP leader	Prof. Dr. Robert Meijer, TNO

Objectives

The goal of WP2 is to give e-Science the capacity to continuously adapt the service of the network and computing resources for an optimal performance of specific virtual laboratory experiments. WP2 addresses the research question R2: *How to program and tune instrument and network resources to realize a functional and sustainable virtual laboratory, with respect to functionality, robustness, predictable and knowable performance and collective behavior*? WP2 applies and extends recent developments in the semantic definition and the programming of distributed systems behavior with respect of network and computational resource usage, which includes the manipulation of network and computing topologies. WP2 has the following activities: A1) develops algorithms for automatic multi-scale ICT optimization, A2) results in architectures and implementations for multi-scale ICT optimization facilities, and A3) creates a domain specific language (DSL) with which the temporal optimization of the behavior (service) of network and computing resources (e.g. topology, capacity) is programmed. Activity A4) creates e-Science workflows that program the behavior of the ICT resources to facilitate an optimal performance of instruments and the e-Science ICT environment itself. Activities A1 and A2 will result also in a PhD thesis, as well as activities A3 and A4. Part of A4 is the dissemination of the results in the 'Early warning system for dike failures' as developed by the IJkdijk Foundation and the UrbanFlood FP7 project where the sensor network, the sensor telecommunication network and the computational infrastructure are resources whose capacities continuously have to match demand. This workpackage relates to P5-WP4 and close collaboration is expected.

Project number P20	
WP title & acronym	Security of (Virtual) e-Science Infrastructure (SeSI)
WP leader	Prof. Dr. Ir. Cees de Laat, UvA-SNE

Objectives

Applications have increasing demands on e-Science infrastructure. This particularly applies for (bio-medical) applications (e.g., P24) which process privacy sensitive information. With the emergence of clouds, infrastructure becomes increasingly dynamic, as different virtual resources may be deployed or migrated over different physical resources. Applications have differing requirements too. In WP3, we study the interaction between application

requirements and the e-Science infrastructure. We will research different methods for creating, finding or allocating (virtual) resources based on an application's requirements. A particular focus is on security and privacy requirements, while keeping an eye on network and "green" requirements. System descriptions (e.g., in RDF) can be used declaratively for on-the-fly creation of virtual machines with a minimal trusted compute base (TCB). A minimal TCB is important for security, as it contains only those components and services needed for a specific application. We will devise methods to compose system descriptions for specific applications. Meanwhile, we need to solve the problem that (outdated or vulnerable) virtual machines and their descriptions can be abused. For this, we require on-the-fly analysis of system descriptions. In this project, we will research the methods needed to create, analyse, use, and possibly share system descriptions for virtual machines in such a way that security is helped, not hampered.

Project number P20		
WP title & acronym	Programming Systems for elastic Cloud applications (ProSysClouds)	
WP leader	Prof. Dr. Ir. Henri Bal, VU-HPDC	

Objective:

Study how to develop flexible programming systems for data-intensive Cloud applications.

Background: to process large data sets, it is attractive to use both local resources (PCs, clusters) and external clouds, which are ideally suited for handling peak loads. Because cloud users pay per resource use, applications should not just be able to scale up, but rather they should adapt themselves to whatever scale is most cost-effective, including scaling down to fewer resources. This is called ``elastic scalability''. We propose to develop programming systems that allow applications to scale elastically between varying numbers of localand cloud resources.

Description: develop a set of programming systems for elastic computing environments, based on Ibis (which supports malleability through Satin, JEL, and Zorilla). Implement applications such as distributed reasoning (with P23) and N-body simulations (astronomy). Extend the work beyond scientific applications by also applying it to mobile systems, for example offloading work from smart phones to Clouds. Do controlled performance experiments on DAS-4 and larger-scale experiments on other grids and cloud systems, taking (a.o.) scalability, efficiency, network stress, ease-of-use, and energy consumption into account.

Project number 20		
WP title & acronym	e-Science applications on large-scale hybrid distributed systems (SALSY)	
WP leader	Dr. Frank Seinstra, VU-HPDC	

Objectives

Study how to develop compute- and data-intensive e-Science applications on large-scale hybrid distributed systems, taking performance and energy-consumption into account.

Background: Future high-performance and distributed systems (clusters, Grids, Clouds) will be far more diverse than today’s systems, as new types of processor accelerators are becoming mainstream. Examples include multi-core processors (e.g. AMD’s 48-core MagnyCours), Graphics Processing Units (GPUs), and Field Programmable Gate Arrays (FPGAs). Apart from the need for speed and scalability for many e-Science applications, there is increasing concern about power consumption.

Description: Develop a distributed programming environment by extending the successful Ibis system with support for accelerators, by integrating multiple equivalent compute kernels (‘equi-kernels’) each targeted at one specific accelerator. Implement applications from e-Science applications in COMMIT and elsewhere using the programming environment. Do controlled experiments on DAS-4 and large-scale experiments on other Grids and Cloud systems. Integrate the scheduling algorithms developed in this Project into the environment, allowing applications to make tradeoffs between performance and energy consumption.

Project number P20	
WP title & acronym	Generic Scheduling and Co-Allocation (SCHOALA)
WP leader	Dr. Ir. Dick Epema, TUD

Objectives

Study how to develop generic scheduling and co-allocation models, mechanisms, and policies for inter-operating federated grids and clouds that is reliable, scalable, cost-aware, and power-aware.

Background: Large-scale distributed systems such as multi-cluster systems, grids, and clouds evolve along multiple dimensions: more heterogeneity in processor types, different ownership and cost models, different types of middleware, larger scale and federations of systems, new faster networks, etc. In order to schedule the resources in these systems, new generic mechanisms and policies for resource discovery, resource selection, and resource (co-) allocation are needed. No single scheduler can provide all this functionality on its own; so multiple cooperating schedulers are required in the large environments of tomorrow.

Description: Develop generic scheduling and co-allocation models, mechanisms, and policies on the basis of the current KOALA scheduler that can deal with resource heterogeneity, multiple types of systems (grids, clouds) in single federations, and that are reliable, cost-aware, power-aware, and can deal with dynamic resource availability. These models and policies will be studied at a more fundamental level, by means of the design of extensions to KOALA, and by means of controlled performance experiments on the DAS and other grids and cloud systems.

Project number P20	
WP title & acronym	Application type-specific Scheduling Support (ATYSUP)
WP leader	Dr. Ir. Dick Epema, TUD

Objective

Study how to develop support for specific application types and application areas in schedulers for large-scale systems such as grids and clouds.

Background: Different application types, both old (e.g., workflows) and new (e.g., mapreduce and data-mining at unprecedented scales), may want to take advantage of the capabilities of evolving large-scale computing infrastructures, such as (inter-operating, federated) grids and clouds with their heterogeneous hardware (processor types) and high-speed networks. No single scheduler can incorporate support for all imaginable types of applications for all types of computing infrastructures. Therefore there is a need to design and build application-support modules in schedulers of large-scale distributed systems.

Description: Develop support in the KOALA scheduler for specific application types and for applications in specific application areas (e.g., life sciences) that are deployed in grids and clouds, and, in particular, that are studied in COMMIT. We will extend the current KOALA scheduler with application-support modules for these application types and areas in such a way that users do not have to deal with details of the underlying systems but can concentrate on application design, and that applications can run (energy-)efficiently and cost-aware.

Project number P20	
WP title & acronym	Workflow Process Modeling& Management (WOPMOM)
WP leader	Prof. Dr. Marian Bubak, UvA-SCS

Objectives

WP20.8 will investigate the challenging issues in modeling and integrating workflow processes of complex experiments, which go beyond execution of installed modules and may involve experimentation with a kind of exploratory programming. These issues are: programming models and tools support for collaborations, achieving reusability cross multiple workflow management systems and provisioning both computing and network resources required sustaining CPU and data intensive applications. The aim is to study and develop a workflow management environment that supports the different phases of the workflow lifecycle: modeling of complex workflows, scheduling and monitoring their execution, and collecting and analyzing data in a collaborative way**.** The main task for the AiO is to develop new methods for modelling and enacting workflow on grid and cloud resources and to allow multi-workflow engine execution. These new methods will be validated using applications developed by the consortium; especially through the collaboration with P24 where workflow support is needed in two WPs: P24.2 and P24.3. This WP will closely collaborate with WP9 in providing methods for complex applications modelling and programming. This WP will also collaborate with P1 on methods of semantic integration and with P2 on collaboration methods and tools.

Project number P20	
WP title & acronym	Workflow Sharing and Reproducibility (WSAR)
WP leader	Dr.Adam Belloum, UvA-SCS

Objectives

WP9 aims at developing methods enabling workflow sharing and reproducibility of results. To enable the sharing of workflows, and make science reproducible, novel methods are needed to enable semantic annotation of workflows and collect provenance information. To achieve this goal methods for supporting cooperative experiments and transfer of knowledge among scientists have to be studied in depth using Web 2.0, and tools developed by the Computer Supported Cooperative Work society. The main task of the AiO is to define generic metadata ontology, schema, and workflow templates that describe the workflow design/pattern and the application and data sources. New methods and tools of domain-oriented semantic descriptions of basic application modules, which, together with knowledge gathered, may result in facilitating an automatic construction of workflow applications. The second task is to apply these methods to improve interoperability among different workflow systems to be achieved through the workflowBus concept. Seamless data integration mechanisms allowing the aggregation of multi-format and multi-source data into a single coherent schema while preserving the autonomy of underlying data sources will also be investigated and developed. The AiO in WP9 will collaborate in particular with WP24.3 in P24, which will provide real requirements and driving use cases. Within P20, the AiO will exploit possible outcomes in WP8, on workflow modeling, and provide generic services that can help other work packages in P20 to capture provenance information.

Project number P20	
WP title & acronym	Workflow and Application Component Integration (WACI)
WP leader	Dr. Thilo Kielmann, VU-HPDC

Objectives, background and description of work in WP

Objective: study how to build e-Science applications and their supporting orchestration environments (workflows, bags-of-tasks, many-task computing environments) based on a consistent set of middleware-independent programming interfaces (API's).

Background: The e-Science community is divided in different middleware camps, focusing on and being bound to their respective middleware environment, for example, gLite, Globus, NorduGrid, Unicore, OMII-UK, or cloud systems like Amazon EC2, Nimbus, Eucalyptus, or OpenNebula. These (and other) middleware platforms provide their native programming models and interfaces; there is no consensus or coordination among these interfaces. Applications and support environments are thus being developed for these particular middleware systems, prohibiting the use of these community-specific codes across different middleware or even middleware releases.

Description: Develop a consistent and cohesive system of API's that supports both e-Science applications themselves as well as supportive programming environments like workflow systems or, more generally, many-task computing systems. Starting point will be the OGF-standardized Simple API for Grid Applications (SAGA) and Open Cloud Computing Interface (OCCI). Build a system of components that implement these API's and facilitate the orchestration of e-Science applications with these components. Evaluate (and possibly refine) these API's using e-Science applications and workflow systems used within the project.

DELIVERABLES

Number of important journal paper
20

Number of important conference contributions
35

Products
-

Software

1. Ontologies for e- Infrastructure description and related software Ontologies for e-Infrastructure description and related software, WP1 (final version Y3).
   The ontology we build in COMMIT must support scheduling and planning decisions within generic e-Infrastructure. Our final OWL ontology fully captures the resources present in a generic e-Infrastructure, and it provides the information needed to support e-Services. We build a functional and complete suite of research and composition tools. These tools are used within P20 and COMMIT to support selection of resources and creation of services. These tools can also be incorporated by external users interested in using the WP1 ontologies in other projects.
   - WP 1 YP 2013

2. Software architecture, algorithms and tools/libraries for programming (sensor) networks and cloud environments.
   Software architecture, algorithms and tools/libraries for programming (sensor) networks and cloud environments, WP2, Q1, Y2. Software. Proof of concept implementations, and robust implementations of software that enables software to program the dynamic (adaptive) service of a network and cloud as a whole as well as individual network elements and virtual machines specifically. Special attention is paid that the technology is able to visualize itself for the benefit of demonstrations and education.
   - WP 2 YP 2012

3. Final code integrated with demonstration application
   Final code integrated with demonstration application, WP2, Q1, Y4 The demonstration application creates a "HelloWorld" / cookbook / templates combining the results of WP2 into a logical bundle able to be distributed and used elsewere. Opensource. Special attention is paid that the technology is able to visualize it self for the benefit of demonstrations and education.
   - WP 2 YP 2014

4. Software to construct or use host descriptions for VM usage

Deliverable: Software (Y3): Software to construct or use host descriptions for VM usage and creation targeted to security/privacy sensitive applications. We will devise tools to create (and search) host descriptions which fully describe a system with security-related configuration details; this is imperative to protect privacy-sensitive data and/or computations in the Grid or cloud. Our tools will allow for searching and deploying existing, pre-configured VM images, and for detection of misconfigurations (outdated) insecure configurations. A particular challenge is that we will try to construct secure, minimal configurations of operating systems on the fly, based on "declarative" host descriptions. The produced software tools should be very useful to manage (dynamically created) resources on the grid.

- WP 3 YP 2013

5. Hybrid distributed reasoner Hybrid distributed reasoner, WP4, Y3

   Both backward and forward reasoning have advantages and disadvantages. The former is computationally expensive over Web-scale data, while the latter cannot eciently deal with datasets which change frequently or when reasoning is needed only on a small portion of the data. We will develop a distributed method to perform hybrid reasoning which balances between the high upfront costs of forward reasoning and the query-time costs of backward reasoning. Hybrid reasoning allows a scalable and distributed architecture that is able to handle large amounts of data with little overhead. Hybrid reasoning will also greatly reduce the complexity of reasoning performed at query time.

   - WP 4 YP 2013

6. Ibis/Constellation

   Deliverable: Software Distribution, Y1-Y4 Package: Ibis/Constellation (as part of the larger Ibis software system) Ibis/Constellation, a software system that allows easy implementation of realistic scientific applications on so-called Jungle Computing Systems. Such systems consist of a simultaneous combination of clusters, grids, clouds, and stand-alone machines, possibly equiped with state-of-the-art multi- and many-core technologies, including GPUs and FPGAs. A new release will be made every year (or more often, if needed), each time with increasing functionality and increased flexibility and integrated intelligence (for example, for the automatic mapping of application tasks on the most appropriate available hardware under QoS constraints).

   - WP 5 YP 2014

7. KOALA releases

   KOALA releases, WP-6, one per year (2011-2014) (see also 1-YEAR) We will (at least) annually release a new version of KOALA. Features to be included are 1) a full MapReduce Runner, 2) an interface of KOALA to one or more cloud resource managers, and 3) application support for applications from the User studies.

- WP 6 YP 2014

8. Workflow Process Modeling and Management System
   Workflow Process Modeling and Management System WP8, prototypes in Y2, Y3, Y4.
   An important challenge is to find such solutions for workflow applications and management systems which enable easy collaboration of scientists and running of such applications on cloud type infrastructures. The preliminary studies show that one of the models to be developed is related to the concept of Workflow as a service (WfaaS). In current scenarios a workflow is submitted to computational resources to process a particular set of data and input parameters; after the processing is finished and results are stored the workflow components finish their execution and the workflow is terminated. When the next set of data and parameters is to be processed the workflow is started again which means that all the workflow components have to be re-submitted, re-scheduled on likely different Grid resources, and the data channel re-established. WfaaS keeps the workflow running on the resources even after a particular data set has been processed. The next data set can be submitted to the workflow that is already instantiated. To better handle the allocated time, the WfaaS approach has to consider the allocated time and other scheduling constraints before starting assigning new execution to ready workflow.
   - WP 8 YP 2014

9. WS-VLAM releases
   WS-VLAM releases, WP9, one per year, including components developed in COMMIT workflow provenance tracking, +WP9, Y1-Y4 WS-VLAM is a workflow management system developed in the VL-e project following a Service Oriented Architecture paradigm; WS-VLAM will be used as a starting point to develop further workflow issues in the context of COMMIT. Workflow provenance services will be integrated and delivered as part of the future WS-VLAM releases. Two types of data provenance are targeted: system level provenance and application level provenance. System level provenance records data at system level, this information will keep track of context in which the experiment has been performed, like when the workflow was executed, on which machines, which libraries and data have been used, been produced etc. Usually this information is available or generated at run time, sometimes saved in logs. Experience shows that users of the workflows can require this information for validation and reproducibility reasons. Thus making this information persistent and searchable is something that is valuable to both workflow developers and users. Application level provenance records application specific data that can help to have insight on the internal behavior of the workflow, and correlate them with system level provenance data to be able to better explain phenomena that occurs at runtime.
   - WP 9 YP 2014

10. Software release of JavaGAT/SAGA, one per year.

These deliverables are software releases of JavaGAT/SAGA, due in Q4, in years 2, 3, and 4. These software releases of JavaGAT/SAGA will continuously implement the results from WP10 throughout the course of the project, continuing from where the first software release from Q4 of year 1 has stopped. The year 2 release will most importantly implement the API package as described in the "OGF standardization document on resource allocation in cloud environments" from Q4, year 1. The releases in years 3 and 4 will reflect the research findings of WP10; mostly aiming at explicit support for dependencies among large sets of tasks as they are typical for workflows and many-task computations. Their precise functionality will depend on the research outcomes by the time of their due dates.

- WP 10 YP 2014

User studies

1. AMUSE, Collaboration: Leiden University (Prof. Simon Portegies Zwart).
   Collaboration: Leiden University (Prof. Simon Portegies Zwart), WP4, Y3 In collaboration with the Leiden University (LU), we will explore the potential of Ibis/Constellation to be applied in the AMUSE software system, which is used for advancing understanding in the field of computational astrophysics. The user study should provide insight in the capabilities and short-comings of Ibis/Constellation, and to allow researchers at LU to gain deeper insight in the potential of effectively using Jungle Computing Systems in their own research efforts.
   - WP 1 YP 2014

2. CineGrid media delivery systems
   CineGrid media delivery systems as usecase for validation of WP1 results, WP1 (Y2 initial study and Y4 final study) We use media delivery applications, such as the ones developed in the CineGrid project, to identify the shortcomings of the e-Infrastructure ontologies produced in WP1. In the second year we perform an initial user study that validates the initial version of the ontologies; in the fourth year we expect the refined ontologies to fully support the delivery applications.
   WP WP 1 YP 2014

3. Multimedia Computing -Collaboration: University of Amsterdam (Dr. Cees Snoek)
   Deliverable: User Study, WP5, Y2 Collaboration: University of Amsterdam (Dr. Cees Snoek). In collaboration with the University of Amsterdam, we will explore the potential of Ibis/Constellation to be applied in the software stack which is applied by the UvA in the yearly TRECVID benchmark evaluation. The user study is intended to provide insight in the capabilities and short-comings of Ibis/Constellation, and to allow researchers at the UvA to gain deeper insight in the potential of effectively using Jungle Computing Systems in their own research efforts.
   - WP 5 YP 2012

4. Remote sensing, Collaboration: University of Extremadura (Prof. Antonio Plaza) and European Space Agency (ESA)

   Description Deliverable: User Study, WP5, Y2, Y3 Collaboration: University of Extremadura (Prof. Antonio Plaza) and European Space Agency (ESA) In collaboration with the University of Extremadura, we will explore the potential of Ibis/Constellation to be applied in the field of remote sensing, using data obbained from the European Space Agency (ESA). The user study should provide insight in the capabilities and short-comings of Ibis/Constellation, and to allow researchers in remote sensing to gain deeper insight in the potential of effectively using Jungle Computing Systems in their own research efforts.

   - WP 5 YP 2013

5. Scheduling medical imaging applications (prof.dr. J Kok)

   We will collaborate with COMMIT project P24 (e-Biobanking with Imaging) on efficient resource management and scheduling for medical applications (e.g., proteomics) in multicluster grids and clouds, and to provide support in KOALA for these types of applications.

   - WP WP 7 YP 2012

6. Scheduling offline social network analysis (prof.dr.ir. M. van Steen)

   We will collaborate with COMMIT project P9 (Very Large Wireless Sensor Networks for Well-Being) on efficient resource management and scheduling for the offline analysis of the data obtained from monitoring sensor-based social networks in multicluster grids and clouds, and to provide support in KOALA for these types of applications.

   - WP 7 YP 2013

7. Studies InterCloud for urgent computing

   InterCloud for urgent computing (prof. Rob Meijer UvA/TNO), WP-9, Y3 InterCloud for urgent computing developed in Y1, will be extended to address the case of early warning systems, this work will build up on the results developed and +in the +context of the UrbanFlood project.

   - WP 9 YP 2013

8. XML-tracing Workflow provenance

   Description XML-tracing Workflow provenance, (prof. V. Sander, fh-aachen, de), WP9, (initial) WP9, Y1 and final Y3 The collaboration with the group of Prof. Sander will focus on validating the outcome of computational processes as the XML tracing model enforces task execution according to deposited restrictions and Service Level Agreements. Assuring the provenance of workflows will likely become even more important with respect to the incorporation of human tasks to standard workflows by emerging standards such as WS-

HumanTask. The works will address this trend by an actor-based workflow approach that actively supports provenance. A framework to track and store provenance information automatically that can be applied to various workflow management systems. In particular, the introduced provenance framework supports the documentation of workflows in a legally binding way.

- WP 9 YP 2013

Other results

1. A set of 'best practices' guidelines

   A set of 'best practices' guidelines to construct safe, secure host descriptions for virtual machine usage and creation. In this WP, tools will be devised to deploy and/or search systems based on detailed host descriptions. These tools are aimed to avoid usage of vulnerable packages on (dynamically created) systems in the cloud. However, to date, deciding what host configuration is secure or not is more an art than a science. Drawing on our experiences with describing and using host descriptions (and, e.g., with vulnerability databases) we aim to not only provide tools to use host descriptions, but we will add a best practices document to construct safe host descriptions for virtual machine usage and creation.

   - WP 3 YP 2014

2. Domain specific language for temporal optimization of ICT

   Domain specific language for temporal optimization of ICT, WP2, Q2 Y2 Other. This DSL allows the adaptive programming of the service of network and cloud as a whole, whilst a compiler will distribute software specifying collective (common) and specific behavior of network and cloud elements. The ultimate goal of the research activity is to develop a compiler that is able to compile itself (and to distribute itself) over the network. Special attention is paid that the technology is able to visualize itself for the benefit of demonstrations and education.

   - WP 2 YP 2012

3. Algorithms for automatic (adaptive) multi-scale optimization of network and cloud services

   Algorithms for automatic (adaptive) multi-scale optimization of network and cloud services, WP2, Q3, Y2 Other. Illustrative algorithms, representative for a certain class of problems, with proof of concept implementations on the testbed.

   - WP 2 YP 2012

4. E-Science workflows

   E-Science workflows that program the behavior of physical resources to facilitate an optimal performance of instruments in the e-Science environment itself, WP2, Q1, Y3 A software framework for the monitoring and control of non-ict resources of cyber physical systems. In sensor networks this is typically a common informations space, a construct that reads out and distributes

(pub sub) sensor data.

- WP 2 YP 2013

5. Application to IJkdijk

   Description Application to IJkdijk early warning system for dike failures, WP2, Q3, Y3 The IJkdijk (now the FP7 UrbanFlood project (UVA, TNO, ...) early warning system for dike failures encorporates (in fact on a general purpose framework for very large scale (environmental, systems, structure) monitoring) a (software) factory for early warning systems (EWSs). The factory instantiates EWSs and adapts the workflow and compute facilities in instantiated ones to different loads. The UrbanFlood technology is cloud based and uses internet based clouds. By advancing the UrbanFlood technology with results of P20, an advanced general purpose monitor and control framework for cyber physical systems (e-science, e-infrastructures, e-systems) is created, an internet of things construct. Special attention is paid that the technology is able to visualize itself for the benefit of demonstrations and education.

   - WP 2 YP 2013

6. Architectures and implementations for multiscale ICT optimization facilities

   Description Architectures and implementations for multi-scale ICT optimization facilities, WP2, Q1, Y4 Vizalisation Environment for the cyber physical systems technologies, testbed and algorithms in

   - WP 2 YP 2014