Project Description
AAA Continuation Activity 2013

Date: 31 August 2012

Project Lead	Project	Institution
Sergio Maffioletti	**Swiss Academic Compute Cloud**	**UZH**

Note: Project will be approved and managed within the project "Information scientifique: accès, traitement et sauvegarde"

Authors: *Placi Flury, Peter Kunszt, Pierre Kuonen, Sergio Maffioletti, Marko Niinimaeki, Heinz Stockinger*

Project start and end: *1 May 2013 – 31 December 2013*

Title of project: *Swiss Academic Compute Cloud*

Project leader: *Sergio Maffioletti, sergio.maffioletti@gc3.uzh.ch, Tel: +41 44 635 4222*

Deputy project leader: *Placi Flury, placi.flury@switch.ch, Tel: +41 44 268 1584*

Project responsible / sponsor and institution: UZH/GC3

Other project participants of the university: *responsible person underlined*

Institution	*Participating people*
SyBIT/SystemsX.ch/ETHZ	*Peter Kunszt, Dean Flanders*
EPFL *	*Vittoria Rezzonico*
Vital-IT/SIB	*Heinz Stockinger*
WSL	*Thomas Wuest*
FMI *	Stefan Grzybek
HES-SO//Ge	*Nabil Abdennadher, Marko Niinimaeki*
HES-SO//Fr	*Pierre Kuonen, Jean-François Roche*
LHEP	*Sigve Haug, Gianfranco Sciacca*
University of Zurich - Grid Computing Competence Center GC3	*Sergio Maffioletti, Riccardo Muri, Tyanko Aleksiev, Antonio Messina*
SWITCH	*Placi Flury, Valéry Tschopp*
ZHAW ICCLAB*	*Thomas Bohnert, Andrew Edmonds,Christof Marti*

** Observing partner*

Vision: We establish a Swiss-wide "computational science" platform composed of resources and services of various type as well as high quality know-how for user and application support. Every Swiss researcher can profit from such a platform to address his/her computational needs.

Mission: support computational research and improve computational science

Objective of the proposal: Consolidate infrastructure and user support effort to enable a Swiss-wide platform for large-scale data analysis.

Project description, goals and benefit:
The project sustains the ongoing grid/cloud activities of the *AAA/SWITCH program* that are essential for the research communities. It bridges the activities that are expected to become relevant to the upcoming *Informationsversorgungsprojekt [crus]*.

The project merges the contributions of the AAA/SWITCH projects, namely the Swiss Multi Science Computing Grid (SMSCG), the Academic Compute Cloud Provisioning and Usage and other related AAA/SWITCH projects (like VM-MAD, GridCertLib and RS-NAS), which conform with our mission. We focus on the sustainability of the computational service for our users and are, therefore, not pursuing any of the activities of these projects that are not directly contributing to this goal, e.g. like for example exploring next generation infrastructure technology.

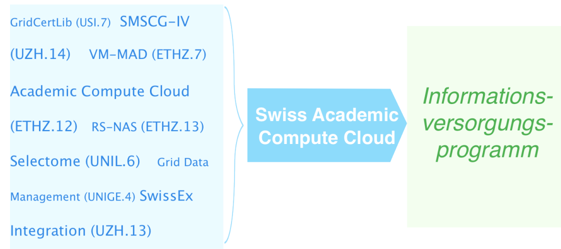

Figure 1 AAA/SWITCH projects upon which Swiss Academic Compute Cloud is building up.

The primary goal of this project is to sustain and keep together the research communities that are currently served by the 'grid track' projects of the AAA/SWITCH program during the Q2-Q4 period of 2013. This is to preserve and update the experience and knowledge of the AAA/SWITCH community and keeping it active and coordinated for the *Informations-versorgungsprojekt*.

The secondary goal is to serve a greater number of research communities. A large community increases the mutual benefits and can profit from economies of scale; we get more experience and more feedback, which helps us to improve the quality of the services provided. Besides, the size of the community determines our long-term sustainability plans. A last goal is to provide a sound survey for decision making on infrastructure models, like in house vs. outsourcing, and on technologies, like running a cluster or going to the cloud (see milestone WP1-SA, page 8).

Users first is the maxim and essence of the project. It is reflected in our mission statement and it is the yardstick for all proposed activities. This proposal aims to establish and consolidate the effort in enabling user communities of above-mentioned projects. The expected outcome is a network of competent support teams (see Collaborative Distributed Support model on page 9) which help/assist scientists in porting/enabling their large-scale data analysis pipelines on a flexible cloud-like platform that can easily be shaped to accommodate their needs. To achieve this objective, this proposal aims in satisfying the following requirements we believe are fundamentals for quality user and application support:

- *Simplicity and flexibility*: simplify the use of the services and of the infrastructure while satisfying the requirements and the complexity of the research problems that are addressed. A user can access an infrastructure consisting of approximately 3k-7k Cores and up to 20TB of project storage space. (For a snapshot of the currently available computational resources visit https://monitor.smscg.ch/public)
- *Responsive support and competent consultancy*: collaborate with users in understanding the infrastructural requirements of their scientific use cases, in designing efficient and scalable solutions, in implementing the solutions on the infrastructure and in providing long-term support (e.g. continuous integration model).
- *Speed and scalability*: be able to deal with large numbers of computational jobs and large amounts of data, while producing results in due time.
- *Supportive interfaces*: provide interfaces, especially to data, which do not impose (major) changes in the data models and data structures in use, e.g. files, databases, etc.
- *Cost effectiveness*: users are cost aware and will carry out a cost-benefit analysis to weigh the quality of the offered service with the efforts and costs of using the service. We, therefore, continuously *innovate* by exploring new and exploiting current technologies[1] and by adapting and optimizing processes and procedures.

Innovation is the concept we use to satisfy the above requirements. We conduct *service innovation* by establishing long-term relationships with the research communities. Their needs define the application portfolio we support and supply. Service innovation doesn't end here. We continuously strive to enhance and simplify the interfaces to our services and to adequately document how to most effectively and efficiently use them; this is a process that has been already initiated in some of the ongoing AAA/SWITCH projects like SMSCG, Academic Cloud and RS-NAS.

We conduct *process innovation,* by periodically reviewing the processes we have put in place to support and consult the research communities we serve, but also those used to manage our project. The immediate implementation of lessons learned and the resulting improvement proposals are prime factors for quality assurance and for cost reduction.

We innovate on the *technological* level, by exploiting the knowledge we acquired in the various AAA/SWITCH projects. We compile a technology portfolio, which displays the various options for the people, who are responsible for the infrastructure[2], and which shall aid them in their decisions. Criteria are, costs, performance, and suitability for target applications.

The two main beneficiary stakeholders of the bridge project are, the research communities and the infrastructure providers (typically the IT services).

- Serviced research communities benefit from continuity of the offered services. They benefit together with newly acquired research communities from an experienced IT consulting group porting their applications and use cases and advising them in their computational research problems.
- Infrastructure providers benefit:
    - If operating local infrastructure: from the close cooperation with the researchers and with the other infrastructure providers of the project. The first ensures a fit between the infrastructure and the use cases and applications that are accommodated. The second strengthens the individual competences through cooperation.
    - If having to ensure access to infrastructure: to be able to base their decisions on a sound foundation. The ties to the local research communities permits to know the requirements and the ties to the bridge project the pros and cons of each available model.

[1] Exploring new technologies is not within the scope of the bridge project. Moreover, we expect that the leading technologies were already assessed within the AAA/SWITCH program.

[2] A person responsible for infrastructure can either operate infrastructure in-house (be a resource provider), or outsource the service (e.g. rent a cluster).

A bridge project is per definition linking the activities of the old program that ought to be sustained until the new program starts. Still, it does not relieve us from taking action. We must create added value to a (larger) user community which shall mandate us to be sustained beyond 2013; in particular for the critical non-operational activities, which shall most desirable be sustained through funding programs like the *Informationsversorgungsprogramm*.

Expected number of users: User Communities and User Groups:

We satisfy the needs of the Swiss research community in progressive steps. We started in the SMSCG project with few yet fault tolerant and representative users. Their use cases helped us validating the procedures and shaping the overall service.

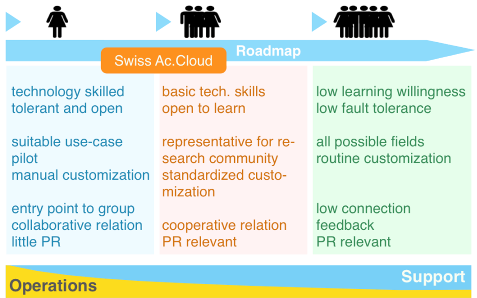

Figure 2 User enabling roadmap

The last SMSCG phases targeted users who already solve research problems with help of computational processing; their IT skills speed up the task of servicing them. Virtualization techniques (from the VM-MAD project and AppPot) and a powerful toolbox for jobs management (GC3Pie) introduced our efforts of standardizing the procedure of enabling new users and their research communities.

The Swiss Academic Cloud project resumes there. It inherits an application portfolio from the SMSCG of 20 applications (see http://www.smscg.ch/WP/applications/ for an updated list) in the fields of life sciences, earth sciences, economics, computer sciences, engineering, cryptography and physics. Research groups in the following projects, experiments or institutions, use these applications:

Institution	Applications	Characteristics (on SMSCG infrastructure)	Research field
IMSB/ETHZ	Rosetta, Proteomics pipeline, imaging pipeline	2-3 users, 152k jobs, 300k walltime hours during last 12 months	life sciences
IBF/UZH	gpremium	1 user, 916k jobs, 238k walltime hours, (1.5.2011 -- 1.7.2012)	economics
EBI/UZH	mhc-coev	About to be certified	evolutionary biology
GEO/UZH	GEOTop	4-8 users, 37k jobs with 11k walltime hours since June 2012	earth sciences

LACAL/EPFL	home-grown	1-4 users, 16k jobs, 32k walltime hours during last 6 months	cryptography
UniL/Vital-IT	Selectome	3 users (development). 152k jobs, 7k walltime during last 12 month	life-sciences
UniGE Geneva	MetaPIGA	About to be certified	life-sciences, engineering
LHEP/UniBe	ATLAS	Several (international) users (200?), ca 27k jobs/month with 239k walltime hours	physics
WSL/SLF	CATS, Alpine3D, SwissEx	CATS: 1 user, 1month, 2300 jobs, 38k walltime hours SwissEx: about to be integrated/certified	earth sciences

The users and their communities should not notice the transition from SMSCG (and the other AAA/SWITCH projects) to the Swiss Academic Compute Cloud. The changes occur under the hood and will be transparent.

Technical Description:

WP1 Operation and Infrastructure

The bridge project inherits two computational and two data management models from the AAA/SWITCH program. In the VM-MAD project we demonstrated the feasibility of an integration of the computational models. Cloud resources were allocated dynamically (termed cloud-bursting) to extend the batch-system pooling model, which predominates the SMSCG infrastructure. This first success needs, however, to be interpreted with care. It is a one-way demonstrator, from grid to cloud. Its scope did consciously not include issues that are essential for a production level integration, such as complete accounting with charging and billing, usage policies, etc.

An integration of the data management models has, on the other hand, not been addressed. There, the user requirements must first be better understood. From our current understanding and experience (e.g. from the RS-NAS project) data management solutions must be designed and customized for each community individually. We can neither target for a generic one-size-fits-all solution, nor are we aiming at achieving full interoperation of the data management models. The complexity of such an endeavor dwarfs any imaginable benefit.

In the bridge project we will:

1. Operate the inherited infrastructure with its computational and storage services
2. Advance the development of a standardized user interface
3. Aim for a uniform authentication and authorization mechanism for shared resources and
4. *4.* Encourage the standardization of the interfaces of the emerging Swiss private clouds.

Access Interfaces

Figure 3 depicts the computational model of the bridge project. The computational resources are classified into local(host), cluster/grid, private cloud, and public (commercial) cloud. Users access the computational service with a programmable user interface (e.g. GC3Pie). The interface is used for large-scale job management and manages the authentication and authorization (AA). It gives the user the illusion of uniform technology/model-agnostic access to the computational and data management service (the latter is not within the scope of the bridge project), while the support team of the Swiss Academic Cloud carries out the mapping to the underlying model. It permits users to set up a development and processing pipeline for their applications. Command line options and configuration parameters permit targeted computational resource selection, e.g. whether to first run local and later on the local cluster, on the grid or on the cloud.

The task of porting the user application is a user support task and is done within WP2. The task of WP1 is to extend the interface to accommodate the AA mechanisms that are imposed by the private and public cloud resources.

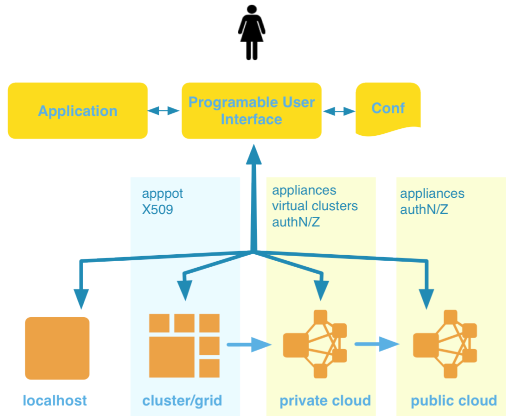

Figure 3 Simplified computational model of the infrastructure of the Swiss Compute Cloud project

Uniform authentication and authorization

Each of the four resource types is currently using its own AA mechanism. A consolidation of the AA is a necessity for effective user management.
The AA of the grid model is based on X509 certificates, for both computational and data management. With the Short Lived Credential Service (SCLS) we provided an intermediate mechanism to support federated AA. The SLCS converts user credentials from AAI into X509 user certificates. The solution ought, in the long term, to be replaced with the more flexible and extensible Security Token Service (STS) [see STS]. By April 2013 the EGI is expected to provide an STS implementation.

The cloud model lacks of a *federated* AA, a vital component for the integration of heterogeneous and distributed computational and data storage resources. Each provider allows individual AA (mostly based on SSH RSA keys); what is required for a Swiss Academic Cloud infrastructure is a common base AA infrastructure that levers, from one side, existing authorization and authentication infrastructures (e.g. AAI) and, allows, from the other side, grouping and shared access to resources (e.g. STS).

Standard interfaces of private clouds

At the time of this writing there are no well-accepted standards for the cloud software stacks. The Swiss Academic Cloud project, therefore limits itself to the observation of the trends and developments in this area with special attention how they relate to and/or depend on the requirements of the user communities.

Shortcomings

The bridge proposal will deliver a consolidated high throughput-computing platform, which satisfies most of the needs of the user communities. Not all aspects of data management -- like a sound data lifecycle management -- and its ramification on AA will be resolved given the scope and resources of the project. Our goal is to provide a solid foundation upon which a finalizing project can resume.

Tasks

- Consolidation of the internal procedures for effective and efficient operation of the infrastructure. We aim for routine. In this task we merge and streamline the existing documentation of the formerly individual projects and set up a common knowledge base. Both shall be maintained and extended in a collaborative way with possible user involvement. **[Estimated effort: 1PM (1 SWITCH), Milestones WP1-OP]**
- Review and adapt central services that are vital for the provision of a reliable quality computational and data storage service. User management, monitoring and accounting must be extended to incorporate the different technologies and models we use. We will piggyback on NGI-CH efforts where possible. **[Estimated effort: 1PM (SWITCH), Milestone WP1-CS]**
- Review usage policies and adapt where necessary. **[Estimated effort: 0.2 PM (SWITCH), Milestone WP1-UP**
- Uniform authorization for computational and data storage services. The Academic Compute Cloud Provisioning and Usage (ETHZ.12) and the RS-NAS (ETHZ.13) projects will have identified the requirements for a production-quality AA. In the bridge project we will define an AA architecture that meets these requirements. We will implement a pilot that provides the most basic functionality and that helps validating our concepts. Notice, the sound implementation and the deployment of the proposed AA architecture exceed the resources of the project. We focus therefore on establishing a solid foundation for subsequent projects and collaborations, e.g. external open source community projects. **[Available PMs: 12PM (2PMs GC3, 2 PMs ETHZ, 8PMs SWITCH, Milestones WP1-AA1, WP1-AA2]**
- Cost benefit assessment and suitability guidelines. We carry out cost-benefit estimations, which shall provide a foundation for investment and purchasing decisions; in more detail, what infrastructure models to select given application requirements and deployment constraints. We will therefore conduct a suitability analysis for application profiles and infrastructure models (e.g. whether applications with XY requirements run efficiently on a private cloud or whether a dedicated infrastructure is economically better suited.) **[Estimated effort: 1PM (SWITCH), Milestones WP1-SA]**
- Virtual Appliances management: create Swiss Academic Cloud appliance templates that shall be used for application-specific customization. The appliances will include components to ease management and visibility on the Swiss Academic Cloud infrastructure. For the seamless integration of existing data analysis pipelines we will provide appliances for on demand instantiation of virtual clusters. The later will be stored, cataloged and made accessible by various cloud providers.
  All appliances will be certified by the support teams, which are porting the user applications. **[Estimated effort: 3PM (1PM SWITCH, 1PM ETHZ, 1 PM GC3]**
- Bridge with the Venus-C project (Virtual Multidisciplinary EnviroNments USing Cloud Infrastructures, http://www.venus-c.eu). This 7th framework project has 14 partners and 15 pilot applications, among them MetaPIGA that was ported to the platform by HES-SO//Ge. Venus-C uses the Microsoft Azure platform that, to our knowledge, has not yet been utilized in Swiss high-performance computing. Within Venus-C, the Azure platform is injected with "generic workers" that download and execute pre-defined tasks that correspond to SMSCG runtime environments.[3] Bridging Venus-C means connecting Azure to SMSCG like a "public cloud" in Figure 3, thereby making its infrastructure (estimated 1000 cores) available for Swiss researchers. **[Estimated effort: 2PM (HES-SO//Ge)]**.

[3] Y. Simmhan: Bridging the Gap between Desktop and the Cloud for eScience Applications, IEEE 3rd International Conference on Cloud Computing, Los Angeles, CA, July 2010.

Information scientifique: accès, traitement et sauvegarde

- Application integration testbed. An integral component of a production cloud infrastructure is the availability of a testbed that can reproduce on small scale the cloud production execution environment. Such a cloud testbed is an important component for the operational infrastructure as updates and external component integration can be tested and evaluated. Similarly, it simplifies the application integration task by allowing pre-production tests. The testbed will be provided by HES-SO//Fr, HES-SO//Ge and Vital-IT and will be part of the accessible infrastructure for all non-production activities. The model aimed here is similar to what provided by the FutureGrid project (https://portal.futuregrid.org/) where an easily accessible test infrastructure can be used for development and benchmarking purposes as well as application integration and for testing technologies in a pre-production environment **[Estimated effort: 6PM (1PM HES-SO//Ge, 3 PM HES-SO//Fr, 1PM GC3, 1PM Vital-IT, Milestones WP1-AI]**

Remarks: The operational overhead for a site contributing with production infrastructure is estimated to be 1.2 PMs per site for the duration of the project. Of these the site has to cover 0.8 PMs by itself. Coordination of operational tasks will be lead by SWITCH.

Milestones:

WP1-OP Knowledge base and consolidated documentation by June 2013
WP1-UP Review policies by June 2013
WP1-AA1 AA architecture defined by July 2013
WP1-AI: Cloud testbed available by August 2013
WP1-CS Review and adapt new central services by September 2013
WP1-SA Technology portfolio analysis document by October 2013
WP1-AA2 AA test/pilot implementation by November 2013

WP2 Application and user support

Starting from the experiences and the activities already initiated in the previous projects, this proposal aims to consolidate application and user support over a distributed computing infrastructure for large scale data analysis.

The minimal granularity level of "Application and User support" activities will be centered around software appliances (application and/or use case specific) that will be created on demand either by the support units of by the experienced users directly. The different use cases will be then supported by integrating and/or developing tools and component to orchestrate such appliances on the available computing and data infrastructure.

There are currently two trends in supporting scientific computing on cloud infrastructure:

1. Extend on-premise infrastructure with cloud resources while maintain the batch-based usage model unmodified. This has the advantage that it brings minimal if no disruption of the services for the end-users that will be able to continue running existing data analysis pipelines indirectly levering cloud resources as an extension or as an augmentation of the local IT infrastructure. This is a short/medium term approach that could be used for supporting established research group to progressively transition from a classical local batch-oriented model to a more flexible and transparent cloud one. The main disadvantage of this approach is that it is not possible to fully lever cloud features like elasticity and load balancing and, at the same time, it still maintain a certain level of perception of the underlying computing infrastructure for those dealing with application integration.
2. Redesign analysis pipelines and/or scientific applications to directly lever cloud features like interactiveness of the allocated nodes. This approach has the advantage of fully exploiting the transparency and the elasticity of cloud features from, potentially, within the algorithmic part of the scientific application. The gain is higher level of scalability, robust resilience and a complete transparency of the underlying infrastructure. The disadvantage is that, as of today, there are no standards in tools to control and coordinate large number of parallel virtual appliances, thus making difficult to define a

uniform support layer that can be transparently used in various use cases and on various cloud instances. Nevertheless, this seems to be clearly the direction scientific computing will follow in the coming future.

As the proposed project aims to sustain the ongoing support activities inherited from the above-mentioned AAA/SWITCH projects both of the presented approaches will be modeled as part of the "Application and User support" activities:

- *Integration of cloud into analysis pipelines*: for the already established use cases, the integration/migration towards cloud-based computing and data infrastructure will be done almost transparently by integrating cloud resources at the infrastructural level and by exposing them to the data analysis pipelines as part of the available resource portfolio. The result of this will be application-specific appliances that will be activated through cluster-like interfaces integrated either within existing computing resources (cloud-bursting) or as individual virtual computational resources (virtual clusters).
- *Integration of analysis pipelines into cloud*: For the new use cases, decided on a case-by-case, we will evaluate the best integration pattern that can, from one side, fully exploit cloud models and, from the other, better support the use case. The result will be a set of appliances and a control logic that will be customized and tuned for the underlying cloud infrastructure and for the specific use case.

Collaborative Distributed Support (CDS)

The Swiss Academic Cloud inherits the user support model of SMSCG (see figure 4), which we term Collaborative Distributed Support (CDS). It has international equivalents in the EGI Champions https://wiki.egi.eu/wiki/VT_EGI_Champions, the UK NGS Campus Champions http://www.ngs.ac.uk/campus-champions, the Strategic Applications Collaborations/Strategic Community Collaborations (SAC/SCC) at SDSC http://www.sdsc.edu/us/sac/ and many more.

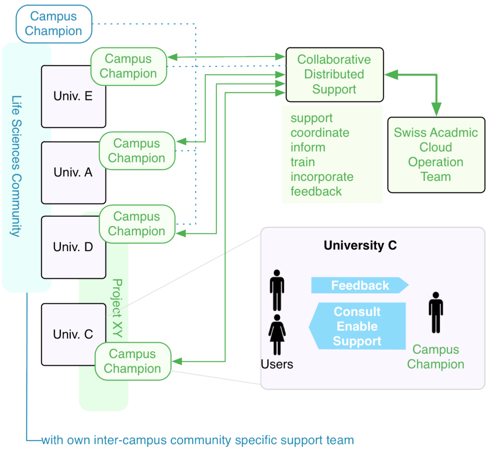

Figure 4 Campus Champion support model.

The campus champion is our man/woman at the campus (for an actual champions list visit the prototype user website: http://jupiter.switch.ch/internals.html). The champion is the link between the local research communities and the Swiss Academic Cloud operations and infrastructure team. It is the local contact point that establishes and cultivates the relation to the users and is

meant to harmonize with the already established local support units. Its responsibilities are to provide support and consulting, to enable the local communities on the infrastructure, to disseminate and possibly train, and to inform the communities about occurrences. Just as important, it collects the feedback of the community, thus providing most valuable input for the improvement of our services. The bridge project addresses the following tasks:

- Expand campus champion outreach to include all partner institutions and more.
- Professionalize engagement with (local) user communities. (Guidelines, procedures, and training of champions)
- Enhance coordination between campus champion and operations. Clear roles model and establish routinized processes (e.g. benchmarking user enabling within one working day as upper limit.).
- Strengthen role of competence center for application enabling.
- User feedback processing. Deduce service extensions and tasks upon clear user request and user need (new capabilities for users that allow for innovative resource usage like workflow technology; include compute, data, and visualization resources.)
- Support research groups and/or communities in their international collaborations (e.g. EGI, etc.)
- Explicit work on public relations: become more visible, create quality reputation, and organize community meetings through SwiNG.

The above-mentioned activity is to be considered as an advanced support that is general enough to be an added value for the whole computational research community. Key aspects of this model are the competence of the people involved and the coordination that ensures consistency and coherence of the support activities. For this a core competence team needs to be formed to guarantee the basic knowledge base and to have a placeholder where the "Application and User support" members can share and discuss their approaches and where such activities can be observed in respect of the project's perspective. Initially the core competence team will be formed by those groups that have been most involved in support activities in the previous projects: GC3 from UZH, Vital-IT, ETHZ and HES-SO//Ge. Such competence core team will be responsible to:

- Organize internal seminar events where every member of the "Application and User support" can participate with his/her own contribution.
- Provide the necessary manpower to complement, when needed, the local support competences.
- Make sure emerging best practices are properly integrated into the project mainstream activities.

Contribution from Vital-IT

Over the last years the user community from SIB/Vital-IT has been extended from mainly based in the Lemanic region (Lausanne, Geneva) to reach out to several cities currently mainly in Western Switzerland including Fribourg and Bern. Additionally, since 2011 Vital-IT deploys HPC clusters on three sites (UNIL, EPFL and UNIGE) as well as life science software and databases on the UBELIX cluster in Berne. This effort has mainly been dictated by the necessity to cope with specialized infrastructure and expertise at each research institution, something that a single institution could not establish alone. Given the distributed user community in different locations, the proposed "Campus Champions" model fits well here, applied at two levels: (1) technical level in terms of computing infrastructure and software to be used (2) scientific level in terms of biological, bioinformatics and computational biology domain knowledge (i.e. how to transform a high impact biological question in the omics area into a scientific workflow ("pipeline") that can then be appropriately used on a cluster and/or cloud infrastructure. This requires both training as well as consulting with local scientific user and expert groups. The activity of the Vital-IT group will focus primarily on consulting and training of user groups in Geneva, Fribourg and Berne. **[Estimated effort: 3PMs (Vital-IT)]**

Contribution from GC3/UZH

GC3 has been one of the most active units in application and user support within SMSCG. It has provided support at the national level by enabling use cases and consulting research groups from different institutions. In this phase, GC3 will continue its consulting activity with a national scope complementing, when possible, the local institution CDS unit. It will also take the responsibility of coordinating the dissemination, in form of technical seminars, to keep the distributed CDS units cohere and up-to-date on the ongoing activities. **[Estimated effort: 4PMs (GC3)]**

Contribution from ETHZ

The ETHZ is in the process of establishing a local Science Support unit, integrating the infrastructure, software and application support needs for all of the ETHZ compute intensive sciences. This unit has not been established yet at the time of writing this proposal, but it is expected that it will be operational by the beginning of this project. The ETHZ Brutus Cluster team, the Center for Information Science and Databases at the ETHZ/Basel and the SystemsX.ch-funded persons at the Institute of Molecular Life Sciences are the starting point for this new effort, which will be consolidated and expanded during the coming years. This new unit will provide the necessary services to not only the ETHZ but all collaborating scientists as well. **[Estimated effort: 2PMs (ETHZ)]**

Contribution from LHEP, HES-SO//Ge, ETH

The other members of the CDS will initiate within their local institution a campus champion support unit that will integrate with the existing user support activities. Such a unit will be responsible to actively identify, contact and support local candidate research groups that will be willing to invest their own effort in exploring whether and how their scientific use cases could be supported on the infrastructure. The main benefit for these institutions will be to participate into an organized national support activity that explores emerging approaches in enabling large scale scientific data analysis. **[Estimated effort: 4PMs (1 HES-SO//Ge, 2 WSL, 1 LHEP]**

[Overall estimated effort: 13PMs (4 GC3, 2 ETH, 3 Vital-IT, 1 HES-SO//Ge, 2 WSL, 1 LHEP]

Milestones

WP2-CDS1: Organize internal seminars (1 every 2 months = 4 seminars)
WP2-CDS2: Organize site visits for reach-out, information and dissemination

Specific use cases

- *User level standardized interfaces*: GC3Pie (http://code.google.com/p/gc3pie/) is a programming library conceived to streamline and simplify the development of complex control logics for large scale data analysis. GC3Pie is one of the main tools used within SMSCG by the support units to develop end-to-end solutions for the various supported use cases. We extend the GC3Pie library to provide uniform technology-agnostic access to the infrastructure. We start by interfacing the computational service, leaving aside the far more complex data storage services. **[Estimated effort: 3PM, Milestones WP2-INT (GC3)]**
- *OSPER integration*: The Open Support Platform for Environmental Research (OSPER) will take forward key data management components developed under the Swiss Experiment Platform project, and integrated in the SMSCG infrastructure through the AAA/SWITCH UZH.17 project, to achieve improved usability and a wider scope. OSPER aims to provide the technology for storage, management and exploitation of environmental research data, focusing on data interoperability and documentation. The goal of the UZH.17 project is/was to implement a showcase as a proof of the concept

that a cloud-like infrastructure could be beneficial for OSPER as it could spare from the basic infrastructure operation management from and it could seamlessly scale towards large commercial providers. The showcase makes a subset of the public data acquired from the distributed GSN (Global Sensor Networks) available and allows simple statistical and analytical pipelines to be run through simplified interfaces. The proposed activity will focus on full integration of the OSPER data lifecycle: all of the data collected by the distributed GSN instances will be stored and archived by a scalable and centralized database as one of the Academic cloud appliances. The appliance will be accessible and visible by the integrated computing infrastructure thus allowing running more complex statistical and analysis pipelines. While the goal of this activity is to provide a scalable solution for a centralized archive of the GSN data, additional statistical and analytical pipelines will be integrated in accordance with the outcome of the UZH.17 project. This effort will be complemented by a direct involvement of the OSPER technical and user team as stated in the attached support letter. **[Estimated effort: 4PMs, Milestone WP2-SwX (4 GC3)]**

- *RS-NAS integration*: The Remote Scalable Network Attached Storage project intends to provide technology for cloud-based academic storage. This technology enables access to central storage that can be located anywhere in the country through a local interface, as if it was local. This enables the provisioning of very large petabyte-scale strategic storage for academia while maintaining the flexibility and locality of the data, without large performance penalties. The important data would be cached locally. The scientific initiation use case comes from SystemsX.ch, where several research groups generate large amounts of data across the country, which are again needed for integration and analysis at other places - between Basel and Zurich mostly. Through this technology both sites can work as if data was local and new data is accessible without delay. SWITCH networking technology is a crucial part of this effort. **[Estimated effort: 3PMs, Milestone WP2-SyX (ETHZ)]**
- *GEOtop* is an environmental simulator supported in the SMSCG project for the group of Dr. Stephan Gruber from the Department of Geography at the University of Zurich (GEO/UZH). This simulator represents the complex interactions between climate, snow cover and frozen soil and is a central research tool in understanding climate impacts on permafrost and related natural hazards. The application has been enabled on the SMSCG infrastructure using the GC3Pie framework and relying on the ARC-based available computational resources. In the proposed activity, GEOtop will be enabled as a virtual appliance that could be executed for large-scale data analysis on an EC2 and S3 compliant cloud infrastructure. This is important because simulators such as GEOtop have a large parameter space for which sensitivities and uncertainties need to be explored if reliable data products or conclusions are to be generated. While the control logic, based on GC3Pie, will remain the same, this approach will make possible to scale the computational requirements depending on which subset of the datasets will be needed to analyze (this impact the number of computational resources that needs to be provided) and on what level of detail (this impact the profile of the computational resources provided). An efficient control logic that will allow automatic scaling and elasticity of the cloud-like resources available will permit to easily and efficiently tune the computational infrastructure to the verifying needs of the datasets. Encapsulating GEOtop into a dedicated appliance will also allow embedding the post-processing steps, a workflow of R scripts, thus further simplifying the data management part. This work will be performed by GC3/UZH in orchestration with Dr. Gruber who will commit 1PM for the scientific evaluation as stated in the attached support letter. **[Estimated effort: 4PMs, Milestone: WP2-GEO (GC3)]**
- High-memory machines with a minimum of 64 GB become the standard in the genomics and large-scale data mining projects but are usually not available in Grid environments. While the grid middleware used in the WLCG, EGI and SMSCG projects allows the uniform connectivity layer, they lack in flexibility in provision, classification

and retrieval of specialized resource instances. The batch-based computing model and the intrinsic overhead of an infrastructure aimed to aggregate heterogeneous resources (especially the information system and the complex brokering) make the grid infrastructure not adequate to suite most of the life science execution models. For those a dedicated system for fast processing is required. Several life science applications are either still in their infancy phase or redesigning their core functionalities thus creating a scenario were short cycle application update is the norm. Another aspect is that several bioinformatics applications make use of novel programming models such as MapReduce (one of the goals of SMSCG-IV as well as partner projects such as academic cloud). Another trend is that there exist more and more predefined "pipelines" for genome analysis that are packaged for running on clouds. This is often due to the reason that the pipelines have several software dependencies which make it difficult for end users to use and/or cluster administrators to efficiently install and configure. It is therefore of great interest for both scientists and resource providers to have easy access to ready-deployable instances which allow for a fast usage of existing and newly developed pipelines. This approach renders as well the maintenance of versioning of the pipeline easier as it is self-contained within a "packaged cloud image". Relevant experience and knowledge from related projects (e.g. academic cloud) is important in this activity to have an infrastructure deployed and usable for life and medical scientists. Testing, and preliminary deployment of cloud provisioning for one selected genomics application (for example: QIIMEE http://www.qiime.org/). **[Estimated effort: 3PMs, Milestone WP2-BIO (Vital-IT)]**

- Western Switzerland has recently launched a CleanCity collaboration http://iig.hevs.ch/switzerland/clean-city.html that uses data collection and software tools to determine the effects of pollution in cities. The project will build methods to evaluate counter-measures to pollution. One of the state-of-the-art methods is to simulate particles of pollution in air flows under different conditions (wind speed, traffic density..). HES-SO//Ge uses OpenFOAM, the open source computational fluid mechanics software for this. In an earlier Multi-Site ARC project, a gridified OpenFOAM was deployed. However, large OpenFOAM simulations require dozens (if not hundreds) of cores and CPU-hours. In Swiss Academic Cloud project, we shall provide an MPI-aware, check-pointing OpenFOAM environment for the CleanCity community. In a related project, HES-SO (GE) will port the SWAT http://swat.tamu.edu/ tool in collaboration with the University of Geneva. A letter of support from UNIGE's EnviroGRIDS group is pending. **[Estimated effort: 4PMs, Milestone WP2-CTY (HES-SO//Ge]**

WP3 Project Management

The project management is responsible for the regular coordination of WP1 and WP2 activities. Special and dedicated involvement will be required in:

- **Reaching out for new research communities**. Extending the Collaborative Distributed Support to new institutions is a task of project management; WP3 establishes the first contact and makes the negotiation with the local support groups on the level of involvement, the activities to be followed. The follow-up of the participation in the CDS activities is coordinated within WP2.
- **Reputation management.** Ensure quality and promote dissemination and awareness.
- **Sustainability beyond bridge project.** Being the link between the community (together with SwiNG) and the *Informationsversorgungsprogramm*. Inform members on developments. Provide sustainability strategy.

Milestones

WP3-KO: Kick-off
WP3-SP: Sustainability plan.
WP3-CR: Closing report.

Project plan / deliverables / milestones:

Milestone	Title	Date Due	Comment
WP3-KO	Kick-off meeting	1 May 2013	
WP2-CDS1	Internal support seminars	June 2013	First seminar by end of June
WP1-OP	Knowledge base and consolidated documentation	June 2013	
WP1-UP	Reviewed usage and security policies	June 2013	
WP1-AA1	AA architecture defined	July 2013	
WP2-CDS2	Site visits (reach-out)	July 2013	First visit by mid July
WP2-INT	User interface for new computational models available	July 2013	
WP1-AI	Cloud testbed available	August 2013	
WP2-SvX	OSPER integrated	August 2013	With new UI of WP2-INT
WP2-SyX	RS-NAS Integration	August 2013	
WP2-GEO	GEOtop integration	August 2013	
WP2-BIO	QIIMEE deployment	September 2013	
WP2-CTY	SWAT tool ported	October 2013	
WP3-SP	Sustainability plan	August 2013	Foundation to submit project proposal.
WP1-CS	Revision and adaptation of central services finished	September 2013	
WP1-SA	Technology portfolio analysis document	October 2013	
WP1-AA2	AA test/pilot implemented	November 2013	
WP3-CR	Closing report	31 December 2013	

Manpower/other costs: *<Estimate the manpower in man months and the estimated costs of the project (manpower, hardware, software).>*

The manpower and hardware costs are estimated as follows:

Table 1 Hardware contribution by institution

Institution	Institutional Contributions	Requested Federal Funds	Comment

	[CHF]	[CHF]	
ETHZ			
SWITCH	Not applicable	Not applicable	
GC3			
WSL			
Vital-IT			
HES-SO//Ge			
LHEP			
HES-SO//Fr			
EPFL			
TOTAL			

Table 2 Manpower contributions in person month by institution

Institution	Institutional Contributions [PM]	Requested Federal Funds [PM]	Comment
ETHZ	4	4	operations
SWITCH	7	7	operations mngt
GC3	12	12	support mngt
WSL	1	1	support
Vital-IT	3.5	3.5	support
HES-SO//Ge	4	4	support, cloud bridging, applications
LHEP	0.5	0.5	support
HES-SO//Fr	1.5	1.5	operations, support Cloud testbed infrastructure
Total	**33**	**33**	**67**

Institution	WP1 [PMs]	WP2 [PMs]	WP3 [PMs]	Comment
ETHZ	3	4	1	
SWITCH	12		2	operations mngt
GC3	4	16	4	support mngt
WSL		2		
Vital-IT	1	6		
HES-SO//Ge	3	5		
LHEP		1		
HES-SO//Fr	3			
Total	**26**	**33**	**7**	

Information scientifique: accès, traitement et sauvegarde

Table 3 Contribution by institution (sum of hardware and manpower contribution)

Institution	Institutional Contributions [CHF]	Requested Federal Funds [CHF]	Comment
ETHZ			
SWITCH	Not applicable	Not applicable	
GC3			
WSL			
Vital-IT			
HES-SO//Ge			
LHEP			
HES-SO//Fr			
EPFL			
TOTAL			

Risks, possible problems:

- Emergence of freemium models for academic users by commercial cloud providers as a result of competition (e.g. 500GB of free storage incl. transfer and 100khours CPU time per researcher).
  Such a development would be beneficial for our mission. We would concentrate on the user support and consulting, reducing the priority of operating an 'own' computational and storage infrastructure.
- Small and distributed infrastructure contributions by partners and/or observers. Effort to operate small infrastructure contributions is prohibitive. We can only accept such if an added value, that is, if it's not part of the production infrastructure and is instead used for experimenting and testing.

Boundaries to other projects and activities:

Prerequisites for success of the project:

Confidentiality of project submission: *<Choose one of the following three options:>*
A) CRUS can communicate project content
B) CRUS can communicate project title, leading institution and name of project leader
C) CRUS should treat project submission confidential

Remarks:

Signature project leader

Signature project responsible of the university

Information scientifique: accès, traitement et sauvegarde

References

[crus] Informationsversorgungsprogramm (fr.Acces a l'information scientifique). See http://www.crus.ch/funktionen/sitemap.html (direct link of call not available at time of this writing).

[STS] https://forge.switch.ch/redmine/projects/sts/wiki/Wiki/diff/6

Glossary

AA	Authentication and Authorization
AAA/SWITCH	e-Infrastructure for e-Science program (April 2008 - April 2013)
AAI	Authentication and Authorization Infrastructure
EC2	Amazon Elastic Computing Cloud
IaaS	Infrastructure as a Service
OCCI	Open Cloud Computing Interface
RS-NAS	Remote Scalable Network Attached Storage
SMSCG	Swiss Multi Science Computing Grid
STS	Security Token Service
SwiNG	Swiss National Grid Association
NGI_CH	National Grid Initiative (for Switzerland)
VM-MAD	Virtual Machines Management and Advanced Deployment