Logical Infrastructure Composition Layer, the GEYSERS Holistic Approach for Infrastructure Virtualisation

Joan A. García-Espín, Jordi Ferrer Riera, Sergi Figuerola
Distributed Applications and Networks Area (DANA), Fundació i2CAT, C./ Gran Capità 2-4, office 203, Nexus 1 building, 08034, Barcelona, Catalunya. Spain
e-mail: joan.antoni.garcia@i2cat.net, jordi.ferrer@i2cat.net, sergi.figuerola@i2cat.net

Mattijs Ghijsen, Yuri Demchemko
University of Amsterdam, Informatics Institute, System and Network Engineering group, Science Park, 904, 1090 Amsterdam, The Netherlands
e-mail: m.ghijsen@uva.nl, y.demchemko@uva.nl

Jens Buysse, Marc De Leenheer, Chris Develder
Ghent University – IBBT, Dept. of Information Technology (INTEC) – IBCN, Gaston Crommenlaan 8 bus 102, BE-9050 Gent, Belgium
e-mail: jens.buysse@intec.ugent.be , marc.deleenheer@intec.ugent.be, chris.develder@intec.ugent.be

Fabienne Anhalt, Sébastien Soudan
Lyatiss, 39 rue du Mail, 69004, Lyon, France
e-mail: fabienne@lyatiss.com, sebastien@lyatiss.com

Paper type

Research Paper

Abstract

This article presents the Logical Infrastructure Composition Layer as a solution produced in the context of the European project GEYSERS for ICT infrastructure virtualisation in the Future Internet. The concepts behind this layer are based on infrastructure resources virtualisation, regardless of their nature (either IT or network); the Resources-Ownership-Roles-Actors model for virtualised resources access and the concept of the Virtual Infrastructure and its management attributes. This article provides an overview on them. After that, the work is focused on the Virtual-to-Physical Infrastructure mapping problem and different options to allocate Virtual Infrastructures over the same substrate. Finally we present the grouped VI mapping strategy and its analysis, depending on the utility function considered and the context conditions. Our findings show that batched VI mapping strategy enhances the amount of VIs to be allocated on the physical substrate. The technological solution and simulations on the potential benefits show a novel ICT infrastructure control and management solution that is able to accommodate the optimisation requirements for the Future Internet (cost, energy, availability, flexibility, etc.) in coordination with application deployments and cloud service models.

Keywords

IT and Optical Network Virtualisation, Virtual Infrastructure, Resource Abstraction, Coordinated IT and Network Resource Provisioning

1 Introduction

The current Internet has become a ubiquitous commodity to provide communication services to both enterprises and residential users [1]. Cloud computing has emerged as a key paradigm in order to provide computing services addressing users' requirements over the Internet. Cloud computing stands for transparent, on-demand access to IT hardware or software resources, which are geographically spread and interconnected by networks. In fact, analyses predict that in 2020 more than 80% of the infrastructure will be outsourced within the Cloud [2]. While there are countless definitions for the Cloud computing term, there seem to be common characteristics that a cloud infrastructure should have: (i) pay-per-use (no on-going commitment, utility prices); (ii) elastic capacity and the illusion of infinite resources; (iii) self-service interface; and (iv) resources that are abstracted or virtualised [3]. Others also argue that broad network access is the fifth essential characteristic [15], but this requires problem analysis and robust integration design of IT and networks resources and their associated service flows, and this is typically not addressed. This fifth characteristic and the shortcomings on its investigation strongly motivate the work we present in this article.

None of the common points (i to iv above) mention the network, its availability or even the network resources composing it. This new IT provisioning paradigm considers that the network is always available and provisioned, which clearly is not necessarily true. Applications or services running in the cloud may be affected by network performance, throughput, delay, or any other QoS parameter: current applications, such as 3D-video streaming, hold high requirements in terms of network performance. Furthermore, bandwidth-provisioning systems typically do not take into account specific characteristics of the IT resources and services connected at the edge of the network. In other words, architectures for coordinated IT and network resource provisioning have been barely investigated [4].

This divergence has been present in the research community for several years, where provisioning optimisation has been scarcely addressed considering both realms at the same time. Coordinated IT and network infrastructure service provisioning is one of the main challenges to be faced. In order to dynamically provision IT resources and gain full benefit of these thanks to Cloud technologies, it is crucial to have control over the quality of the network connections used.

In parallel to the emergence of this novel coordinated provisioning, ideas and concepts behind virtualisation have matured enough after being on the research arena for a while. The fact of sharing a common good in order to improve its efficiency, usage, and productivity has been a key goal in the research community, especially given the fact that the good, in our case, is an expensive and operationally costly ICT infrastructure. Therefore, since both IT and network realms have been totally independent, there have not yet been many approaches in the community considering combined resource virtualisation, using resources from both domains.

Current trends in the telecom realm are moving towards such holistic architectures, capable of handling both IT and network resources in a converged manner. The physical infrastructure, the common good to be shared, is totally decoupled from the services that may be offered on top of it, through the combined resource virtualisation. The Generalised Architecture for Dynamic Infrastructure Services (GEYSERS) [16] project aims at addressing this coordinated IT and network resource virtualisation in order to address some of the Future Internet challenges [5]. The combination of the Infrastructure as a Service (IaaS) model for both optical network and IT resources with a novel resource access model, namely Resource-Ownership-Role-Actor (RORA) [6], enables this separation.

The Logical Infrastructure Composition Layer (LICL) is the element in GEYSERS responsible of acting as a middleware. It aims at decoupling infrastructure resource management from the actual service provisioning, regardless of resources' nature. The LICL utilizes a specific semantic resource description model, namely the Information Modelling Framework (IMF), in order to provide a Virtual Infrastructure (VI) that hides technological and/or vendor specific details of the underlying physical infrastructure from the operators. At the same time, the IMF provides the LICL with a uniform language to manage both types of resources. Hence, the IMF allows the LICL to perform flexible, coordinated resource virtualisation with the different infrastructure resources.

Following with the previously commented issues, the paper addresses the challenges on coordinated provisioning of virtualized IT and optical network resources through the LICL, contextualized within the GEYSERS project. The paper presents the LICL concepts and design for manipulating physical and virtual infrastructure resources, both from IT and network domains. It elaborates in the LICL concept as the basis for the Future Internet infrastructure management, in the context of the GEYSERS project. Furthermore, we discuss the different service lifecycles and lifetimes for the virtual infrastructure resources.

The article is structured as follows. First, in Section 2, we introduce the GEYSERS project, its main components and its functional layered architecture. Then, we present in detail the LICL and provide a brief summary of the state of the art regarding resource virtualisation. We also explain how the VI concept is handled in the FP7 SAIL project compared to the GEYSERS approach and provide a brief overview of the RORA model in GEYSERS. Next, Section 3 contains the whole specification and characterisation of the VI concept in GEYSERS: we present the model, the lifecycles and the timelines of the VI provisioning service. After that, in Section 4, we provide the functional architecture of the LICL and the specification of the IMF used to virtualise the physical infrastructure. We present the VI mapping problem and provide in the subsequent Section 5 some results on a grouped VI mapping approach, exploring the benefits of simultaneously handling groups of VI.

2 Related Work

2.1 The GEYSERS project

The GEYSERS architecture presents an innovative approach by adopting the concepts behind the Infrastructure as a Service (IaaS) servicing model from cloud computing and service-oriented networking to enable infrastructure operators offering new IT and network converged services. In the GEYSERS layered architecture physical devices populating the bottom layer – physical infrastructure layer – are abstracted and partitioned or grouped into virtual resources that can be selected to form the virtual infrastructures. This process takes place in the LICL, the key element of the GEYSERS stack in order to provide converged infrastructure services. On top of the virtual infrastructures, there is the Service Middleware Layer (SML) and the Network Control Plane (NCP+), responsible for configuring and managing virtual resources. Furthermore, the SML is responsible for translating the application requests and Service Level Agreements (SLAs) into technology specific requests in order to trigger the provisioning procedures at the NCP+ level [7].

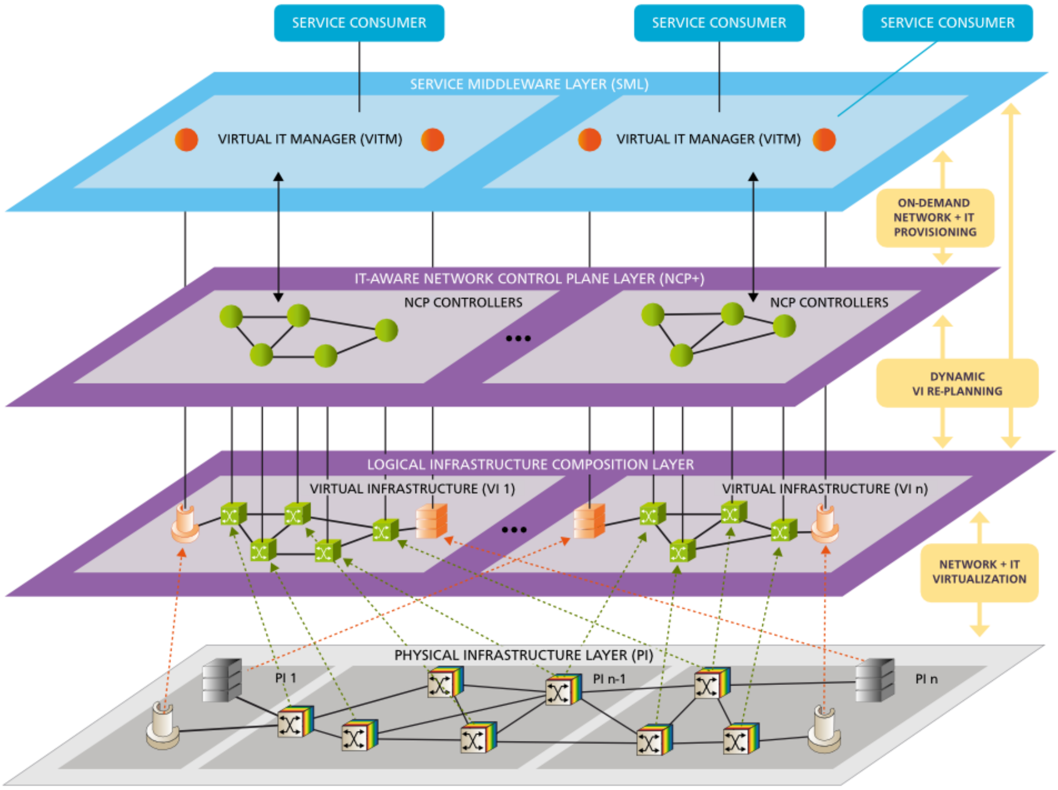

Figure 1: GEYSERS global architecture

The SML is a convergence layer for coordinating the management of IT resources that belong to an aggregate service. The SML contains the Virtual IT Manager, which is the element responsible of the end-to-end IT service management and the virtual IT resource configuration. The GEYSERS NCP+ performs all control and management functions necessary to operate the virtual network resources within the virtual infrastructure. The NCP+ also offers a set of functionalities towards the SML, in support of on-demand and coupled provisioning of the IT resources and associated network connectivity. The LICL is a key component in the GEYSERS architecture and represents one of its architectural innovations. It is responsible for the creation and maintenance of virtual resources as well as virtual infrastructures composed of those virtual resources. Such a layer acts as a middleware on top of the physical infrastructure and offers a complete toolset to the involved entities in the infrastructure service provisioning workflow. More information of each layer presented can be found in [4,6].

The LICL provides the infrastructure services within GEYSERS. This layer is the component in the GEYERS architecture responsible for abstracting and virtualising the physical resources, and thus offering them as a service to the upper layers on the GEYSERS stack, being those layers the NCP+ and the SML. Although having such other components and elements (e.g., security), we focus on this article into the three pillars on which the main LICL functionalities rely: (i) resource abstraction, (ii) the IMF, and (iii) the synchronisation mechanisms.

Abstraction is the first step for resource virtualisation, since it provides the foundations for partitioning and mapping the resources populating the underlying infrastructure. It is defined as the process representing physical resources as a set of uniform attributes, characteristics and functionalities while hiding unnecessary characteristics from the resource itself [8] The IMF specifies and provides a detailed description of this abstracted model, and all the additional parameters that are necessary in order to provide infrastructures services. Finally, the synchronisation mechanisms keep abstracted information on the IMF consistent with the underlying resources. The functional architecture of the LICL and detailed description of these three pillars is provided in Section 4.

2.2 Virtual Infrastructures in the SAIL project

Network and computing infrastructure are composed of various devices operating at different layers. It is nowadays common to virtualise servers and easily possible to virtualise the network layer 2 and layer 3 networks, using technologies such as VLAN, router virtualisation, e.g., Virtual Router Forwardings (VRFs), or even OpenFlow technology [11]. Difficulty remains in managing virtual infrastructures as a whole, e.g., creating, configuring, or interconnecting virtual networks.

Addressing this issue, a virtual network management paradigm is currently being investigated by EU FP7 SAIL project [17]. Virtual network partitions can be reserved and deployed on demand as ‘Flash Network Slices’ (FNS), which interconnect resources located on different sites by a layer-2/layer-3 network that is seen by the user as a simplified network interconnection. This approach is interesting for a maximum level of abstraction and automation.

VIs in the GEYSERS context differ from this FNS approach by two main criteria: (i) operators are given access to partitions of optical network devices for reconfiguration during runtime, and (ii) VIs are composted of virtual network resources and virtual IT resources (computing and storage). This cross-layer approach allows virtual network and application providers to jointly provision network, computing and storage resources for providing optimally adapted services to applications.

2.3 The RORA Model in GEYSERS

The GEYSERS architecture facilitates the emergence of business entities that implement new behaviours depending on how they interact with the infrastructure. Such an architecture, decoupling the traditional infrastructure from the service provided on top of it, facilitates the appearance of new challenges on how to model interactions among the entities involved in the service provisioning operations and on how to match them with the current infrastructure operators workflows.

Considering these challenges, the aforementioned architecture, and current trends towards vertical disintegration of the telecom businesses models, we have defined the RORA model, which aims at providing the required tools and utilities in order to define and specify any business model. The RORA model takes its name from the four components it is based on: Resources, Ownership, Roles, and Actors. Resources represent the first component of the model. Within GEYSERS, a resource can be any physical resource populating the physical infrastructure layer, as well as a virtual resource or even a virtual infrastructure. Then, the model considers the ownership scheme over a given resource, which determines the different set of actions that can be performed by an actor over that resource. The ownership scheme is based on previous research of Dijkstra et al. [12,13]. Under the GEYSERS umbrella we consider four different types of ownership: (i) the legal or economic ownership, (ii) the administrative ownership, (iii) the operative ownership, and (iv) the usage ownership. The model comes to the roles, which help in the chain value description. Basically, a role names the behavior of an entity participating in a particular context and generally is used to identify it [14]. Within GEYSERS we identify four different roles, associated to the whole stack previously presented: the Physical Infrastructure Provider (PIP), the Virtual Infrastructure Provider (VIP), the Virtual Infrastructure Operator (VIO), and the Service Consumer (SC). The toolset offered by the LICL is used by two of these roles: the VIP and the PIP, while the VIO is responsible for consuming the virtual infrastructure service. Finally, the model considers an actor, which is a materialisation of one or various roles. Thus, a single actor may have different roles. The RORA model considers each entity involved in a given use case as an actor, but depending on the role or roles that the actor implements, it has different responsibilities and duties.

3 Virtual Infrastructure

3.1 Concept and Model

One of the key components of any on-demand cloud provisioning system is the Service Delivery Framework (SDF) [18]. The SDF, originally proposed by the TeleManagement Forum, is a conceptual framework, intended for static systems, managed by human operators. The GEYSERS SDF extends this conceptual framework by connecting it to a specific architecture (i.e., the GEYSERS architecture) and applying it to a dynamic and automated system.

The GEYSERS Service Delivery Framework (SDF) supports the provisioning of a Virtual Infrastructure (VI) as a service. The workflow of the VI provisioning is shown in Figure 2, which comprises five phases, namely, (1) Service requests and SLA negotiation; (2) Planning/design; (3) Deployment/Configuration; (4) Operation and (5) Decommission. The GEYSERS SDF is compliant with the TMF Service Delivery Framework [19] with the necessary extensions to facilitate the combined network + IT services.

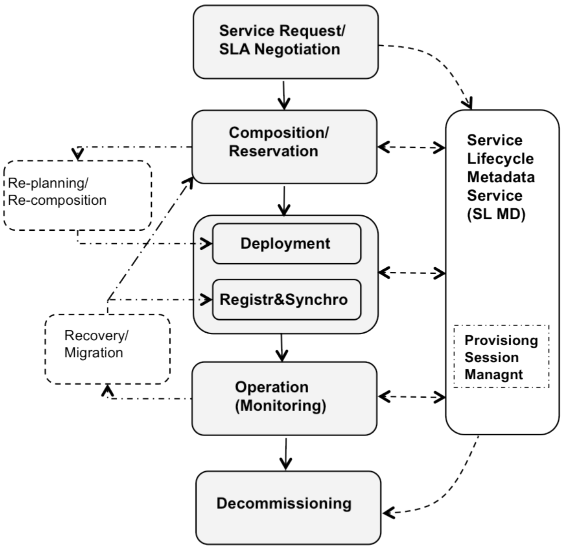

Figure 2: GEYSERS Service Delivery Framework workflow chart

The GEYSERS SDF imposes a number of requirements on the management and operation of a VI, which will be discussed in the following sections.

3.2 Management

Virtual infrastructure management consists of different operations involving the creation, configuration and destruction of virtual resources (VR). In the context of GEYSERS, we consider computing, storage and network resources down to layer one, organized in a specific topology. This makes the VI entity a complex construct with many configuration parameters. In addition, several actors on different layers of the GEYSERS model (cf. Fig. 1) have different roles in the management of VIs. Application providers manage computing and storage resources by requesting their instantiation and decommissioning to virtual infrastructure providers (VIP). VIPs in turn must manage the configuration of these resources as well as the network interconnecting them, and the provisioning of the requested services.

A first problem in the management of VIs is their parameterization, i.e., the description of each of their VRs and their interconnection topology. In GEYSERS, each resource can be configured with different attributes using the IMF (cf Section 3.1), hence allowing the management and service configuration of the VI. In addition, such a precise parameterization allows selecting among the available physical resources the ones where each VR can be hosted. A second problem is the selection of physical resources to map VRs. To ease this process, the physical resources are abstracted to IMF descriptions. Based on their capabilities and the VR-attributes, the allocator[1] used by GEYSERS performs the mapping. This allocator performs sub-graph isomorphism detection; incorporating the VR parameters and topology constraints given by the IMF description, to ensure the VI will provide the requested service levels. Finally, another issue is the management of the physical substrate considering the heterogeneity of resources, e.g., optical cross connects, layer 2 access network, computing and storage resources and Cloud management systems. In GEYSERS, all these devices are configured through device-specific adaptors, but managed in a common way, as these adaptors expose the devices as abstracted resources to the management layer.

3.3 VI Service lifecycle

Virtual infrastructures are de-materialized resource aggregates, which provide a service over a time-limited period. Their lifecycle involves their planning and creation phases, the service delivery phase, and finally their

---

[1] The allocator currently used by the LICL is part of Lyatiss CloudWeaver (www.lyatiss.com).

decommissioning. Figure 3 shows the different states of a virtual infrastructure lifecycle within GEYSERS.

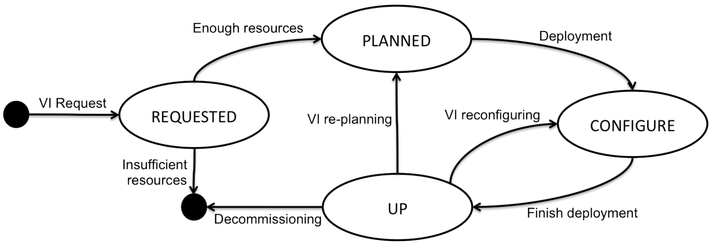

Figure 3: Lifecycle of a Virtual Infrastructure

When a user requests a VI, such request is submitted to an allocator. The state of a virtual infrastructure always depends on the different states of the virtual resources that compose it. Hence only once all resources of the virtual infrastructure could have been allocated, the VI moves to the *planned* state, otherwise the request is rejected. Once the VI planned, its different virtual resources are configured with the different parameters given in the VI request, and it moves to the *configure* state. The virtual infrastructure becomes *up* once all of its virtual resources have been instantiated on the physical resources to which they had been assigned. It is then ready to be handed over to the users (e.g., network operators, application providers), who can proceed with operating the network resources, installing applications or even create new virtual nodes. At the end of the reserved lifetime of a virtual infrastructure, all its virtual resources are decommissioned.

3.4 VI Provisioning Timescales

Infrastructure services timescales typically are confused with connectivity services timescales or even IT-resource provisioning timescales, even being those services different in nature. It is quite important to distinguish between operation and reservation or management phases of the delivery framework. Within GEYSERS, we have developed the infrastructure service, we have characterized the virtual infrastructure, with the model and the lifecycle associated to the service; furthermore, we provide of the comparison for the timescales when this service is valid. Figure 4 depicts clearly the range where the infrastructure service lives.

If we consider operation services — such as connectivity services, which basically consist of configuration commands on the resources — we can see how the LICL introduces certain delay, due to the additional virtual layer added on top of the physical resources. This is one of the main constraints of the virtualisation approach, the delay introduced in the operation phase. However, the delay in the timescale is minimized for the reservation phase, in other words, for the management stage. Therefore, the figure clearly depicts how the planning and automated re-planning phases introduce a significant enhancement of the provisioning time with regards with current infrastructure provisioning practices. Although not usually mentioned, such an enhancement in the infrastructure provisioning time, we consider that it represents one of the key advantages of the virtual infrastructures. To sum up, the extra delay in the operation commands takes place in different orders of the timescale, and thus it is not comparable to the benefits of the whole infrastructure provisioning time.

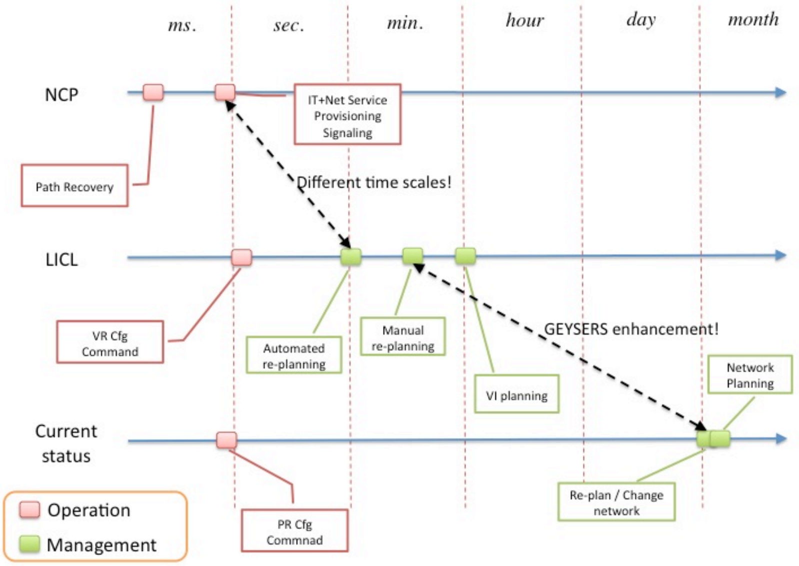

Figure 4: GEYSERS ecosystem timescales

4 Logical Infrastructure Composition Layer

As stated before, the LICL is the key element in the GEYSERS architecture in order to provision infrastructure services. This section provides a detailed description of the functional architecture of the LICL, divided into two major components, the upper and the lower LICL, depending on the functionalities covered. It also contains the detailed description of the IMF and how it enables the virtualisation process. Furthermore, this section introduces the problem of mapping the virtual infrastructures requests on top of the physical resources.

4.1 Functional Architecture

The LICL has been divided into two main sub-systems depending on the functionalities implemented in each sub-system and also depending on the role that uses such functionalities. On the one hand, there is the upper-LICL, which is responsible mainly for the virtual infrastructure management and satisfies the needs and requirements of the virtual infrastructure provider. On the other hand we have, the lower-LICL, which is responsible for physical resource virtualisation and management and which satisfies the requirements of the physical infrastructure provider.

The upper-LICL is composed of different modules. The functionalities covered at this level are the virtual infrastructure creation, management and re-planning, and the Service Level Agreement (SLA) enforcement. The virtual infrastructure creation is done as a composition of different virtual resources available from one or multiple PIPs. Such a virtual infrastructure is provisioned towards the virtual infrastructure operator as a unit. Furthermore, the upper-LICL offers dynamic re-planning functionalities as a response to the changing requirements of the VIO. Such dynamic re-planning may involve the inclusion of new resources to the virtual infrastructure, the release of un-used resources, or even the resizing of some of them (e.g., increase or decrease the total bandwidth capability of a virtual link). As a part of the system oriented to provide dynamic infrastructure services, the upper-LICL provides capabilities to ensure SLA levels are met during the whole service lifecycle.

The lower-LICL covers the functionalities regarding physical resource abstraction and resource virtualisation. The tools offered by the lower-LICL are used by the PIP in order to manage its own infrastructure. The lower-LICL is responsible for the physical resource abstraction that basically comprehends all the necessary steps to create a logical resource representing the physical resource. It also is in charge of the virtual resource creation and management, as well as the resource monitoring and configuration. The lower-LICL also offers an information service, which is used by the PIP to send information about its domain capabilities towards the different VIPs.

Figure 5 depicts the functional architecture of the LICL, split into the two aforementioned components. It also shows the different interfaces in each component in order to communicate with the outer world. In the case of the upper-LICL, it has the Management-to-LICL (MLI) interface, which offers all the virtual infrastructure management operations (e.g., request, re-planning, decommission) and then the SML-to-LICL (SLI) interface and the Call Controller Interface (CCI), used to offer operation capabilities over the virtual infrastructure. In detail, the SLI offers operations over the virtual IT resources and the CCI over the virtual network resources. However, it is remarkable that this is a logical differentiation, since the implementation of the system offers one interface and handles the virtual resources in a converged manner independently of its nature. Finally, the lower-LICL offers the VR request service, used to request for single virtual resources, the Resource Operation Service, that represents the operation interfaces for the virtual resources, and the information service, which is used to exchange information with the different physical infrastructure providers.

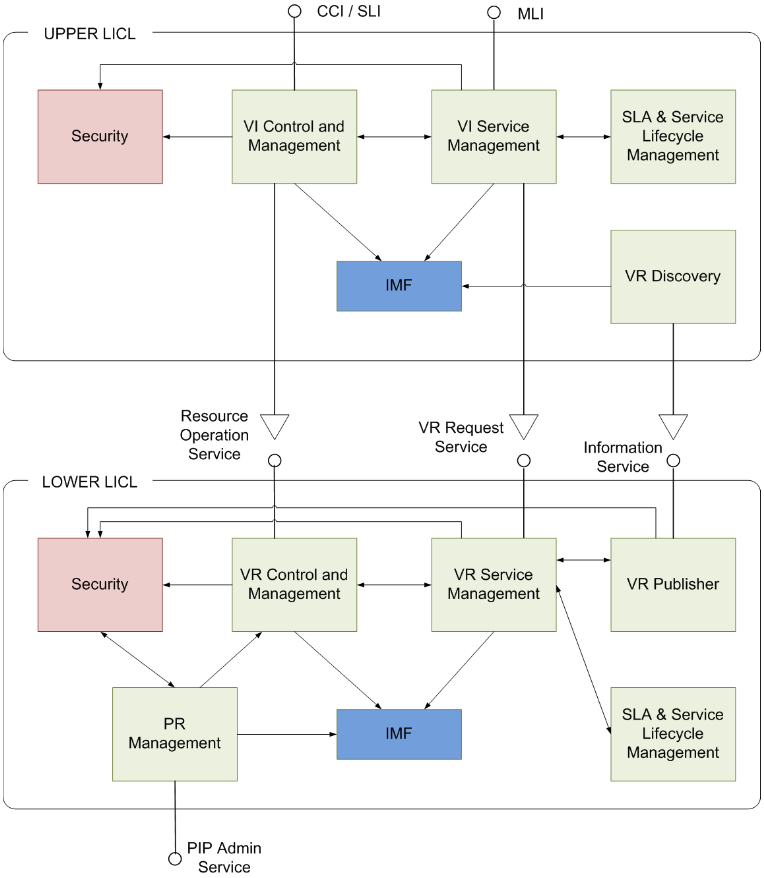

Figure 5: LICL functional architecture overview

4.2 Virtualisation through the Information Modelling Framework

In order for the different LICL components to interact and exchange information, a common information model is needed. The GEYSERS Information Modelling Framework (IMF) draws upon VXDL [20] and the Network Description Language (NDL) [21] and also adopts the latter's semantic approach. As a result the IMF was used as one of the inputs for the Infrastructure and Network Description Language (INDL) [22].

The main resource hierarchy of the IMF is shown in Figure 6. The main concept hierarchy contains the three main resource types: *Node*, *NodeComponent* and *NetworkElement*.

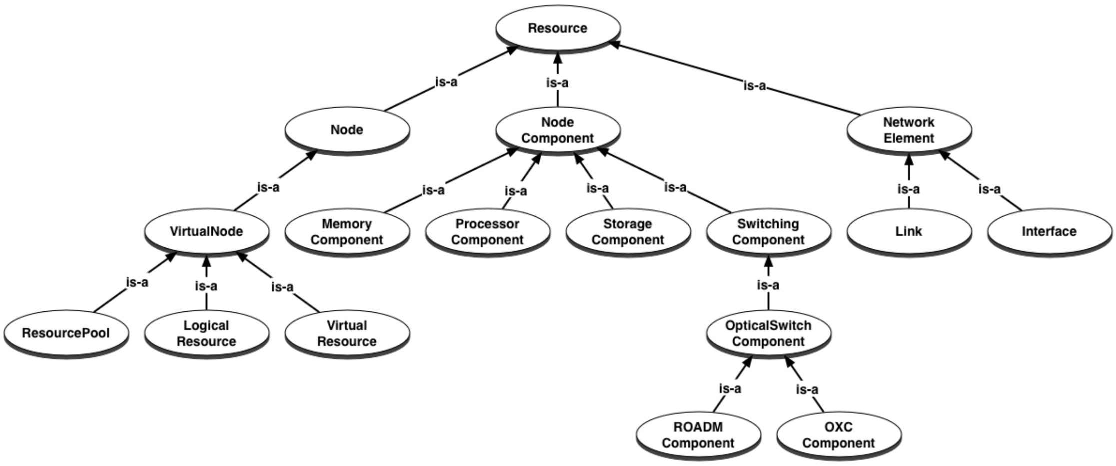

Figure 6: IMF Resource Hierarchy

Virtualization is modelled using the *VirtualNode* concept, which is an abstract class (i.e., it cannot have any instances). The IMF defines different types of *VirtualNodes* that can be implemented on top of other nodes. The subclasses of a virtual node are: *LogicalResource*, *ResourcePool* and *VirtualResource*. A logical resource is used to model an abstracted physical resource, a resource pool is used to model a reserved capacity for future use, and a virtual resource is used to represent an instantiated virtual machine.

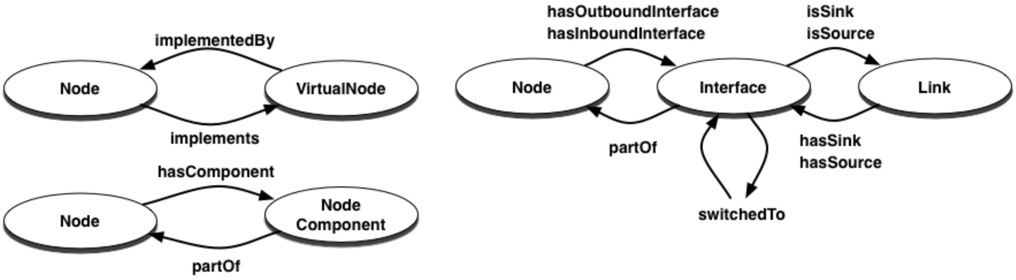

Figure 7: IMF main resource properties

Figure 7 shows the main properties of the resources in the IMF. Important for this section are the *implementedBy* and *implements* relations between *Node* and *VirtualNode*. Using these two relations, we can indicate on which node, a virtual node is implemented. Furthermore, because *VirtualNode* is a subclass of *Node*, we can also implement virtual nodes on top of other virtual nodes. Thus we can create an arbitrary number of virtualization layers.

To illustrate the use of the *implementedBy* relation we show how the different layers of virtualization are modelled in the REQUESTED state and PLANNED state of the VI lifecycle (see Figure 3 for the complete VI lifecycle states). As an example we discuss a single-domain case but the approach for multi-domain is identical. For clarity we do not show the *implements* relation and only show the *implementedBy* relation.

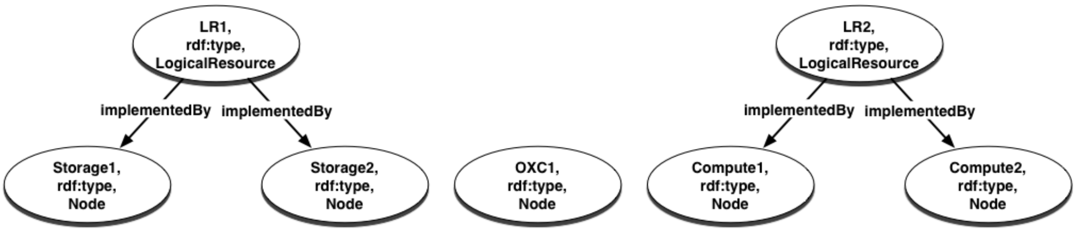

Figure 8: VI in REQUESTED state

When the VI is in the requested state (see Figure 8), the lower LICL already has a number of PR nodes instantiated. These nodes are modelled using the *Node* concept. In the case of IT nodes, these can be aggregated into a single LR node that is managed for example by OpenNebula. The LR is modelled using a *LogicalResource* concept. When a VI containing a certain storage and compute capacity connected by a switch is requested, and the request can be satisfied, the VI will move to the planned state as shown in Figure 9.

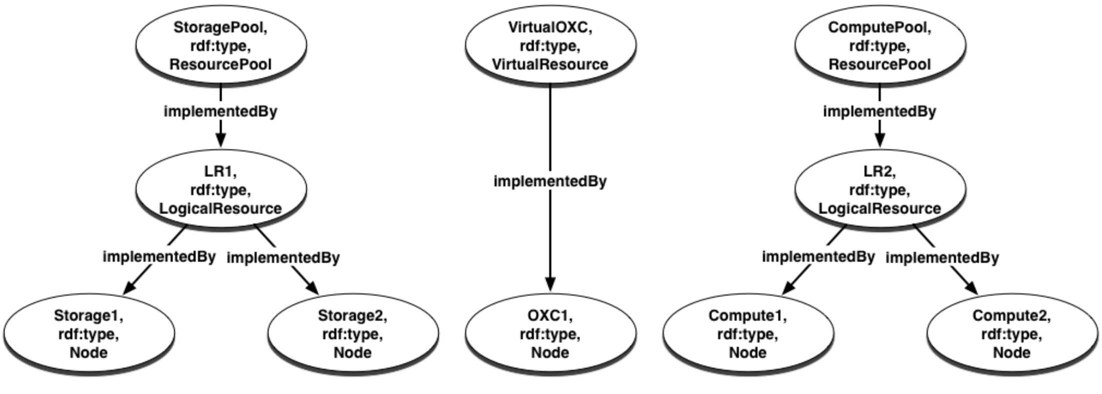

Figure 9: VI in PLANNED state.

In the planned state, the reserved capacity for storage and computing is modelled using the *ResourcePool* concept. The virtual OXC is modelled using the *VirtualResource* concept. In later stages of the VI lifecycle, virtual machines for computing or storage can be deployed on the logical resources. In that case, the virtual machine will be modelled using a *VirtualResource* concept and will be implemented on top of the logical resource.

4.3 Virtual Infrastructure Mapping

In the previous sections we have described the GEYSERS architecture, which is able to provide such an infrastructure-on-demand service. When multiple virtual infrastructures need to be provided on the same physical substrate, selecting which virtual resource to be mapped to which physical resource in order to make efficient use of the underlying physical infrastructure becomes a major issue, which we denote as the "VI mapping problem". A VI mapping problem may have several objectives of which some have been investigated in the scope of the GEYSERS project.

The authors of [26] have investigated the VI mapping problem in transparent optical networks, where intermediate optical switches keep the traffic in the optical domain. In that case, physical link impairments (PLI) inherent to WDM networks may lead to signal degradations, where multiple transparent connections can mutually impact each other. Since multiple VIs can coexist on the same physical substrate, the resulting interferences from PLI between VIs will impact the performance of each VI. Their key message is that in transparent optical networks, when considering a classical VI mapping problem from a PLI perspective, more VI requests mappings are rejected, as their signal quality is unacceptable. Therefore they propose a PLI-aware VI mapping solution, which drastically decreases this blocking ratio (up to 55% in some cases).

In [24], two studies are presented that try to minimise energy consumption of concurrent VIs over a shared IT and network infrastructure, for which the authors have provided a Mixed Integer Linear Programming model. Their results indicate that such an energy aware mapping solution can decrease energy consumption up to 40% compared to a scheme where the distance between IT end points is minimized. However, this energy saving brings along the fact that longer paths are taken between these IT end points, which leads to higher network utilization and higher delay.

A. Pagès et al. [23] studied the impact of the transport technology in the amount and characteristics of the virtual infrastructures that can be built on top of a physical optical network infrastructure. For this purpose they state the so-called Virtual Optical Network Allocation (VONA) problem and focus on two cases: wavelength switching and spectrum switching. The authors conclude that the spectrum-switching case maximises the amount of Virtual Optical Networks than can be mapped to the physical infrastructure, for demands requiring fine bandwidth granularity.

Below, we focus on potential benefits in terms of resource requirements that can arise from grouping several VI requests and provision them jointly. Intuitively, we expect to achieve some network capacity advantages (similar to the effect of grooming in more classical multilayer networks). We will focus on VI mapping in a WDM scenario, where the effect is investigated of grouping VI requests (that are expressed as a set of nodes, and the pair wise traffic they are expected to exchange) in a cluster wherein bandwidth (i.e., wavelength circuits) can be shared, whereas different clusters are still properly isolated.

5 Grouped VI Mapping Approach

In a (D)WDM or Dense Wavelength Division Multiplexing infrastructure virtual topologies consist of lightpaths (i.e., end-to-end wavelength connections) and each virtual topology can be managed independently. However, isolation usually leads to an increase in required capacity with respect to the physical optimal infrastructure, since each virtual network is allocated its own network resources. Given the coarse bandwidth granularity in current commercial (D)WDM products (each wavelength offers 10, 40 or 100 Gbps), total network capacity may be very high while resource utilization unacceptably low. Therefore, in [25] we have proposed to cluster VI requests and introduce traffic grooming in these clusters. As such, we do not introduce de facto isolation within each cluster although full isolation is enforced between different clusters.

A small number of isolated virtual infrastructures maximizes the opportunities of statistical multiplexing and as such will lead to the highest resource utilization. However, this will lead to large isolated virtual networks that in turn degrade control plane scalability, since the number of control plane messages is directly influenced by the number of nodes in a network. As such, we show the trade-off between resource utilization and control plane scalability. More formally, we solve the following VI mapping and clustering problem:

Given

- Physical infrastructure topology
- Set of virtual infrastructure requests, each specified as a traffic matrix.
- The number of isolated virtual infrastructures that should be mapped on the physical topology. Each

isolated virtual infrastructure is composed of one or more virtual network requests.

Find

- The composition of the isolated VIs, i.e., which VI requests jointly form what isolated virtual network.
- The mapping of the isolated virtual networks on the physical topology.

Our solution to this problem comprises a two-step algorithm: first we perform clustering to group individual VI requests in groups of virtual networks, after which we determine the mapping of these virtual networks onto the physical network, basing the exact topology of each virtual network on the aggregate network demand of all involved VIs.

For clustering we have an ILP-based solution, which provides optimal results for the clustering of virtual network requests, which we benchmark against a random clustering approach. For the actual mapping onto the physical infrastructure, we considered two alternatives: (i) a *FullMesh* strategy which minimizes hop distance between in the virtual network nodes, and (ii) a *MaxUtil* strategy that aims at filling the available link capacity as efficiently as possible, by maximally exploiting statistical multiplexing. For more details on these variants we refer to [25].

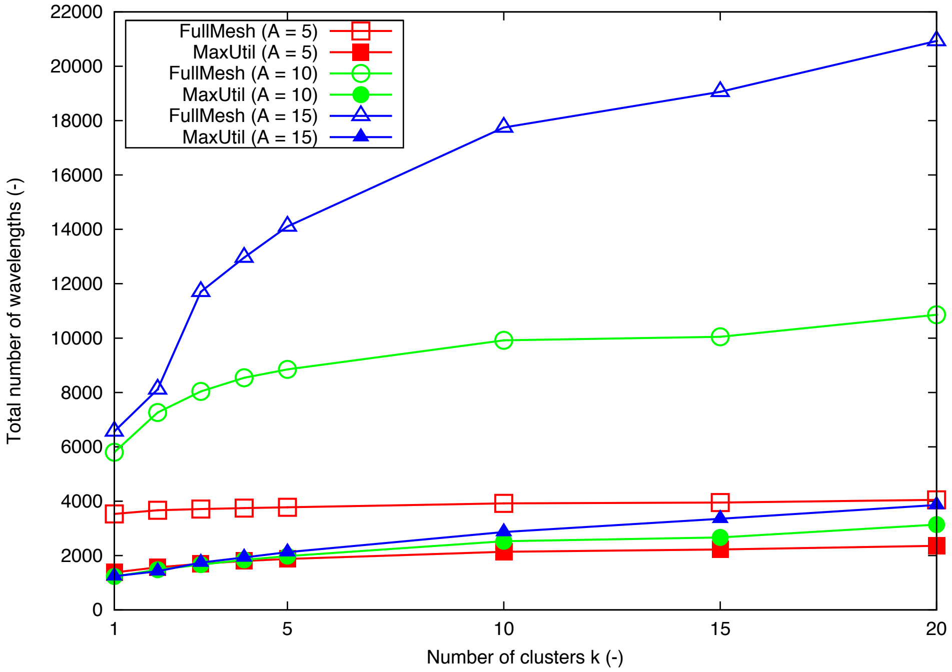

Figure 10: Relatively stable wavelength usage for *MaxUtil* virtual network design, and fast growth for the *FullMesh* approach, for varying sizes of the VI requests (A = number of active nodes, i.e., nodes that generate traffic, in a VI request).

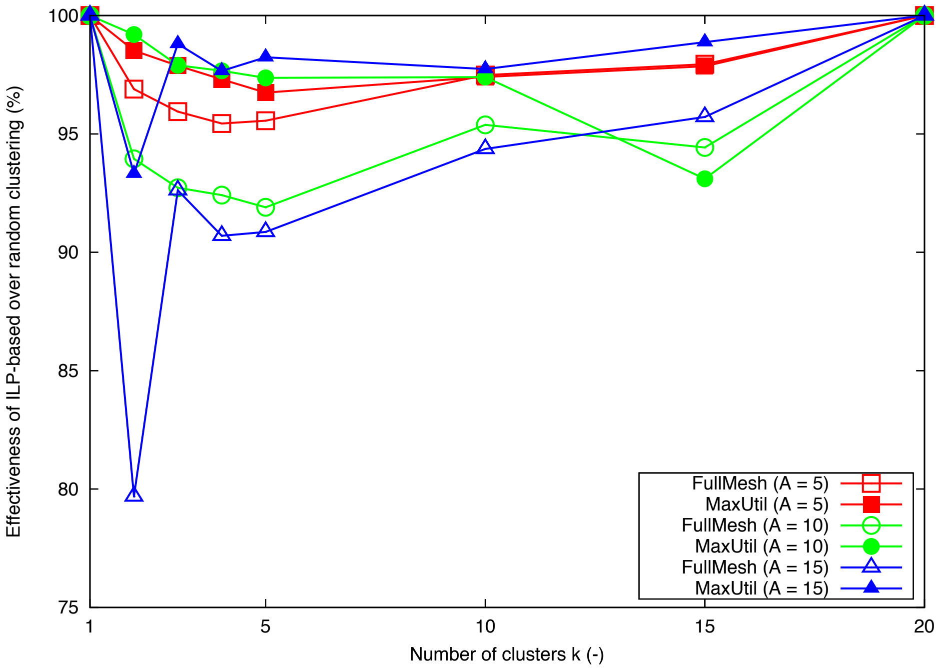

Figure 11: Ratio of total wavelength capacity (= effectiveness) of ILP-based over random clustering

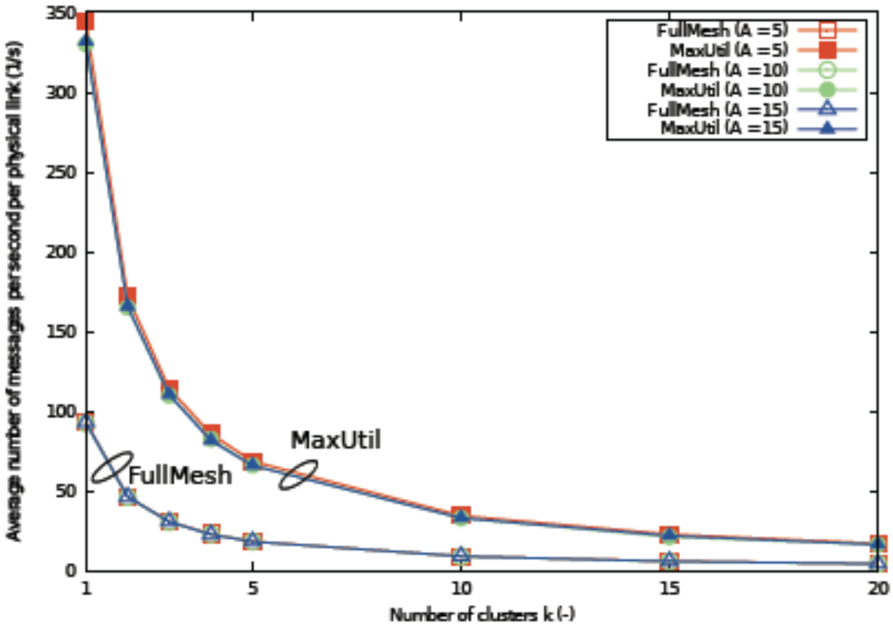

Figure 12: Convergence of average message exchange rate for large number of virtual network clusters

Figure 10 shows the total number of wavelengths necessary to instantiate a varying number of virtual network clusters (here denoted as k), using the ILP-based clustering algorithm. We observe the relatively slow growth in wavelength usage for the *MaxUtil* approach, which is in stark contrast to the behaviour of the *FullMesh* virtual network design.

The effectiveness of the ILP-based clustering algorithm is studied in Figure 11, showing the ratio of the total number of wavelengths for the ILP-based over random clustering. We note that the ILP algorithm requires 5 to 10% less wavelength capacity compared to random clustering. Also, the effectiveness of the random clustering reaches a minimum around 2-3 clusters, indicating the region where intelligent clustering is most relevant. However, the relatively low improvement of ILP-based over random clustering indicates that more advanced clustering should be developed. Indeed, our approach only incorporated node activity of virtual network requests, whereas the potential for network grooming was not considered. Also, since we study completely random traffic matrices between VI requests, their correlation is limited and such opportunities for intelligent clustering are most likely limited.

Within each provisioned VI, connection requests are assumed to be issued at certain rates associated with the traffic matrix used for the VI mapping. In Figure 12 the control plane message exchange rate associated with those dynamic connection requests) is averaged over all virtual networks. We only consider connection signalling traffic (RSVP-TE), as this forms the majority of control plane traffic (especially true when introducing flooding reduction techniques for OSPF). Note that the hop count within the VI always equals 1 in the *FullMesh* case, while the distance between any two VI nodes in the *MaxUtil* case depends on the number of intermediate GMPLS controllers. Figure 12 shows that both design techniques converge to approximately the same average message exchange rate, although the *MaxUtil* approaches a very high control plane load for a small number of virtual networks. The size of the virtual network requests does not influence the message exchange rates at all. Note that the average message exchange rate is a hyperbolic function and thus the total message exchange rate (sum over all virtual networks) remains constant. However, the reduction in control traffic within each cluster indicates that virtualization offers virtual network operators the compelling advantage of control plane scalability (since the associated controllers can run independently from each other).

6 Conclusions

Cloud computing in essence has emerged thanks to the increased availability of network connectivity and bandwidth. However, despite the crucial role that networks play in making cloud services possible, network resource provisioning to date is not an integral part of the cloud service provisioning process. To alleviate this, and thus assure that network performance is satisfactory to meet the specific characteristics of the (cloud-based) applications, the GEYSERS project proposes a holistic architecture, handling both IT and network resources in a converged manner, while exploiting virtualization of both of them to maximize their efficient utilisation in an Infrastructure-as-a-Service (IaaS) model.

In this paper we present the Logical Infrastructure Composition Layer (LICL), which acts as middleware to decouple infrastructure resource management from actual provisioning. We outlined how it is uses a RORA model to handle the complexity of virtualizing resources owned, managed and operated by different actors. The workflow of the provisioning process through GEYSER's service delivery framework (SDF) enables us to manage the lifecycle of different types of resources (e.g. virtual infrastructures) in a uniform way. Furthermore we explained the semantic resource description framework called IMF to formally describe (virtualization of) network and IT resources and we show how the IMF relates to the lifecycle stages of the GEYSERS SDF. The IMF is exploited in the resource mapping/allocation problem to match possible resources to allocate to particular virtual infrastructure (VI) requests, for which we detailed the various steps. For the mapping problem, we provided illustrative results of advanced mapping approaches that our framework allows to implement. In particular, we showed the effects of clustering multiple VI requests to provision them together in a single virtual network: results showed that intelligently clustering them can attain non-negligible advantages in network capacity needed (order of 10%).

7 Acknowledgements

The European Commission, under ICT-2009.1.1 Project No. 248657 Generalised Architecture for Dynamic Infrastructure Services (GEYSERS), has funded this work. We thank our project partners for their valuable contributions. Part of this work has also been partially funded by the TIN2010-20136-C03-03 grant from the Spanish Ministry of Science and Innovation (MICINN).

8 References

[1] Vicat-Blanc, P., Figuerola, S., et al: "Bringing Optical Networks to the Cloud: an Architecture for a Sustainable Future Internet". The Future Internet, vol. 6656. Lecture Notes on Computer Science, pp.307-320. Springer Berling/Heidelberg 2011. 10.1007/978-3-642-20898-0-22.

[2] Nelson, M.R.: The Next Generation Internet, E-Business, and E-everything, Available online at: http://www.aaas.org/spp/rd/ch20.pdf [Retrieved 2011-11-17].

[3] Rajkumar Buyya, James Broberg, and Andrzej Goscinski. Cloud computing: principles and paradigms. Wiley, Hoboken, New Jersey, 2011.

[4] Escalona, E., Nejabati, R., et al: "GEYSERS Deliverable D2.1: Initial GEYSERS Architecture & Interfaces Specification". EC-FP7 GEYSERS project. Available online at: http://www.geysers.eu/images/stories/deliverables/geysers-deliverable_2.1.pdf [Retrieved 2011-11-18].

[5] Figuerola, S, García-Espín, J.A., Ferrer Riera, J., and Ciulli, N.: "An Optical network and IT infrastructure virtualisation framework". Optical Internet (COIN), 2010. 9th International Conference on. 11-14 July, Korea (South). 978-1-42448221-4. DOI: 10.1109/COIN.2010.5546328.

[6] Escalona, E., et al: "GEYSERS Deliverable D2.2: GEYSERS overall architecture, interfaces specification, and service provisioning workflow. EC GEYSERS Project. Available online at: http://www.geysers.eu/images/stories/deliverables/geysers-deliverable_d2.2_update_final.pdf [Retrieved 2011-11-18].

[7] Tzanakaki, A., Anastasopoulos, M., Georgakilas, K., Buysee, J., De Leenheer, M., Develder, C., Peng, S., Nejabati, R., Escalona, E., Simeonidou, D., Ciulli, N., Landi, G., Manfredi, A., Lopez, E., Ferrer Riera, J., Garcia-Espin, J.A., Donadio, P., Parladori, G., and Jimenez, J. Energy Efficiency in integrated IT and Optical Network Infrastructures: The GEYSERS Approach. Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on; 10-15 April 2011, pp. 343 - 348, 978-1-4577-0249- 5, 10.1109/INFCOMW.2011.5928835

[8] Abosi, C., et al., A Novel Service Composition Mechanism for the Future Optical Internet. Journal of Optical Communication and Networking IEEE/OSA, vol. 1, issue 2, pp. 106-120, July 2009.

[9] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. A view on cloud computing. Communications of the ACM, 53(4):50–58, April 2010.

[10] A. Willner, C. Barz, J. A. Garcia-Espin, J. Ferrer Riera, S. Figuerola, and P. Martini. Harmony - advance reservations in heterogeneous multi-domain environments. In Luigi Fratta, Henning Schulzrinne, Yutaka Takahashi, and Otto Spaniol, editors, NETWORKING 2009, volume 5550 of Lecture Notes in Computer Science, pages 871–882. Springer Berlin / Heidelberg, 2009. 10.1007/978-3-642-01399-7

[11] McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rex- ford, J., Shenker, S., and Turner, J. (2008). OpenFlow: enabling innovation in campus networks. SIGCOMM CCR, 38(2):69–74.

[12] Dijkstra, F., and de Laat C.: "Optical Exchanges". In proceedings of the 1st International Workshop on Networks for Grid Applications GridNets04. San José, CA, USA.

[13] Dijkstra, F., et al.: "A terminology for control models at Optical Exchanges". Inter domain management. Lecture notes in Computer Science, 2007, Vol. 4543/2007, 49-60, DOI: 10.1007/978-3-540-72986-0_5.

[14] Booch, G., Rumbaugh, J., and Jacobson, I., The UML Reference Manual. AW 1999. ISBN 1931777446

[15] Zack, B.: "Patterns for Cloud Computing", presentation at the Cloud Computing Expo 2011, New York. Available on line at: https://skydrive.live.com/?cid=f522894cef160b44&id=F522894CEF160B44%21318#!/view.aspx?cid=F522894CEF160B44&resid=F522894CEF160B44%21319 [Retrieved 2012-04-15]

[16] European 7th Framework Programme project GEYSERS. Available online at: http://www.geysers.eu

[17] European 7th Framework Programme project SAIL. Available online at: http://www.sail-project.eu

[18] Y. Demchenko, J. van der Ham, M. Ghijsen, M. Cristea, V. Yakovenko, C. de Laat. On-Demand Provisioning of Cloud and Grid based Infrastructure Services for Collaborative Projects and Groups, The 2011 International Conference on Collaboration Technologies and Systems (CTS 2011), May 23-27, 2011, Philadelphia, Pennsylvania, USA.

[19] TMF Service Delivery Framework. http://www.tmforum.org/servicedeliveryframework/4664/home.htm

[20] G. Koslovski, P. Vicat-Blanc Primet and A. Schwertner Charão. VXDL: Virtual Resources and Interconnection Networks Description Language. In GridNets 2008, Beijing, China.

[21] J. van der Ham, P. Grosso, R. van der Pol, A. Toonk, and C. de Laat, "Using the network description language in optical networks," in Tenth IFIP/IEEE Symposium on Integrated Network Management, May 2007.

[22] M. Ghijsen, J. van der Ham, P. Grosso and C. de Laat. Towards an Infrastructure Description Language for Modeling Computing Infrastructures, in 10th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA 2012), Leganés, Madrid, 10-13 July 2012.

[23] A. Pagès, J. A. García-Espín, J. Perelló, J. Ferrer Riera, S. Spadaro, and S. Figuerola, "Optimal Allocation of Virtual Optical Networks for the Future Internet", in Proceedings of the 16th Conference on Optical Network Design and Modeling (ONDM 2012), April 2012.

[24] Anastasopoulos, M., Tzanakaki, A., Georgakilas, K., and Simeonidou, D.: "Energy Aware Planning of Multiple Virtual Infrastructures over Converged Optical Network and IT Physical Resources". Proc. of 37th European Conf. Optical Communication ECOC, Geneva, Switzerland, September 2011.

[25] De Leenheer, M., Buysse, J., Develder, C., and Mukherjee, B.: "Isolation and Resource Efficiency of Virtual Optical Networks". Proc. Int. Conf. Computing, Networking and Communications ICNC, Maui, Hawai,

USA, 2012.

[26] Peng, S., Nejabati, R., Azodolmolky, S., Escalona, E., and Simeonidou, D.: “An impairment-aware virtual optical network composition mechanism for future Internet”. Proc. of 37th European Conf. on Optical Communication ECOC. Geneva, Switzerland, 2011.

9 Biographies

Joan A. García-Espín is a research project manager at the Distributed Applications and Networks Area of the i2CAT Foundation in Barcelona, Spain. He received his MSc degree from the Technical University of Catalonia (UPC) in 2007 for a thesis on design and implementation of TE-enabled, DiffServ-aware MPLS networks providing end-to-end QoS. He is currently the work package leader of the Logical Infrastructure Composition Layer in the EU-FP7 GEYSERS. He contributed to create the bandwidth on demand tool named Harmony in the EU FP6 Phosphorus project. He also contributed to the Network Services Framework and Interface recommendations in the Open Grid Forum.

Jordi Ferrer Riera received his MSc degree in Computer Science at the Technical University of Catalonia (FIB-UPC). He joined the i2CAT foundation as a software engineer in early 2007 for developing the MANTICORE project. In 2008, he started his collaboration in the PHOSPHORUS project. He also collaborates with the GLIF Generic Network Interface Technical Group (GNI-TG), adapting the Harmony Service Interface (HSI) to the GNI specification. In 2010, he started collaborating actively in the GEYSERS project. He is currently a PhD candidate of the Telematics Engineering Department of the Technical University of Catalonia (UPC).

Mattijs Ghijsen received his MSc in Computer Science from the University of Twente in 2004. He has a background in Artificial Intelligence with a focus on designing distributed intelligent and adaptive systems. Previously he has applied his research in the domains of crisis management and smart electricity grids. He is currently working as a researcher and software developer at the System and Network Engineering group at the University of Amsterdam. His main activities are carried out in the EU funded GEYSERS project.

Yuri Demchenko received his M.Sc. degree and later Ph.D. from the Kiev Polytechnic Institute, National Technical University of Ukraine. He is a Senior Researcher with the System and Network Engineering (SNE) Research group at the University of Amsterdam. His main research areas include security architecture and distributed authorisation service infrastructure, complex resource provisioning and manageable security services. He is involved in two European projects GÉANT3 and GEYSERS where he takes part in the development of the GEMBus Composable Services middleware and Authentication and Authorisation Infrastructure (AAI) for on-demand Infrastructure Services provisioning. Yuri is actively contributing to OGF, in particular to Infrastructure and Security areas.

Jens Buysse is a Phd. researcher at Ghent University, sponsored by the government agency for Innovation by Science and Technology. He received his Licentiate degree in Computer science in 2007 at Ghent University. His main interests lie in the field of photonic networks, distributed computing, virtualization techniqiues in optical network and Energy Efficient design principles. He participated in one National Project GEISHA and in the European project PHOSPORUS. He is currently involved in GEYSERS, where the goal is to qualify optical infrastructure providers and network operators with a new architecture, to enhance their traditional business operations.

Marc De Leenheer received the M.Sc. and Ph.D. degrees in computer science engineering from Ghent University, Ghent, Belgium, in June 2003 and December 2008, respectively. He holds a postdoctoral fellowship from the Research Foundation - Flanders, and is currently a visiting Postdoctoral Scholar at Stanford University, CA, USA. His research focuses on the design and performance analysis of optical networks and cloud computing systems. He is author or co-author of over 50 publications in peer-reviewed journals or international conference proceedings.

Chris Develder received the M.Sc. degree in computer science engineering and a Ph.D. in electrical engineering from Ghent University (Belgium), in July 1999 and Dec. 2003 respectively. In Oct. 2007 he obtained a part-time, and since Feb. 2010 a fulltime associate professorship at Ghent University. He was and is involved in national and European research projects (IST David, IST Phosphorus, IST E-Photon One, BONE, IST Alpha, IST Geysers, etc.). His research interests include dimensioning, modelling and optimizing optical (Grid) networks and their control and management, smart grids, as well as multimedia and home network software and technologies.

Fabienne Anhalt is currently a research engineer at Lyatiss, working mainly on network virtualisation and software defined networks. She received her MSc in Computer Science from Institut National des Sciences Appliquées de Lyon in 2008 and obtained her PhD funded by Institut National de Recherche en Informatique in 2011 from Ecole Normale Supérieure de Lyon, France. Since 2010, she is actively involved in the EU FP7 GEYSERS project.

Sébastien Soudan co-founded Lyatiss in 2010 and brings his technical expertise in bandwidth sharing, network scheduling, optical networks, network economics and game theory. Prior to founding Lyatiss, Sebastien Soudan graduated from Ecole Centrale de Lyon in 2005 and obtained his PhD in 2009 from Ecole Normale Supérieure de Lyon. He won the Marconi Prize "Young Scholars" in 2009 for his research on dynamic bandwidth provisioning.

Sergi Figuerola is the Director of the Distributed Applications and Networks Area at i2CAT. He is graduated in Telecommunications Engineering by the UPC, and holds a Master in Project Management from La Salle-URL. He is currently the Technical Manager of the FP7 IP GEYSERS project, and participates in the FP7 I3 FEDERICA among other initiatives/projects at national level. He participated as WPL in the FP6 PHOSPHORUS project. From 2000 to 2003, he was involved in research at UPC's Optical Communications Group (GCO), and he is a PhD candidate at the UPC. His main interests are network resource provisioning and IaaS.