Coordinating IT and Network Virtualisation to Provide Infrastructure as a Service: A GEYSERS Approach for the Future Internet

Sergi Figuerola[1,*], Joan Antoni García-Espín[1], Jordi Ferrer Riera[1], Ester López[1], Eduard Escalona[2], Reza Nejabati[2], Shuping Peng[2], and Dimitra Simeonidou[2]

[1] Distributed Applications and Networks Area, i2CAT Foundation
C/ Gran Capità 2, office 203, Nexus 1 building, 08034, Barcelona, Catalonia, Spain
{sergi.figuerola,jage,jordi.ferrer,ester.lopez}@i2cat.net
[2] High Performance Networks Group, University of Essex,
Colchester, CO4 3SQ, United Kingdom
{eescal,rnejab,pengs,dsimeo}@ac.essex.uk

Abstract. In this article we propose the GEYSERS project approach for Future Internet based on provisioning full infrastructures as a service by using IT and transport network resource virtualization. An overview of the project's architecture is provided that shows the layering scheme empowering the definition of Virtual Infrastructure service. The Logical Infrastructure Composition Layer in the architecture allows for operator-tailored planning of infrastructures, which also allows for a dynamic re-planning of the infrastructure depending on service needs. The rest of the article elaborates on the concept, characteristics and variants for modifying Virtual Infrastructures, as well as their implications.

Keywords: Virtualisation, Future Internet, Infrastructure Planning.

1 Introduction

Current network operators are focused mainly on providing and selling network services on top of the infrastructures they own and manage while keeping the resource control and the service provisioning process hidden to the service consumer. However, the evolution of the Internet usage and the exponential growth of its traffic raise critical issues associated with telecom infrastructure design principles and operational models that need to be addressed. Moreover, emerging applications are bringing into the arena new requirements in terms of requesting network and computational resources simultaneously. These new requirements are difficult to accommodate with the existing telecom operational models because telecom companies have full control over the infrastructure [1]. Infrastructure services and virtualisation are key enablers for allocating isolated instances from networking and information technology (IT) devices to different users or applications. As an example, applications such as cloud computing and 3D-video streaming require optimisation and combined provisioning of different infrastructuresincluding both IT and network.

K. Pentikousis et al. (Eds.): MONAMI 2011, LNICST 97, pp. 244–257, 2012.

Emerging trends within the Future Internet research community are developing new architectures that decouple the traditional roles into independent entities in order to lower market entry barriers. The flexibility and dynamism enabled bythe new models provide support to new services, which typically demand high performance and high capacity networks and IT appliances.

On the other hand, these new infrastructure operation models and the related applications provoke that demands become more sporadic and variable. This variability, at the same time, creates the necessity of making infrastructure service provisioning dynamic and manageable. Since physical infrastructures that are interconnected through these new dynamic infrastructure services present a high level of heterogeneity, generalised services are required in order to successfully face all the challenges arising in the Future Internet ecosystem.

With these requirements, the GEYSERS project aims to build a future internet architecture that enables new service paradigms and business models for today's telecom operators, carriers and service providers. The dynamism introduced by GEYSERS enables a flexible partitioning of physical infrastructure resources (network and IT) and dynamic composition of Virtual Infrastructures (VIs), which, in turn, are offered as a service to operators. These operators will be able to operate a virtual infrastructure with an enhanced Network Control Plane (NCP) that facilitates-the provisioning of coupled, optimized and dynamic on-demand 'net+IT' services.

In GEYSERS scenario, ownership and operation are split and assumed by different players: a Physical Infrastructure Provider (PIP) owns the physical infrastructure, which can be partitioned and tailored to different customers depending on their demands; a Virtual Infrastructure Provider (VIP) that acts as a virtualresource broker, and offers to a Virtual Infrastructure Operator (VIO) a customized VI composed by resources coming from PIPs; VIOs will operate these dynamic and adaptable virtual networks that will fit their needs, avoiding both under- and over-provisioning.

GEYSERS infrastructure resources will be delivered as a service, reducing the CapEx related to initial investment in hardware and fixed maintenance fees, and enabling a "pay-as-you-go" model.

2 Related Work

GEYSERS research focus is related to existing initiatives that are taken into account when defining the project's architecture. Some of these initiatives and research projects are the basis for some of GEYSERS requirements and research work.

- **NGN Open Service Environment (OSE):** The NGN reference model, according to ITU-T Y.2011 Recommendation, suggests the separation of the transport network and application services and defines them as NGN service stratum and NGN transport stratum consisting of User plane, Control plane and Management plane. The NGN Y.2012 architecture defines also the Application Network Interface (ANI) that provides an abstraction of the network capabilities and is used as a channel for applications to access network services and resources.

- **Composable Services Lifecycle Management:** The Service Oriented Architecture-based technologies provide a good basis for creating composableservices that, in case of advancing to dynamically re-configurable services rely on the well-defined Services Lifecycle Management (SLM) model. The GEYSERS framework considers dynamic provisioning as a major issue, thus, dynamically provisioned and re-configured services will require re-thinking of existing models and propose new security mechanisms at each stage of the typical provisioning process.
- **4WARD Project:** EU-FP7 4WARD's goal is to make the development of networks and networked applications faster and easier, leading to both more advanced and more affordable communication services. According to 4WARD outcomes,network virtualisation is not only an enabler for the coexistence of multiple architectures, but also provides a path for the migration towards more evolutionary approaches to the Future Internet. In 4WARD's vision, virtualisation can help to keep the Internet evolvable and innovation-friendly, particularly since it can mitigate the need to create broad consensus regarding the deployment of new technologies among the multitude of stakeholders that compose today's Internet. In 4WARD, the goal is to develop a systematic and general approach to network virtualisation. The problem space is divided into three main areas: virtualization of network resources, provisioning of virtual networks and virtualization management. The 4WARD roles model has been used as state of the art for GEYSERS.
- **RESERVOIR Project:** The FP7 RESERVOIR project's goals were defined as to enable massive scale deployment and management of complex IT services across different administrative domains, IT platforms and geographies.The project considered virtualization technologies to transparently provision distributed resources and services on-demand with the specified QoS-based on Service Level Agreement (SLA). Compared to GEYSERS approach, RESERVOIR was limited to server virtualisation, regardless of the transport networks, since it was focused to build the foundations of clouds services.The expertise generated on RESERVOIR for the management of IT resources has been very useful and so considered along the project when creating virtual infrastructures.

3 GEYSERS Architecture

GEYSERS aims at designing and specifying an architecture [3] which main objective is to support dynamic infrastructure services and unified network and IT resource provisioning. This architectureallows opening new business models not only to service providers but also allows infrastructure providers to offer infrastructure services on demand.This is achieved through a novel Logical Infrastructure Composition Layer (LICL) that offers a framework for abstracting, partitioning and composing virtual infrastructures from a set of physical resources in an automated way. On the other hand, service providers will be able to provide advanced dynamic end-to-end services on top of these virtual infrastructures by deploying an enhanced Network

Control Plane (NCP), capable of seamlessly controlling the virtualized network and IT resources.

In terms of infrastructure management, the main objective of the LICL is to provide a mechanism for virtualisation and associated techniques such as uniform resource description, resource abstraction and composition. The LICL relies on a solid resource description framework, which allows applying a common set of procedures and signals to both network and IT resources. This is strengthened by a synchronisation sub-system that ensures resource state information to be coherent between physical and virtual instances of the resources. Moreover, the LICL allows novel planning and re-planning actions to be invoked over the virtual infrastructure in a coordinated action between the VIO and the LICL, for better performance, optimal resource usage and efficient service provisioning.

The provisioning and control of the end-to-end reservation of IT and network resources is performed by the NCP, which provides optimized path computation able to minimize the energy consumption, not only realizing the energy efficient routing at the network level, taking into account a variety of "green TE-parameters" as additional constraints, but also achieving the optimal energy consumption at the IT level by selecting the most efficient IT end-points.

The GEYSERS architecture deliversan overall service provisioning solution that addresses the following requirements:

- **Scalability:** Applications designed for cloud computing need to scale with workload demands so that performance and compliance with service levels remains on target. Moreover, the elasticity concept must be ported to transport networks, in order to allow combined IT and network scalability.
- **Availability:** Through a synchronised combination of virtualisation, control and management techniques applied to network and IT domains.
- **Reliability:** Application data has to be processed properly and delivered with the minim losses.
- **Security:** Applications need to provide access only to authorized, authenticated users, and those users need to be able to trust that their data is secure. GEYSERS ports this approach to infrastructure services provisioning, especially when both IT and network resources are considered.
- **Flexibility and agility:** In the GEYSERS approach, the dynamic virtual infrastructure planning and re-planning processes introduce a new level of flexibility to the virtualization services.
- **Serviceability:** Once a virtual infrastructure is deployed, it needs to be maintained. Proper synchronisation mechanisms must be inherent to the infrastructure provisioning service, which GEYSERS foresees based on mixed synchronous-asynchronous models.
- **Efficiency:** In terms of energy or resource utilisation, GEYSERS proposes a set of techniques for achieving efficient resource and service operation at different layers, ranging from physical infrastructure up to control plane [2].

4 VI Provisioning Service

As abovementioned, GEYSERS introduces a new layer of abstraction and virtualization between the physical layer and the control plane, the LICL. This layer relies on IT and network physical resources to create virtual infrastructures which will be composed of virtual IT resources (at the edges) interconnected by virtual network resources (for connectivity provisioning between IT resources) with a virtual topology based on circuit-switched connectivity. Moreover, virtual resources can be created by aggregation or partitioning of the physical resources, allowing the co-existence of multiple virtual infrastructures over the same physical substrate and thus increasing the resource usage efficiency.

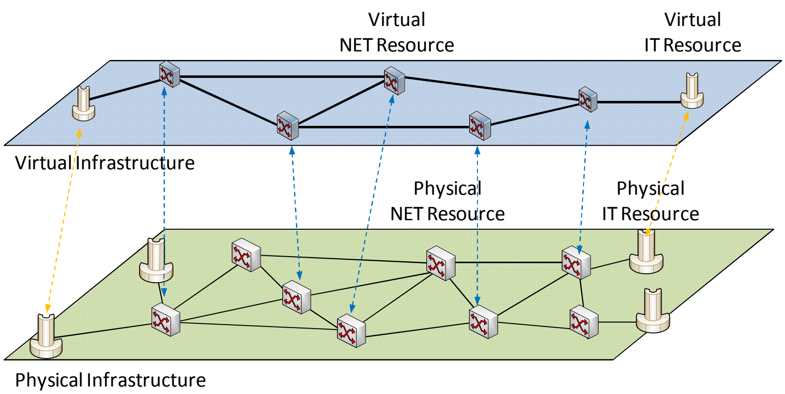

Fig. 1. VI to physical infrastructure mapping

The GEYSERS architecture generically supports any circuit-oriented network, including L1 (e.g. Optical), L2 (e.g. Ethernet) or L2,5 (e.g. MPLS LSP); nevertheless, the concept of virtualisation in GEYSERS focuses in optical L1 switching. The reason is that L2 and L2,5 virtualisation services/mechanisms are widely known and have been developed for many years now, while optical network virtualisation opens a whole new area of research and poses new challenges imposed by the analogue characteristics of optical networks.

In GEYSERS, network resources are considered as optical switching nodes and the links interconnecting these nodes. Ideally, the IT resources would be directly connected to a port of an edge node by an optical link. However, IT resources are typically located in data centres in a separate L2 or L3 network and connected to the core network via L2/L3 switches with optical interfaces. The GEYSERS architecture allows the virtualisation of L2 and L3 devices but for the reasons exposed above and to focus the developments, the support of L2/L3 virtualisation will be associated with the IT resource virtualisation and will rely on the capabilities of the cloud management systems (e.g. OpenNebula, OpenStack).

4.1 Resource Virtualisation

In the context of network and computing infrastructure, virtualisation is the creation of a virtual version of a physical resource (e.g. network, router, switch, optical device or computing server), based on an abstract model of that, which is often achieved by partitioning (slicing) and/or aggregation.

Resource virtualisation is a critical enabler for the LICL that is closely related to the subsequent VI provisioning and operation. Resource virtualisation in LICL can be categorized into four paradigms: aggregation, partitioning, abstraction and transformation, as shown in figure 2, GEYSERS contemplates these four paradigms and how they can be supported for different type of resources, e.g. IT and optical network resources.

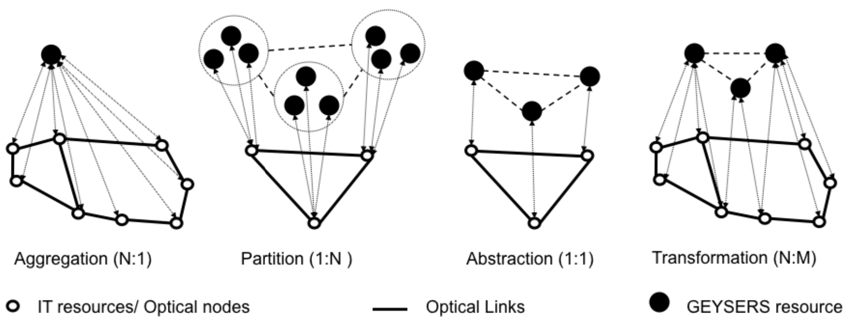

Fig. 2. Virtualisation paradigms considered within GEYSERS

4.2 IT Resource Virtualisation

The IT resources considered in GEYSERS are computing and storage nodes, running user applications, and being interconnected by a virtualized network infrastructure. Users can request such IT resources that are in reality partitions or aggregations of real physical resources.

An IT resource can be partitioned into N virtual IT resources using common virtualisation technologies, such as Xen [4], KVM [5], VMware [6], VServer [7], etc., where each partition is represented as a virtual machine (VM) with computing and storage resources. These technologies use different types of virtualisation to partition the resources. While OS-level virtualisation (e. g. VServer) offers interesting performance [8], it allows only limited isolation and customization. On the contrary, performing emulation and hardware virtualisation (e. g. KVM, Xen), each VM has its own isolated execution environment where any OS can run.

As opposed to partitioning, aggregation consists in exposing a set of physical IT resources as a single virtual IT resource to the user. Such aggregation is for example possible with vSMP [9](Versatile SMP) which aggregates many physical servers and makes them appear to the OS like one giant machine with many cores. Regarding only storage nodes, it is possible to aggregate different disks into a common logical storage pool. This can also be done using SNIA technology, allowing not only sharing a device into several ones, but also to aggregate several physical devices and make them appear as one single virtual device.

4.3 Network Resource Virtualisation

Network virtualisation has brought the concept of server virtualisation to the framework of communication networks. The first attempts for network virtualisation come from the IP routing world (e.g. [10], [11]), where routers are sliced into virtual routers and interconnected by virtual links such as VLANs. However, virtualising optical networks has major differences to such technologies. Optical networking traditionally relies on a strong involvement of manual planning, engineering and operation, due to the fundamental impact of physical impairments in network creation and service provisioning. Even with the widespread of GMPLS-based control planes, optical networks are still manually operated by most network operators. The first attempts to automate such networks are facing the challenge of how to automate the impairment-aware service provisioning and integrate it with the concept of dynamic control planes and zero-touch networking. Thus, optical network virtualisation has to address all these challenges inherent to optical networking.

Optical network virtualisation is the creation of virtual instances of optical network resources the behaviour of which is the same to that of their corresponding physical optical network resources. It relies on the abstraction of heterogeneous network resources, including nodes, links and segments comprising both nodes and links. Optical network virtualisation techniques depend on the type of optical element to be virtualized and should enable the representation of the virtual optical resources inheriting the critical characteristics of the physical ones.

In GEYSERS, the basic optical elements to be considered are optical nodes and optical links. Each virtual instance of an optical node has its own ports and switching capability and the separation and isolation between the control of each virtual instance depends on the virtualisation capabilities of the device itself. Regarding optical link virtualization, it consists of abstracting optical data links as virtual instances by partitioning or aggregation. The partitioning of optical data links is introduced by dividing the link capacity into smaller units, resulting in the granularities of sub-wavelength and wavelength while the aggregation results in a granularity of waveband (or fibre or group of fibres). Optical fibre partitioning is easily achieved in DWDM where the optical links (i.e. fibres) can be inherently split into individual wavelength channels. Highest bandwidth granularity allowing for more efficient bandwidth utilisation can be achieved by having access to even lower bandwidth units at the sub-wavelength level. The virtualization capability of a link is related to the optical port characteristics of the associated optical node.

4.4 Service Delivery

The GEYSERS VI Provisioning service consists of several phases that include both automated and engineer/human assisted procedures (figure 3).

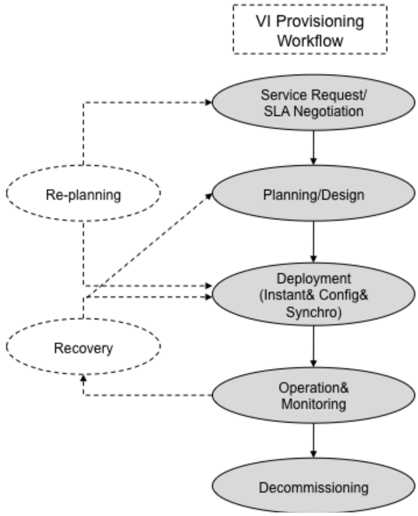

Fig. 3. GEYSERS Virtual Infrastructure provisioning workflow

In the figure, circles represent the different stages composing the whole workflow and arrows represent transitions between the different phases. The VI service provisioningrequest starts with a Service Request/SLA Negotiation. In this phase, the VIO defines the requirements for the desired VI and initiates the SLA negotiation with the VIP. The SLA defined in this phase provides a set of basic requirements, which includesQoS requirements, security policies and robustness requirements among some others.The security policies defined in this phase will be used in the planning, deployment and operation phases. Additionally, SLAs may also contain trust anchors in a form of public key certificates. It is an initiation point of the VI lifecycle.

Just after the service request phase, we identify the Planning/Design phase. This service can be decoupled into the following sub-phases:

- Virtual Infrastructure design: The Virtual Infrastructure design is carried out by a VIP based on the requirements received from the service request and SLA negotiation phase. These requirements are in form of SLAs that are expected to be fulfilled for the service provisioning. These SLAs are decomposed into more technical constraints in the SML in a semi-automated fashion. The SML is a rule-based expert system which can also support a human interaction for planning. Fully automated planning schemes are also considered in GEYSERS SML architecture.
- Virtual resources selection and composition: In this sub-phase, the VIP searches/negotiates the virtual resources offered by one or multiple PIPs. Dynamic Algorithms are used to composing IT and network resources and to produce as a result, a blueprint of the virtual network is available, and ready to be included in a contract between the VIO and the VIP.

- Virtual resource reservation: In this sub-phase, each selected virtual resource is associated with a common reservation ID (to be hereafter referred to as Global Reservation ID (GRI)) that also binds the reservation session/instance with the SLA initiated at the provisioning process. The reserved resources need to be configured and initiated in the deployment phase.

Once the Planning/Design phase is finished, the process enters into the deployment phase**.** During the VI deployment phase, the reserved infrastructure instances are instantiated, configured, registered, and initialised. This phase should allow the review and approval from the network/IT engineers.The deployment phase can also be decoupled into the following sub-phases:

- NCP and IT controller'sinstantiation and deployment: in this sub-phase, the VIO deploys its NCP by taking consideration of the VI specifications. The software modules are then deployed/installed to control the different virtual network nodes composing the VI. Similarly the ITcontrollers related to the IT virtual resources are deployed/installed.
- Configuration of the NCP and IT controllers: The network controllers and PCE modules deployed in the NCP are configured with the network topology information and policies. Similarly, the IT controllers are configured with information about virtual IT resource availability and properties.
- NCP initialization: The NCP modules are initialized and started, and the network auto-configuration process takes place (e.g. neighbour/UNI discovery, initial flooding of TE parameters and routing protocol convergence). In this phase, the IT controller also injects the capabilities of the IT site under its control into the NCP.
- Instant Network+IT service/infrastructure registration and initialisation: this sub-phase allows a new service to be registered in the VIO and put into operation. It also allows binding security and provisioning sessions with the service ID and (underlying/implementing) platform runtime environment. The importance of specifying this phase is defined by the need to address such scenarios as infrastructure re-planning and failure restoration.

As a result of the instantiation phase, the VIO has configured the virtual resources and has deployed its control plane over the virtual infrastructure. The virtual infrastructure is up and running. At this point we enter into the Operation and Monitoring phase.

At this point we enter the Operation and Monitoring phase, whichincludes all the processes for the provisioning of network + IT services (NIPS) to users. During the operation phase, the VIO runs its own virtual infrastructure provisioning service that is targeted to deliver the necessary infrastructure resources (both network and IT) to users, project or applications. It is intended that this provisioning process is automated and allows using the same business model as traditional physical operators although behaving under different roles, depending on the model role. The on-demand service provisioning happens in this phase.

Along the whole provisioning service the Re-planning and Recoveryphases may take place. These are additional phases triggered by special events during operation, or on the request process of any of actors. Re-planning is a special VI stage in which the LICL implements a change in the VI. This phase is further detailed in the next section.. A recoveryphase/process takes place when the running virtual/provisioned service fails (e.g. because of hardware failure).

Depending on the type of failure, restoration may require just restarting/re-deploying the virtual service or will involve new planning/design/reservation processes.

The last phase of the VI provisioning service is the Decommissioning, which is triggered whenever a virtual infrastructure is no longer in operation and must be terminated. This usually happens when the leasing contract between VIP and VIO ends and the VI is no more suitable for other VIO customers of the VIP. The termination phase ensures that all the authorization right of the VIP for access to the PIP resources are inactivated as well as the authorization right of the VIO for access to the virtualized physical resources. Once a VI is decommissioned, the physical resources of the PIPs become available for planning and instantiation of new VI.

In essence, the VI provisioning service consists in creating virtual infrastructures upon request and in on-demand basis. From a business perspective the VI provisioning service involves the participation of several of the mentioned GEYSERS roles. In GEYSERS we consider the virtual infrastructure operator (VIO) as the entity generating the request. Nevertheless, anyone in need of a virtual infrastructure can issue a VI request (e.g. application provider).shows the most basic workflow diagram for the VI provisioning service. The service starts with the VIO requesting the creation of a VI to a virtual infrastructure provider (VIP). The VIP processes the requests and interacts with the Physical Infrastructure Provider (PIP) to request the creation of virtual instances of the physical resources (by partitioning or aggregation). Once the required virtual resources (VR) have been created, the VIP uses them to compose a virtual infrastructure and offer it to the VIO.

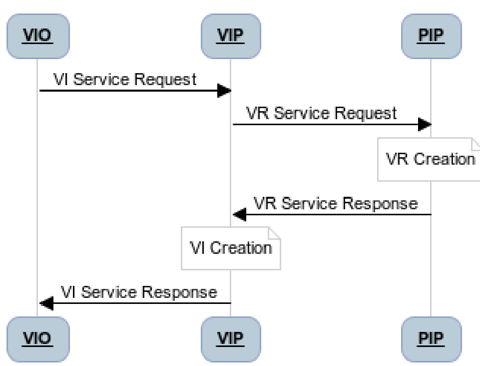

Fig. 4. Service provisioning from a business perspective

During this process, negotiation between the different roles is required when different actors carry them out.

5 VI Re-planning

Re-planning is a VI provisioning phase that is triggered by special events or upon request by a GEYSERS actor during the operation phase of an already planned and provisioned VI exists. Therefore, in the scope of GEYSERS, VI re-planning is a stage in which the LICL will be responsible for implementing the required changes in the virtual infrastructure. The requirement for VI re-planning can be triggered at any time once the virtual infrastructure has been instantiated.

In the scope of GEYSERS the VI re-planning procedure may involve a modification of a network infrastructure, IT infrastructure or both network and IT infrastructures. When the procedure describes modifications applied to the network infrastructures, it is referred to as VI network re-planning. Analogously, when the procedure affects IT infrastructures only, it is referred to as VI IT re-planning.

From an operator perspective, infrastructure re-planning usually takes place in long timescales and is usually human-driven, although supported by dimensioning tools, as in some cases it may involve the investment of CapEx for the acquisition of new physical equipment and its installation in the infrastructure. However, in the context of GEYSERS, it is relevant to consider a VI re-planning that takes place in short timescales, useful for infrastructure operators to adapt their infrastructures dynamically with the aim to improve the efficiency of resource utilisation and increase service availability for end users.

We must also take into account how the whole process of modification of the existing virtual infrastructure should be performed. It has been identified as a key functionality for the VI re-planning to be performed in short timescales and be supported by a dynamic procedure for the implementation of changes in the virtual infrastructure. In this article this procedure is referred to as automatic VI re-planning. The automatic procedure allows applying the required changes immediately, with a limited or preferably no human intervention with supporting tools. When the process of VI re-planning takes place in long timescales, there is no need from an infrastructure operator point of view to have full automation of the whole VI re-planning procedure. This procedure is named manual VI re-planning.

5.1 VI Re-planning Requests

Following the initial VI planning request, the VI re-planning requests can vary with regards to the amount of specific information that is provided as part of the request itself. In this context, VI planning requests can be generated by the VIO and communicated to the PIP through the VIP. It is important to note that to perform VI planning and hence re-planning the VIP has to rely on the physical resources information provided by the PIP, being the entity that has full knowledge of the physical infrastructure. The types of re-planning requests are:

- Service-driven VI request: This type of request allows the VIO to request a VI with the maximum degree of flexibility. The information provided as part of the request involves the prediction of the volume and type of infrastructure expected to support the required services, as well as other service related information as determined by the associated SLAs, including availability etc.
- Constrained VI request: This type of request follows in general terms the "service driven VI request" described above, but may also impose some additional constraints associated e.g. with the location of some or all involved IT or network resources in the space of an area/country/continent or possible energy consumption requirements.
- Specific VI request: This type of VI request will include specific information regarding the IT resource requirements and processing capability in addition to the usual service specific requirements including availability etc. Taking into consideration the above, the "specific VI request" will therefore result in a request for a specific virtual topology that is already capacitated with regards to the IT and network resources required.

In the specific VI requests, the virtual topologies of the VI need to be indicated, including the virtual nodes and the virtual links. The virtual nodes are partitioned from a single physical node in GEYSERS, and the geo-location can also be specified to get the location of the physical node to be mapped. Finding the optimum mapping between the VRs and the PRs is part of the VI planning that applies optimisation with specific objectives.

5.2 Service Delivery

The manual VI re-planning workflow is shown in figure 5. The whole procedure involves the three roles: VIO, VIP and PIP. Once the VIO issues a request for VI re-planning, the VIP checks the consistency of the request. Once the request is positively validated, the VIP sends a confirmation (VI Re-planning response) back to the VIO. In the next step the VIP initiates the procedure of implementation of the request in the VI to identify a list of virtual resources (VRs) capable of satisfying the VIO request. Once the list of VRs is ready, the VIP issues the request to the PIP to instantiate them and attach to the existing VI. Once the VI is changed, the notification is sent to the VIO and VIO instantiates (or re-configures) relevant NCP and VITM controllers (located at the SML), to manage the network and IT resources, respectively. During the initialization of the new NCP controllers, the status of particular virtual resources is exchanged between VIO and VIP, and finally, based on received information, the NCP controllers are synchronized with relevant information.

In GEYSERS, the manual VI re-planning request is considered a management operation. Thus, the functional module within the VIO responsible for the manual re-planning is located at the SML.

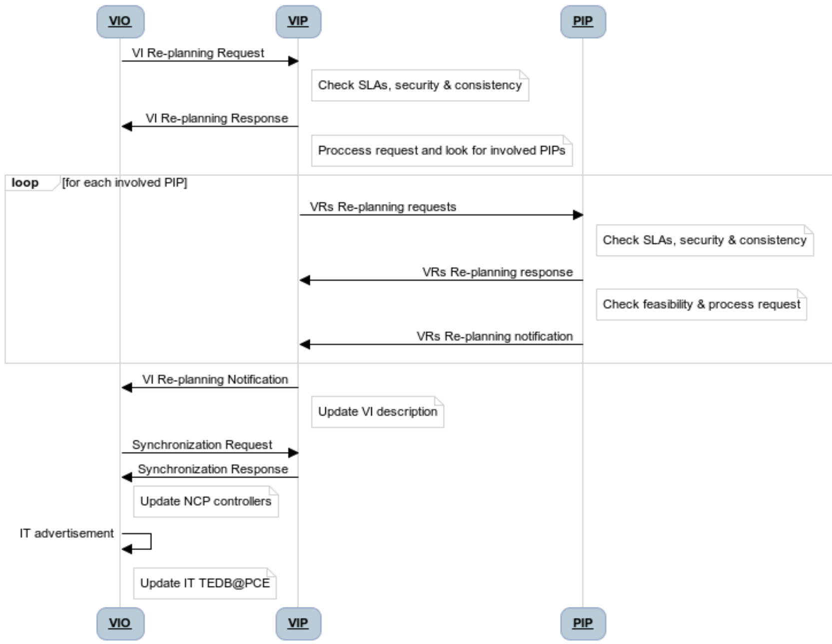

Fig. 5. Manual VI re-planning workflow

6 Conclusion

Some common goals when conceiving Service oriented architectures lead to the delivery of business services with an integrated IT strategy supported by a set of linked services and information systems. In this architecture, infrastructure services are of paramount importance not only for the IT resources but also for the network resources required to interconnect them. Infrastructure services allow the possibility of leasing physical resources, releasing the burden of having to purchase physical infrastructure for application providers.

The GEYSERS project proposes and implements an architecture which main objective is to support dynamic infrastructure services and unified network and IT resource provisioning based on the virtualization of infrastructure resources. The dynamicity and flexibility of the proposed architecture allow providing specific virtual infrastructures on demand while promoting the emergence of new business roles such as the virtual infrastructure provider. This architecture also enables the online modification of the virtual infrastructure through an on demand re-planning service.

Acknowledgments. This work has been supported by the FP7 GEYSERS through the European Commission under Grant Agreement no. 248657.

References

1. Figuerola, S., Lemay, M.: Infrastructure Services for Optical Networks [Invited]. Journal of Optical Communications Networks (JOCN) 1(2), 247–257 (2009)
2. Tzanakaki, A., Anastasopoulos, M., Georgakilas, K., Buysse, J., Leenheer, M.D., Develder, C., et al.: Energy Efficiency in integrated IT and Optical Network Infrastructures: The GEYSERS approach. In: INFOCOM 2011 Workshop on "Green Communications and Networking", Shanghai, China (2011)
3. Escalona, E., Peng, S., Nejabati, R., Simeonidou, D., Garcia-Espin, J.A., Ferrer, J., Figuerola, S., Landi, G., Ciulli, N., Jimenez, J., Belter, B., Demchenko, Y., de Laat, C., Chen, X., Yukan, A., Soudan, S., Vicat-Blanc, P., Buysse, J., De Leenheer, M., Develder, C., Tzanakaki, A., Robinson, P., Brogle, M., Michael Bohnert, T.: GEYSERS: A Novel Architecture for Virtualization and Co-Provisioning of Dynamic Optical Networks and IT Services. Submitted to Future Network and Mobile Summit (2011)
4. Xen hypervisor, `http://www.xen.org`
5. Kernel-based Virtual Machine, `http://www.linux-kvm.org`
6. VMware, `http://www.vmware.com`
7. VServer, `http://linux-vserver.org`
8. Bavier, A., Bowman, M., Chun, B., Culler, D., Karlin, S., Muir, S., Peterson, L., Roscoe, T., Spalink, T., Wawrzoniak, M.: Operating system support for planetary-scale network services. In: Proceedings of the 1st Conference on Symposium on Networked Systems Design and Implementation, San Francisco, California, March 29-31, vol. 1, p. 19. USENIX Association, Berkeley (2004)
9. vSMP: versatile SMP, `http://www.scalemp.com/smp`
10. EU FP7 4WARD project, `http://www.4ward-project.eu`
11. EU FP7 Mantychore project, `http://www.mantychore.eu`