GENERALISED ARCHITECTURE FOR DYNAMIC INFRASTRUCTURE SERVICES

Large Scale Integrated Project
Co-funded by the European Commission within the Seventh Framework Programme
Grant Agreement no. 248657
Strategic objective: The Network of the Future (ICT-2009.1.1)

Start date of project: January 1st, 2010 (36 months duration)

Deliverable D2.2 (update)
GEYSERS overall architecture & interfaces specification and service provisioning workflow

Deliverable leader: University of Essex

Author list: Eduard Escalona (UEssex), Reza Nejabati (UEssex), Javier Jimenez (TID), Anna Tzanakaki (AIT), Kostas Georgakilas (AIT), Markos Anastasopoulos (AIT), Florian-Alexandru Antonescu (SAP), Giada Landi (NXW), Joan A. Garcia-Espin (i2CAT), Sergi Figuerola (i2CAT)

Dissemination Level

☒	**PU:**	Public
☐	**PP:**	Restricted to other programme participants (including the Commission Services)
☐	**RE:**	Restricted to a group specified by the consortium (including the Commission Services)
☐	**CO:**	Confidential, only for members of the consortium (including the Commission Services)

Project:	GEYSERS (Grant Agr. No. 248657)
Deliverable Number:	D2.2 update
Date of Issue:	14/05/11

Abstract

This document is a complement to deliverable D2.2 to provide updates to specific topics relevant to the overall architecture of GEYSERS. The main aspects included in this report are a refined version of the architectural requirements and a detailed description of how GEYSERS addresses each of them. Specifically, optical network virtualisation characteristics, re-planning mechanisms, a complete specification of the interfaces and workflows, security framework and an analysis of multi-domain issues are provided.

Table of Contents

Figure Summary

Table Summary

0 Executive Summary

This document extends and complements deliverable D2.2 [2] to provide a simpler and evolutionary view of the GEYSERS architecture, globalizing the set of functionalities presented in a more comprehensive way.

Section 1 presents the GEYSERS architecture as an evolution of current Next Generation Network (NGN) architectures, showing its two main supported services with basic usecase workflows and functionalities while Section 2 describes the energy efficiency approach tackled by GEYSERS and the solutions proposed for optimized energy consumption at both the virtual infrastructure composition and provisioning level.

1 GEYSERS Architecture

Service oriented architectures can be conceived in many different ways, since it is difficult to find an agreed definition. However, some common goals 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.

GEYSERS aims to design and specify an architecture which main objective is to support dynamic infrastructure services and unified network and IT resource provisioning. The roadmap evolution of the architecture during the lifetime of the project involves several phases:

1. Initial specification of the architecture: The initial GEYSERS architecture conceived in deliverable D2.1 [1] presents a basic layered structure with its main functionalities and interfaces and some possible scenarios.
2. Overall architecture specification: The next step in the GEYSERS architecture evolution has been delivered in deliverable D2.2 [2] and is extended by this document. It comprises a more comprehensive view of GEYSERS supported services including the overall service workflows, detailed interface specifications and the GEYSERS approach for security management. Specifically, this document describes the architecture in a generic way for simple case scenarios and presents an evolutionary approach from the Next Generation Networks (NGN) architecture (Section 1.1).
3. Refined architecture specification: The architecture will be refined and evolved to support the different multi-domain case scenarios identified in D2.2 [2] and after feedback received from other WP2 tasks (T2.5 and T2.6) and from WP3 and WP4 developments. These refinements will be presented in next deliverables (D2.6 and D2.7).

1.1 Infrastructure Virtualisation

The GEYSERS architecture introduces a new layer of abstraction and virtualization between the physical layer and the control plane (Figure 1), the LICL.

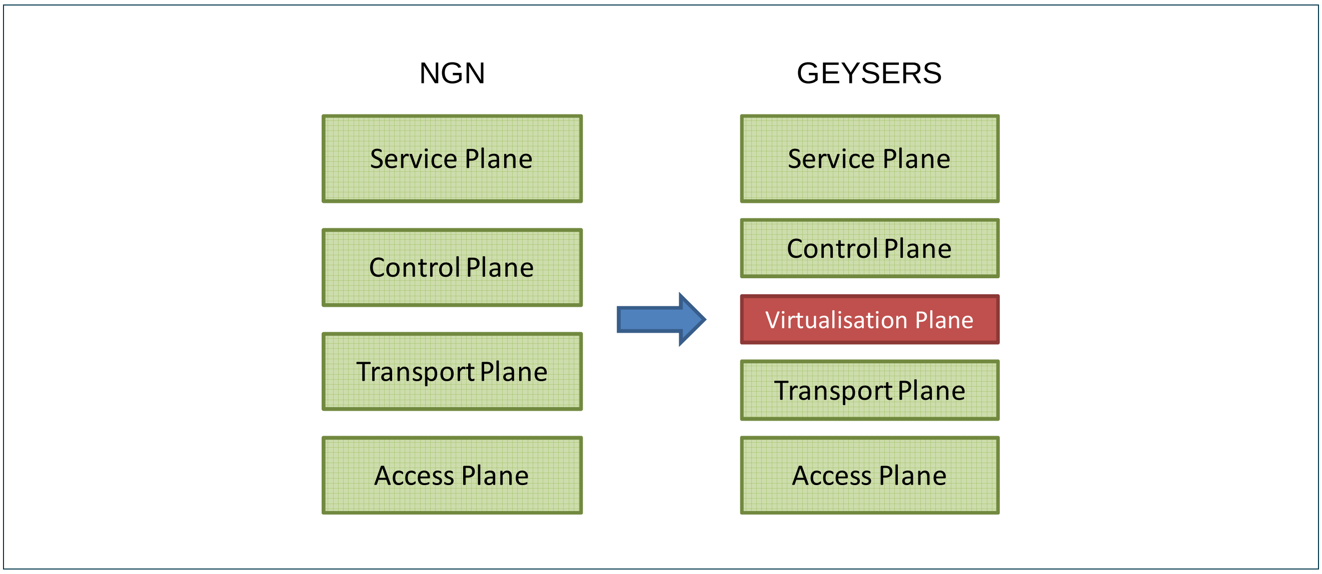

Figure 1: GEYSERS virtualisation layer

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 between IT resources) (Figure 2) 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.

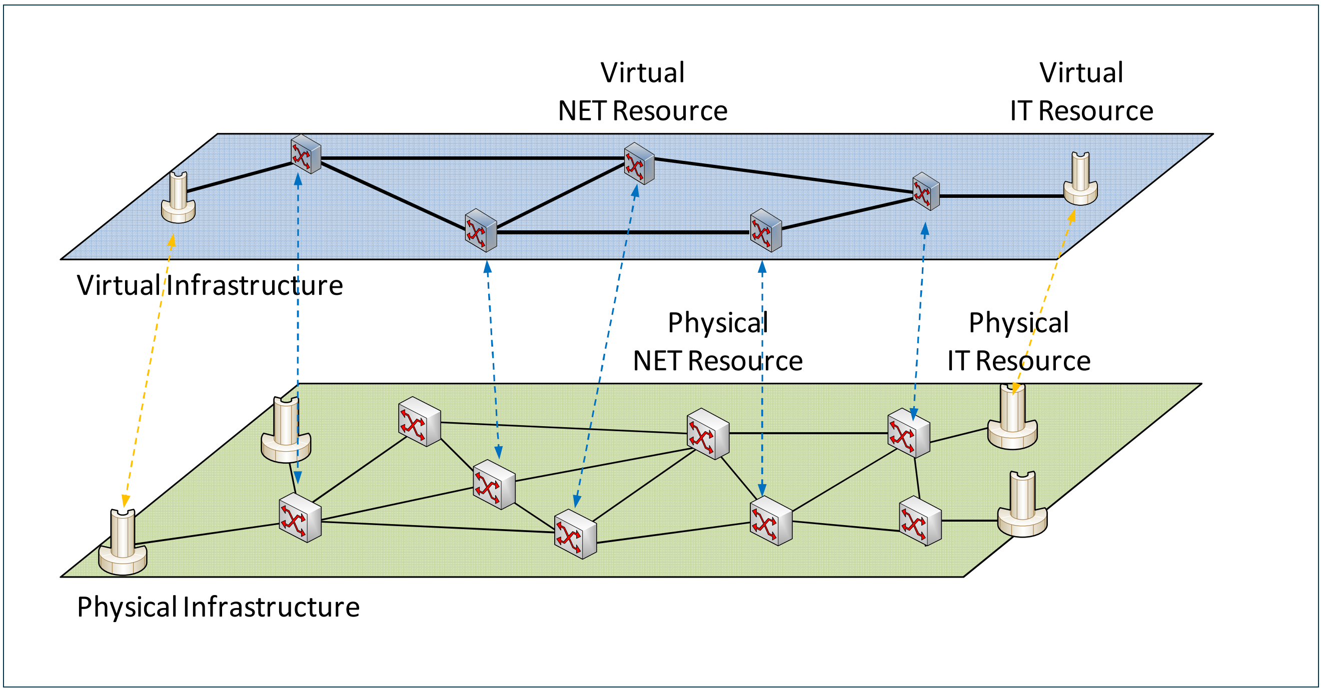

Figure 2: Infrastructure virtualisation

The GEYSERS architecture generically supports any circuit oriented network, including L1 (e.g. Optical), L2 (e.g. Ethernet) or L3 (e.g. MPLS LSP); nevertheless, the concept of virtualisation in GEYSERS focuses in optical L1 switching. The reason is that L2 and L3 virtualisation 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 physical 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).

1.2 Virtual Infrastructure Control

The next step in the GEYSERS architecture is to effectively provision the network and IT resources over the created virtual infrastructures. GEYSERS offers the possibility to deploy any control plane or management system since the LICL exposes the virtual infrastructure as if it was a physical one.

Typically, IT and network resources are provisioned separately using different control and management systems. However, GEYSERS deploys an enhanced Network Control Plane that implements most of the functionalities required for the co-allocation and combined provisioning of IT and network resources, including advertisement of IT resource information and selection of IT resources (Figure 3).

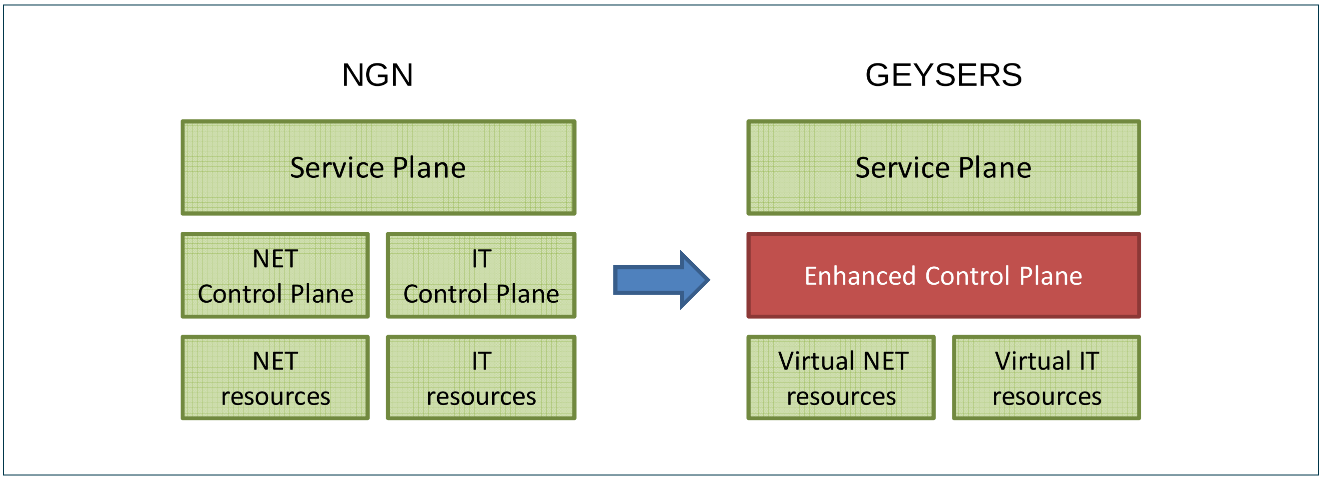

Figure 3: GEYSERS enhanced control plane

The only functionality not implemented by the GEYSERS control plane is the actual configuration of the IT resources since already existing systems can be leveraged and used for this purpose in collaboration with the control plane.

1.3 Services offered by the GEYSERS architecture

The GEYSERS architecture as described in D2.1 [1] offers two main services, the virtual infrastructures (VI) provisioning service and the network and IT provisioning service (NIPS). These two services are independent in execution but dependent in operation since the NIPS service can just be offered over already created virtual infrastructures.

1.3.1 VI provisioning service

The VI provisioning service supported by the GEYSERS architecture consists in creating virtual infrastructures upon request. 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, researchers ...). The triggering of a VI request is considered a management operation, thus, independently of its origin, it is included as one of the functionalities of the vertical management layer described in D2.2 Section 4.1 [2], and will be specifically implemented in the Service middleware layer (SML).

The VI provisioning service can be described from a business perspective or from an architectural perspective. The former should consider the interaction between the different roles involved in the VI provisioning while the latter considers the interaction between the different involved layers. The workflows described next are a simple version of the workflow presented in D2.2 Section 4.2.2, [2] simplified here for the sake of clarity.

1.3.1.1 *Business perspective*

D2.2 Section 3 [2] described generically the atomic services that constitute the interaction between the different roles in GEYSERS. However, the VI provisioning service includes several of the mentioned services and roles. Figure 4 shows the most basic workflow diagram for the VI provisioning service. The service starts with a VIO requesting the creation of a virtual infrastructure 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.

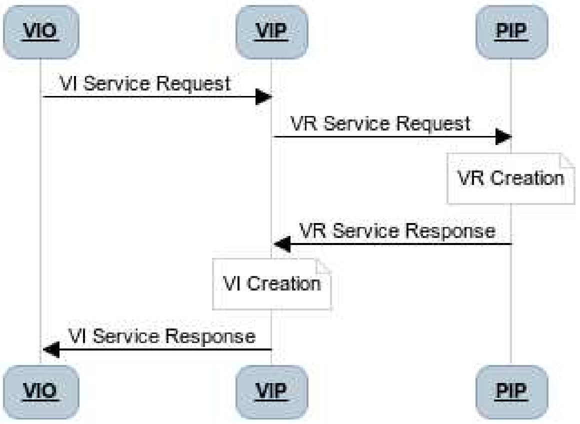

Figure 4: VI provisioning service (business perspective)

During this process, negotiation between the different roles is required when they are carried out by different actors.

1.3.1.2 *Architectural perspective*

Architecturally, the layers involved in the VI provisioning are the Physical Infrastructure Layer, the LICL and the Service middleware Management System (SMS) of the vertical management layer located within the SML. Upon a VI Service request from the SML, the LICL performs the required operations to gather information about the physical resources, virtualise them and construct the virtual infrastructure.

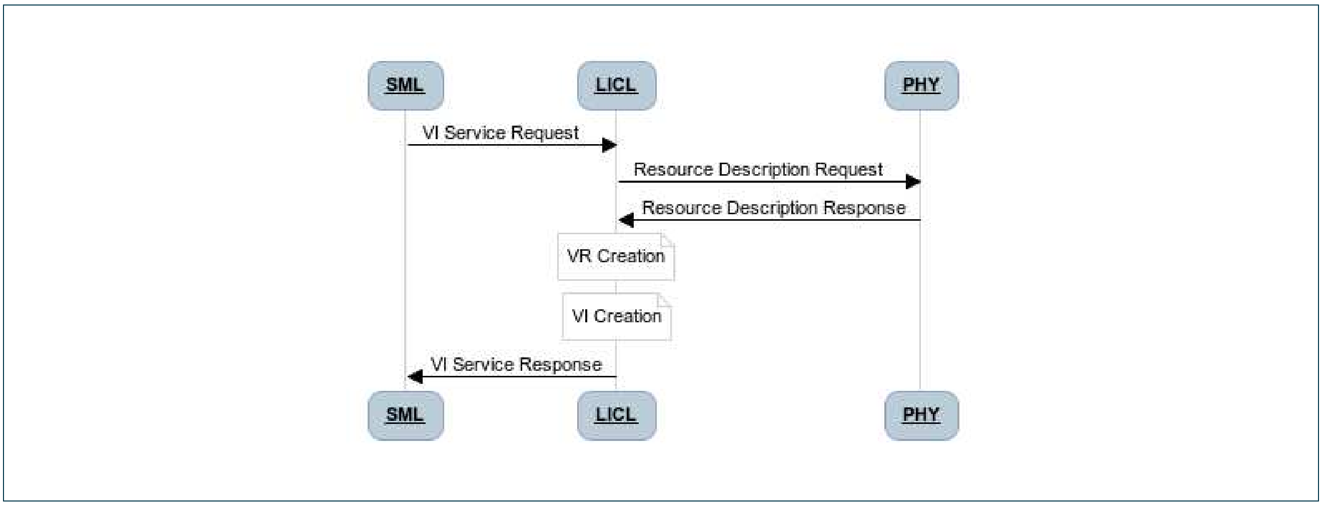

Figure 5: VI provisioning service (architectural perspective)

1.3.2 Network and IT provisioning Service (NIPS)

The second main service offered by the GEYSERS architecture is the NIPS service, which involves the selection and configuration of virtual resources over a virtual infrastructure to provide end-to-end IT resource provisioning including connectivity.

A pre-condition to enable the NIPS service is the deployment and configuration of the NCP controllers by the VIO according to the provided VI. This step is a management operation which in GEYSERS implementation will be performed manually.

1.3.2.1 *Business perspective*

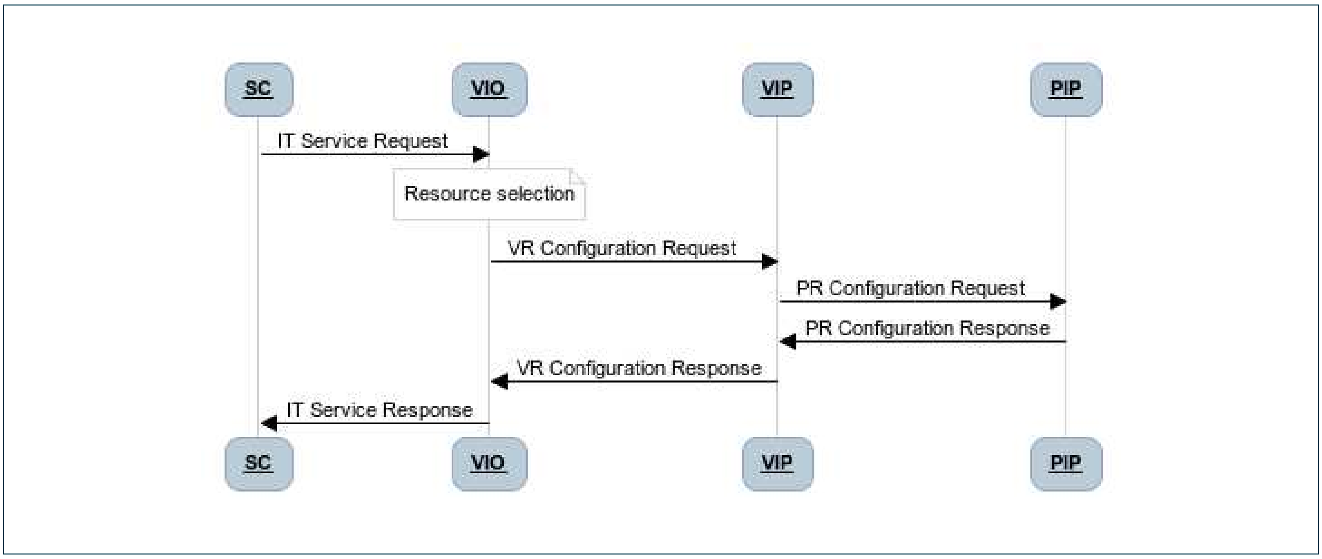

Figure 6: Network and IT provisioning Service (business perspective)

1.3.2.2 *Architectural perspective*

The resource selection at the VIO combines both network and IT resources and is performed on-demand depending on the specification of the requested IT service. This selection requires the cooperation between the SML and the NCP+ and, in the most advanced scenarios, can be performed jointly at the NCP+ in order to optimize the usage of the overall set of heterogeneous resources. The basic scenario is depicted in Figure 7, where the selection of the IT end-points and the network path between them is performed separately following the traditional unicast model. In the first step the SML chooses the IT end-points to be used for the application and asks the NCP+ to dynamically establish a network connectivity service between them. Based on the results of the network path computation, the NCP+ interacts with the LICL requesting the configuration of the virtual optical resources. Once the end-to-end path is established, the SML interacts with the LICL in order to configure the IT resources.

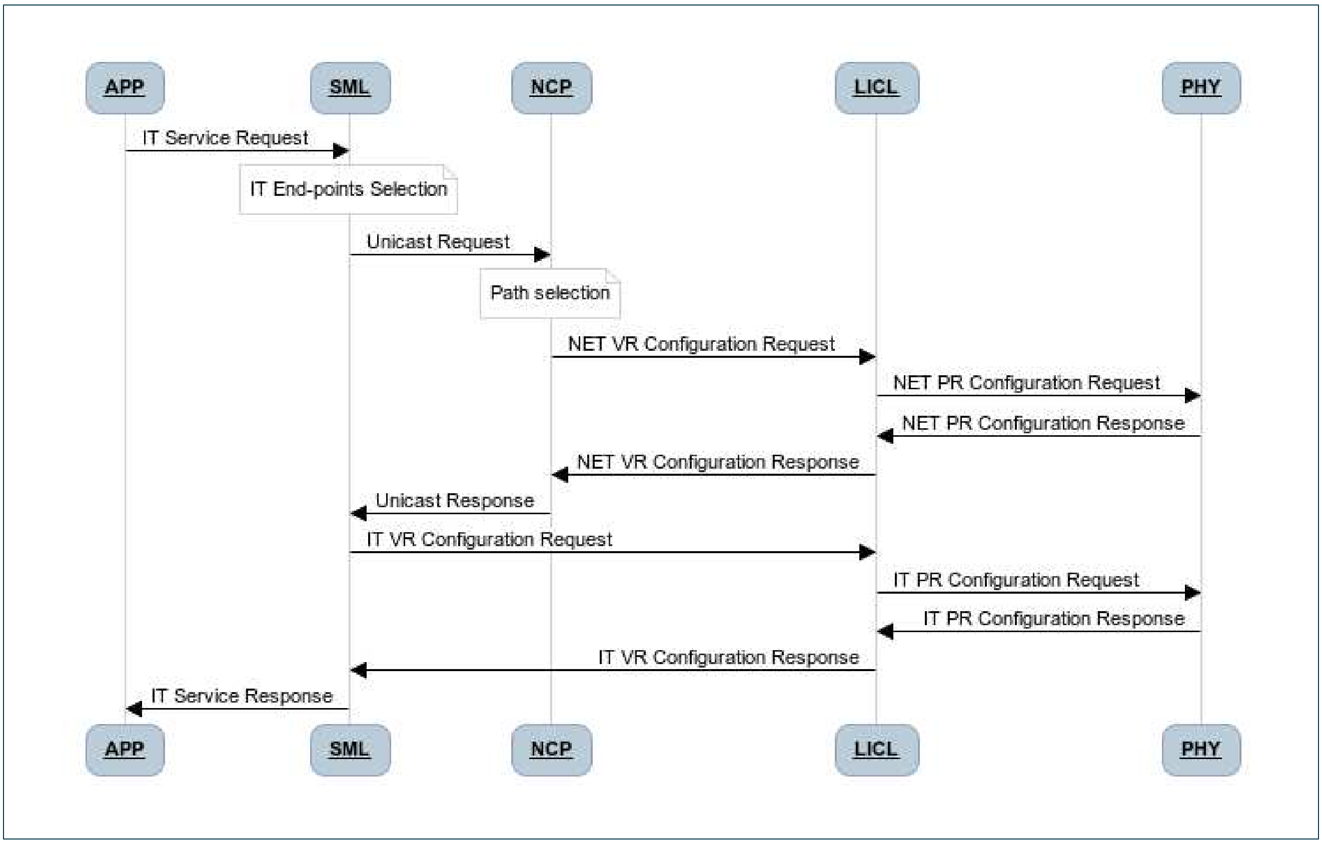

Figure 7: Network and IT provisioning Service – Unicast (architectural perspective)

In the full anycast scenario (Figure 8), the choice of the IT end-points is completely delegated to the NCP+. In the anycast request, the SML specifies the requirements for the IT end-points, in terms of e.g. amount of available storage or CPU, and the NCP+ runs some enhanced routing algorithms to jointly compute the best combination of IT end-points and network path. Such algorithms are able to take into account constraints for both the network (e.g. required bandwidth) and the IT resources (e.g. required computing power, memory or storage). Once this computation is performed, the selected IT end-points are notified to the SML. It should be noted that SML and NCP+ remain responsible for the configuration of the resources in their own segment (IT and net respectively). This choice is flexible enough to support different business models where the VIO role can be covered by a single actor or split between two different actors (VIO-IT and VIO-N), each of them able to control a specific part of the VI (VI-IT and VI-N respectively) through the mechanisms exposed by the LICL.

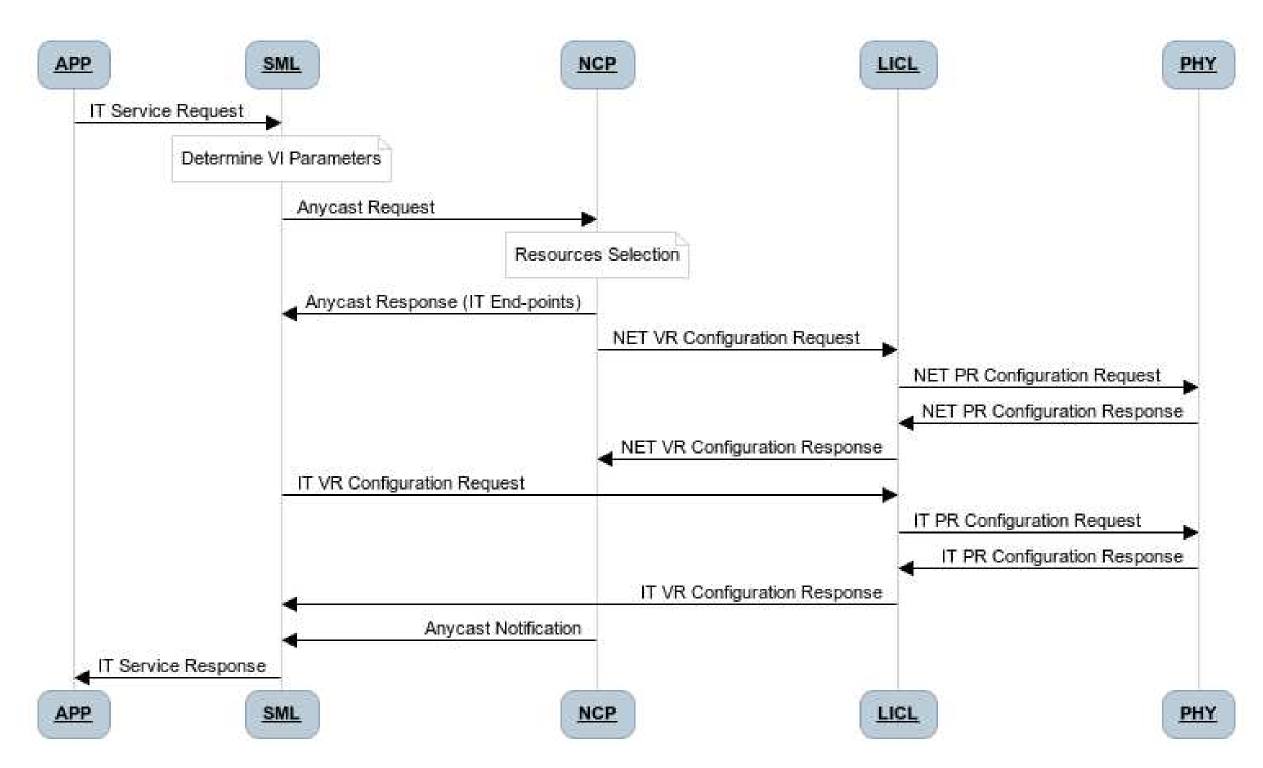

Figure 8: Network and IT provisioning Service – Anycast (architectural perspective)

As an example, an application can request a Virtual Desktop Infrastructure (IT Service Request) that should support a given number of users (e.g., 100). This request is translated at the SML into virtual resource requirements (e.g., number of virtual cores, amount of memory and storage, as well as total network bandwidth and maximum latency) and then sent to the NCP as an Anycast request. After receiving the IT endpoints, SML will begin configuring Virtual Machines in order to configure requested Virtual Desktop Infrastructure.

The previous example is described in the following diagram (Figure 9), emphasizing the interactions between the Service Consumer SC, SML, NCP and LICL. Here, the request coming from SC and response of SML are shown. The request will be used by SML to look up its knowledge-base that consists of:

- Concrete object-type requirements. For example, VDI would be registered as an asset in the SML, where its min CPU, MEM and Network requirements are provided.
- Target-Type information that narrows the interpretation of “Cloud” as a target type given the Object type. In this case it can be determined that “Cloud” refers to a provider of Virtual Machine instances.
- VM image and binary repository for performing the actual installation of the VDI.

- Scripts for invoking the LICL API with a set of parameters including CPU, MEM, Storage and NIC.

The VR request does not include information that a VDI has been requested by the SC, but only parameters specific to VR (in this case VM), creation and configuration details.

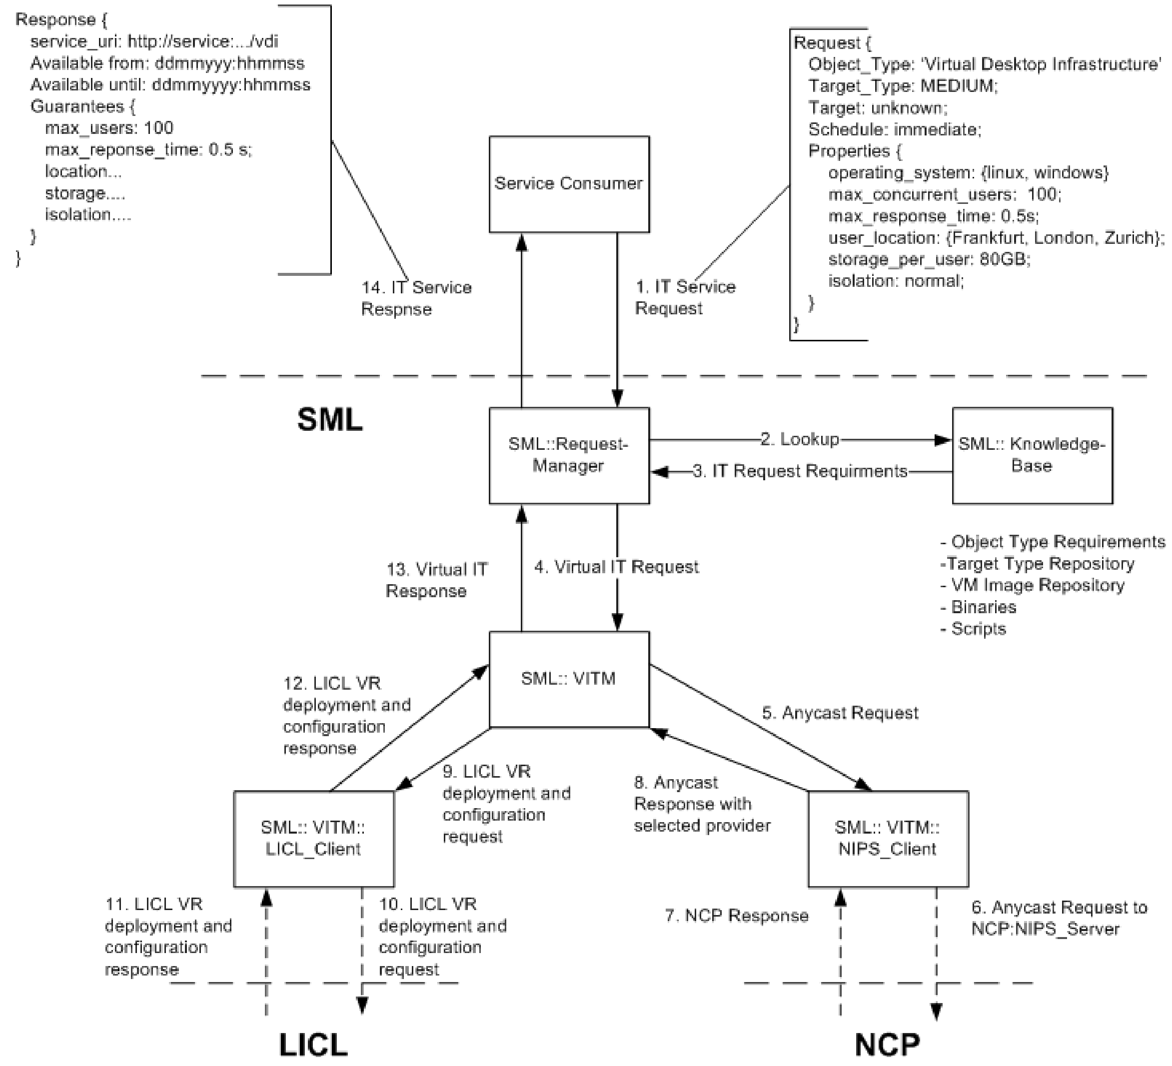

Figure 9: SML interactions in anycast scenario

1.3.3 VI re-planning service

The VI re-planning service allows the modification of already created virtual infrastructures during its operation phase. It consists in adding or removing virtual IT and network nodes or links without disrupting the provisioning services running in the virtual infrastructure. The different re-planning aspects and modes considered by GEYSERS have been detailed in deliverable D2.2 [2].

2 GEYSERS energy efficiency approach

In GEYSERS, energy efficiency is targeted, in an integrated IT and optical network infrastructure, taking a multilayer approach. GEYSERS is inherently addressing the issue of energy efficiency given its focus on optical network technology that has been proven to be significantly more energy efficient than alternative network technologies. However, additional energy efficiency considerations are addressed in both the VI planning and the VI operation/service provisioning phases taking into consideration the combined energy consumption of both optical network and IT resources.

Regarding the VI planning phase the responsible layer is the LICL. The LICL defines and implements a resource information model to uniquely identify and uniformly abstract physical resources, including energy efficiency properties of physical infrastructure elements. The objective of VI planning is to implement a dynamically reconfigurable virtual network that not only meets customer's specific needs, but also satisfies the virtual infrastructure provider's-driven requirements, while maintaining cost-effectiveness and in our case the specific requirement of energy efficiency. Through this process the least energy consuming VI that can support the required services is identified in terms of both topology and resources. To identify this least energy consuming VI the detailed power consumption models and figures of the underlying physical infrastructure, including joint consideration of optical network and IT resources, are taken into consideration. Mapping the virtual resources to the physical resources is also part of the VI planning phase. Therefore the VI planning phase is also responsible to define the energy consumption parameters of the VI itself. Through the VI planning process, virtual resources will be effectively abstracted from the physical devices and will be marked with "green parameters". Using these parameters the virtual infrastructures can be operated in an energy-efficient manner, through the implementation of energy-aware provisioning mechanisms at the NCP+ layer.

The GEYSERS NCP+, based on the ASON/GMPLS and Path Computation Element (PCE) architectures and protocols, is aiming at playing a key role in the operational energy efficiency of VI infrastructures. This can be achieved as the NCP+ is responsible for the path computation, required as part of the service provisioning, and therefore the corresponding resource allocation of both network and IT resources. In the effort to maximize the energy efficiency of operating VIs, the NCP+ can compute end-to-end paths with an objective different to these considered in conventional approaches (e.g. length-based, or hop-based shortest path etc.), such as the minimization of the energy consumption associated with the corresponding services. This computation is performed at the PCE(s) and, in case of inter-domain paths, requires the cooperation of multiple PCEs following a hierarchical model. The routing algorithms implemented at the PCEs take into account the “green” parameters, describing the power consumption of the heterogeneous resources involved. To further increase the energy saving during the operation of the VI the option of switching-off of unused resources (network and IT) is also proposed. This approach involves the introduction of a sleep or stand-by mode during which unused resources can be set to, and from which they can be awaken in a fast and seamless manner when required to be used [3], [4]. Moreover, in order to enable energy awareness at the routing level, information about energy consumption has to be included in the data model that describes both network and IT virtual resources and that is provided by the LICL to the VIO at the VI provisioning stage.

In addition, in the GEYSERS architecture the LICL dynamically and continuously synchronizes the energy status of physical resources and their associated virtual resources. Therefore, sudden energy changes of the infrastructure can be compensated by initiating VI re-planning.

Initial modelling results indicate that the proposed energy efficient approach addressing both the design and operation of integrated optical network and IT infrastructures achieves significant energy savings [8] and these become more substantial when service resilience is also taken into account.

2.1 Energy Models for Network and IT resources

The estimation of energy consumption of the PI network resources is highly sensitive to the network architecture employed and the network technology used. As GEYSERS is focusing on optical network technologies utilizing wavelength division multiplexing (WDM) the energy consumption of the corresponding technologies is taken into account. The overall network power consumption model is based on the power-dissipating (active) elements of the network and can be classified as switching nodes, referred to as Optical Cross-Connect (OXC) nodes in the context of WDM networks, and transmission line related elements. More specifically the OXCs assumed are based Micro-Electrical Mechanical Systems (MEMS), while for the fibre links a model comprising a sequence of alternating single mode fibre and dispersion compensating fibre spans together with optical amplifiers to compensate for the losses is employed. The details of these models are described in [5] with the only difference being that unlike [5] the current work assumes wavelength conversion capability available at the OXC nodes.

The physical IT infrastructures in this approach are considered to be data centres i.e. facilities used to house computer systems and associated components. Data centres are primarily used to process data (servers) and data storage (storage equipment). The collection of this processing and storage equipment is referred to as the "IT equipment". The rest of the IT infrastructure facilities, such as cooling mechanisms etc correspond to "infrastructure equipment". A large part of the energy consumption of a data centre resides in the IT equipment and the cooling [6], making infrastructure equipment (especially cooling) an important factor when trying to reduce the energy consumption in data centres.

In this approach the analysis is based on a linear power consumption model that mainly concentrates on the power consumption associated with the CPU load of IT resources. More specifically p_s is referred to the CPU resources and E_s to the energy consumption for utilizing a portion $u_s = p_s / p_s^{\max}$ of the maximum CPU capabilities $p_s^{\max}$ of server s. For simplicity, the following linear energy consumption model has been adopted [7], [8]:

$$E_s(u_s) = P_{idle}^{s} + \left(P_{busy}^{s} - P_{idle}^{s}\right)\cdot u_s$$

where P_{idle}^{s} and P_{busy}^{s} is the energy consumption of IT server s at idle state and under full load, respectively. For further details regarding the technical specifications of the IT servers the reader is referred to [7].

2.2 Energy Aware VI Planning

The VI planning is the process of using historical operating data and estimated future virtual resource (VR) requirements, to determine the optimal design of the VI. The objective of VI planning is to identify the topology and determine the virtual resources required to implement a dynamically reconfigurable VI based on both optical network and IT resources. This VI not only will meet customer's specific needs, but will also satisfy the virtual infrastructure provider's (VIP's) requirements for minimum energy consumption in our case. We formulate the Energy Aware VI planning problem through a Mixed Integer Linear Programming (MILP) model that aims at minimizing jointly the total energy that is consumed by optical network components including WDM transponders, amplifiers, switches and IT resources.

2.2.1 Problem Formulation

The problem is formulated using a network that is composed of one resource layer that contains the physical infrastructure and will produce as an output a virtual infrastructure [9]. The physical infrastructure is described through an eleven-node topology corresponding to the Pan-European optical network in which randomly selected nodes generate demands d (d = 1, 2,…, D) to be served by a set of IT servers s (s = 1, 2, … S). The granularity of demands is the wavelength. The IT locations (demand destinations) at which the services will be handled, are not specified and are of no importance to the services themselves. Therefore, the demand destinations will be identified through the optimization performed by the proposed model.

The objective of the current problem formulation is to minimize the total cost of the resulting network configuration as this cost consists of the following components:

- k_g is the cost of the capacity of link g of the PI. It consists of the energy consumed by each lightpath due to transmission and reception of the optical signal, optical amplification at each fibre span and switching according to the model described in [5].
- E_s is the cost for using capacity p_s of the IT servers. The linear energy consumption model [7], [8] described previously has been adopted where:

It should be noted that in addition to the power consumption due to data processing, a 100% power overhead due to cooling has been incorporated to the energy consumption model described above.

In this context, minimum energy consuming VI is obtained by minimizing the following cost function:

$$\textit{Minimize } F = \sum_g k_g u_s + \sum_s E_s(u_s)$$

Subject to the following constraints:

- Every demand d has to be processed to a single IT server. This allocation policy reduces the complexity of implementation and increases the reliability of the resulting VI.
- The planned VI must have
    - Adequate IT server resources such as CPU, memory, disk storage to support all requested services.
    - Sufficient optical link capacity in order all demands to be transferred to their destinations.
- The capacity of each link in the VI should be realized by specific PI resources.

2.2.2 Numerical Results

To investigate the energy efficiency of the proposed VI design scheme in the GEYSERS architecture, Figure 10 is considered: the lower layer depicts the PI and the layer above depicts the VI. For the PI the COST239 Pan-European reference topology [10] has been used in which four randomly selected nodes generate demands to be served by two IT servers located in Luxemburg and Milan. Furthermore, we assume a single fibre per link, 40 wavelengths per fibre, and wavelength channels of 10Gb/s each. It is also assumed that each IT server can process up to 2Tb/s of and its power consumption ranges from 6.6 to 13.2KW, under idle and full load, respectively.

An example of the optimal VI topology design is depicted in Figure 10. In this scenario, the generated VI topology consists of 7 virtual links and 6 virtual nodes.

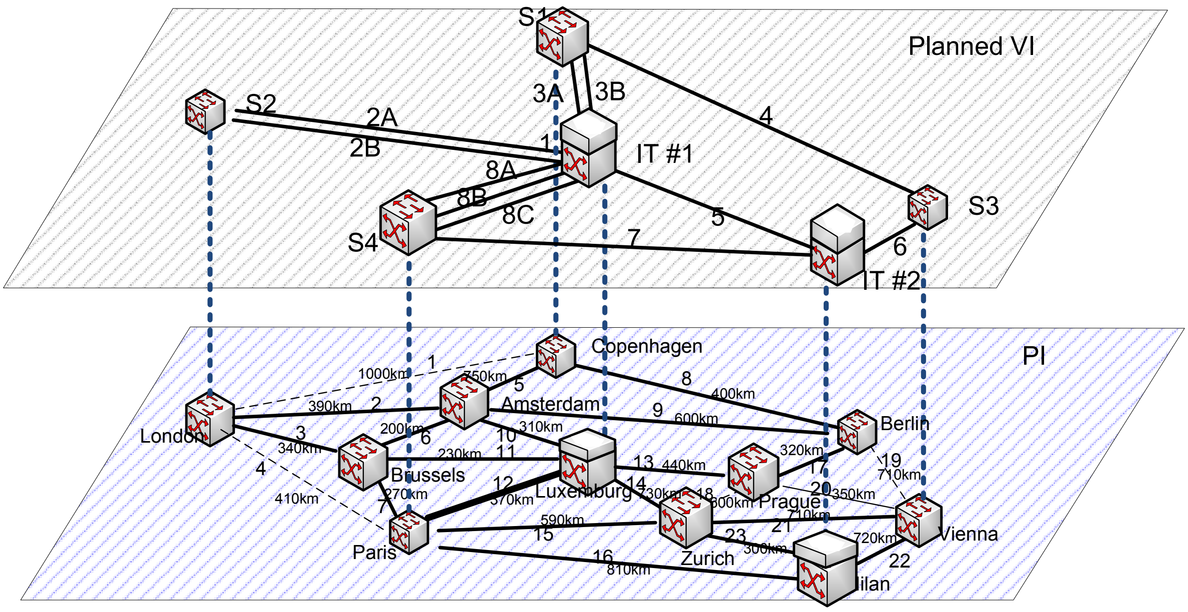

Figure 10: Example of the virtualization of a physical infrastructure

The capacity of each virtual link along with its mapping to the PI is given in Table 1.

Virtual link	Capacity (wavelengths)	Physical Layer Paths realizing virtual links	Capacity of PI paths (wavelengths)	Average cost / wavelength for Virtual link
Y2	50	Path A: u2-u10	15	21,97
		Path B: u3-u11	35	18,97
Y3	65	Path A: u5-u10	25	20,96
		Path B: u8-u17-u13	40	29,69

Y4	15	u5-u6-u7-u15-u21	15	51,86
Y5	15	u14-u23	15	20,83
Y6	40	u22	40	11,25
Y7	40	u16	40	11,62
Y8	70	Path A: u12	40	9,83
		Path B: u15-u14	25	22,01
		Path C: u7-u11	5	18,68

Table 1: Sample Virtual to Physical Mapping

In Figure 11, the performance of the proposed energy aware VI design is compared to the demand allocation scheme presented in [11] where demands from each source node are assigned to its closest IT server. Note that “closest” refers to the shortest distance between a source node and a data centre. Comparing these two schemes, it is observed that the energy aware VI design consumes significantly lower energy for serving the same amount of demands compared to the closest IT scheme in the order of 30%: in the former approach fewer IT servers are activated to serve the same amount of demands. Given that the power consumption required for the operation of the IT servers is dominant in this type of networks, switching-off the unused IT resources achieves significant reduction of energy consumption. Furthermore, it is observed that in both schemes the average power consumption increases almost linearly with the number of demands. However, the relative benefit of the energy aware design decreases slightly with the number of demands, as we get closer to full system load.

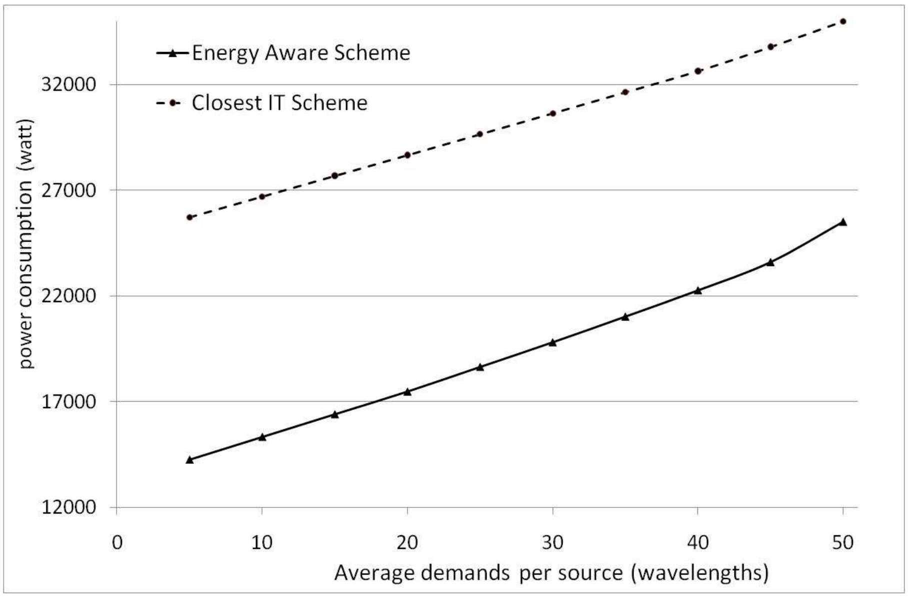

Figure 11: Comparison of the energy aware scheme with the closest IT server demand allocation scheme

2.2.3 Multiple Virtual Infrastructure Planning

In this section, the analysis presented is extended to address the problem of multiple VI planning over an integrated IT and optical network infrastructure. Taking into account the detailed power consumption models and figures of the underlying converged IT and optical network physical infrastructure [5][6], the least energy consuming VIs are identified. The problem is formulated using again an MILP model in which for each VI_i $(i=1,2,...,I)$, there is a set of demands d_i $(d_i=1_i,2_i,...,D_i)$ to be served by a set of IT servers s $(s=1,2,…,S)$. The scope of the MILP problem formulation is to minimize the total energy consumption of the planned VIs' under optical links' capacity and IT servers' processing constraints.

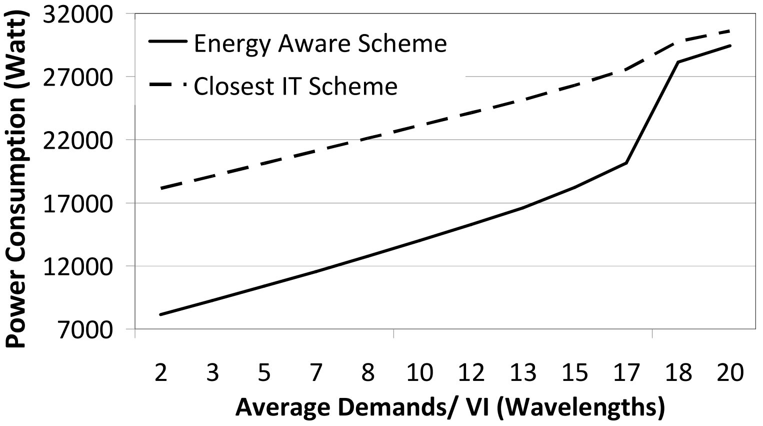

Figure 12: Comparison of the energy aware with closest IT scheme (3 VIs)

Preliminary results on the total consumed power of the infrastructure (optical network and IT resources) when applying the proposed MILP approach optimizing for energy or distance between sources and IT servers are illustrated in Figure 12. Comparing these two schemes, it is observed that the energy aware VI design consumes significantly lower energy for serving the same amount of demands compared to the closest IT scheme in the order of 40%: in the former approach fewer IT servers are activated to serve the same amount of demands. Given that the power consumption required for the operation of the IT servers is dominant in this type of networks, switching-off the unused IT resources achieves significant reduction of energy consumption. Furthermore, it is observed that in both schemes the average power consumption increases almost linearly with the number of demands. However, the relative benefit of the energy aware design decreases slightly with the number of demands, as we get closer to full system load.

2.2.4 Virtual Infrastructure Planning with Resilient Considerations

Uninterrupted service provisioning is of crucial importance in the deployment of transport optical networks. To this end, countermeasures against failures of the IT servers and optical links of the PI that may lead to service disruption should be taken into account during the VI planning process. In traditional planning algorithms, a possible failure of the primary IT server is treated by forwarding its demands to a secondary IT server. On the other hand, in case of failure of an optical link, demands are routed to their destination via alternative paths.

Within the context of GEYSERS, energy aware virtual infrastructure planning algorithms with resilient consideration have been implemented quantifying significant energy saving compared to traditional approaches. Preliminary results are depicted in Figure 13, where the performance of two variations of the proposed energy aware VI design with and without (w/o) protection mechanisms is compared to the demand allocation scheme aiming at using the closest IT resources . Note that “closest” refers to the shortest distance between a source node and a data centre. Comparing these two schemes, it is observed that the energy aware VI design without protection mechanism consumes significantly lower energy for serving the same amount of demands compared to the closest IT scheme: in the former approach only one IT server is activated to serve the same amount of demands. Given that the power consumption required for the operation of the IT servers is dominant in this type of networks, switching off the unnecessary IT resources achieves significant reduction of energy consumption. Furthermore, the energy aware scheme enhanced with protection mechanisms achieves significantly lower power consumption or for high traffic demands the same, with the closest IT scheme without protection.

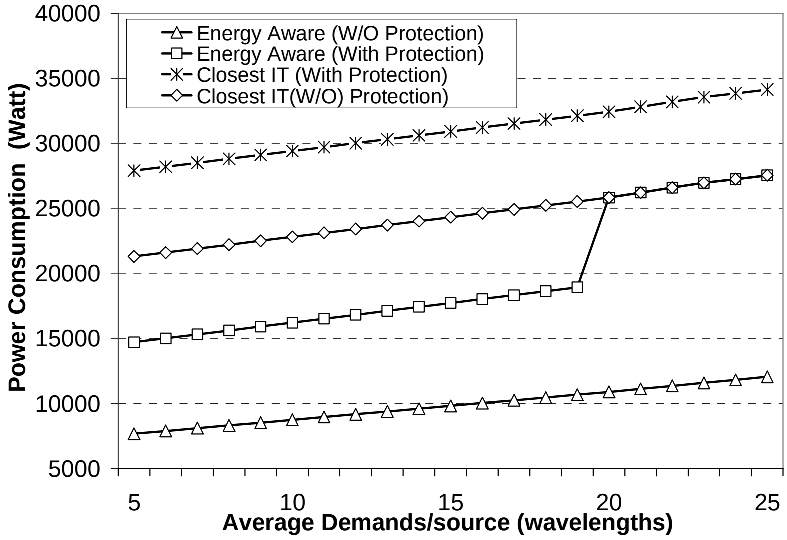

Figure 13: Comparison of the energy aware scheme with the closest IT server demand allocation scheme

2.3 Energy Aware Service Provisioning over VIs

As already discussed the multi-layer GEYSERS architecture allows for energy efficiency considerations both at the VI planning phase as well as the operation phase of the VI through the use of energy efficient service provisioning supported by the NCP+.

In order to facilitate energy efficient operation of the VI we propose an energy efficient routing algorithm. This algorithm aims at provisioning IT services that are originating from specific source sites and need to be executed by suitable IT resources (e.g. data centres). Candidate IT resources reside at different geographical locations and connectivity of the source site to these IT resources is provided through the underlying optical network as dictated by the already designed VI, available as the output of the VI planning phase. The routing scheme used is anycast routing [13], since the requirement for the IT services is the delivery of results, while the exact location of the execution of the job is of no interest. In this context energy consumption is achieved by identifying the least energy consuming IT and network resources required to support the services, and switching-off of any unused network resources (such as links, nodes, etc) and IT resources (such as servers) [3].

To evaluate the performance of energy awareness in the service provisioning phase, we have created an Integer Linear Program (ILP) based on the model described in [13] modified appropriately. This ILP assumes a capacitated VI and an energy consumption model where a data centre's energy increases linearly with processing load. The energy consumption of the optical network is calculated through the VI planning phase previously described. Our model takes as inputs the: (a) parameters specifying the energy consumption of the VI IT and network resources, (b) a set of IT resource sites which are able to handle IT requests, (c) a set of lightpath requests per source, which need to end in one of the proposed server sites, and (d) the capacity of the links and the Data Centres. The output of the model includes: (a) the IT resource allocation (i.e. which server site serves which source node's light path request) and (b) a route to the allocated IT site allowing switching of server sites, OXCs and links.

In order to compare the performance of the proposed energy efficient algorithm, we have also implemented three alternative strategies: (i) one where the objective is to minimize the number of used wavelengths, which corresponds to conventional shortest path (SP) calculation; (ii) one where we merely try to minimize the energy consumed by the network elements; and (iii) a third one which aims at minimizing the energy consumed by the IT resource sites only. We have run the calculations for 10 random demand vectors for each demand size (15 to 285), so the results which are plotted in Figure 12 actually represent averages over those 10 demand vectors.

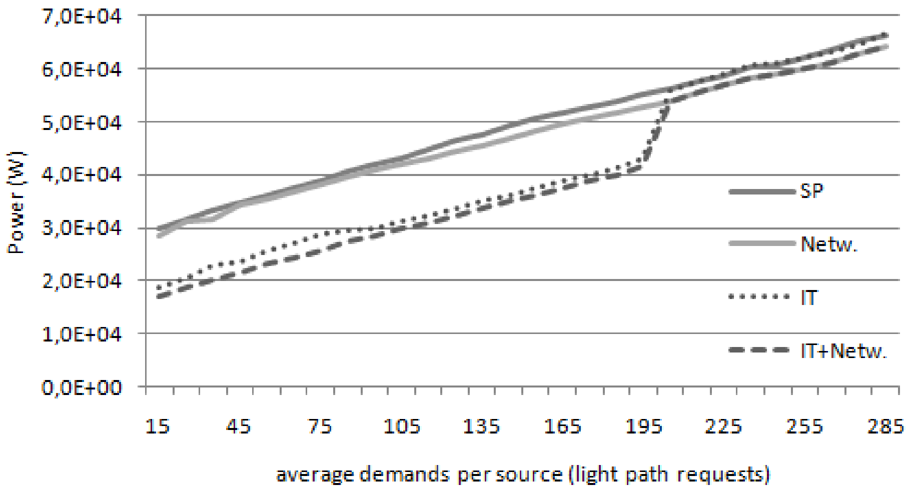

Figure 14: Energy aware service provisioning schemes vs SP routing

Figure 14 illustrates the energy consumption of the VI infrastructure when applying the proposed energy aware algorithm compared to SP. It can be seen that the energy saving achieved through the energy aware scheme reaches 40% for low service loads and reduces with the load. This effect of the energy saving reduction with the service load is due to that at high loads, most of the VI resources are utilized and the overall energy consumption reaches its maximum value. Among all approaches under consideration the highest VI energy consumption is obtained as expected, in the case where no energy consideration is taking place. In addition, the approach that focuses on the network power consumption provides very similar performance to that of the SP, with a small difference (3,37% on average). On the other hand, the proposed scheme considering jointly the energy consumption of IT and network resources always provides optimum performance. However, it should be noted that the routing approach taking into consideration only the power consumption of IT resources, demonstrates only a small additional energy penalty compared to the proposed solution (a difference ranging from 2,76% up to 11,05% depending on the offered load). These observations clearly indicate that in the VI under consideration the IT resources have the most dominant contribution in the overall power consumption compared to the optical network resources. Figure 12 shows that when applying the proposed routing approach as well as the approach considering the energy consumption of IT resources, a step like increase in energy consumption is observed at 195 requests. This is due to that a resource site’s maximum capacity is 200 requests. Hence, the step like increase indicates the start-up cost for powering on a second server site: in low loading conditions the preferable solution for energy minimization involves the use of only one server site. However, in high load conditions powering on a second server site is inevitable to accommodate the load.

This approach is extended in [14] taking into account the granularity with which a data centre is able to switch on/off servers. Simulation results show significant energy savings that can reach up to 55% compared to energy-unaware schemes, depending on the granularity with which a data centre is able to switch on/off servers.

3 References

[1] GEYSERS deliverable D2.1: Initial GEYSERS architecture and interfaces specification

[2] GEYSERS deliverable D2.2: GEYSERS overall architectire & interfaces specification and service provisioning workflow

[3] Energy Star, “Small network equipment, www.energystar.gov/index.cfm?c=new specs.small network equip.

[4] A. Muhammad, P. Monti, I. Cerutti, L. Wosinska, P. Castoldi, A. Tzanakaki, “Energy-Efficient WDM Network Planning with Protection Resources in Sleep Mode”, Proc. Globecom 2010, Miami, Florida, USA, 6-10 Dec. 2010.

[5] A. Tzanakaki et al, “Power Considerations towards a Sustainable Pan-European Network”, JWA061, OFC2011

[6] Davis, Z. (2005). Power Consumption and cooling in the data center: A survey.

[7] Oracle Data Sheet, “Sun Oracle DataBaseMachine”, http://www.oracle.com/us/products/database/database-machine-069034.html

[8] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. 2007. Power provisioning for a warehouse-sized computer. *SIGARCH Comput. Archit. News* 35, 2 (June 2007), 13-23

[9] A.Tzanakaki et al., "Energy Efficiency in integrated IT and Optical Network Infrastructures: The GEYSERS approach", in proc. of IEEE INFOCOM 2011, Workshop on Green Communications and Networking (2011)

[10] A.Tzanakaki et.al, "Dimensioning the future Pan-European optical network with energy efficiency considerations," JOCN 3, 272-280 (2011)

[11] P. Batchelor et al., “Study on the implementation of optical transparent transport networks in the European environment-Results of the research project COST 239”, Photonic Network Communications, vol. 2, no. 1, pp. 15-32, 2000.

[12] K. Bouyoucef, I. Limam-Bedhiaf, and O.Cherkaoui,“Optimal Allocation Approach of Virtual Servers in Cloud Computing”, in proc. Of 6th Euro-NF conference on Next Generation Internet,July 2010.

[13] J. Buysse, M. De Leenheer, B. Dhoedt, C. Develder, “On the impact of relocation on network dimensions in optical grids”, Proc. Int. Conf. Optical Netw. Design and Modeling (ONDM 2010), Kyoto, Japan, 31 Jan.-3 Feb. 2010.

[14] J.Buysse, K.Georgakilas, A.Tzanakaki, M. Leenheer, B Dhoedt , C.Develder and P. Demeester, “Calculating the minimum bounds of energy consumption for cloud networks”, in proc. Of ICCCN 2011

4 Acronyms

AAA	Authentication, Authorization and Accounting
AAI	Authentication, Authorization Infrastructure
ASON	Automatically Switched Optical Networks
BoD	Bandwidth on Demand
CCAMP	Common Control and Management Protocol
CCI	Connection Controller Interface
CSSI	Common Security Service Interface
ENNI	External Network to Network Interface
ERO	Explicit Route Object
FA	Forwarding Adjacency
FEC	Forward Error Correction
GMPLS	Generalized Multi-Protocol Label Switching
GRI	Global Reservation Identifier
IACD	Interface Adjustment Capability Descriptor
LICL	Logical Infrastructure Composition Layer
LRI	Local Reservation Identifier
LSP	Label Switched Path
MLN	Multi-Layer Network
MRN	Multi-Region Network
NCP	Network Control Plane
NIPS	Network + IT Provisioning Service
NMS	Network Management System
NNI	Network to Network Interface
OBS	Optical Burst Switching
OPS	Optical Packet Switching
OTN	Optical Transport Network
PAP	Policy Administration Point
PCC	Path Computation Client
PCE	Path Computation Element
PCEP	PCE Communication Protocol
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PIP	Physical Infrastructure Provider

PIP	Policy Information Point (Security)
PR	Physical Resource
QoS	Quality of Service
RD	Routing Domain
RORA	Resources, Ownership, Roles, Actors
RSVP	Resource reSerVation Protocol
RWA	Routing and Wavelength Assignment
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SSLM	Security Services Lifecycle Management
TE	Traffic Engineering
TLV	Type, Length, Value
UNI	User to Network Interface
VI	Virtual Infrastructure
VIO	Virtual Infrastructure Operator
VIP	Virtual Infrastructure Provider
VR	Virtual Resource
WDM	Wavelength Division Multiplexing
WG	Working Group