UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Bringing optical networks to the Cloud: an architecture for a sustainable future Internet

Vicat-Blanc, P.; Figuerola, S.; Chen, X.; Landi, G.; Escalona, E.; Develder, C.; Tzanakaki, A.; Demchenko, Y.; García Espín, J.A.; Ferrer, J.; López, E.; Soudan, S.; Buysse, J.; Jukan, A.; Ciulli, N.; Brogle, M.; van Laarhoven, L.; Belter, B.; Anhalt, F.; Nejabati, R.; Simeonidou, D.; Ngo, C.; de Laat, C.; Biancani, M.; Roth, M.; Donadio, P.; Jiménez, J.; Antoniak-Lewandowska, M.; Gumaste, A.

DOI
10.1007/978-3-642-20898-0_22

Publication date
2011

Document Version
Final published version

Published in
The Future Internet

License
CC BY-NC

Link to publication

Citation for published version (APA):
Vicat-Blanc, P., Figuerola, S., Chen, X., Landi, G., Escalona, E., Develder, C., Tzanakaki, A., Demchenko, Y., García Espín, J. A., Ferrer, J., López, E., Soudan, S., Buysse, J., Jukan, A., Ciulli, N., Brogle, M., van Laarhoven, L., Belter, B., Anhalt, F., ... Gumaste, A. (2011). Bringing optical networks to the Cloud: an architecture for a sustainable future Internet. In J. Domingue, A. Galis, A. Gavras, T. Zahariadis, D. Lambert, F. Cleary, P. Daras, S. Krco, H. Müller, M.-S. Li, H. Schaffers, V. Lotz, F. Alvarez, B. Stiller, S. Karnouskos, S. Avessta, & M. Nilsson (Eds.), *The Future Internet: Future Internet Assembly 2011: achievements and technological promises* (pp. 307-320). (Lecture notes in computer science; Vol. 6656). Springer. https://doi.org/10.1007/978-3-642-20898-0_22

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

Download date:09 Nov 2026

Bringing Optical Networks to the Cloud: An Architecture for a Sustainable Future Internet

Pascale Vicat-Blanc[1], Sergi Figuerola[2], Xiaomin Chen[4], Giada Landi[5], Eduard Escalona[10], Chris Develder[3], Anna Tzanakaki[6], Yuri Demchenko[11], Joan A. García Espín[2], Jordi Ferrer[2], Ester López[2], Sébastien Soudan[1], Jens Buysse[3], Admela Jukan[4], Nicola Ciulli[5], Marc Brogle[7], Luuk van Laarhoven[7], Bartosz Belter[8], Fabienne Anhalt[9], Reza Nejabati[10], Dimitra Simeonidou[10], Canh Ngo[11], Cees de Laat[11], Matteo Biancani[12], Michael Roth[13], Pasquale Donadio[14], Javier Jiménez[15], Monika Antoniak-Lewandowska[16], and Ashwin Gumaste[17]

[1] LYaTiss
[2] Fundació i2CAT
[3] Interdisciplinary Institute for BroadBand Technology
[4] Technische Univeristät Carolo-Wilhelmina zu Braunschweig
[5] Nextworks
[6] Athens Information Technology
[7] SAP Research
[8] Poznan Supercomputing and Networking Center
[9] INRIA
[10] University of Essex
[11] Universiteit van Amsterdam
[12] Interoute
[13] ADVA
[14] Alcatel-Lucent
[15] Telefónica I+D
[16] Telekomunikacja Polska
[17] Indian Institute of Technology, Bombay

Abstract. Over the years, the Internet has become a central tool for society. The extent of its growth and usage raises critical issues associated with its design principles that need to be addressed before it reaches its limits. Many emerging applications have increasing requirements in terms of bandwidth, QoS and manageability. Moreover, applications such as Cloud computing and 3D-video streaming require optimization and combined provisioning of different infrastructure resources and services that include both network and IT resources. Demands become more and more sporadic and variable, making dynamic provisioning highly needed. As a huge energy consumer, the Internet also needs to be energy-conscious. Applications critical for society and business (e.g., health, finance) or for real-time communication demand a highly reliable, robust and secure Internet. Finally, the future Internet needs to support sustainable business models, in order to drive innovation, competition, and research. Combining optical network technology with Cloud technology is key to addressing the future Internet/Cloud challenges. In this con-

J. Domingue et al. (Eds.): Future Internet Assembly, LNCS 6656, pp. 307–320, 2011.
© The Author(s). This article is published with open access at SpringerLink.com.

text, we propose an integrated approach: realizing the convergence of the IT- and optical-network-provisioning models will help bring revenues to all the actors involved in the value chain. Premium advanced network and IT managed services integrated with the vanilla Internet will ensure a sustainable future Internet/Cloud enabling demanding and ubiquitous applications to coexist.

Keywords: Future Internet, Virtualization, Dynamic Provisioning, Virtual Infrastructures, Convergence, IaaS, Optical Network, Cloud

1 Introduction

Over the years, the Internet has become a central tool for society. Its large adoption and strength originates from its architectural, technological and operational foundation: a layered architecture and an agreed-upon set of protocols for the sharing and transmission of data over practically any medium. The Internet's infrastructure is essentially an interconnection of several heterogeneous networks called Autonomous Systems that are interconnected with network equipment called gateways or routers. Routers are interconnected together through links, which in the core-network segment are mostly based on optical transmission technology, but also in the access segments gradual migration to optical technologies occurs. The current Internet has become an ubiquitous commodity to provide communication services to the ultimate consumers: enterprises or home/residential users. The Internet's architecture assumes that routers are stateless and the entire network is neutral. There is no control over the content and the network resources consumed by each user. It is assumed that users are well-behaving and have homogeneous requirements and consumption.

After having dramatically enhanced our interpersonal and business communications as well as general information exchange—thanks to emails, the web, VoIP, triple play service, etc.—the Internet is currently providing a rich environment for social networking and collaboration and for emerging Cloud-based applications such as Amazon's EC2, Azure, Google apps and others. The Cloud technologies are emerging as a new provisioning model [2]. Cloud stands for on-demand access to IT hardware or software resources over the Internet. Clouds are revolutionizing the IT world [11], but treat the Internet as always available, without constraints and absolutely reliable, which is yet to be achieved. Analysts predict that in 2020, more than 80 % of the IT will be outsourced within the Cloud [9]! With the increase in bandwidth-hungry applications, it is just a matter of time before the Internet's architecture reaches its limits.

The new Internet's architecture should propose solutions for QoS provisioning, management and control, enabling a highly flexible usage of the Internet resources to meet bursty demands. If the Internet's architecture is not redesigned, not only mission-critical or business applications in the Cloud will suffer, but even conventional Internet's users will be affected by the uncontrolled traffic or business activity over it.

Today, it is impossible to throw away what has made the enormous success of the Internet: the robustness brought by the datagram building block and the end-to-end principle which are of critical importance for all applications. In this context, we argue that the performance, control, security and manageability issues, considered as non-priority features in the 70s [3] should be addressed now [6]. In this chapter, we propose to improve the current Internet's architecture with the advanced control and management plane that should improve the integration of both new optical transport network technologies and new emerging services and application that require better control over the networking infrastructure and its QoS properties.

In this chapter we explore the combination of Cloud-based resource provisioning and the virtualization paradigm with dynamic network provisioning as a way towards such a sustainable future Internet. The proposed architecture for the future Internet will provide a basis for the convergence of networks—optical networks in particular—with the Clouds while respecting the basic operational principles of today's Internet. It is important to note that for several years, to serve the new generation of applications in the commercial and scientific sectors, telecom operators have considered methods for dynamic provisioning of high-capacity network-connectivity services tightly bundled with IT resources. The requirements for resource availability, QoS guarantee and energy efficiency mixed with the need for an ubiquitous, fair and highly available access to these capacities are the driving force of our new architecture. The rest of this chapter exposes the main challenges to be addressed by the future Internet's architecture, describing the solution proposed by the GEYSERS[18] project and how it is a step towards the deployment and exploitation of this architecture.

2 Challenges

There are various challenges that are driving today's Internet to the limit, which in turn have to be addressed by the future Internet's architecture. We consider the following six challenges as priorities:

1. **Enable ubiquitous access to huge bandwidth:** As of today, the users/applications that require bandwidth beyond 1 Gbps are rather common, with a growing tendency towards applications requiring a 10 Gbps or even 100 Gbps connectivity. Examples include networked data storage, high-definition (HD) and ultra-HD multimedia-content distribution, large remote instrumentation applications, to name a few. But today, these applications cannot use the Internet because of the fair-sharing principle and the basic routing approach. As TCP, referred to as the one-size-fits-all protocol, has reached its limits in controlling—alone—the bandwidth, other mechanisms must be introduced to enable a flexible access to the huge available bandwidth.

[18] http://www.geysers.eu

2. **Coordinate IT and network service provisioning:** In order to dynamically provision external IT resources and gain full benefit of these thanks to Cloud technologies, it is important to have control over the quality of the network connections used, which is a challenge in today's best-effort Internet. Indeed, IT resources are processing data that should be transferred from the user's premises or from the data repository to the computing resources. When the Cloud will be largely adopted and the data deluge will fall in it, the communication model offered by the Internet may break the hope for fully-transparent remote access and outsourcing. The interconnection of IT resources over networks requires well-managed, dynamically invoked, consistent services. IT and network should be provisioned in a coordinated way in the future Internet.
3. **Deal with the unpredictability and burstiness of traffic:** The increasing popularity of video applications over the Internet causes the traffic to be unpredictable in the networks. The traffic's bursty nature requires mechanisms to support the dynamic behavior of the services and applications. Moreover, another important issue is that the popularity of content and applications on the Internet will be more and more sporadic: the network effect amplifies reactions. Therefore, the future Internet needs to provide mechanisms that facilitate elasticity of resources provisioning with the aim to face sporadic, seasonal or unpredictable demands.
4. **Make the network energy-aware:** It is reported in the literature [10], that ICT is responsible for about 4 % of the worldwide energy consumption today, and this percentage is expected to rapidly grow over the next few years following the growth of the Internet. Therefore, as a significant contributor to the overall energy consumption of the planet, the Internet needs to be energy-conscious. In the context of the proposed approach, this should involve energy awareness both in the provisioning of network and IT resources in an integrated globally optimized manner.
5. **Enable secured and reliable services:** The network's service outages and hostile hacks have received significant attention lately due to society's high dependency on information systems. The current Internet's service paradigm allows service providers to authenticate resources in provider domains but does not allow them to authenticate end-users requiring the resources. As a consequence, the provisioning of network resources and the secure access of end users to resources is a challenge. This issue is even more significant in the emerging systems with the provisioning of integrated resources provided by both network and IT providers to network operators.
6. **Develop a sustainable and strategic business model:** Currently, the business models deployed by telecom operators are focused on selling services on top of their infrastructures. In addition, operators cannot offer dynamic and smooth integration of diversified resources and services (both IT and network) at the provisioning phase. Network-infrastructure resources are not understood as a service within the value chain of IT service providers. We believe that a novel business model is necessary, which can fully integrate the network substrate with the IT resources into a single infrastructure. In

addition, such business model will let operators offer their infrastructures as a service to third-party entities.

3 Model

In order to address the aforementioned challenges and opportunities, the proposed architecture introduces the three basic concepts featured by the future Internet:

- The **Virtual Infrastructure concept** and its operational model as a fundamental approach to enable the on-demand infrastructure services provisioning with guaranteed performance and QoS, including manageable security services.
- A **new layered architecture** for the Control and Management Plane that allows dynamic services composition and orchestration in the virtual infrastructures that can consistently address the manageability, energy-efficiency and traffic-unpredictability issues.
- The **definition of the new Role-Based operational model** that includes the definition of the main actors and roles together with their ownership and business relations and the operational workflow.

3.1 Virtual Infrastructures

A Virtual Infrastructure (VI) is defined as an interconnection of virtualized resources with coordinated management and control processes. The virtual infrastructure concept consists in the decoupling of the physical infrastructure from its virtual representation, which can either be aggregating or partitioning a physical resource. This concept has little to do with the way data is processed or transmitted internally, while enabling the creation of containers with associated non-functional properties (isolation, performance, protection, etc.). The definition of a VI as an interconnection of a set of virtual resources provides the possibility of sharing the underlying physical infrastructure among different operators, and granting them isolation. At the same time, dynamic VI-provisioning mechanisms can be introduced into the infrastructure definition, creating the possibility to modify the VI capabilities in order to align them with the VI usage needs at any given instant.

As stated above, optical network technologies are among the key components for the future Internet. They inherently provide plenty of bandwidth and, in particular, the emerging flexible technology supported by the required control mechanisms enable a more efficient utilization of the optical spectrum and on-demand flexible bandwidth allocation, hence addressing challenge #1. IT resources comprise another important category of future Internet shared resources aggregated in large-scale data centers and providing high computational and storage capacities. In order to build VIs, these resources are abstracted, partitioned or grouped into Virtual Resources (VRs), which are attached to the VI

and exposed for configuration. Each VR contains a well-defined, isolated subset of capabilities of a given physical entity. Taking the example of an optical cross-connect, the VRs are built on its total amount of ports or its wavelengths, and control switching exclusively between these allocated ports and wavelengths.

3.2 A Novel Layered Architecture

The proposed architecture features an innovative multi-layered structure to re-qualify the interworking of legacy planes and enable advanced services including the concepts of Infrastructure-as-a-Service (IaaS) and service-oriented networking [4]. We aim to enable a flexible infrastructure provisioning paradigm in terms of configuration, accessibility and availability for the end users, as well as a separation of the functional aspects of the entities involved in the converged service provisioning, from the service consumer to the physical ICT infrastructure.

Our architecture is composed of three layers residing above the physical infrastructure as illustrated in Fig. 1. They are referred to as the **infrastructure-virtualization layer**, the **enhanced control plane**, that corresponds to the network management layer, and the **service middleware** layer. Each layer is responsible for implementing different functionalities covering the full end-to-end service delivery from the service layer to the physical substrate.

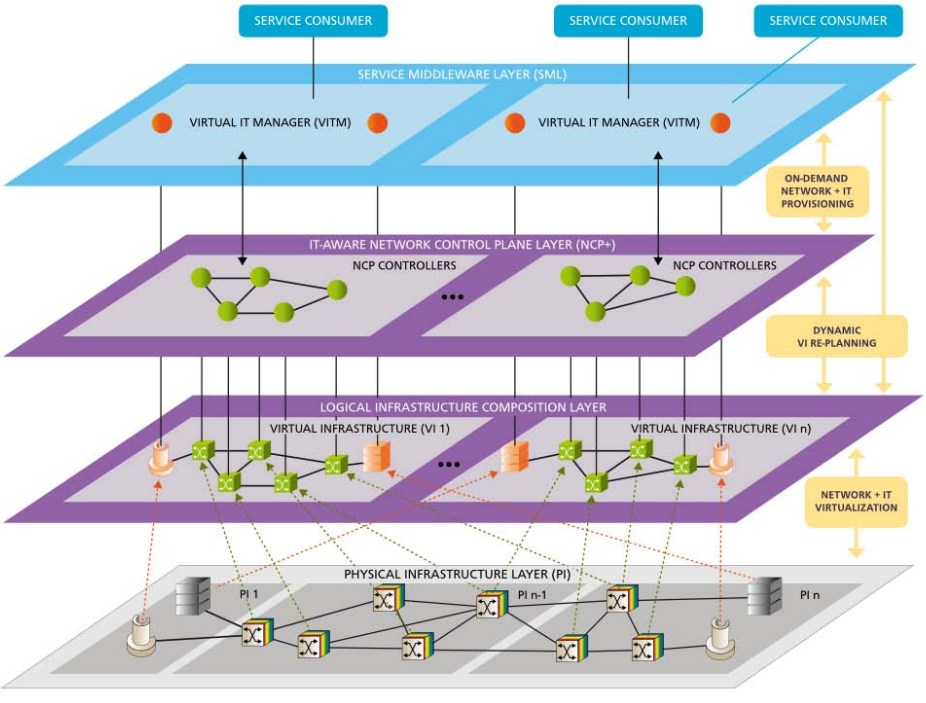

Fig. 1. Reference model as proposed in the GEYSERS project.

1. At the lowest level, the physical infrastructure layer comprises optical-network, packet network and IT resources from physical-infrastructure providers. Central to this novel architecture is the **infrastructure virtualization layer** which abstracts, partitions and interconnects infrastructure resources contained in the physical infrastructure layer. It may be composed of physical resources from different providers.
2. An **enhanced control plane** responsible for the operational actions to be performed over the virtual infrastructure (i.e., controlling and managing the network resources constituting the Virtual Infrastructure) is closely interacting with the virtualization layer.
3. Finally, a **service middleware layer** is introduced to fully decouple the physical infrastructure from the service level. It is an intermediate layer between applications running at the service consumer's premises and the control plane, able to translate the applications' resource requirements and SLAs into service requests. These requests specify the attributes and constraints expected from the network and IT resources.

The proposed architecture also addresses two major security aspects: secure operation of the VI provisioning process, and provisioning dynamic security services, to address challenge #5. Fig. 1 shows the reference model of our architecture as it has been modeled in the context of the GEYSERS project. The infrastructure-virtualization layer is implemented as the Logical Infrastructure Composition Layer (LICL) and the enhanced control plane as the NCP+. As an example, the figure shows two virtual infrastructures, each of them controlled by a single control-plane instance. The different notations used in the rest of this chapter are summarized in Tab. 1.

Table 1. Abbreviations

LICL	Logical Infrastructure Composition Layer
NCP	Network Control Plane
NCP+	Enhanced Network Control Plane
NIPS	Network + IT Provisioning Services
PIP	Physical Infrastructure Provider
SML	Service Middleware Layer
VI	Virtual Infrastructure
VIO	Virtual Infrastructure Operator
VIO-IT	Virtual IT Infrastructure Operator
VIO-N	Virtual Network Infrastructure Operator
VIP	Virtual Infrastructure Provider
VR	Virtual Resource

3.3 New Roles and Strategic Business Model

Given this virtualized network and IT architecture, new actors are emerging: the traditional carriers role is split among Physical Infrastructure Providers

(PIP), Virtual Infrastructure Providers (VIP) and Virtual Infrastructure Operators (VIO). PIPs own the physical devices and rent partitions of them to VIPs. These, in turn, compose VIs of the virtual resources rented at one or several PIPs, and lease these VIs to VIOs. VIOs can efficiently operate the rented virtual infrastructure through the enhanced Control Plane, capable of provisioning on-demand network services bundled with IT resources to meet challenge #2. New business relationships can be developed between Virtual IT Infrastructure Operators (VIO-IT) and Virtual Network Infrastructure Operators (VIO-N), but they can also converge to a single actor, with traditional operators moving their business towards higher-value application layers. This creates new market opportunities for all the different actors addressing challenge #6: infrastructure providers, infrastructure operators and application providers cooperate in a business model where on-demand services are efficiently offered through the seamless provisioning of network and IT virtual resources.

To illustrate this, we now describe a sample use case. A company hosts an Enterprise Information System externally on a Cloud rented from a Software-as-a-Service (SaaS) provider. It relies on the resources provided by one or more IT and network infrastructure providers. It also connects heterogeneous data resources in an isolated virtual infrastructure. Furthermore, it supports scaling (up and down) of services and load. It provides means to continuously monitor what the effect of scaling will be on response time, performance, quality of data, security, cost aspect, feasibility, etc. Our architecture will result in a new role for telecom operators that own their infrastructure to offer their optical network integrated with IT infrastructures (either owned by them or by third-party providers) as a service to network operators. This, on the other hand, will enable application developers, service providers and infrastructure providers to contribute in a business model where complex services (e.g., Cloud computing) with complex attributes (e.g., optimized energy consumption and optimized capacity consumption) and strict bandwidth requirements (e.g., real time and resilience) can be offered economically and efficiently to users and applications.

4 Virtual Infrastructures in Action

4.1 Virtual Infrastructure Life Cycle

The definition of the VI's life cycle is an important component of the proposed architecture that provides a common basis for the definition of the VI-provisioning workflow and all involved actors and services integration.

The VI life cycle starts with a VI request from the VIO, as represented on Fig. 2. A VI request may be loosely constrained, which considers the capabilities that the VI should provide but does not require a specific topology or strict resource attributes; or constrained to a topology, specifying the nodes that must be included and their capabilities and attributes. If a VI request is loosely constrained, the Virtual Infrastructure Provider (VIP) is responsible for deciding the final infrastructure topology. The VI request is not limited to a static definition of a VI, but temporal constrains can also be included. To specify VI requests,

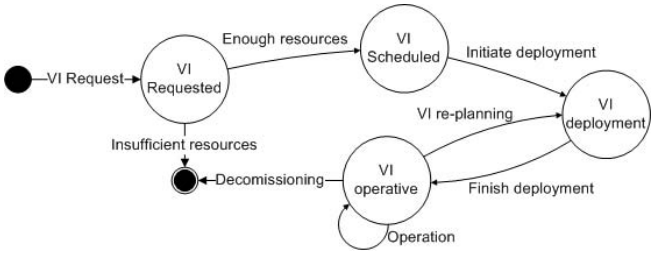

Fig. 2. VI Life Cycle

the *Virtual Infrastructure Description Language (VXDL)*[19] [7] can be used. It allows to formulate the topology and provisioning of a VI for its whole lifecycle [1]. After receiving the VI request, it is VIP's responsibility to split it into VRs request to different Physical Infrastructure Providers (PIPs), detailing the VRs that will be needed and their characteristics. Finally, each corresponding PIP will partition the physical resources into isolated virtual resources to compose and deploy the VI. Once the VI is in operation, the VIO can request for a VI re-planning to update its infrastructure so that it is optimized for a new usage pattern, involving for example traffic changes as predicted in challenge #3. A VI re-planning considers modifying the attributes of a node from the infrastructure (VR resizing), as well as including or excluding nodes into the VI. Depending on the initial SLAs agreed between the infrastructure operator and provider, a VI re-planning may involve a re-negotiation of the corresponding SLA or the current one may include it. The final stage in the VI life cycle is decommissioning, which releases the physical resources that were taken by the VIO.

4.2 Controlling the Virtual Infrastructures

Cloud applications are characterized by high dynamic and strict requirements in terms of provisioning of IT resources, where distributed computing and storage resources are automatically scaled up and down, with guaranteed high-capacity network connectivity. The enhanced Network Control Plane (NCP+) proposed in our architecture (Fig.1) offers integrated mechanisms for Network + IT Provisioning Services (NIPS) through the on-demand and seamless provisioning of optical and IT resources. These procedures are based on a strong inter-cooperation between the NCP+ and the **service middleware layer (SML)** via a service-to-network interface, named NIPS UNI during the entire VI service life cycle. The NIPS UNI offers functionalities for setup, modification and tear-down of enhanced transport network services (optionally combined with advance reservations), monitoring and cross-layer recovery. In particular, the NIPS UNI is based on a powerful information model that allows the SML to specify multiple aspects of the application requirements in the service requests. These requirements describe not only the characteristics of the required connectivity in terms

[19] http://www.ens-lyon.fr/LIP/RESO/Software/vxdl/home.html

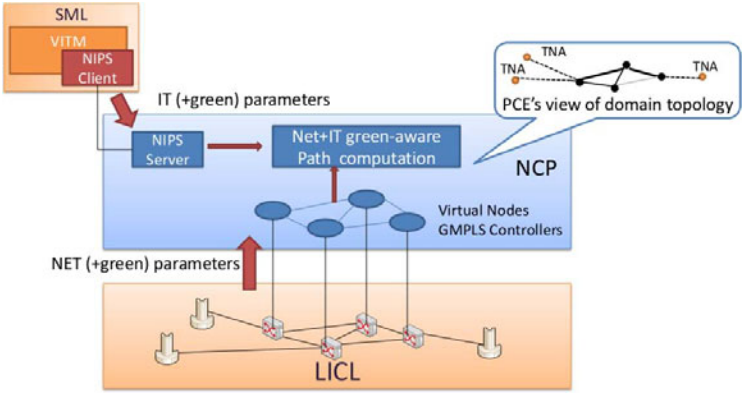

Fig. 3. Enhanced NCP—Network + IT energy-efficient service provisioning

of bandwidth, but could also specify further parameters, like the monitoring information required at the service layer or the mechanisms for cross-layer recovery. Following the service specification included in the requests, the network connectivity services are automatically tailored to the cloud dynamics, allowing for an efficient utilization of the underlying infrastructure. The integrated control of network and IT resources is achieved through different levels of cooperation between service plane and Network Control Plane, where part of the typical service plane's functionalities for the selection of the IT end points associated to an aggregate service can be progressively delegated to the NCP+. In anycast services the SML provides just a description of the required IT resources (e.g. in terms of amount of CPU), while the NCP+ computes the most efficient combination of end-points and network path to be used for the specific service. This feature requires a NCP+ with a global knowledge of the capabilities and availabilities of the IT resources attached to the network; this knowledge is obtained pushing a subset of relevant information from the SML to the NCP+, using the NIPS UNI. The NCP+ in the proposed architecture is based on ASON/GMPLS [8] and PCE [5] architectures and is enhanced with routing and signaling protocols extensions and constraints designed to support the NIPS. In particular, it is enhanced with the capability of advertising the energy consumption of network and IT elements, as well as availability, capabilities and costs of IT resources. The network-related parameters are collected from the LICL and flooded through the OSPF-TE protocol among the GMPLS controllers of the associated domain; on the other hand IT parameters retrieved through the NIPS UNI are communicated to a set of PCEs and not flooded among the GMPLS controllers. The path computation is performed by dedicated PCEs that implements enhanced computation algorithms able to combine both network and IT parameters with energy-consumption information in order to select the most suitable resources and find an end-to-end path consuming the minimum total energy (see Sec. 5). Figure 3 shows a high-level representation of the control plane: the routing algo-

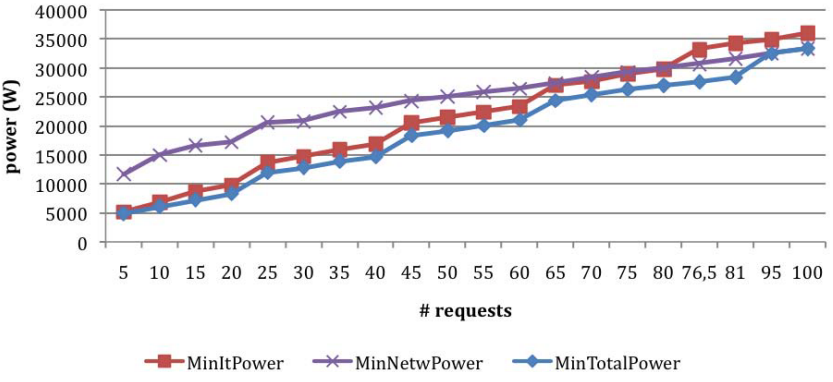

Fig. 4. Total Power consumption.

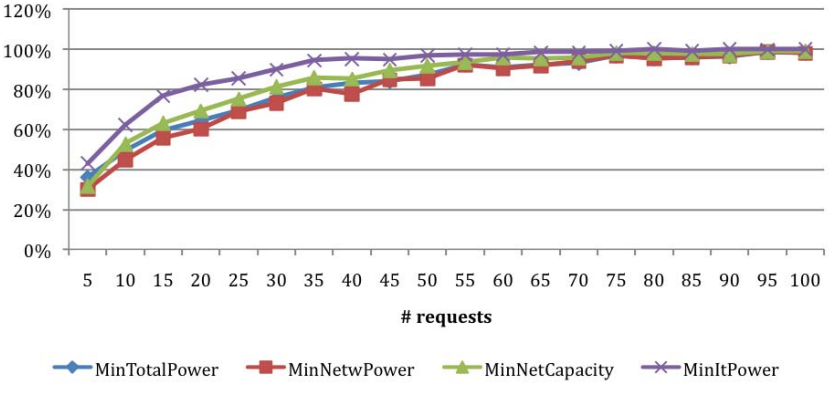

Fig. 5. Number of activated OXCs.

rithms at the PCE operate over a topological graph created combining network and IT parameters with "green" parameters, retrieved from the SML (IT side) and the LICL (network side).

Finally, another key element for the control plane is the interaction with the infrastructure-virtualization layer, in order to trigger the procedures for the Virtual Infrastructure's dynamic re-planning on the network side, besides the IT re-planning. In case of inefficiency of the underlying infrastructure, the control plane is able to request the upgrade or downgrade of the virtual resources, in order to automatically optimize the VI's size to the current traffic load.

5 Energy-Consumption Simulation Results

In this section we illustrate the potential of the proposed architecture in terms of energy-awareness, thereby addressing challenge #4. We evaluated an energy efficient routing algorithm(due to space limitations, the detailed algorithm is not

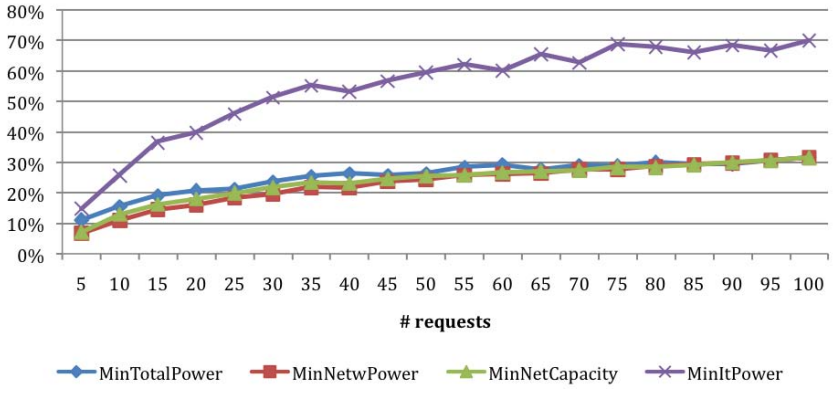

Fig. 6. Number of activated fibers.

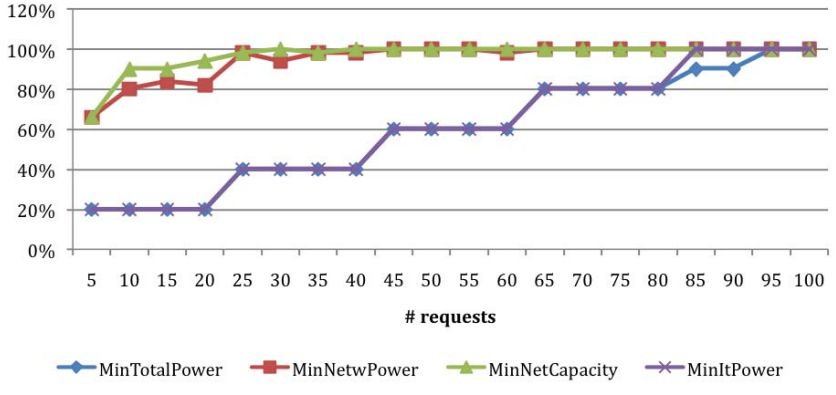

Fig. 7. Number of activated data centers.

shown here) from a networked IT use case: each source site has certain processing demands which need to be satisfied by suitable IT resources (in a data center). Note that we assume anycast routing, implying that the destination IT resource can be freely chosen among the available ones, since the requirement for the particular service is the accurate delivery of results, while the exact location of the processing is of no interest. 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. Simulation results (see Fig. 4-6) indicate that our proposed algorithm can decrease the energy consumption by 10% compared to schemes where only IT infrastructure is considered and up to 50% when taking only the network into account, in an initial sample case study on an European network where we have run the calculations for 10 random demand vectors for each demand size going from 5- up to 100 connections

for which, we have averaged the results[20]. (Please note that the savings are illustrative only, and depend on the case study assumptions.)

6 Conclusion

In this paper we have proposed a new way of configuring optical and IT resources within a new, layered architecture enabling the configuration of a virtual network infrastructure as a whole. Our goal is to address the six most critical challenges the Internet has to face urgently to support emerging disruptive applications and continue to grow safely. This chapter has presented the building blocks of this new architecture: the virtual infrastructure concept, a layered control and management architecture to complement the existing TCP/IP protocol stack and role-based operational model. Each of these building blocks has been further detailed and illustrated by test cases. The overall architectural blueprint complemented by the detailed design of particular components feeds the development activities of the GEYSERS project to achieve the complete software stack and provide the proof of concept of these architectural considerations. One approach of the project is to validate some concepts in a simulation environment. To make this proposal a viable solution for future production networks, a realistic and powerful test-bed will carry promising experiments. One of the test-bed's goals is to serve as a platform to implement and evaluate prototypes of the different software components creating and managing optical virtual infrastructures. The other goal is to evaluate the performance and functionality of such a virtualized infrastructure in a realistic production context.

Acknowledgement. This work has been founded by the EC ICT-2009.1.1 Network of the Future Project #248657.

Open Access. This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Anhalt, F., Koslovski, G., Vicat-Blanc Primet, P.: Specifying and provisioning virtual infrastructures with HIPerNET. Int. J. Netw. Manag. 20(3), 129–148 (2010)
2. Buyya, R.: Market-Oriented Cloud Computing: Vision, Hype, and Reality of Delivering Computing as the 5th Utility. In: Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid, Washington, DC, USA. CCGRID '09, p. 1. IEEE Computer Society Press, Los Alamitos (2009), doi:10.1109/CCGRID.2009.97

[20] **MinTotalPower**: minimizing both network- and IT power; **MinNetworkPower**: only minimizing the network power; **MinNetCapacity**: minimizing the number of wavelengths needed to establish all requested lightpaths; **MinItPower**: only minimizing the energy consumed by the data centers.

3. Clark, D.: The design philosophy of the DARPA internet protocols. SIGCOMM Comput. Commun. Rev. 18, 106–114 (1988), doi:10.1145/52325.52336
4. 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., Leenheer, M.D., Develder, C., Tzanakaki, A., Robinson, P., Brogle, M., Bohnert, T.M.: GEYSERS: A Novel Architecture for Virtualization and Co-Provisioning of Dynamic Optical Networks and IT Services. In: ICT Future Network and Mobile Summit 2011, Santander, Spain (June 2011)
5. Farrel, A., Vasseur, J.P., Ash, J.: A Path Computation Element (PCE)-Based Architecture. RFC 4655 (Informational) (Aug 2006), http://www.ietf.org/rfc/rfc4655.txt
6. Handley, M.: Why the Internet only just works. BT Technology Journal 24, 119–129 (2006), doi:10.1007/s10550-006-0084-z
7. Koslovski, G., Vicat-Blanc Primet, P., Charão, A.S.: VXDL: Virtual Resources and Interconnection Networks Description Language. In: Vicat-Blanc Primet, P., Kudoh, T., Mambretti, J. (eds.) GridNets 2008. LNICST, vol. 2, Springer, Heidelberg (2009)
8. Mannie, E.: Generalized Multi-Protocol Label Switching (GMPLS) Architecture. RFC 3945 (Proposed Standard) (Oct 2004), http://www.ietf.org/rfc/rfc3945.txt
9. Nelson, M.R.: The Next Generation Internet, E-Business, and E-everything, http://www.aaas.org/spp/rd/ch20.pdf
10. Pickavet, M., Vereecken, W., Demeyer, S., Audenaert, P., Vermeulen, B., Develder, C., Colle, D., Dhoedt, B., Demeester, P.: Worldwide energy needs for ICT: The rise of power-aware networking. In: Advanced Networks and Telecommunication Systems, 2008. ANTS '08. 2nd International Symposium on. pp. 1–3 (2008)
11. Rosenberg, J., Mateos, A.: The Cloud at Your Service, 1st edn. Manning Publications Co., Greenwich (2010)