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)

D5.1.1
Test-bed implementation update

Version 1.0

Due date: 31/06/12
Submission date: 15/08/12
Deliverable leader: TID
Author list: José Ignacio Aznar, Luis Miguel Contreras, Juan Rodríguez, Amanda Azañón, Sergio Martínez de Tejada (TID), Bartosz Belter, Damian Parniewicz, Łukasz Łopatowski, Jakub Gutkowski, Artur Binczewski (PSNC), Eduard Escalona (UESSEX), Monika Antoniak-Lewandowska, Łukasz Drzewiecki (TP), Giada Landi, Giacomo Bernini, Nicola Ciulli, Gino Carrozo, Roberto Monno (NXW), Attilio Borello, Sébastien Soudan (Lyatiss), Jens Buysse, Chris Develder (IBBT), Jordi Ferrer Riera, Joan A. Garcia-Espin, Steluta Gheorghiu, 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)

Abstract

This document describes the current status of the testing environment deployed in GEYSERS, whose high level description has been provided in deliverable D5.1 [REF 7] and which comprises the validated architecture and the prototypes developed within the project. This document updates the status of the local test-beds deployed by each involved partner and the integration of the prototypes developed in WP3 and WP4 which aim to validate the whole architecture through tests that evaluate its features and operability.

This document also provides the specification of the GEYSERS Demonstrators and matches them to the test-cases described in deliverable D1.5 [REF 1]. This matching covers most of the test-cases demonstrating the benefits of GEYSERS architecture. Moreover, this document describes the development plan for each of the GEYSERS Demonstrators and a detailed description of their physical environment and the virtualization planning process of such physical resources available at the end of M30.

Table of Contents

Figure Summary

Table Summary

Acronyms and Abbreviations

AAI	Authentication and Authorization Infrastructure
API	Application Program Interface
C-VLAN	Client VLAN
CAPEX	Capital Expenditure
CCI	Connection Controller Interface
CPE	Customer Premises Equipment
CPU	Central Processing Unit
DCVM	Datacentre Virtualization Management
DWDM	Dense Wavelength Division Multiplexing
EIS	Enterprise Information System
FS	Fibre Switching
FSC	Fibre Switching Capable
GLIF	Global Lambda Infrastructure Facility
GMPLS	Generalized Multi-Protocol Label Switching
GRE	Generic Routing Encapsulation
IP	Internet Protocol
L2VPN	Layer 2 VPN
LAN	Local Area Network
LICL	Logical Infrastructure Composition Layer
LSC	Lambda Switching Capable
LSP	Label Switched Path
MEMS	Micro Electro-Mechanical System
MLI	Management to LICL Interface
MPLS	Multiprotocol Label Switching
NAT	Network Address Translation
NCP	Network Control Plane
NFS	Network File System
NIPS	Network + IT Provisioning System
OPEX	Operational Expenditure
OSGi	OSGi Alliance (formerly Open Services Gateway Initiative)
PCE	Path Computation Engine
PIP	Physical Infrastructure Provider
PIP-IT	Physical Infrastructure Provider – IT
PIN-N	Physical Infrastructure Provider – Network
Q-in-Q	Ethernet networking standard formally known as IEEE 802.1ad
RAM	Random Access Memory
ROADM	Reconfigurable Optical Add Drop Multiplexer
ROI	Return of Investment
S-VLAN	Service VLAN
SCN	Signalling Communication Network
SDH	Synchronous Digital Hierarchy
SLI	SML to LICL Interface
SML	Service Middleware Layer
TDM	Time Division Multiplexing
UNI	User to Network Interface
VI	Virtual Infrastructure

VIO	Virtual Infrastructure Operator
VIO-IT	Virtual Infrastructure Operator – IT
VIO-N	Virtual Infrastructure – Network
VIP	Virtual Infrastructure Provider
VLAN	Virtual Local Area Network
VM	Virtual Machine
vNIC	Virtual Network Interface Card
VPN	Virtual Private Network
VR	Virtual Resource
WSDL	Web Services Description Language

1 INTRODUCTION

This document describes the current status of the test-bed deployed in GEYSERS as introduced in deliverable D5.1 [REF 7] by presenting the status of the local test-beds deployed by each involved partner and the related prototypes integration developed in WP3 and WP4. The focus of these 2 work packages is on validating the entire architecture through specific tests used to evaluate respective features and operability.

The main goal of this deliverable is to provide a complete specification of test-bed facilities available for GEYSERS partners in order to implement the GEYSERS architecture in the test environment and validate it with a set of well-defined procedures, as part of the GEYSERS Demonstrators. The deliverable reports how the test-bed resources have been allocated and partitioned to satisfy all Demonstrators' needs.

The document is composed of the following sections:

Section 2 presents the details of local test-beds provided by each partner. This section includes the deployment of local control planes for test-bed infrastructures, which are responsible for the configuration and provisioning of particular infrastructure layers. Local test-bed infrastructures include optical layer 1 hardware, layer 2 switching capabilities, and various IT hardware, including computational and storage facilities.

In Section 3 a reference test-bed is explained. The section provides an explanation how a joint provisioning of network and IT resources is performed. It conceptually shows how PIP resources are virtualized by means of a smooth collaboration among GEYSERS specific modules and the agreed upon virtualization tools (e.g. OpenNebula, KWM, etc.). The technology offered by each telecoms Service Provider, the tools used by each IT provider and the role of the SML/LICL and NCP+ in the provisioning of IT and network services are also explained in this theoretical use case. The joint virtualization of resources constitutes a key item from the lower layer (PIP) to the upper layers of the value chain (VIP and VIOs) with the requested services performing GEYSERS prototypes features and operation actions.

Section 4 provides information on the GEYSERS Demonstrators, which have been selected to validate and demonstrate the key features of the GEYSERS architecture. The matching of GEYSERS Demonstrators to the test-cases described in deliverable D1.5 [REF 1] covers most of the test-cases demonstrating the benefits of GEYSERS architecture. This section also provides a detailed description of the Demonstrator physical environment and how the project plans to use these physical resources. The description of each of the four identified Demonstrators includes details of the physical test-bed provided, the joint virtualization plans carried out for the IT and network resources and the virtualized resultant scenarios.

2 GEYSERS GLOBAL TEST-BED DESIGN

2.1 Introduction

This section describes the physical local test-bed infrastructure (including both network and IT resources) deployed by each partner, in order to perform the validation procedures on the GEYSERS prototypes developed in WP3 and WP4. The local test-bed facilities have been distributed among the GEYSERS Demonstrators which match the use cases and scenarios presented in deliverable D1.5 [REF 1]. Local test-beds include optical layer 1 hardware, layer 2 switching capabilities, and various IT hardware, including computational and storage facilities. This section updates the test-bed implementations described in deliverable D5.1 [REF 7] providing a higher level of detail of the specific devices, ports and connectivity, etc.

Figure 1 shows the global picture of the local test-beds integration. A more detailed description of the infrastructure in each local test-bed is given in the remainder of this section. This section concludes with a description of the interconnections between the local test-beds which includes both data plane and control/management plane capabilities.

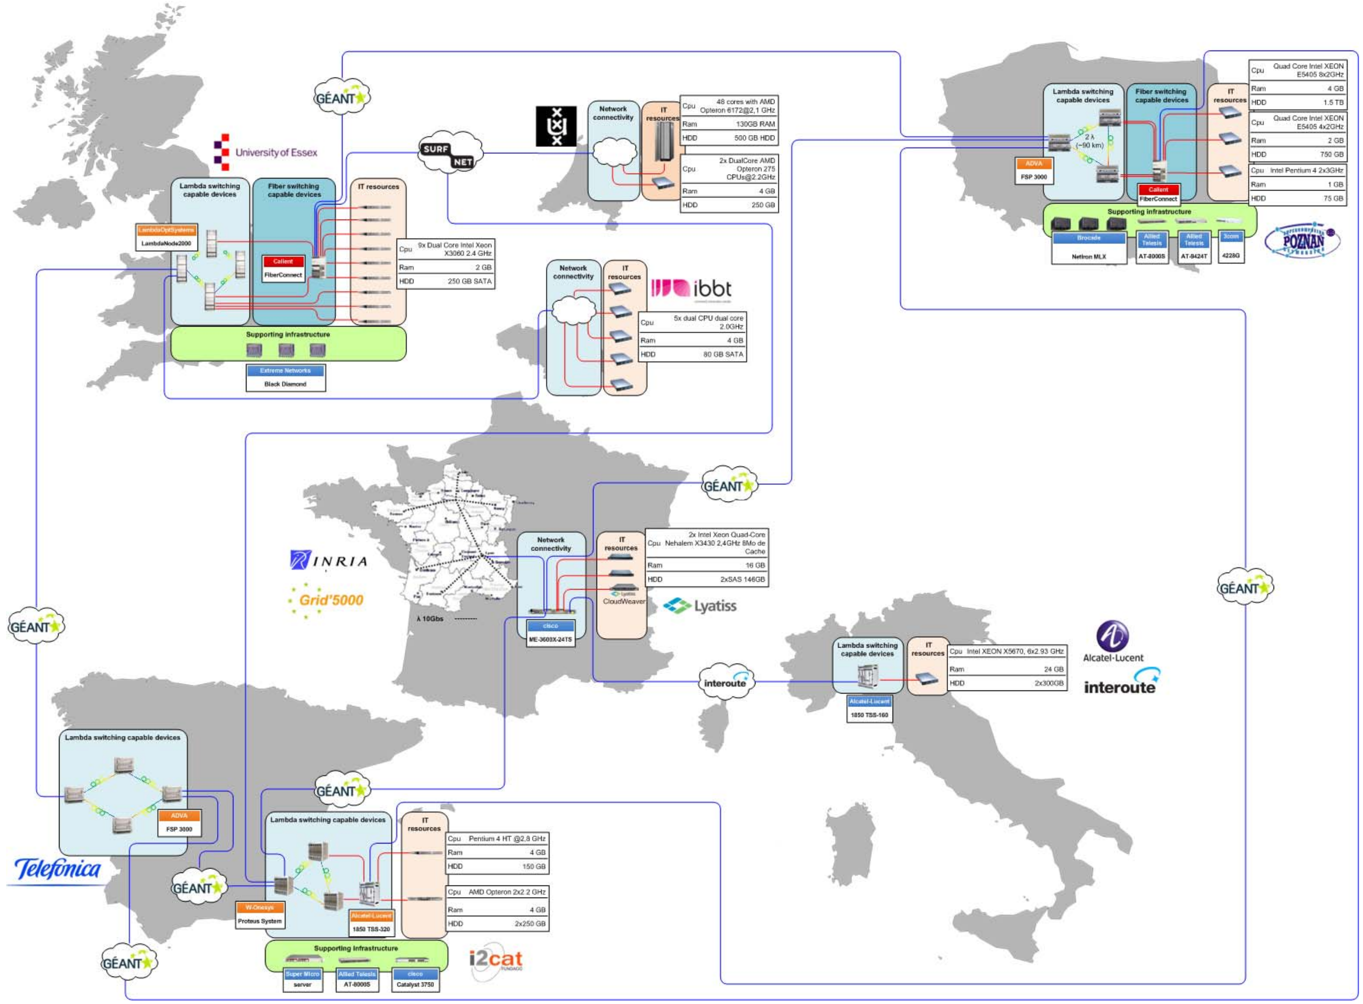

Figure 1: General local test-beds overview

2.2 GEYSERS local test-beds updates

2.2.1 Lyatiss local-test-bed

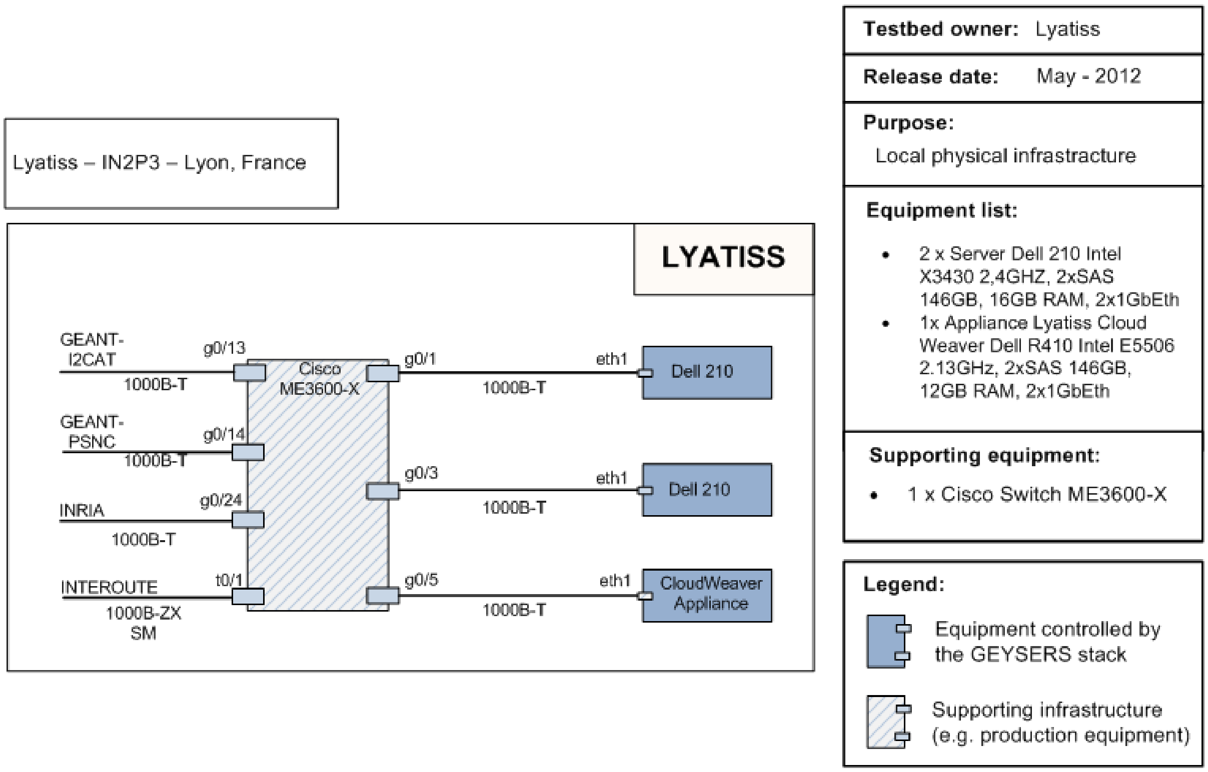

Figure 2: Lyatiss local test-bed details

The Lyatiss test-bed is composed of IT resources, two Dell R210 servers are installed in the IN2P3 (Lyon) Datacenter and KVM is the Hypervisor. Demonstrators can deploy Virtual Machines (VMs) and the GEYSERS stack software needed on these servers.

The technical details of each server are as follows:

- Qt.2 Dell 210 Intel Xeon Quad-Core Nehalem X3430 2,4GHz 8MB Cache, 2xSAS de 146GB 15Ktpm, SAS6/iR (Raid 0 or Raid 1), 16GB DDR-3 à 1066MHz, 2x1GbEth 1000 base-T.

The test-bed provides also a switch Q-in-Q capable for local facilities and also to interconnect the Lyon test-bed with the others GEYSERS partners.

As mentioned in the GEYSERS deliverables D2.6 [REF 3] and D3.3 [REF 5], the Lower-LICL provisions IT resources requested by the Upper-LICL. Lyatiss' CloudWeaver does the planning and provisioning in the Lower-LICL, which includes the states of the physical resources, the allocation of the virtual resources, the triggering of the instantiations and the synchronisation of their states. In this test-bed, Lyatiss CloudWeaver [REF 9] is installed as a VM hosted in a dedicated physical appliance. OpenNebula [REF 8] is deployed as the hypervisor controller. The appliance's technical details are:

- 1x Appliance Lyatiss Cloud Weaver Dell R410 Intel E5506 2.13GHz, 2xSAS 146GB, 12GB RAM, 2x1GbEth.

2.2.2 University of Essex local-test-bed

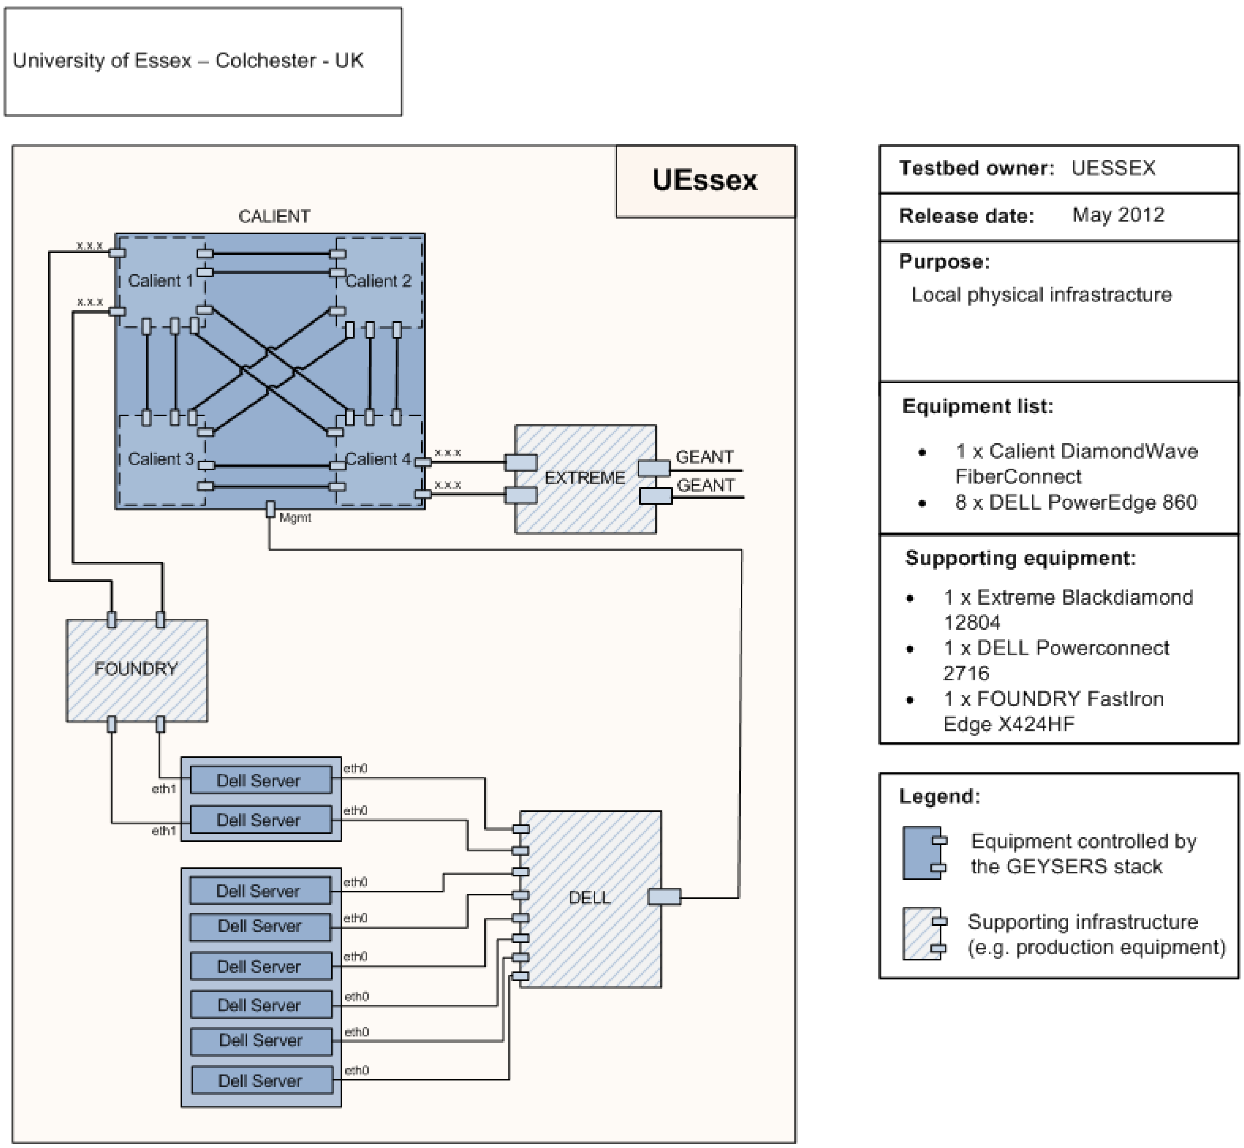

Figure 3: University of Essex local test-bed details

The devices of the UEssex test-bed that will be used for the Demonstrators consist of a Calient DiamondWave Fiber Connect optical switch that provides the FSC data plane infrastructure. The Calient switch is manually partitioned into 4 sub-switches to form a simple mesh network topology. Regarding IT resources, UEssex provides 8 DELL PowerEdge 860 servers, 6 of which will be used for the deployment of the LICL and NCP+ controllers and the other 2 for the actual implementation of the PIP-IT part of the test-bed. In addition, three supporting devices will be used in the test-bed: an Extreme Blackdiamond 12804 will connect the Calient switch to the rest of the test-beds through GEANT. A DELL Powerconnect 2716 will be used to provide SCN connectivity to the LICL and NCP+ controllers, and finally a FOUNDRY FastIron Edge X424HF switch will be used to connect the IT resources to the Calient switch.

2.2.3 IRT/ALU-I local-test-bed

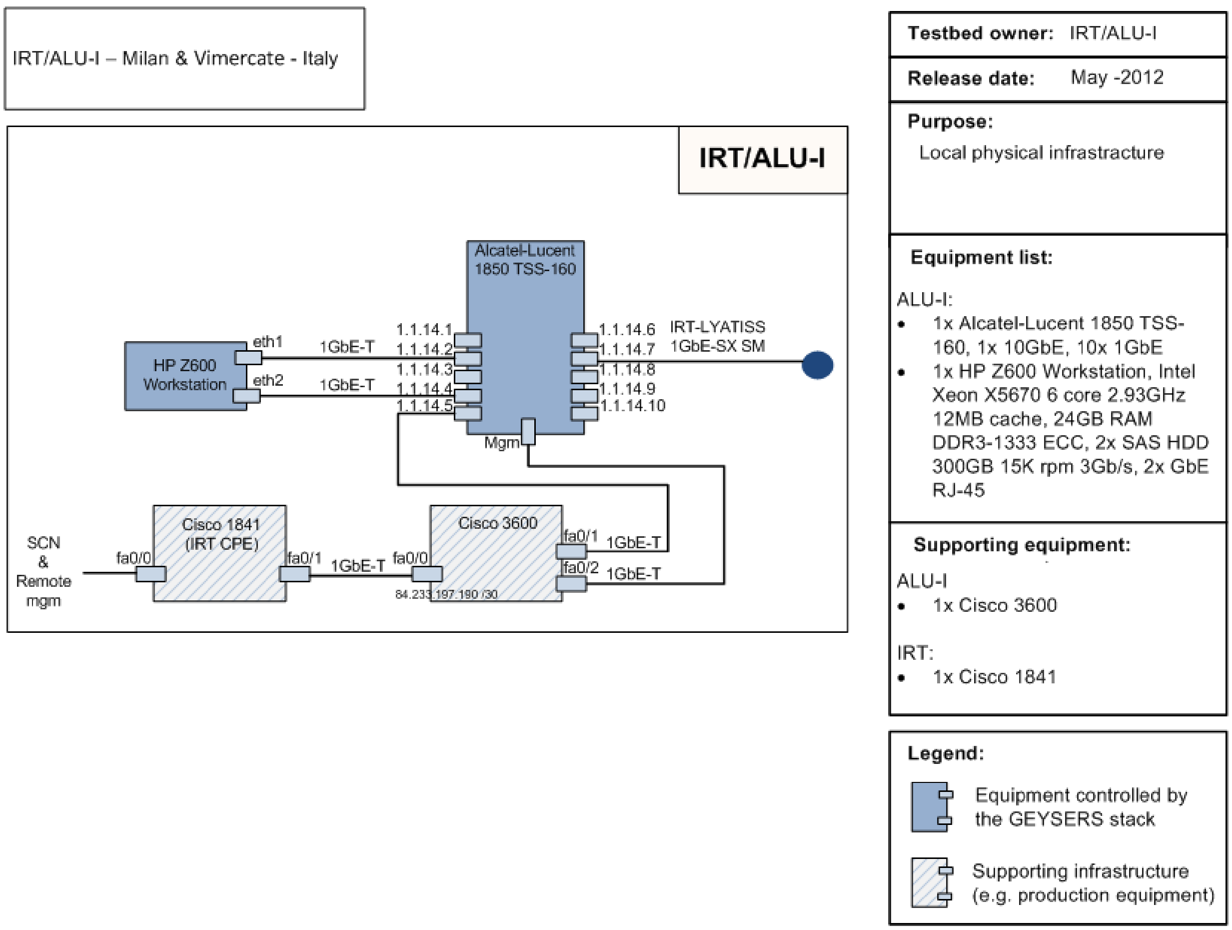

Figure 4: Interoute and ALU-I local test-bed details

IRT and ALU-I jointly deploy a local test-bed for the GEYSERS project, composed of an optical network node and an IT server, plus other devices that act as supporting infrastructure. Figure 4 depicts the physical infrastructure of this local test-bed.

IRT provides connectivity towards the Lyatiss test-bed in Lyon via a 1Gbps Ethernet over an MPLS circuit provisioned over its optical network. Moreover, it offers co-location services for ALU-I equipment and public IP Internet access. In fact, all the equipment is installed at IRT premises in Milan Caldera, Building C. A Cisco 1841 router acts as a CPE and it is deployed as an IRT standard solution for managed public IP Internet access.

The other equipment is provided by ALU-I. The optical network node is an Alcatel-Lucent 1850 TSS-160, a Packet-Optical Transport switch of the Alcatel-Lucent TSS product family. The IT server is an HP Z600 Workstation, which will run several VMs, plus a Lower-LICL instance. ALU-I also deployed a Cisco 3600 router, as a support device for remote management purposes: it terminates VPN tunnels at the ALU-I premises in Vimercate, and performs related NAT functionalities.

2.2.4 PSNC local-test-bed

PSNC's local test-bed (Figure 5) is composed of optical network resources. The test-bed is located at PSNC's optical laboratory and enables the creation of several VIs using both Fiber (FSC) and Lambda (LSC) Switching Capable devices.

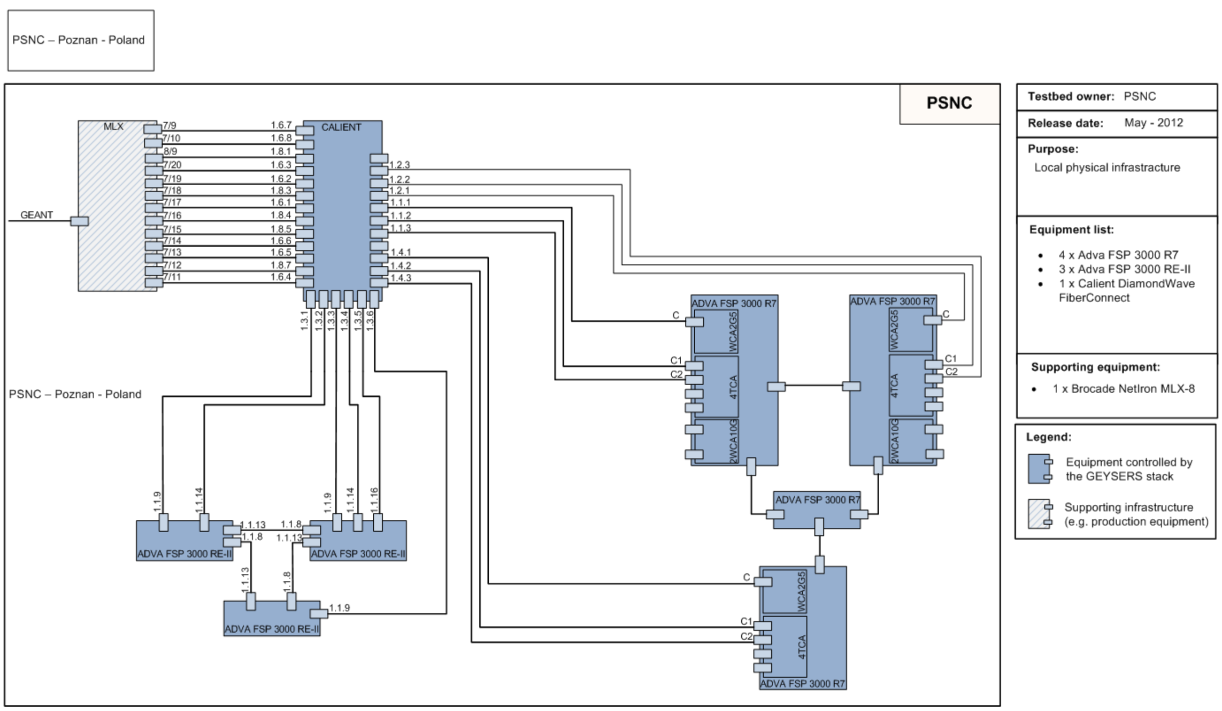

Figure 5: PSNC's local test-bed details

The former resource type is represented by the Calient DimondWave FiberConnect switch manually partitioned into four isolated logical Calient devices (for more details see the following subsection). The main part of this optical switch is a non-blocking matrix with MEMS mirrors technology which, when fully equipped, can handle up to 640 fiber terminations per shelf.

The latter resource type is represented by ADVA equipment. PSNC's test-bed consists of three legacy ADVA FSP 3000 RE-II and four ADVA FSP 3000 R7 devices. Both types are DWDM devices based on Reconfigurable Optical Add/Drop Multiplexing (ROADM) technology.

Additionally PSNC's test-bed provides one Brocade NetIron MLX-8 router which enables the connection with the remote GEYSERS test-beds and the creation of several demo test-beds using Q-in-Q technology. For more technical details on the equipment located at PSNC's test-bed, please refer to deliverable D5.1 [REF 7] Appendix A.1.5.

LOCAL TEST-BED SPECIFIC DETAILS

Figure 6 presents the Calient DimondWave FiberConnect device which was partitioned into four isolated logical Calient devices. In order to achieve this, several ports were interconnected. This procedure allows for a more complicated test-bed topology using only one available physical box. The benefits of this solution are mainly consumed by the Demonstrator 1.

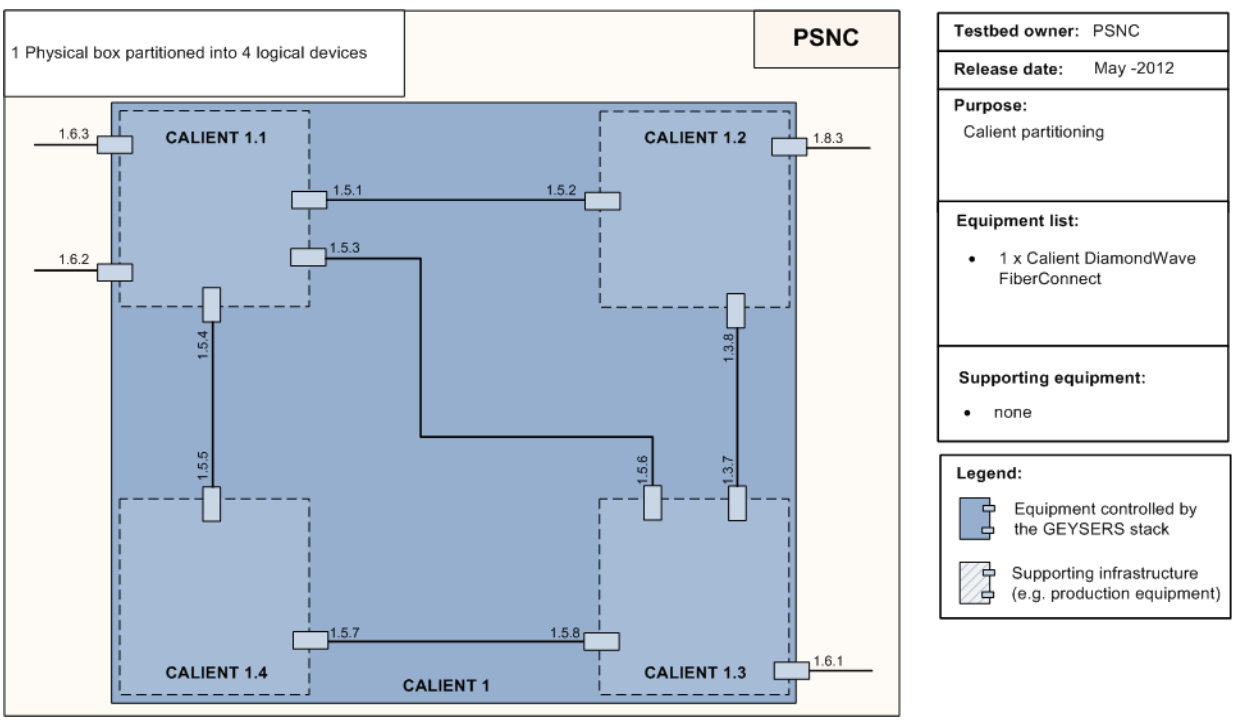

Figure 6: Calient DiamondWave FiberConnect partitioning

The two following figures show the LSC domain of the PSNC's test-bed. Figure 7 presents four interconnected ADVA FSP 3000 R7 devices. Three of them are equipped with client cards which are connected through the Calient and MLX-8 boxes to the remote test-beds. The fourth box, located in the middle, consists of three 8ROADM modules (one for each direction). This configuration enables up to three 1Gbit unprotected tunnels using two lambdas, one for WCA2G5 card and one for 4TCA (TDM card).

Figure 8 introduces the remaining part of the LSC domain at the PSNC test-bed. The three legacy ADVA FSP 3000 RE-II devices are connected to the whole GEYSERS test-bed in the same manner as described above. Available resources allow for the creation of a DWDM ring with three lambdas as shown in the picture.

The former technology will be mainly used by the Demonstrator 2, while the latter technology will be used by the Demonstrator 4.

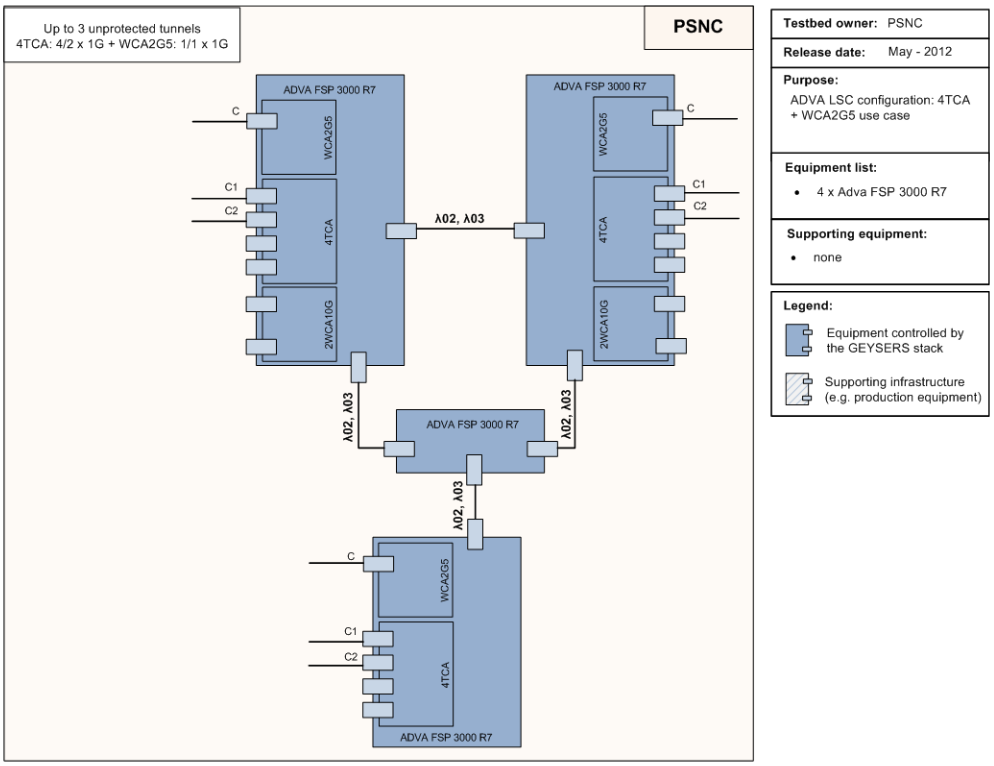

Figure 7: ADVA FSP 3000 R7 DWDM ring

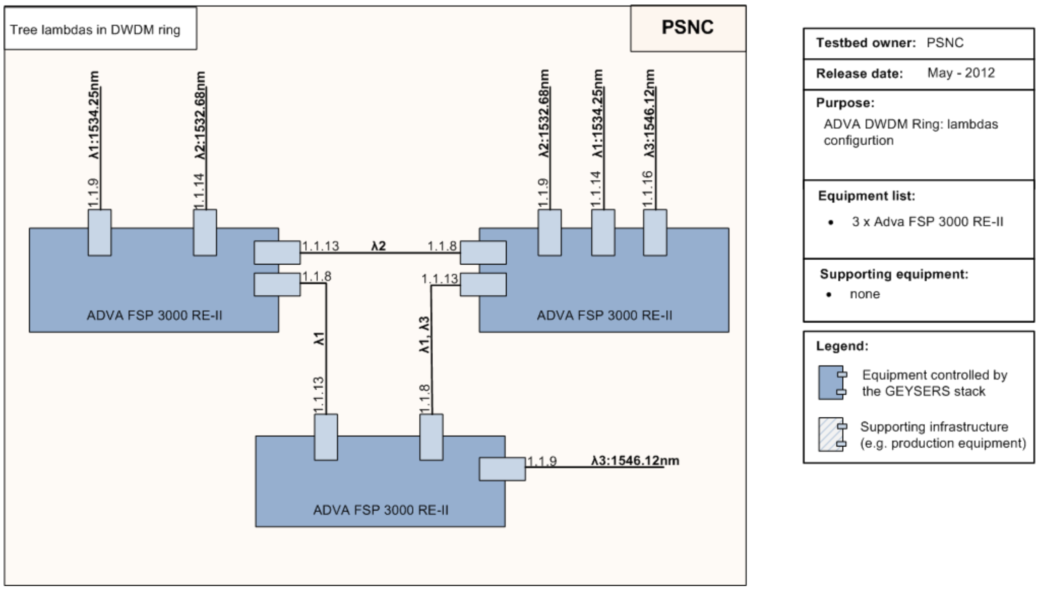

Figure 8: ADVA FSP 3000 RE-II DWDM ring

2.2.5 TID local-test-bed

The TID local test-bed comprises both IT and network resources. IT resources consist of two servers: DELL POWER EDGE and HP DL380. The first one is dedicated to IT resources, while, the second one is dedicated to Parent-PCE and Child-PCE servers, GMPLS+ controllers, NIPS server, and Lower LICL modules. OpenNebula software will be installed in the Front-End and there is also an image repository where VM templates are stored. Network resources consist of four ADVA FSP 3000 ROADMs. Several 1 Gbit LSPs have been established among them in the data plane. The Lower LICL module will manage these network resources to provide virtualized network resources to the global test-bed. Juniper routers enable the access of clients to the ROADMs. Riverstone switch routers constitute the gateways to the global GEYSERS test-bed for both data and signalling planes.

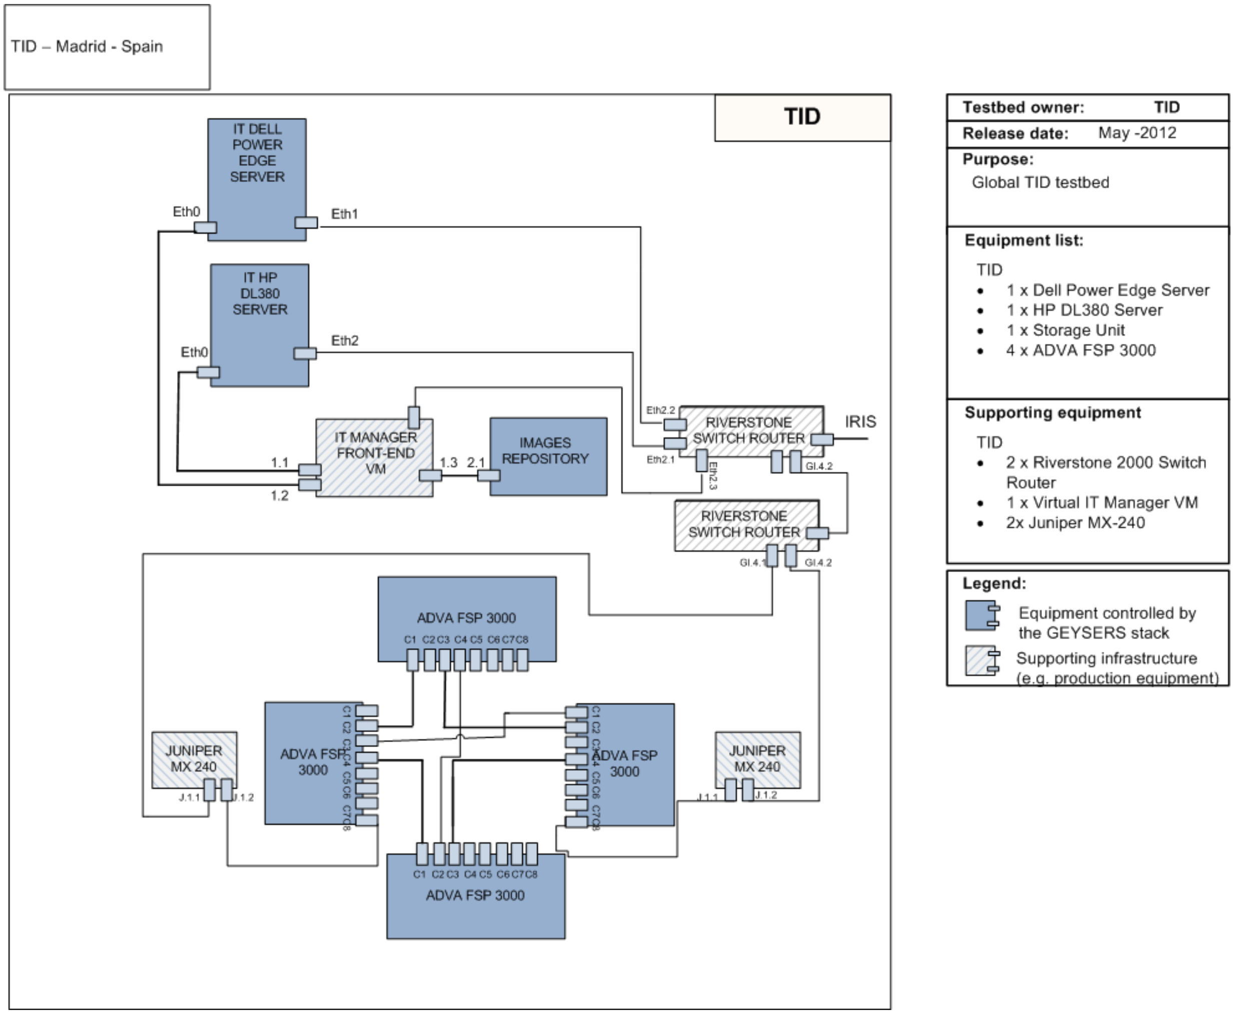

Figure 9: TID local test-bed details

2.2.6 i2CAT local-test-bed

The test-bed in i2CAT is composed of both IT and network resources. One SuperMicro Server will be used for deploying the LICL software, and a SunFire x2200 server will be used for installing OpenNebula. An additional SunFire x2200 together with a Dell PowerEdge 850 server will act as IT resources controlled by the LICL. The network resources consist of three W-Onesys Proteus devices, which make use of two Transmode TM-301 boxes as supporting equipment to provide connectivity to the test-bed.

Several network devices (2 Cisco Catalyst switches, one Alcatel 6850 and two SMC Aggregator devices) are providing the data plane connectivity and the connectivity to other GEYSERS test-beds, while the SCN connectivity is ensured by an Allied Telesis switch.

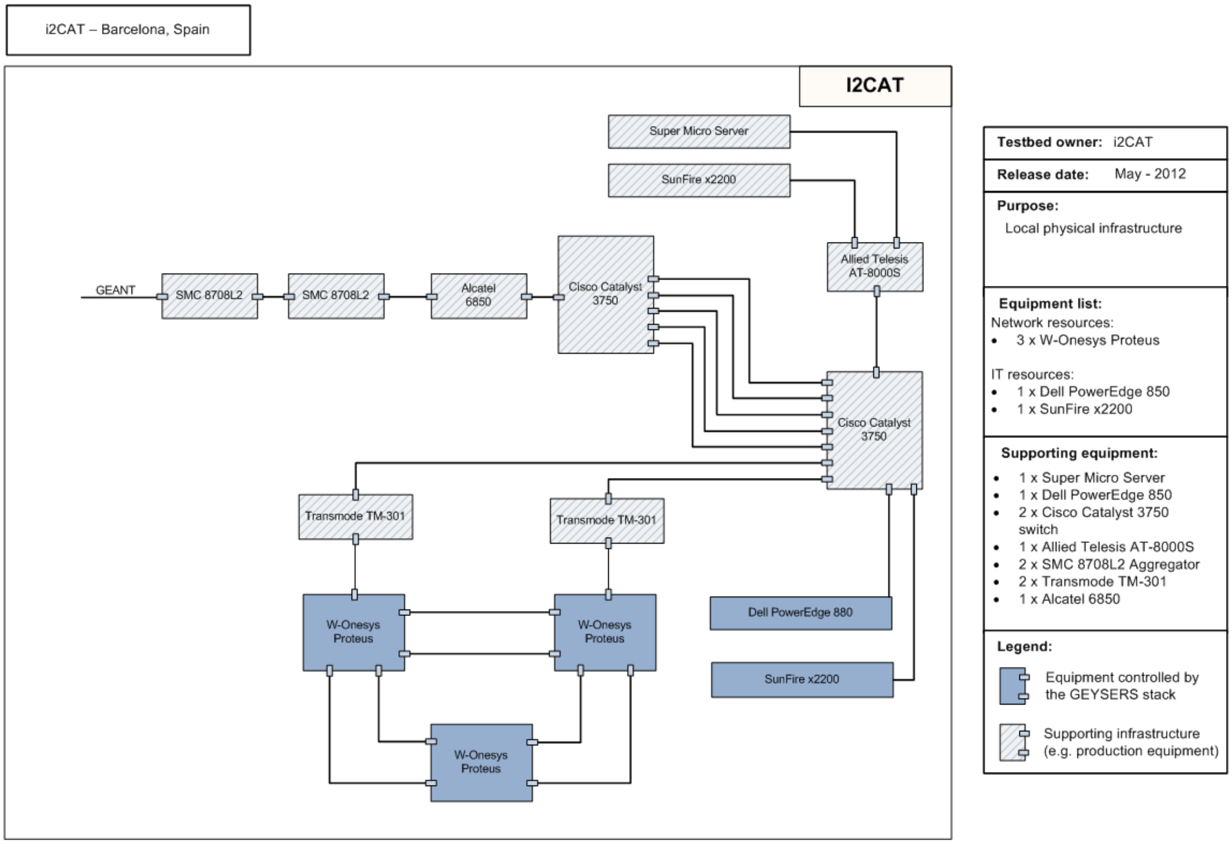

Figure 10: i2CAT local test-bed details

2.2.7 IBBT local-test-bed

IBBT provides 6 nodes with the following specifications:

- 12GB memory.
- 2x Intel Xeon CPU E5620 @ 2.40GHz: a total of 8 cores and hence 16 threads per machine.
- 160GB local storage.
- A shared NFS of 500 GB between the nodes, which is located on a 1Gb linked iSCI server.

One of the servers will have both an OpenNebula installation and act as the NFS server. All other nodes will have a KVM installation.

IBBT – Gante - Belgium

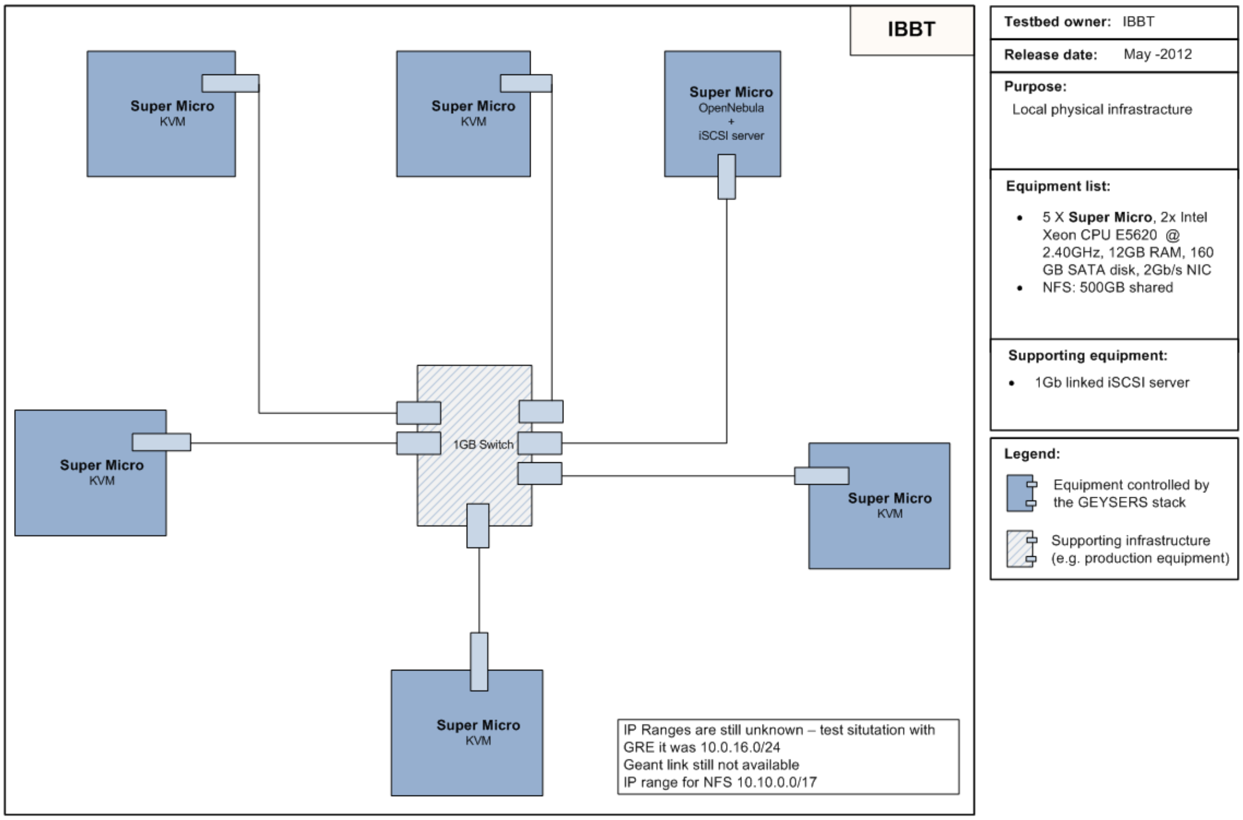

Figure 11: IBBT local test-bed details

2.2.8 UVA local test-bed

The UvA test-bed (Figure 12) provides a cloud-machine (Dell R810: 48cores, 128GB-RAM, 600GB storage) for on-demand computation and storage resources. These resources are provided with the aid of OpenNebula installed and the necessary software to deploy the Lower-LICL. This test-bed provides access via two GLIF lightpaths, to i2CAT and to UEssex. The Dell R810 will be used to provide IT resources (computing and storage) for Demonstrators 1, 3 and 4.

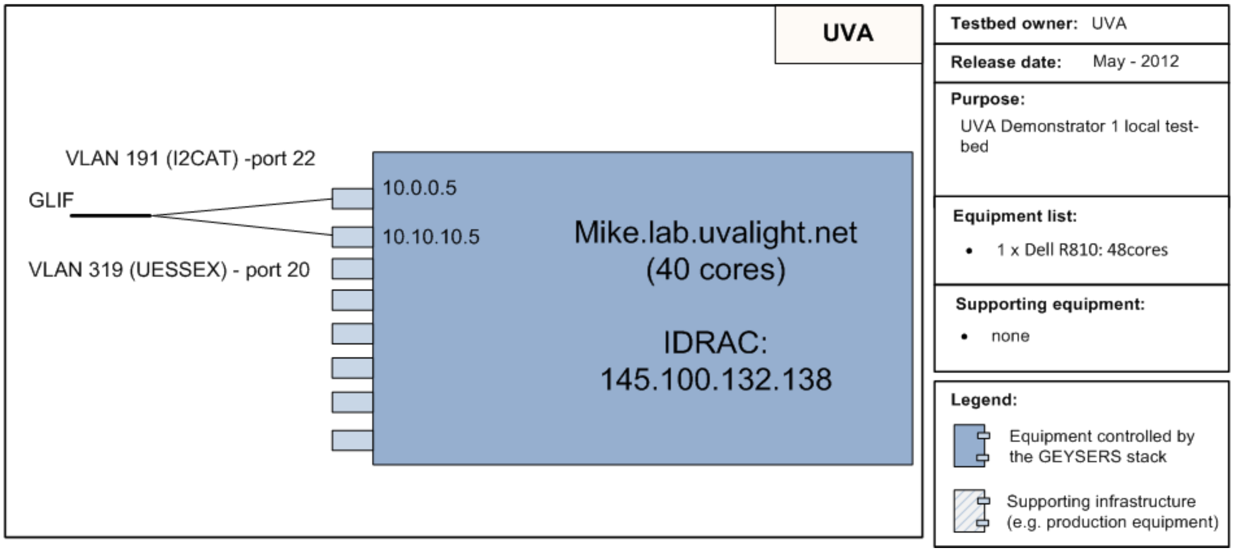

Figure 12: UVA local test-bed details

2.3 Interconnection capabilities between local test-beds update

2.3.1 Data plane inter-connections update

In this sub-section an update in terms of data plane inter-connections is presented. Only the new updates which have taken place in terms of the connectivity between partners' local test-beds with respect to deliverable D5.1 [REF 7] have been included. The remaining data plane inter-connections can be found in the same document.

IRT/ALU – Lyatiss	
Technology	Gigabit Ethernet over SDH
Infrastructure	Interoute private fiber
Status	Operative
Capacity	1Gbps
VLAN transmission	Yes, 1 VLAN, (QinQ available on Lyatiss site)
Notes	Waiting for end-to-end test between IT resources.

Table 1: IRT/ALU – Lyatiss

Local test-bed i2CAT – Local test-bed TID	
Technology	10GbE
Infrastructure	RedIRIS
Status	Operative
Capacity	10 Gbps
VLAN transmission	Yes, 1 VLAN, 688
Notes	-

Table 2: i2CAT – TID

Local test-bed UESSEX – Local test-bed TID	
Technology	GEANT Plus Service
Infrastructure	UESSEX - Janet – GEANT – IRIS - TID
Status	Operative
Capacity	1 Gbps
VLAN transmission	Yes, 1 VLAN, 931
Notes	Working from April 2012

Table 3: UESSEX – TID

Local test-bed TP – Local test-bed PSNC	
Technology	L2VPN
Infrastructure	TP core IP/MPLS network, private fiber
Status	Operative
Capacity	1Gbps
VLAN transmission	Yes, 4096 VLAN's
Notes	Working from June 2012

Table 4: PSNC – TP

2.3.2 Control Plane/Management Plane Update

The control plane to signal among partners' test-beds has been updated along with the data plane. As was previously explained in deliverable D5.1 [REF 7], the control and management plane are based on a "hub-star" topology as observed in Figure 13.

The SCN Main Router is hosted at PSNC's premises. Each partner deployed a VM which acts as a router for local control and management purposes. To make the deployment easier the VM image has been prepared. The person responsible for the VM installation has only to set the configuration file which is used during the booting sequence in order to set a correct network configuration. The communication between all domains is based on GRE tunnels over the Internet. The entire configuration is set as static entries to the routing table of each host so there is no routing protocol involved.

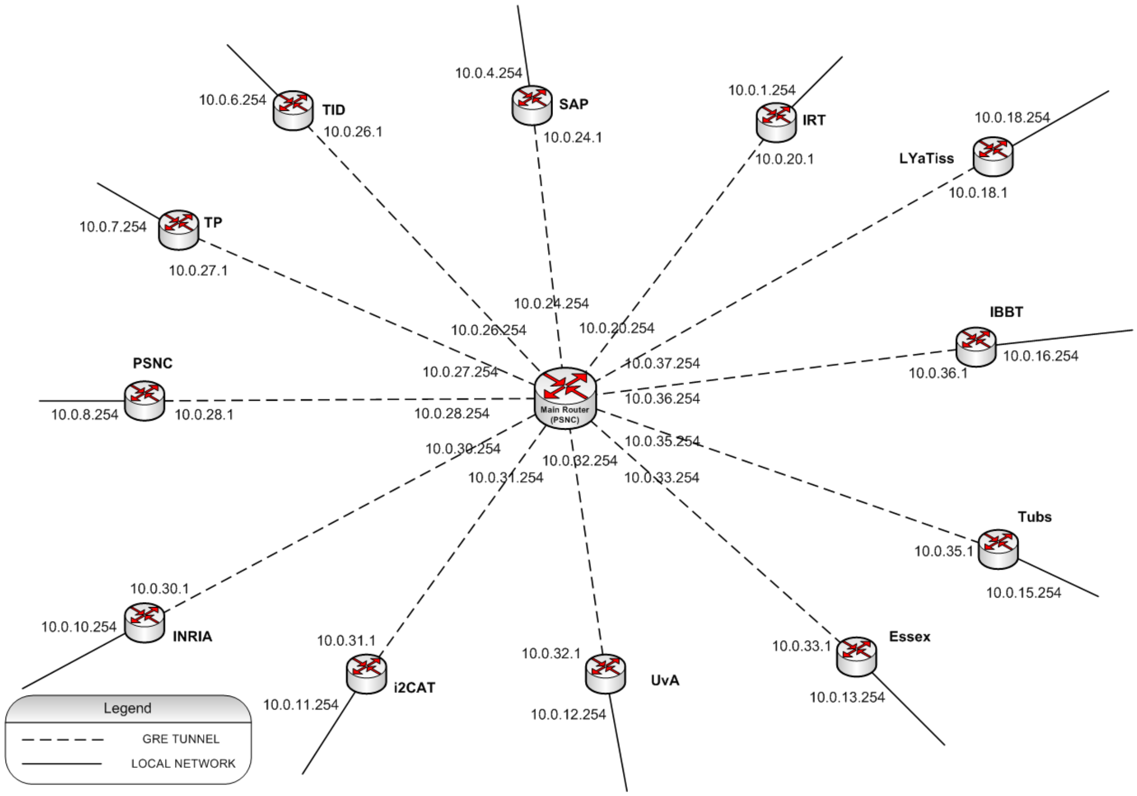

Figure 13: The Control/Management Plane topology in the GEYSERS test-bed

<u>IMPLEMENTATION STATUS</u>

After each partner finalised their local router (SCN Virtual Machine) deployment, the connection status was tested. The tests consisted of checking pings between the SCN addresses. The following example shows which addresses were checked in the case of i2CAT:

- From i2CAT side - 10.0.31.254,
- From Mine Router side - 10.0.11.254.

Table 5 summarises the current status of GEYSERS control plane based on the SCN access routers.

Partner Name	GEYSERS SCN address pool	Public IP Address (physical interface)	Implementation status	Comments
IRT/ALU	10.0.1.0/24	84.233.197.190	to be done	HP Workstation not yet installed in IRT facility in Milan.
SAP	10.0.4.0/24	212.41.199.78	done	
TID	10.0.6.0/24	193.145.240.13	done	
TP	10.0.7.0/24	217.96.70.199	done	
PSNC	10.0.8.0/24	150.254.171.142	done	
i2CAT	10.0.11.0/24	84.88.40.61	done	
UvA	10.0.12.0/24	145.100.132.145	done	
UEssex	10.0.13.0/24	155.245.93.14	done	Not active but solution was checked.
IBBT	10.0.16.0/24	193.191.148.159	done	
Lyatiss	10.0.18.0/24	192.168.100.1/24 STag=235 (Geant) CTag=100	done	
Main Router (PSNC)	x.x.x.x/24	150.254.171.141	pending	The configuration of the Main Router will be finalised when all end points are deployed and connected to the infrastructure.

Table 5: SCN implementation update

3 REFERENCE SCENARIO FOR JOINT PROVISIONING OF NETWORK AND IT RESOURCES

As previously stated in deliverables D2.2 [REF 2] and D3.2 [REF 4], IT resources are typically located in datacentres connected to the optical network infrastructure through a separate L2 or L3 network. The virtualization of this separate network in GEYSERS is out of the scope of development, however, it is effectively carried out by the LICL using OpenNebula functionalities and the actual resource provisioning is also supported by OpenNebula interfaces. This section describes in detail how GEYSERS solves the unified network and IT resource provisioning through the planning and operational phases of the virtual infrastructure lifecycle. With this goal, a reference test-bed scenario will be used to show the virtualization and provisioning of IT resources interconnected to the optical network infrastructure.

3.1 Reference test-bed

The scenario depicted in Figure 14 will be used as the reference to describe how IT resources are connected to the backbone network and the unified provisioning of IT and network resources takes place.

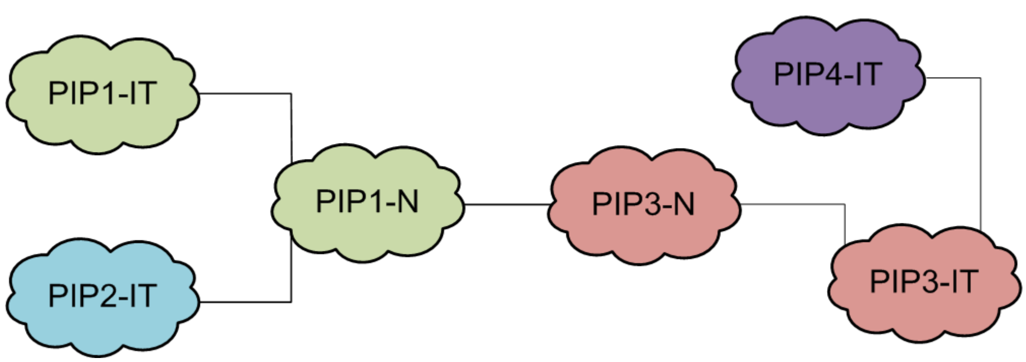

Figure 14: Reference test-bed relative to GEYSERS' adopted roles.

This scenario is composed by four interconnected PIPs. Both PIP1 and PIP3 own an infrastructure with optical network and IT resources while PIP2 and PIP4 are just IT providers (e.g. datacentres). Additionally, PIP4 does not have direct connectivity with the optical network owned by PIP1 or PIP3. It uses a VPN connection (e.g. MPLS L2 VPN) towards the PIP3 IT facilities to make its own resources available. Therefore, logically it can be considered as subset of PIP3-IT. The following figure shows how the GEYSERS test-bed implements the described scenario.

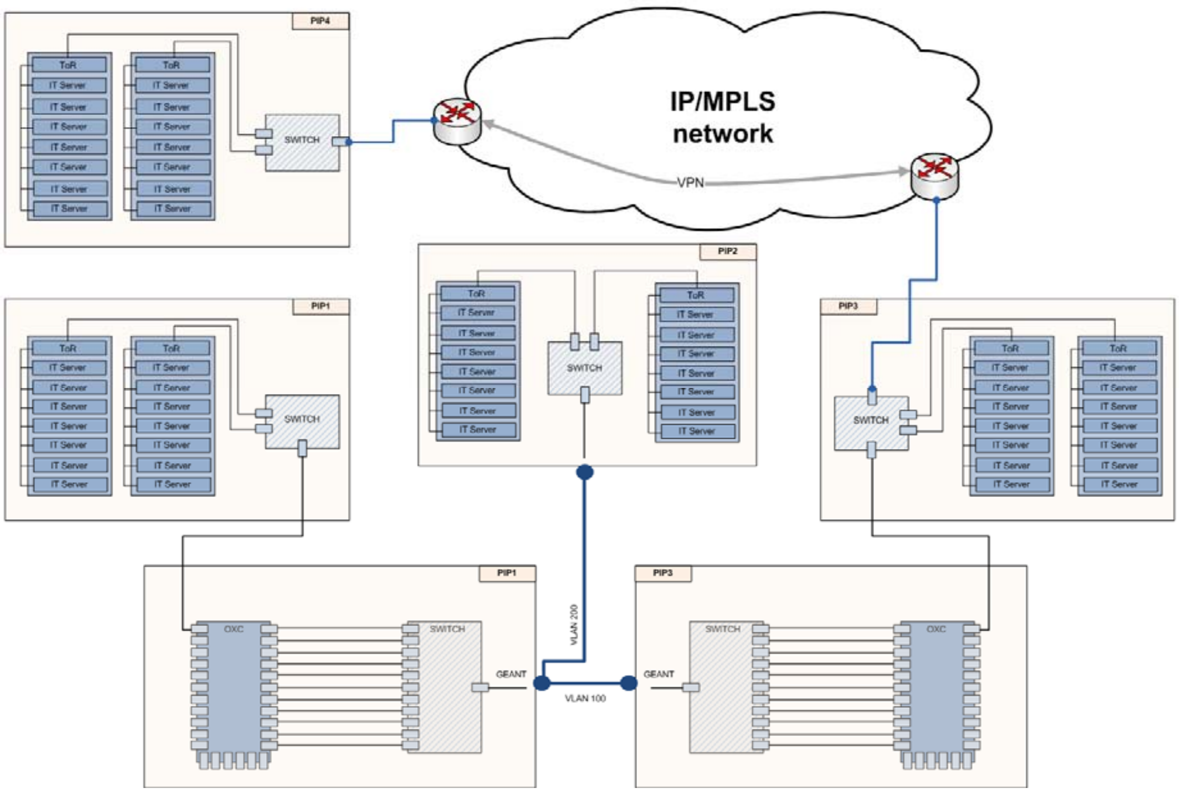

Figure 15: GEYSERS reference test-bed implementation.

In this reference scenario, only the PIPs are depicted. The VIP and VIO roles can be represented by any of the actors in the scenario (any PIP) or different actors with the capabilities and business drivers to implement the functionalities of VIP and VIO roles. These functionalities will be supported by the ability to deploy the Upper-LICL or the NCP+ respectively and the capability to establish business operations with the PIPs and communicate with its respective Lower-LICL systems.

3.2 Network and IT resource virtualization

One of the emergent capabilities of the GEYSERS framework is to enable virtualization technology openness with the approach of physical resource adapters. This means that the lower levels of the stack that implement virtualization functions can be selected from a wide range of technologies existing in the marketplace and as open source today. The physical resource adapters provide a uniform interface that abstracts the details of the underlying approach of virtualization and virtualization management. However, for demonstration and proof of concept purposes, we have committed to specific technologies and assumed certain capabilities in the PIPs implementing resource virtualization, given that they are sufficient to cover the needs of the Demonstrators.

3.2.1 IT Resource Virtualization

IT resource virtualization refers primarily to server or machine virtualization, where it is possible to segment and isolate portions of CPU, memory and storage on a physical server, such that a (different) guest operating system can be run. There are also more specific classes of IT resource virtualization, namely memory and storage.

Figure 16 shows how a server has a Host operating system with multiple VMs running, each representing guest domain 1, ..., n. Domain 0 is reserved for the hypervisor, which translates instructions from VMs to actual physical machine instructions. The figure also shows that each VM only has access to portions of the disk, CPU and RAM on the physical server. As stated before, memory and storage can also be virtualized using other methods besides the single VM instance. Memory virtualization is the decoupling of RAM from single servers in the infrastructure such that they can be made available as a pool, while storage virtualization is the combination of one or more physical storage devices that appear to applications and users as a single device. Virtualization is hence used either for the segmentation or clustering of resources. We typically refer to IT resource virtualization from the segmentation perspective within GEYSERS and the Demonstrators. Figure 16 shows the general layout of a PIP-it with the minimum capabilities expected for virtualization.

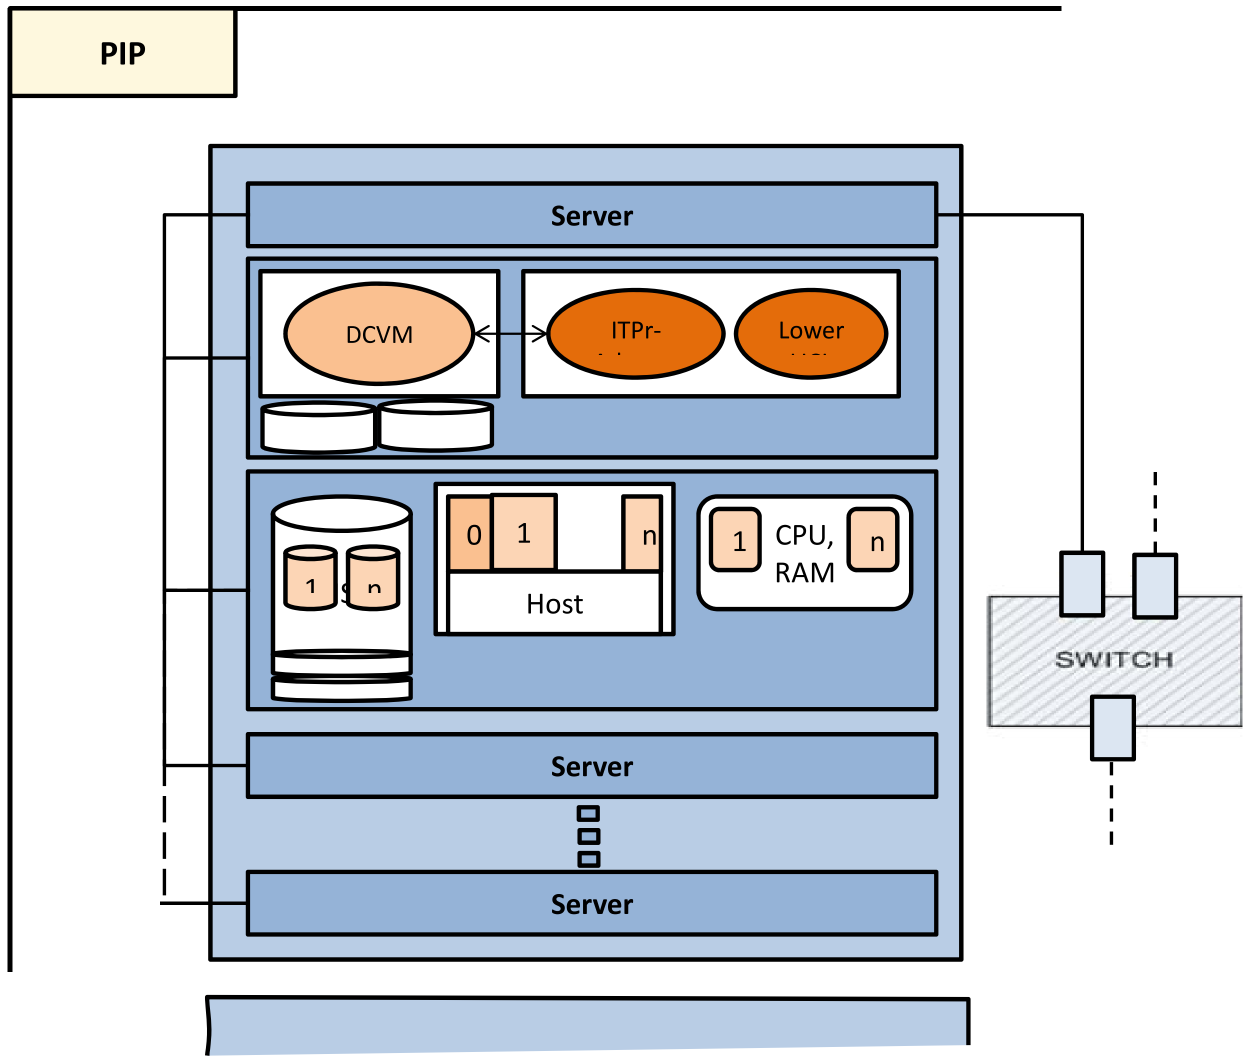

Figure 16: A Virtualization-enabled PIP with a Datacentre Virtualization Management (DCVM) solution deployed

The point of interaction between the PIP's native resource management and the GEYSERS Lower-LIC is the Datacentre Virtualization Management (DCVM). This is interfaced with the IT PR Adapter (an adapter for physical resource control), such that the virtualization operations and attributes of the LICL are constrained by the DCVM. The reference DCVM used and for which we have implemented and tested an IT PR Adapter is OpenNebula, an open source offering now managed by C12G labs. OpenNebula orchestrates hypervisor control, VM images distribution, network isolation, monitoring, and security technologies that enable the dynamic placement of multi-tier services (groups of interconnected VMs) on distributed infrastructures, combining both datacentre resources and remote cloud resources, according to allocation policies. Besides OpenNebula, some partners have chosen CloudWeaver as it extends some OpenNebula features.
DCVMs such as OpenNebula manage multiple hosts, which are physical servers or clusters with virtualization capability. Each host is registered with the DCVM to provide information about host to access, how to monitor, manage storage and execute virtualization mechanisms. For example, with OpenNebula the hostname, information driver (for monitoring), storage driver (e.g. nfs) and virtualization driver are used to register a physical host. Based on the type of virtualization

driver a set of operations are possible, such as start, stop, resume, delete and migrate. The libvirt API [REF 10] is a standard for virtualization operations. There is need for only one type of support and protocol in the DCVM, if all physical hosts implement a libvirt driver for managing the VM lifecycle. OpenNebula has drivers for KVM [REF 11], XEN [REF 12] and VMWare [REF 13]. All partners providing an infrastructure have agreed to use OpenNebula 3.2 as the DCVM and KVM as the hypervisor.

Nevertheless, in some Demonstrators, Lyatiss CloudWeaver will be used as an alternative to OpenNebula since it extends OpenNebula's features and adds the control of the datacentre network between the IT resources of a given datacentre and reaches the optical network as provisioned by the LICL. This enables users to have an end-to-end control of the network needed to communicate between the IT resources. The control of the datacentre network is provided and network resources are dynamically allocated to improve the performance of the applications, thus extending GEYSERS VI concept end-to-end. For instance, waste of performance and time occurs if a VM machine is deployed for a fixed amount of time and the network is not controlled and provisioned after its expiration. CloudWeaver solves these issues by controlling the datacentre cloud network and enables users to add, remove, and change virtual links (e.g. bandwidth) in order to optimise the datacentre resources utilisation and maintain the network in an optimal state in accordance with the VM/Application state.

Another useful feature is the traffic isolation and capacity sharing within the datacentre. CloudWeaver dynamically segments the network (L2 in the Demonstrator context) to build different and isolated Virtual Infrastructures (Vis) and extends the GEYSERS VI concept to the datacentre. Each VI has a predetermined amount of bandwidth between VMs and for every VI the QoS and QoE are preserved during the deployment. Consider, for example, a 1Gbps link, where CloudWeaver provisions and deploys two VIs: one with 600 Mbps and the other with 300Mbps. The two VIs are then completely isolated and the bandwidth is granted according to their requests. CloudWeaver can dynamically re-deploy the parameters if the conditions change due to the virtualized application or the administrator demands more (or less) bandwidth.. When the deployment ends the allocated resources are automatically released and they are available again to provision new VIs.

3.2.2 Network Resource Virtualization

In the context of GEYSERS, network resources are optical nodes and optical links. They compose a network infrastructure that interconnects the IT resources located at the edges in order to support a certain service. In contrast to IT resources, where storage and computing are treated in a similar way, different optical network virtualisation paradigms have different particularities depending on the bandwidth granularity or switching technology to be used. The two main virtualisation elements in GEYSERS for optical networks are at the optical node and optical link (fibre and wavelength) level. Optical node virtualisation can be achieved by partitioning or aggregation. Optical node partitioning consists of the division of the optical device into smaller units by assigning different ports to each virtual instance. The separation and isolation between the controls of each instance depends on the virtualisation capabilities of the device itself. The aggregation of optical nodes consists of the presentation of several physical nodes as a super-node exposing a single management interface towards the upper layers (NCP). This aggregation also includes the links required to interconnect the aggregated physical nodes. The LICL hides the complexity of the virtual node’s internal connections and shows it to the NCP as a unique switching matrix using the external interfaces.

3.2.3 Connecting Virtual IT Resources and Virtual Networks in a Virtual Infrastructure

In order for the virtual IT resources (VMs) controlled by OpenNebula to be connected to the actual network infrastructures, several requirements need to be met beginning with the physically related ones (like connecting the servers to VLAN-trunk enabled switches) and continuing with software requirements for the configuration of the servers which will be used for hosting the VMs. The picture below summarises how the VMs are actually connected to the physical network.

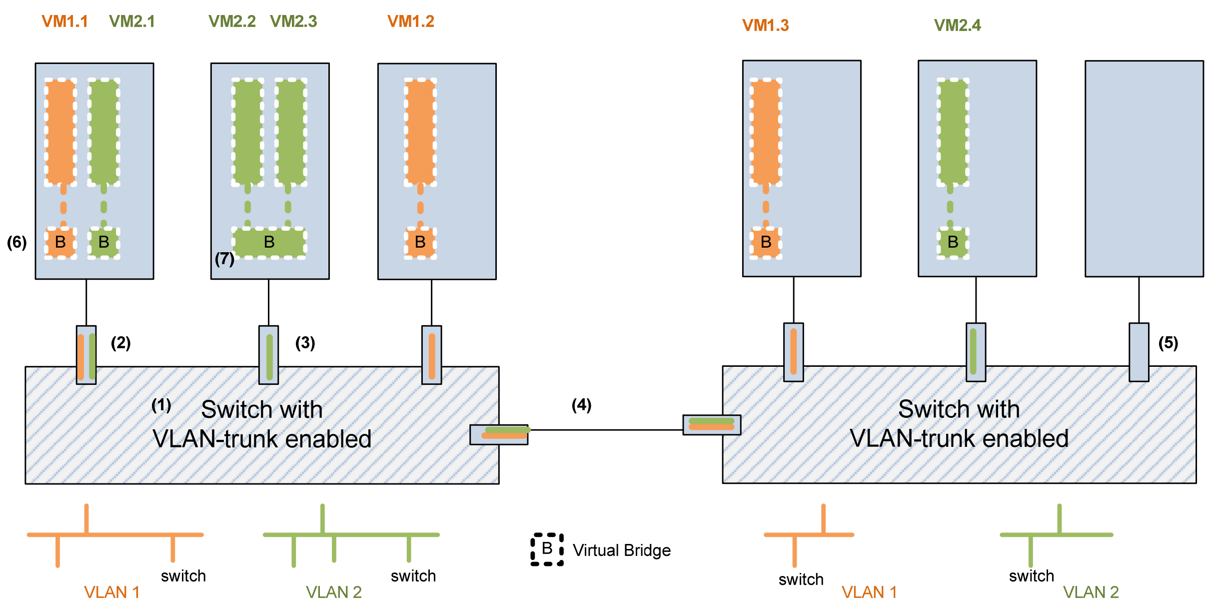

Figure 17: VLAN enabled infrastructure across multiple switches

(1) Each switch used to connect VMs in a VI must be VLAN-trunk enabled
(2) Ports on switch are tagged with VLAN Identifiers such that broadcast messages intended for specific VLANs are only sent to those ports.
(3) Ports may be tagged for a single VLAN, enabling high isolation of network traffic and management broadcast messages
(4) Ports used for connecting switches must also be tagged with the respective VLAN Identifiers
(5) Only ports that are active in VIs need to be enabled, unless used for non-VLAN traffic
(6) VMs are attached to virtual bridges via their virtual network interface (vNIC). Note that a VM may have multiple vNICs each connected to different virtual bridges, if the VM is to be connected to multiple VLANs
(7) VMs on the same physical host and VLAN share the same virtual bridge

Each VM uses at least one virtual network interface for connecting to a virtual network bridge. In case the virtual network bridge does not exist, OpenNebula will create one and attach it to a physical network interface tagged with a VLAN-ID. The OpenNebula Adapter ensures that for each logical resource a different VLAN-ID will be used, ensuring the VMs belonging to the same VI will be able to communicate between them, but not with the VMs belonging to different VIs.

As an IEEE 802.1q [REF 14] mechanism is being used, certain configurations need to be performed on the servers hosting the VMs, such as installing the kernel module 802.1q and the vconfig package, and also allowing the user under which OpenNebula connects to execute vconfig, brctl and IP commands in the host. The physical switch to which the host is being connected must support the VLAN-trunk protocol.

It is recommended to ensure that two different network interface exist in the hosts for handling remote management operations (e.g. configuration and monitoring), so as to allow the VM to connect to the Internet (if desired) given the fact that one of the physical network interfaces in the hosts will be used for carrying VLAN-tagged traffic.

In the case of Virtual Infrastructures distributed across different PIPs, extra attention should be taken in ensuring that the VMs belonging to the same VI use the same VLAN tags, together with using the same technologies for connecting the different VLAN segments,(i.e., using L3 routing for connecting the VI segments).

3.3 Virtual infrastructure provisioning

In order to enable converged network and IT service provisioning, the LICL performs abstraction and virtualisation of the physical resources
[REF 4].

The LICL functionalities are used by two GEYSERS actors – the Physical Infrastructure Provider (the owner of the physical resources) and the Virtual Infrastructure Provider (the infrastructure broker), which is directly involved in the process of VI provisioning. Consequently, the LICL has been implemented as two software modules [REF 5]:

- The Lower-LICL is responsible for resource abstraction, resource publishing, VR creation and management, and VR operation. In the case of the reference scenario described above, this module should be deployed at each of the PIPs;
- The Upper-LICL is in charge of VI creation, management and re-planning, and should be deployed at the VIP.

In the following, the workflow of VI provisioning, with a focus only on the operations related to network and IT resource virtualization at the Lower-LICL layer, leaving aside other steps, such as authentication and authorisation is explained. For a complete description of the workflow, we refer the reader to deliverable D3.2 [REF 4]. In the reference scenario described above, all PIPs should deploy the CloudWeaver appliance, in addition to the Lower-LICL software. Moreover, both PIP1 and PIP3, as IT resource providers, should deploy OpenNebula and KVM hypervisor in order to employ the mechanisms they provide for managing virtualization operations.

In particular, once the VIP receives a request for a VI, the planning phase begins, which involves the Upper-LICL splitting the request into sub-VIs, and further sending these requests to the Lower-LICL at the underlying PIPs. Next, the Lower-LICL makes use of the OpenNebula and/or CloudWeaver functionalities to allocate a VI composed of network and IT resources.

During the planning phase, only the network resources are reserved; booking IT resources at this phase is not possible, because the needed information is lost during the abstraction of the logical resources. Therefore, for IT resources, only a resource pool (i.e. a set of logical storage and computing resources) is reserved in this phase.

Once the VIO requests the instantiation of the VI, the VIP consequently requests the instantiation of the composing VRs. Then, the Lower-LICL proceeds with the VR creation, and requests OpenNebula to allocate the required IT resources at PIP1 and PIP3 from the reserved resource pool, and the supporting datacentre network is configured as described in sub-section 3.2.3.

On the other hand, during the operation phase, the NCP+ deployed by the VIO performs a joint optimised selection of IT and network resources upon NIPS requests as described in deliverable D4.1 [REF 6]. Once the chosen resources have been reserved, the actual allocation of the network resources is signalled to the LICL from the NCP+ through the CCI interface and the activation of the IT resources is signalled from the SML through the SLI interface.

4 GEYSERS TEST-BED DEMONSTRATORS

4.1 Introduction

This section explains the four Demonstrators according to the proposed test-bed and the test-cases described in deliverable D5.1 [REF 7]. The Demonstrators aim, on one hand, to test and validate particular GEYSERS components with a set of well-defined test procedures and, on the other hand, to validate the GEYSERS architecture with proof-of-concept prototypes. The experiments evaluate the model both for all the key technical dimensions and from the usage/business perspectives of the involved actors.

According to the technical constraints of the test-bed offered by the GEYSERS partners, the demonstration value of the use-cases and the temporal constraints of the project, a subset of the test-cases proposed by WP1 in deliverable D1.5 [REF 1] have been selected and are detailed in this document. Each of the following Demonstrators focuses on different components and/or functionalities offered by the GEYSERS architecture.

- Demonstrator 1, Virtual Infrastructures for the VIO, focuses on the LICL functionalities to demonstrate the network virtualization service offered by an Infrastructure Provider to a Virtual Infrastructure Operator (VIO), and how the involved players interact. This scenario aims to better use the available resources through the virtualization of those resources, and to reduce the capital expenditure (CapEx) and operational costs (OpEx).

- Demonstrator 2 shows the on-demand provisioning of enhanced network connectivity services, tailored to the specific requirements of the applications, over a Virtual Infrastructure. The technical focus of this Demonstrator is on the NCP+ functionalities to support the different types of specialised connectivity services that a VIO-N is able to provide to its customer (typically a VIO-IT).

- The objective of Demonstrator 3 is to dynamically adjust the available compute, storage and network resources for an Enterprise Information System (EIS) based on the demand from the users of the EIS. The intention is to motivate methods of synchronising the scaling of IT infrastructure and network capacity dynamically. In this Demonstrator we assess the benefits for cost and overall “satisfaction” of the EIS provider’s operational objectives, without disrupting the application users’ service level objectives and experience.

- Demonstrator 4 is aimed at the demonstration of advanced network and IT management functionalities, with a specific focus on network infrastructure re-planning. The virtual infrastructure re-planning service allows the VIO to request the modification, up-scaling or down-scaling (e.g. upgrade of link capabilities, modification of network topology) of the leased virtual infrastructure in order to optimise the network resource utilisation. The Demonstrator will evaluate and compare manual and dynamic methods for the VI re-planning, by providing usage examples for both of them.

4.2 DEMONSTRATOR 1

4.2.1 Introduction and description

Virtualization has gained sufficient momentum as one of the key paradigms for future networks. Demonstrator 1, Virtual Infrastructures for the VIO, focuses on the network virtualization service offered by an Infrastructure Provider to a Virtual Infrastructure Operator (VIO), and how the involved players interact. This scenario is based on two premises: i) that infrastructure owners will use virtualization to provide third parties with access to their physical infrastructures, thus making a better use of the available resources, and ii) that a VIO prefers to rent part of physical infrastructure (which in turn instantiates a virtual network) instead of deploying a network infrastructure on its own, thus reducing the capital expenditure (CapEx) and operational costs (OpEx).

From the technical point of view, this Demonstrator is focused on showing the functionalities of the LICL layer from the GEYSERS architecture, such as VI request, VI planning, VI instantiation, and VI decommissioning. The Demonstrator involves the participation of three GEYSERS components: the Physical Infrastructure Provider (PIP), the Virtual Infrastructure Provider (VIP) and the Virtual Infrastructure Operator (VIO), and aims at showing several benefits that each of them would gain by using the GEYSERS framework. Firstly, the PIP achieves a better utilisation of the available resources and obtains a faster return on investment (ROI) from the roll-out of its physical infrastructure by being able to virtualize and rent its physical resources. Secondly, the VIP can handle virtual resources from multiple domains, thus overcoming the resource limitations or geographical locations of a single domain, and is able to provide custom, on-demand VIs. Thirdly, the VIO is able to reduce capital expenditure and operational costs, by operating over virtual infrastructures leased from a VIP. In this context, two scenarios have been identified, as explained in Table 6 below:

- Scenario 1: Only IT resources.
- Scenario 2: Network and IT resources

Scenario	Description	Architectural layers	LICL components
1. Only NET resources	• Shows LICL functionalities. • VI composed of only network resources, organised across multiple domains. • Two VI-Ns are created. • SML and NCP+ are out-of-scope. • Dummy stubs: CCI and MLI interfaces • Scope: LICL	LICL	Upper-LICL Lower-LICL
2. NET+IT resources	• Shows LICL functionalities. • VI composed of both network and IT resources, organised across multiple domains. • One VI is created • SML and NCP+ are out-of-scope. • Dummy stubs: CCI and MLI interfaces • Scope: LICL	LICL	Lower-LICL Upper-LICL

Table 6: Demonstrator 1 – scenarios

4.2.2 Physical test-bed provided for this Demonstrator

Demonstrator 1 will be deployed over a subset of the GEYSERS test-bed spanning different sites, in particular using resources from i2CAT, UEssex, PSNC, UvA and IBBT sites, as shown in Figure 18 below.

The physical resources required to implement the scenarios of Demonstrator 1 are defined by the Virtual Infrastructure designed for each case, as well as by the LICL components and the additional components involved in the Demonstrator scenarios. Preliminary considerations about the required resources for each scenario are presented below.

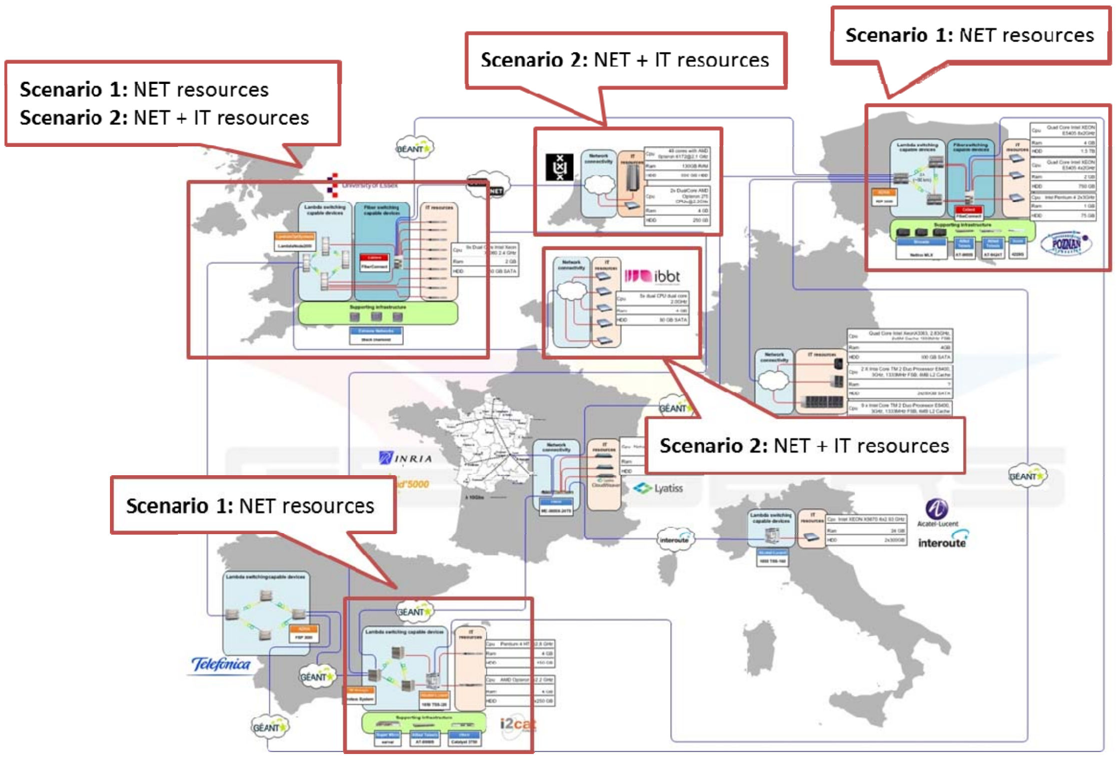

Figure 18: Physical test-bed for Demonstrator 1

Scenario 1: The goal of this scenario is to create Virtual Infrastructures (Vis) composed of only network resources (VI-N), deployed on top of several PIPs. The switching technology adopted for this VI-N is Fibre Switching (FS). Therefore, the devices required at each local test-bed are just the FSC ones and this will also be the switching capability of the virtual network resources created by the LICL. For this scenario, only network resources from UEssex, i2CAT and PSNC sites will be used.

Scenario 2: This scenario goes a step further than the first one, by considering the creation of a VI containing both network and IT resources, with the final goal to demonstrate the LICL's ability to support the joint virtualization of network and IT resources. As in the previous case, FS technology will be used. This scenario will be implemented over the test-beds located at IBBT, UvA and UEssex.

4.2.3 Physical resources planning and virtualization

Table 7 provides an overview of the physical resources used for Demonstrator 1, for each of the considered scenarios.

For Scenario 1, the partners assuming the role of PIP-N (PSNC, i2CAT, UEssex) will deploy the Lower-LICL and the CloudWeaver appliance. PSNC will additionally deploy the Upper-LICL, as it carries the role of VIP as well.

For Scenario 2, as in the previous case, the PIP-N (UEssex) will deploy the Lower-LICL and the CloudWeaver appliance, while the VIP will deploy the Upper-LICL. The PIP-IT (UEssex, UvA, IBBT) will also deploy the OpenNebula v3.2 and the KVM hypervisor. The OpenNebula enables the functionality of instantiation, monitoring, configuration and decommissioning operations of virtual resources (VMs, virtual networks and virtual storage).

The Upper- and Lower- LICL could be deployed on the same host, but it is recommended that they are deployed on different hosts. Moreover, connectivity between CloudWeaver and the Lower-LICL and between the Lower-LICL and the physical resources must be ensured.

	Scenario 1 – Only NET resources	Scenario 2 – NET+IT resources
NET PHY Infrastructure	PSNC 4 x ADVA FSP 3000 R7 1 x Calient DiamondWave FiberConnect UEssex: 1 x Calient DiamondWave FiberConnect i2CAT: 3 x W-Onesys Proteus	UEssex 1 x Calient DiamondWave FiberConnect
IT PHY Infrastructure	N.A.	UEssex 5 x DELL PowerEdge 860 UvA 1 x DELL R810 IBBT 5 x SuperMicro, 2 x Intel Xeon CPU E 5620 @ 2.40 GHz, 12 GB RAM, 160 GB SATA disk, 2 Gb/s NIC NFS: 500GB shared

Table 7: Demonstrator 1 - overview of the physical resources

Table 8 gives details on the LICL components that will be deployed for each of the scenarios:

	Scenario 1 – Only NET resources	Scenario 2 – NET+IT resources
LICL controllers	• Upper LICL • Lower LICL	• Upper LICL • Lower LICL
Other	• CloudWeaver appliance	• CloudWeaver appliance • OpenNebula • KVM hypervisor

Table 8: Demonstrator 1 - summary of the components deployed

PSNC local test-bed resources: PSNC's local test-bed (Figure 19) for this Demonstrator is composed of a Calient DiamondWave FiberConnect that provides the FSC data plane infrastructure. The Calient box is manually partitioned into 4 sub-switches to be able to compose a simple mesh network topology. The Calient switch is connected via a Brocade NetIron MLX-8 router (supporting equipment) to the remote test-beds involved in this Demonstrator.

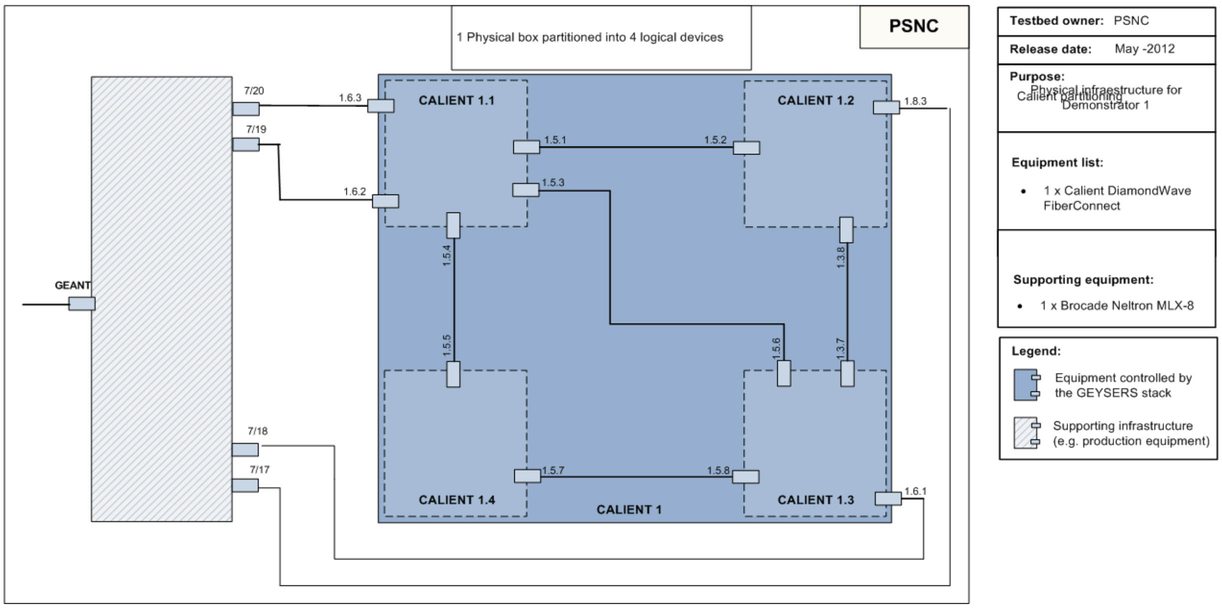

Figure 19: PSNC local test-bed

i2CAT local test-bed resources: The local test-bed at i2CAT is composed of 3 W-Onesys Proteus devices. The LICL will be deployed over the Super Micro Server, while the OpenNebula and KVM hypervisor will be deployed over the SunFire and Dell PowerEdge machines. Additionally, the Lyatiss CloudWeaver will be used as an alternative to OpenNebula in i2CAT test-bed for this Demonstrator.

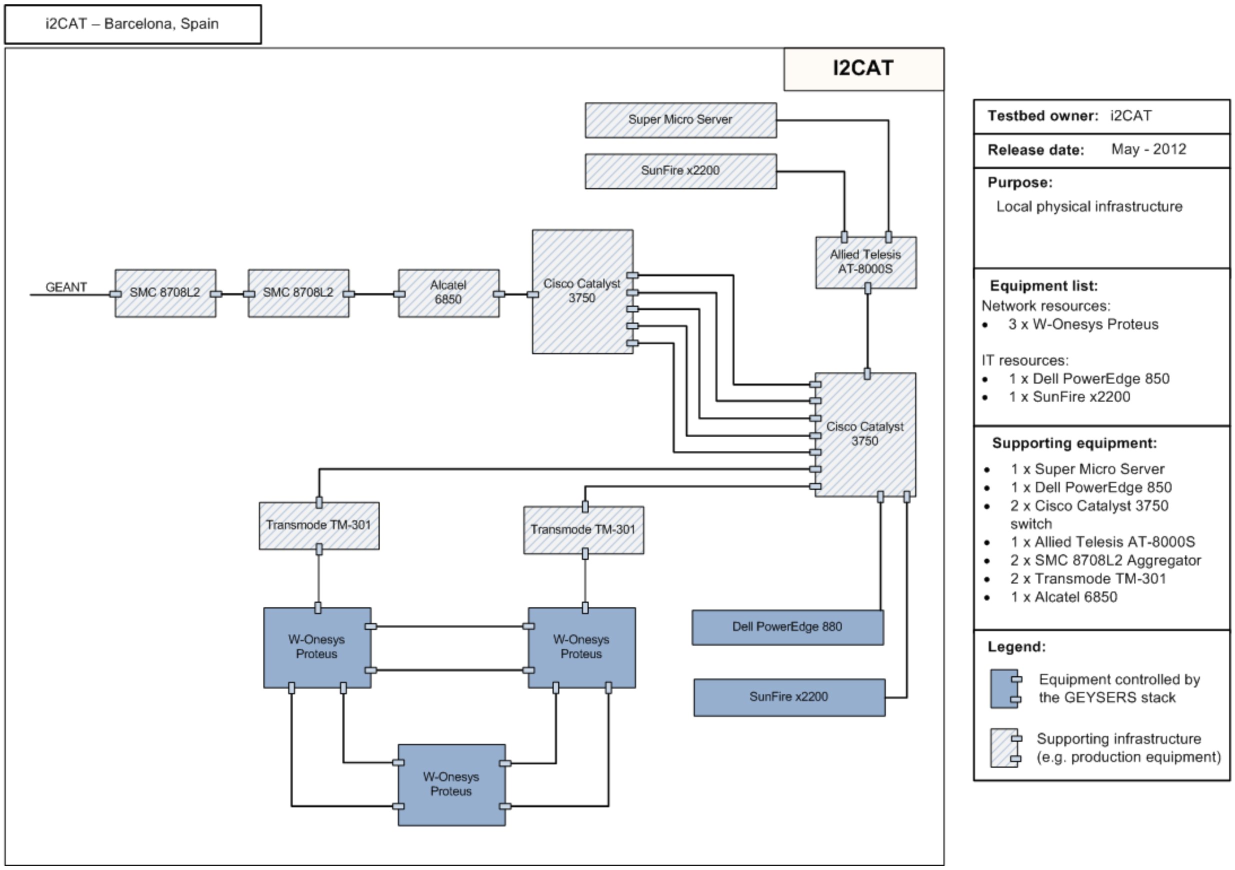

Figure 20: i2CAT local test-bed

UEssex local test-bed resources: The UEssex local test-bed (Figure 21) for this Demonstrator is composed of a Calient DiamondWave FiberConnect that provides the FSC data plane infrastructure. The Calient switch is manually partitioned into four sub-switches to be able to compose a simple mesh network topology. Three DELL PowerEdge 860 servers connected through a DELL PowerConnect Switch will be used to deploy the VMs of the LICL controllers, the CloudWeaver software, OpenNebula and the KVM hypervisor. In addition, two more DELL PowerEdge 860 servers will be used as IT resources, which will be connected to the Calient switch via a FOUNDRY FastIron Edge Switch.

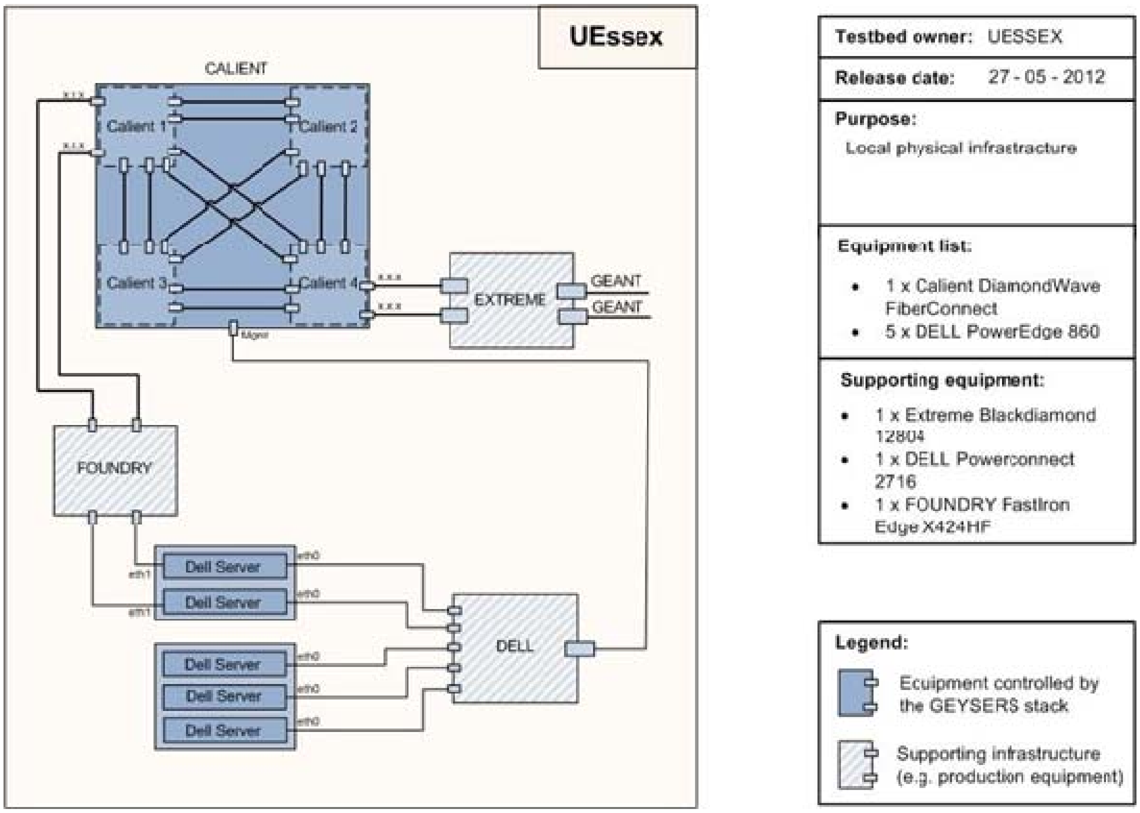

Figure 21: UEssex local test-bed

UvA local test-bed resources: The UvA test-bed (Figure 22) provides a cloud-machine (Dell R810: 48cores, 128GB-RAM, 600GB storage) for on-demand computation and storage resources. These resources are provided with the aid of OpenNebula installed and the necessary software to deploy the Lower-LICL. This test-bed provides access via two GLIF light-paths, to i2CAT and to UEssex.

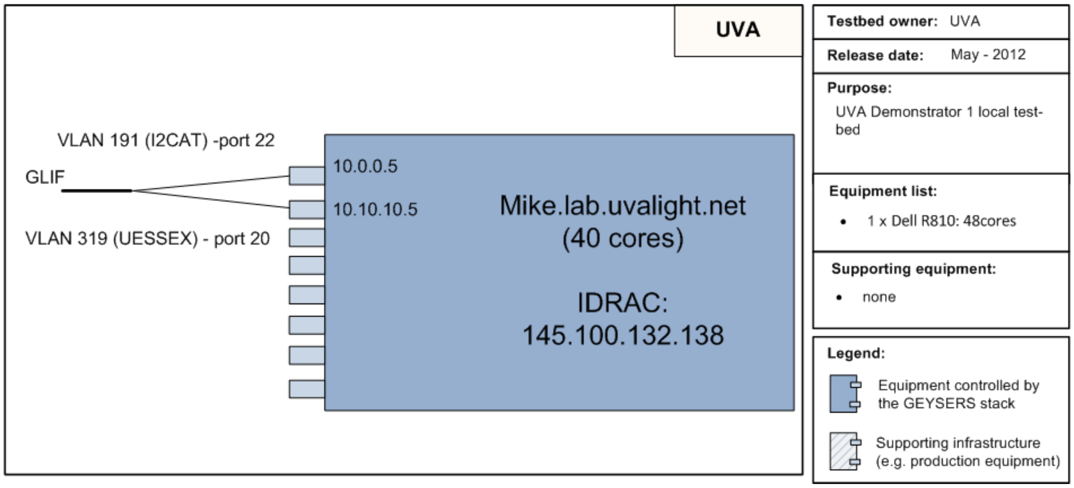

Figure 22: UvA local test-bed

IBBT local test-bed resources: The IBBT local test-bed (Figure 23) for this Demonstrator is composed by 5 Super Micro Intel Xeon nodes, which are interconnected by a 1 GB switch. One of these servers will be used for the CloudWeaver software and OpenNebula and the LICL modules. The other nodes will have the KVM hypervisor.

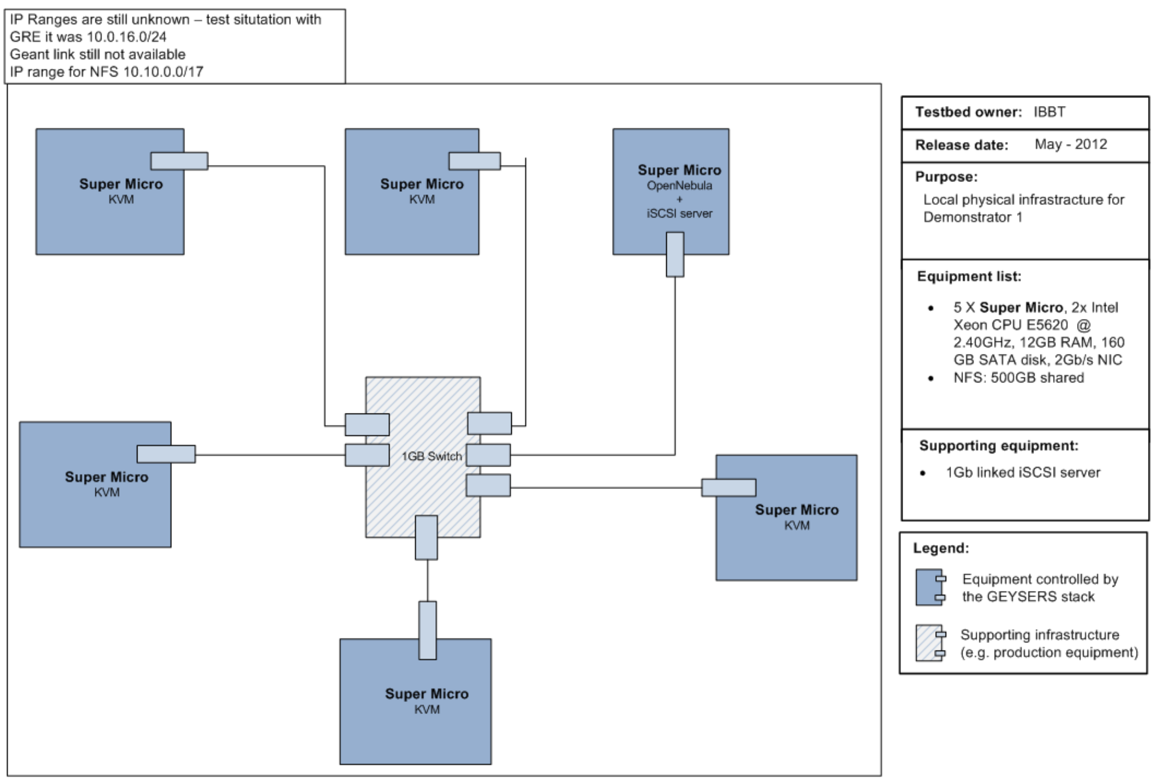

Figure 23: IBBT local test-bed

The resulting test-bed is illustrated in Figure 24 for Scenario 1 and Figure 25 for Scenario 2.

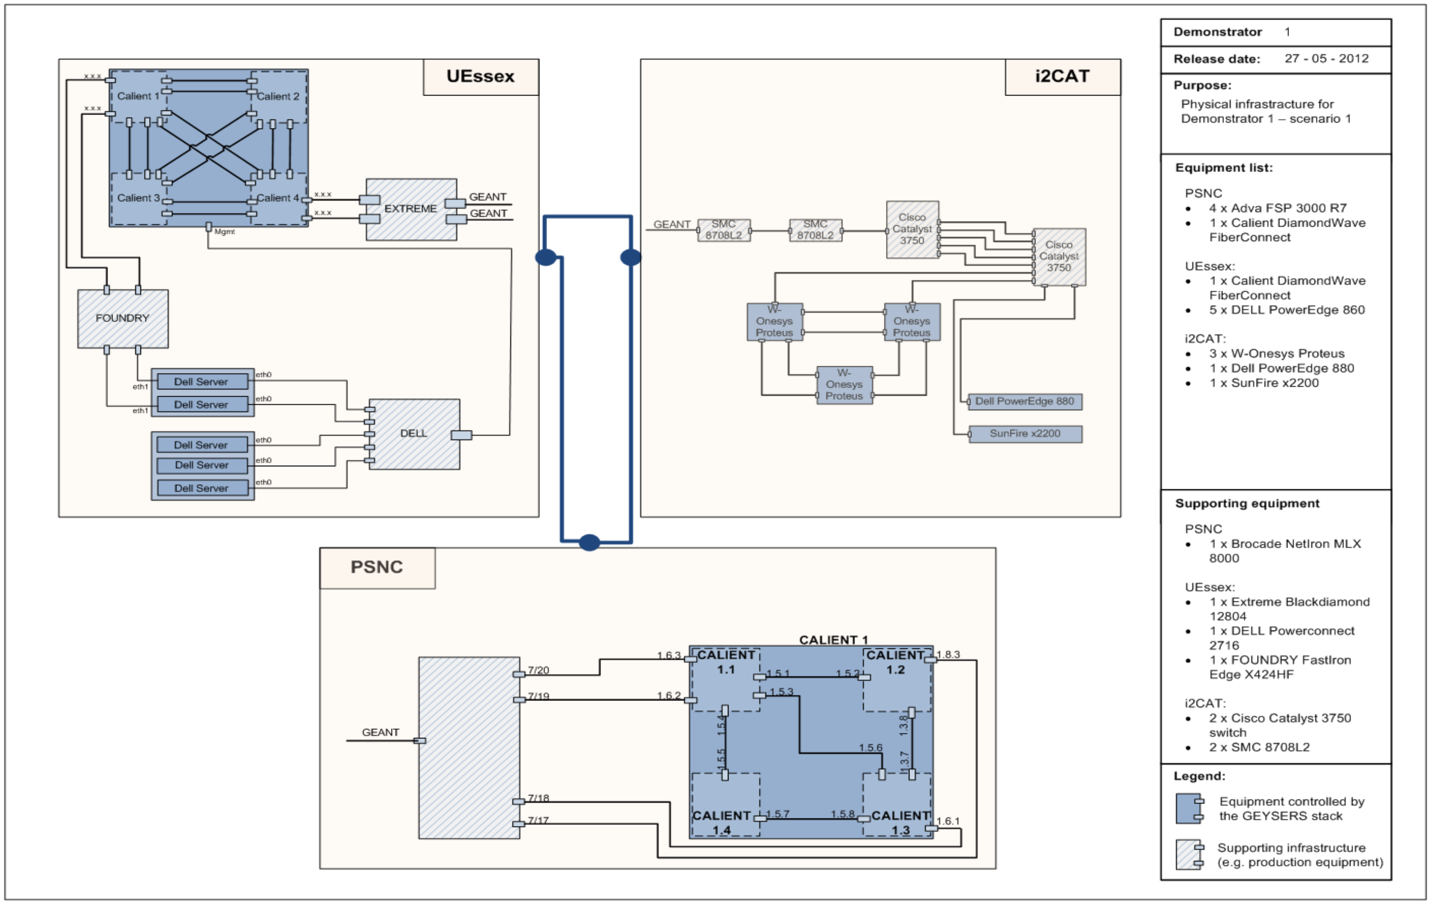

Figure 24: Demonstrator 1 - Scenario 1

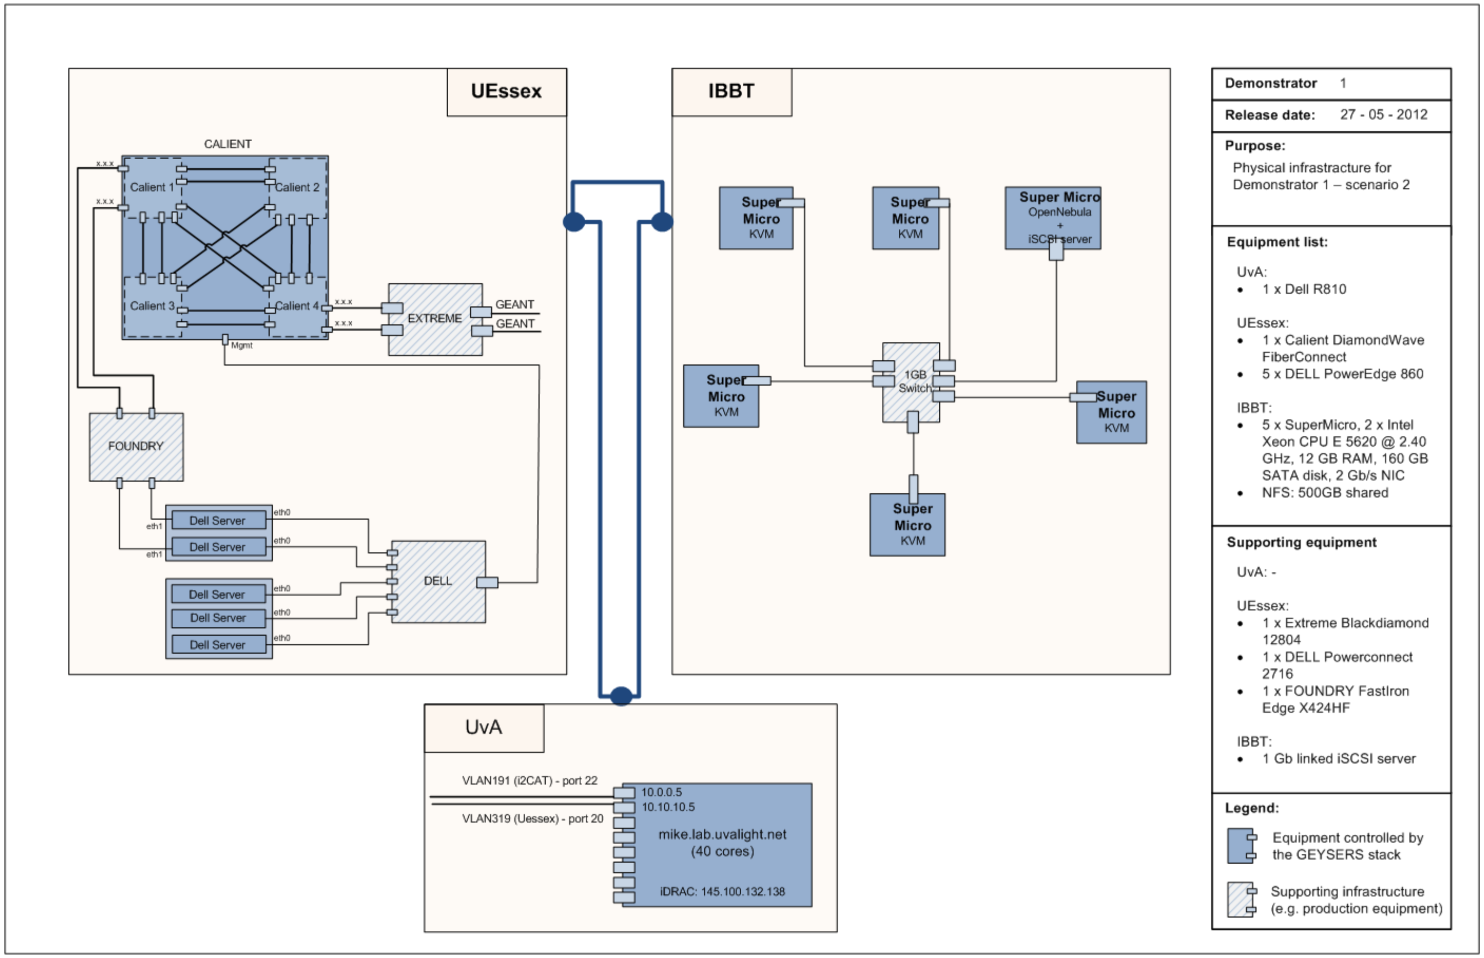

Figure 25: Demonstrator 1 - Scenario 2

4.2.4 Virtualized infrastructure of this Demonstrator

As explained in the previous sub-sections, for Scenario 1 of the Demonstrator 1, two Virtual Infrastructures will be created on top of the physical infrastructure consisting of only network resources, provided at the UEssex, PSNC and i2CAT sites. The high level topology of the virtualized infrastructure is represented in Figure 26.

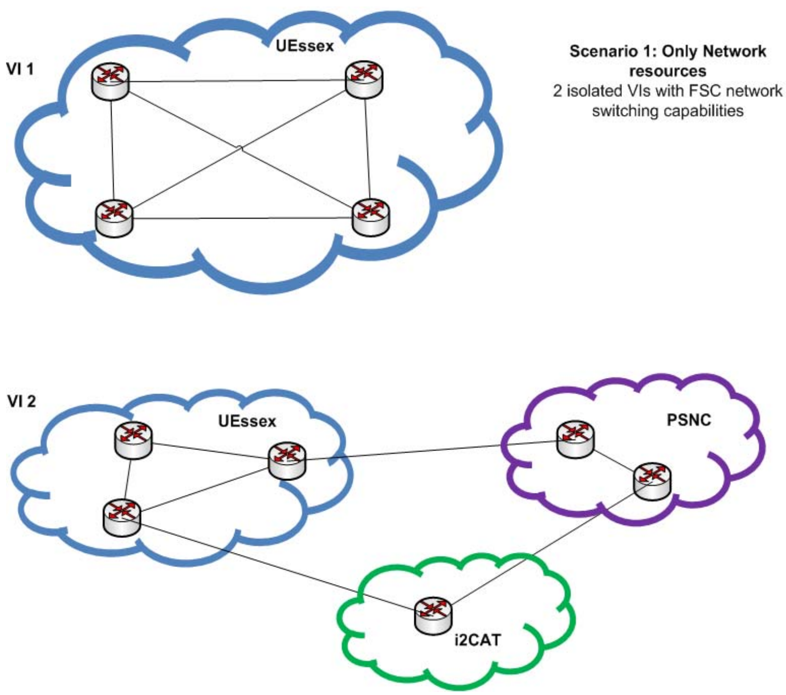

Figure 26: Demonstrator 1 - Virtualized infrastructure for Scenario 1

For Scenario 2 of the Demonstrator, only one virtual infrastructure will be created, that will contain both Network and IT resources, from the sites at UEssex, IBBT, and UvA. Figure 27 below illustrates the high-level overview of the topology in this case.

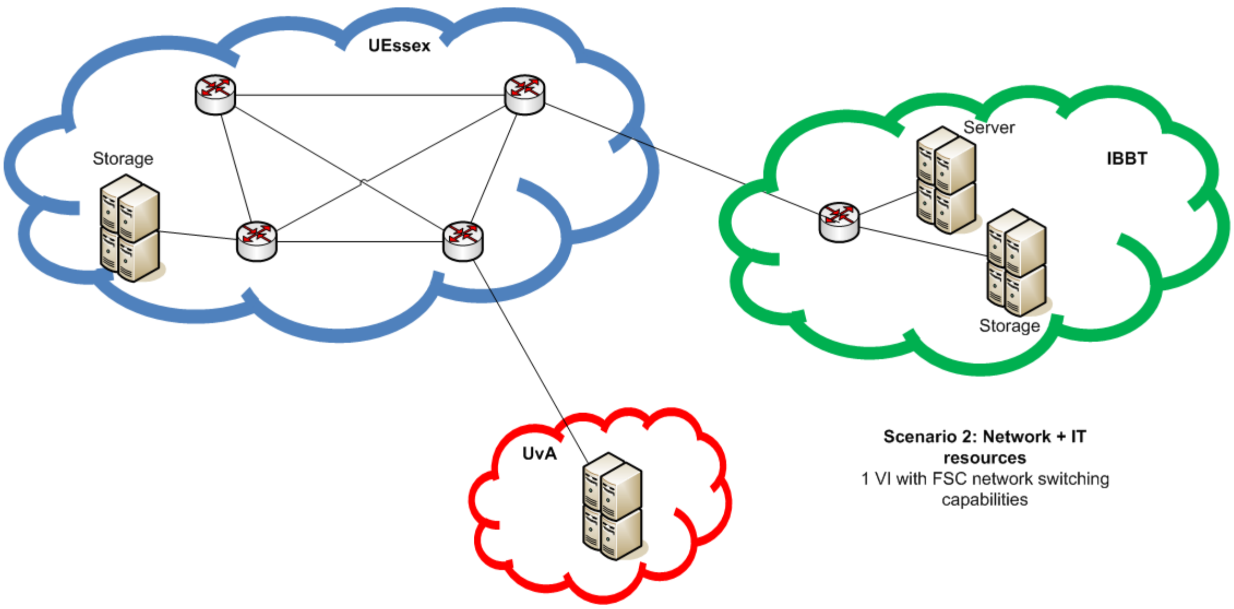

Figure 27: Demonstrator 1 - Virtualized infrastructure for Scenario 2

4.3 DEMONSTRATOR 2

4.3.1 Introduction and description

This Demonstrator will show over a Virtual Infrastructure, the on-demand provisioning of enhanced network connectivity services, tailored to the specific requirements of the applications. The technical focus of this Demonstrator will be on the NCP+ functionalities to support the different types of specialised connectivity services that a VIO-N is able to provide to its customer (typically a VIO-IT). These services are enabled through the cooperation between the NCP+ and the SML over the NIPS UNI, leading to progressive steps towards network and IT convergence and the cross-layer optimisation of the heterogeneous resources composing the entire Virtual Infrastructure. This approach provides multiple benefits for the VIO-N, since it is able to maximise the utilisation of the leased infrastructure and, on the other hand, to offer customised and on-demand connectivity services to its customers with enhanced guarantees in terms of QoS and service resiliency.

The Demonstrator has been designed around three different and progressive technical scenarios, all of them transversal to usage scenarios such as on-demand provisioning of unicast, assisted unicast or anycast services, automatic recovery of network services and advanced service reservation, showing the main service provisioning functionalities offered by the GEYSERS architecture. Special focus will be placed on the NCP+ procedures in support of Network + IT Provisioning Services and on the LICL mechanisms in support of Virtual Infrastructure operation. These scenarios, described in detail in

Table 9 will allow the integration and testing of these components in subsequent steps; therefore, the Virtual Infrastructure and the GEYSERS components will be deployed over the physical resources of the GEYSERS test-bed reserved for this Demonstrator following an incremental approach. Further details about the test-bed deployment for Demonstrator 2 can be found in sub-sections 4.3.2 - 4.3.4.

Scenario	Description	Architectural layers	NCP+ components
Single domain NCP+	• Limited to NCP+ internal functionalities. • VI-N organised in a single routing domain. • SML and LICL out-of-scope. • Dummy stubs to test inter-layer procedures at NIPS UNI and CCI interfaces. • Scope: NIPS UNI <-> NCP+ <-> CCI 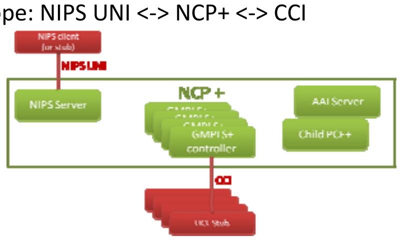	NCP+	NIPS server GMPLS+ controllers Child PCE+ server AAI server
NCP+ and LICL integration	• Includes real operations over the virtual network infrastructure, integrating the LICL. • VI-N organised in a single routing domain. • Two concurrent VI-Ns. • SML out-of-scope. • Dummy stubs to test inter-layer procedures at NIPS UNI. • Scope: NIPS UNI <-> NCP+ <-> CCI <-> LICL 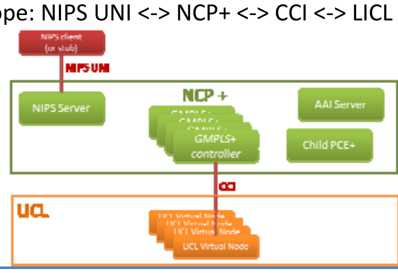	NCP+ LICL	NIPS server GMPLS+ controllers Child PCE+ server AAI server

Multi-domain	• Includes real operations over the virtual network infrastructure, integrating the LICL. • VI-N organised in multiple routing domains (hierarchical path computation). • SML out-of-scope. • Dummy stubs to test inter-layer procedures at NIPS UNI. • Scope: NIPS UNI <-> NCP+ <-> CCI <-> LICL 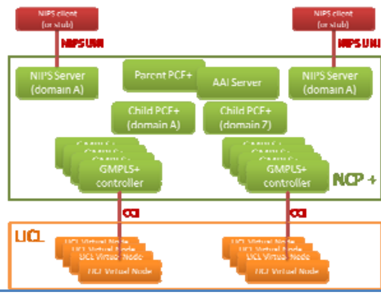	NCP+ LICL	NIPS server GMPLS+ controllers Child PCE+ server Parent PCE+ server AAI server

Table 9: Demonstrator 2 – scenarios

It should be noted that this Demonstrator will consider only a part of the GEYSERS architecture, specifically the NCP+ and the LICL, as it concerns the on-demand network service provisioning during the VI operation phase. The other architectural features and functionalities, as well as the SML mechanisms, are considered out-of-scope since they are covered in detail by other Demonstrators.

4.3.2 Physical test-bed provided for this Demonstrator

Demonstrator 2 will be deployed over a subset of the GEYSERS test-bed spanning three different sites, in particular using resources from UEssex, TID and PSCN sites, as shown in Figure 28.

The physical resources required to implement the scenarios of Demonstrator 2 are defined by the Virtual Infrastructure designed for the different phases of this Demonstrator, as well as by the GEYSERS components progressively involved in the Demonstrator scenarios. Preliminary considerations about the required resources for each scenario are presented below.

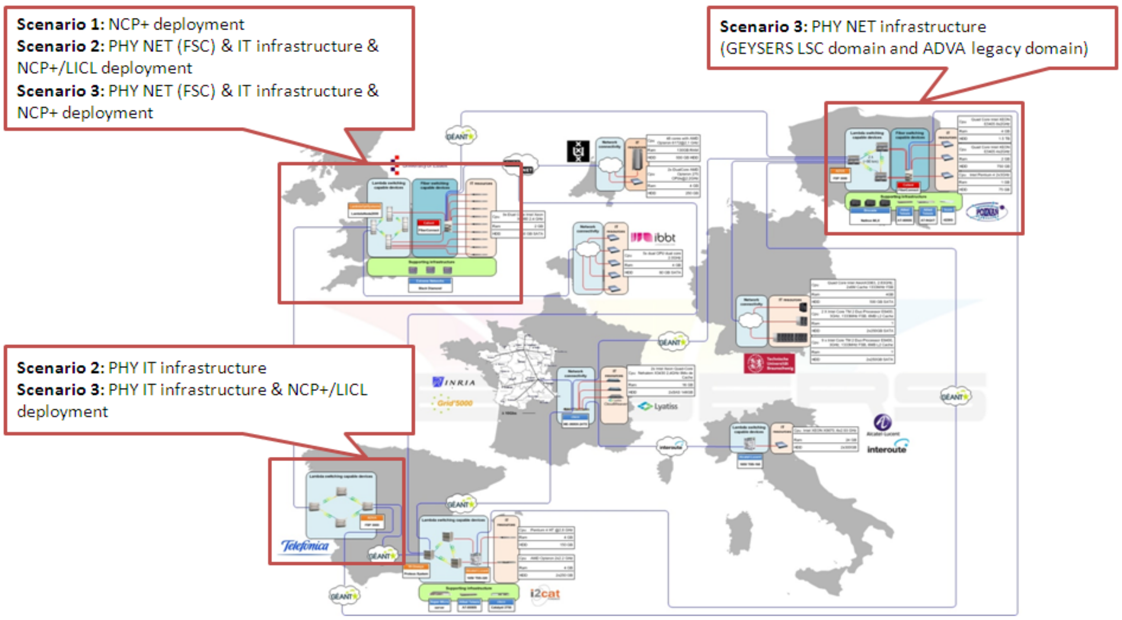

Figure 28: Demonstrator 2 test-beds

Single domain NCP+: This scenario will use dummy stubs to emulate interconnection to the LICL. Therefore, the underlying network technology used for this scenario will not be limited by the actual test-bed devices. This means that both FSC and LSC switching technologies can be considered indistinctly in the Virtual Infrastructure design and tested by properly configuring the LICL stubs. Moreover, since the configuration of the physical infrastructure is out-of-scope for this scenario, the only required test-bed resources are related to the machines utilised for installing and deploying the NCP+ components used to control a single network routing domain (NIPS Server, GMPLS+ controllers, child PCE server, AAI server). A set of machines from the UEssex test-bed will be reserved for this scenario.

NCP+ and LICL integration: The goal of this scenario is to augment the scope of the previous scenario by introducing real IT and network elements virtualized by the LICL and composed in Virtual Infrastructures controlled by the NCP+ as single routing domains on the network side. Consequently, both network and IT resources from UEssex test-bed (FSC network resources) and UEssex and TID test-bed (IT resources) will be dedicated to the construction of the physical infrastructure. In order to show the isolation of multiple virtual infrastructures over the same physical infrastructure, two different VIs sharing the same underlying resources will be considered in this scenario. This means that two instances of single-domain NCP+ (one for each VI) and one instance of LICL (Upper- and Lower- LICL components) must be deployed. The machines dedicated to the NCP+ and LICL deployment will be reserved in the UEssex and TID test-beds.

Multi-domain: This scenario will cover a single multi-technology Virtual Infrastructure, organised in two/three network routing domains. In the simplest configuration, two domains with different switching capabilities will be considered: an FSC domain (derived from physical network resources from UEssex test-bed) connected to an LSC domain (derived from physical network resources located in the PSNC test-bed). IT virtual resources will be created over physical resources provided by TID, UEssex and TP, while the NCP+ and LICL controllers will be installed in machines located in UEssex, PSNC and TID test-bed. It should be noted that, for what concerns the NCP+, a dedicated NIPS Server and child PCE+ server is required for each domain, while a single parent PCE+ server and AAI server will cover the overall VI. A more complex configuration of the multi-domain scenario involves also an internal "legacy" domain, based on ADVA equipment installed in the PSNC test-beds, controlled by their own GMPLS and PCE-based control plane, and supporting the sub-set of the GEYSERS features dedicated to energy efficiency. This deployment has been designed to test the interoperability of the GEYSERS NCP+ with legacy control planes.

4.3.3 Physical resources planning and virtualization

Table 10 provides an overview of the physical resources used for Demonstrator 2 in each scenario. Part of these resources will be dedicated to the creation of the physical infrastructure to be virtualized by the LICL and controlled by the NCP+ once the associated Virtual Infrastructures are instantiated. On the other hand, some machines will be reserved for the deployment of the NCP+ and LICL components, delivered as a set of VMs, as detailed in Table 11.

	Scenario 1 – Single domain	Scenario 2 – NCP+/LICL	Scenario 3 – Multi-domain
NET PHY Infrastructure	N.A.	**UEssex** test-bed network devices: FSC domain, based on CALIENT Diamond Wave	**UEssex** test-bed network devices: FSC GEYSERS domain, based on CALIENT Diamond Wave **PSNC** test-bed – ADVA equipment: LSC GEYSERS domain, based on ADVA FSP 3000R7 – 4TCA cards **PSNC** test-bed – ADVA equipment: LSC legacy domain, based on ADVA FSP 3000R7 – WCA2G5 cards
IT PHY Infrastructure	N.A.	**TID** test-bed server: DELL PowerEdge **UEssex** test-bed server: DELL PowerEdge 860	**TID** test-bed server: DELL PE **UEssex** test-bed server: DELL PowerEdge 860 **TP** – test-bed server: HP Proliant DL580 G7
Machines to deploy NCP+ controllers	**UEssex** test-bed servers: DELL PowerEdge 860 (6 VMs)	**UEssex** test-bed servers: DELL PowerEdge 860 (12 VMs)	**UEssex** test-bed servers – DELL PowerEdge 860 (6 VMs for FSC domain controllers) **TID** test-bed servers – HP DL380 (1 VM for Parent PCE server) **PSNC** test-bed servers – IBM System x3550 M3 (6 VMs for LSC domain controllers)
Machines to deploy LICL controllers	N.A.	**UEssex** test-bed servers: DELL PowerEdge 860 (3 VMs for the Upper- and Lower-LICL) **TID** test-bed servers: HP DL380 (2 VMs for the Lower-LICL)	**TID** test-bed servers – HP DL380 (3 VMs for the Upper- and Lower-LICL) **UEssex** test-bed servers – DELL PowerEdge 860 (2 VMs for the Lower-LICL) **PSNC** test-bed servers - IBM System x3550 M3 (2 VMs for the Lower-LICL)

Table 10: Demonstrator 2 – overview of physical resources

:

<table>
<tr><th></th><th>Scenario 1 – Single domain</th><th>Scenario 2 – NCP+/LICL</th><th>Scenario 3 – Multi-domain</th></tr>
<tr><td>NCP+ controllers</td><td>• GMPLS+ controllers (4 VM)
• NIPS Server + AAI server (1 VM)
• Child PCE+ server (1 VM)
<u>Total:</u> 2 servers required (UEssex)</td><td>VI 1
• GMPLS+ controllers (4 VM)
• NIPS Server + AAI server (1 VM)
• Child PCE+ server (1 VM)

VI 2
• GMPLS+ controllers (4 VM)
• NIPS Server + AAI server (1 VM)
• Child PCE+ server (1 VM)
<u>Total:</u> 4 servers required (UEssex)</td><td>FSC GEYSERS domain
• GMPLS+ controllers (4 VM)
• NIPS Server + AAI server (1 VM)
• Child PCE+ server (1 VM)
<u>Total:</u> 2 servers required (UEssex)

LSC GEYSERS domain
• GMPLS+ controllers (4 VM)
• NIPS Server (1 VM)
• Child PCE+ server (1 VM)
• Parent PCE+ server (1 VM)
<u>Total:</u> 1 server required (PSNC) + 1 server required (TID)</td></tr>
<tr><td>LICL controllers</td><td>N.A.</td><td>• Upper LICL (1 VM)
• Lower LICL and CW appliance (2 VM)
<u>Total:</u> 1 server required (TID)
• Lower LICL and CloudWeaver appliance (2 VM)
Total: 2 servers required (UEssex)</td><td>• Upper LICL (1 VM)
• Lower LICL and CloudWeaver appliance (2 VM)
<u>Total:</u> 1 server required (TID)
• Lower LICL and CloudWeaver appliance (2 VM)
<u>Total:</u> 2 server required (UEssex)
• Lower LICL and CloudWeaver appliance (2 VM)
<u>Total:</u> 1 server required (PSNC)</td></tr>
</table>

Table 11: Demonstrator 2 – GEYSERS LICL and NCP+ components deployment

In the following, details about the physical resources allocated for Demonstrator 2 from each partners' test-bed are presented.

TID local test-bed resources - The TID local test-bed (Figure 29) consists of two servers (i) DELL POWER EDGE and (ii) HP DL380. The first is dedicated to build up, manage and destroy on-demand VMs which will represent virtual IT resources ready to be used by Virtual Infrastructure providers. The second is dedicated to allocate Parent-PCE and Child-PCE servers, GMPLS+ controllers, NIPS server, and Lower LICL modules. OpenNebula software will be installed in this Front-End and, in case necessary, it may allocate additional virtual VMs as well. The global local test-bed running process consists of the following: IT virtual resources will be created within (i) and will be managed from (ii) by means of the different controllers that have been installed on it. The Front-End Manager will also interact with an image repository where VM templates are stored. The decision is still pending whether the image repository will be a separate module or embedded in the Front-End machine. A Riverstone switch router constitutes the gateway to the global GEYSERS test-bed for both data and signalling.

PSNC local test-bed resources - The PSNC local test-bed is composed of four ADVA FSP 3000 R7 devices and partitioned into two independent LSC domains; the first domain using 1GbE WCA25G cards, and the second domain using 1GbE 4TCA cards). All tributary ports of the ADVA DWDM system are connected to the Calient DiamondWave fiber switch which performs an important role supporting Demonstrator 2 scenarios. A reconfiguration of the Calient switch allows easy migration between different Demonstrator 2 scenarios. Additionally, the Calient switch connects the PSNC LSC domains to another supporting Brocade NetIron MLX 8000 switch responsible for establishing VLAN Q-in-Q connections with the UEssex and TP test-beds. The last supporting equipment within the PSNC test-bed is an IBM System x3550 M3 server, which will allow deploying VMs with Lower-LICL and NCP+ software. Figure 30 summarises this description.

UEssex local test-bed resources - The UEssex local test-bed (Figure 31) for this Demonstrator is composed of a Calient DiamondWave FiberConnect that provides the FSC data plane infrastructure. The Calient switch is manually partitioned into four sub-switches able to compose a simple mesh network topology. The VMs of the NCP+ and LICL controllers will be deployed in five DELL PowerEdge 860 servers. Up to three VMs will be installed in each server which will be connected to the SCN via a DELL PowerConnect Switch. Finally, a FOUNDRY FastIron Edge Switch will be used to connect the IT resources to the Calient switch.

TP local test-bed resources - The TP local test-bed (Figure 32) includes an HP Proliant DL580 G7 that will be used to provide a set of physical IT resources. These resources will be considered in combination with the physical network resources within the PSNC test-bed as a physical domain owned by a single PIP. This means that the overall set of resources located in the PSNC and TP test-beds will be managed by a single instance of a Lower-LICL (instantiated on the PSNC servers).

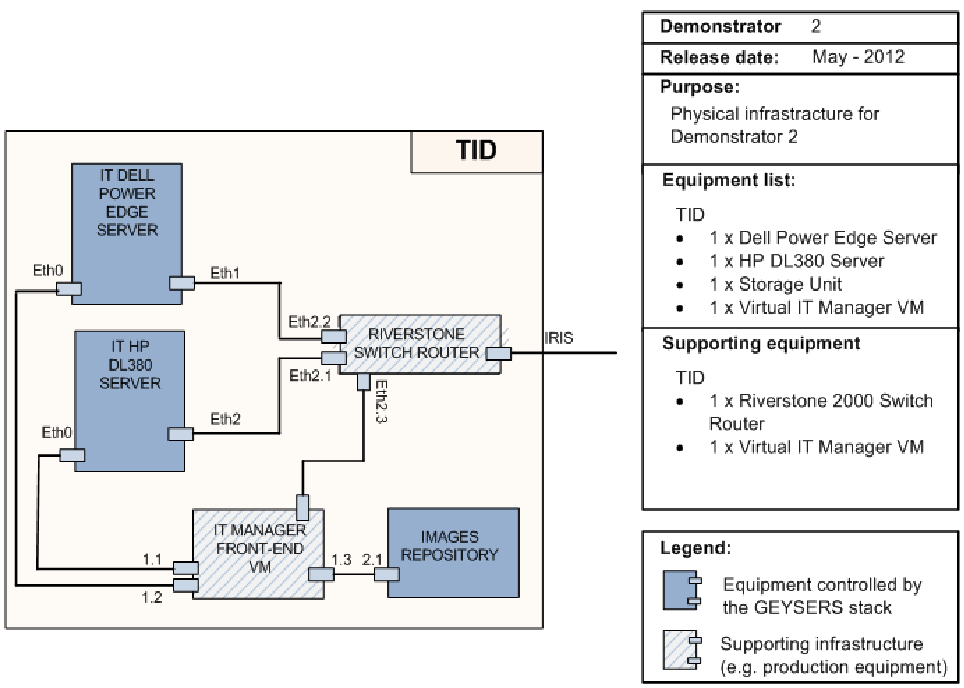

Figure 29: TID test-bed for Demonstrator

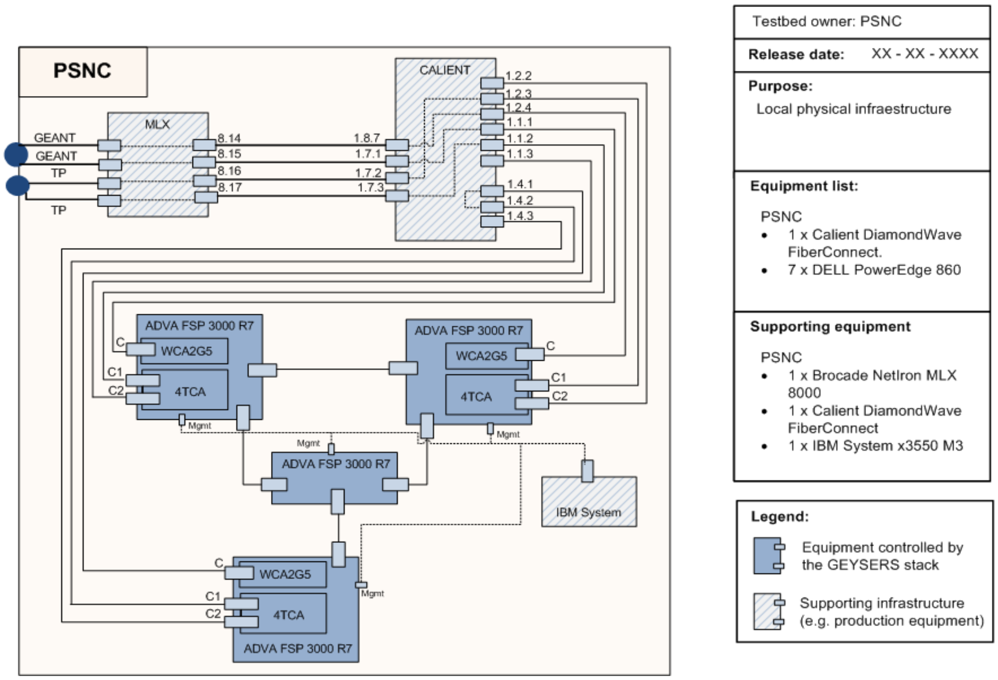

Figure 30: PSNC test-bed for Demonstrator 2

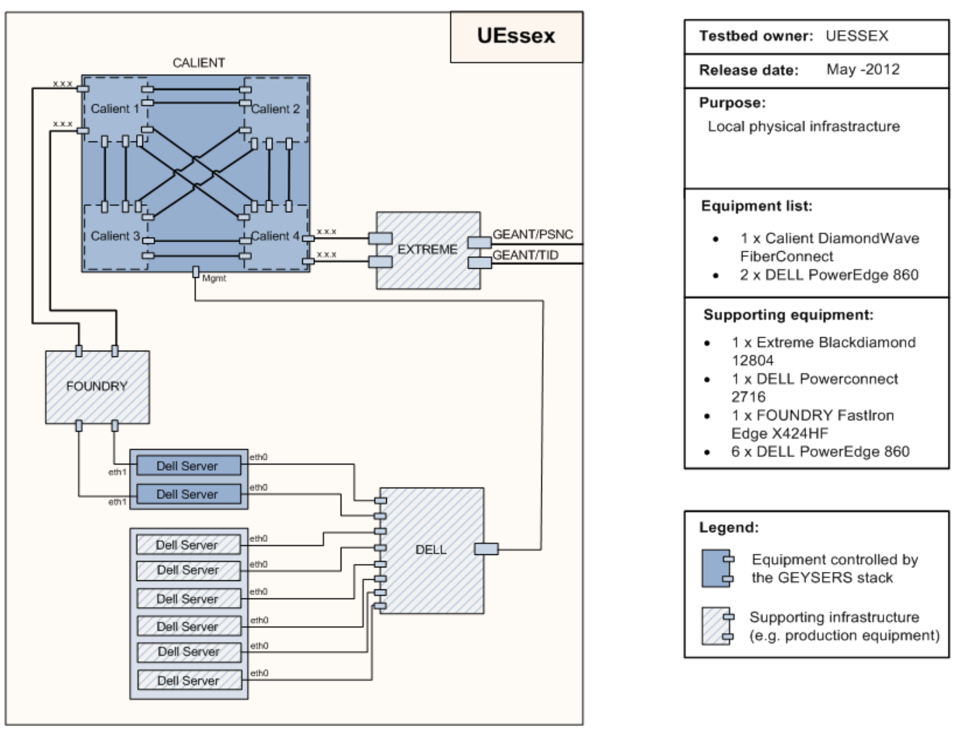

Figure 31: UEssex test-bed for Demonstrator 2

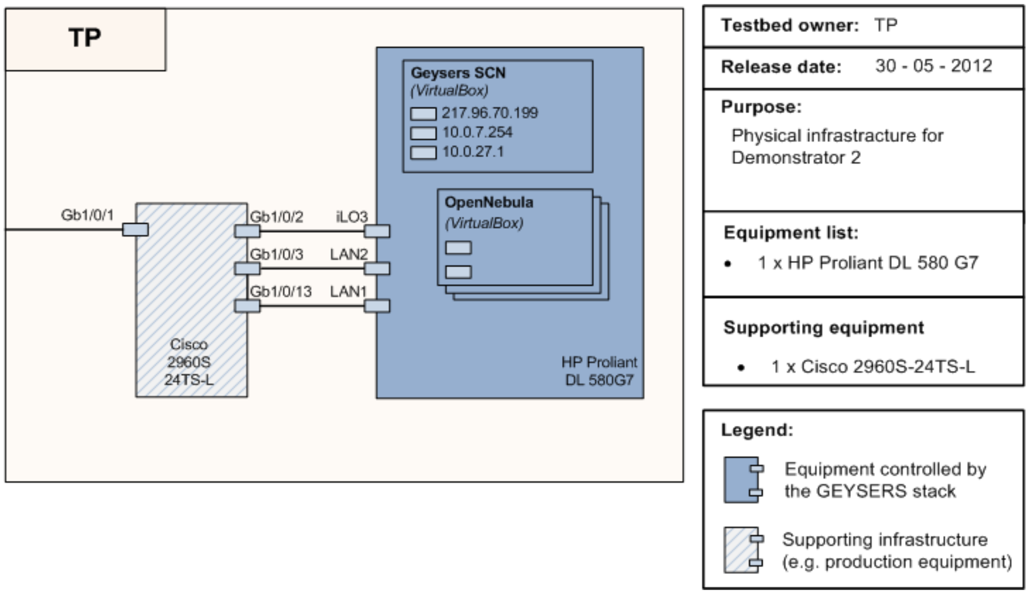

Figure 32: TP test-bed for Demonstrator 2

The resulting test-bed for Demonstrator 2 is shown with details about the role of the different equipment in the Demonstrator.

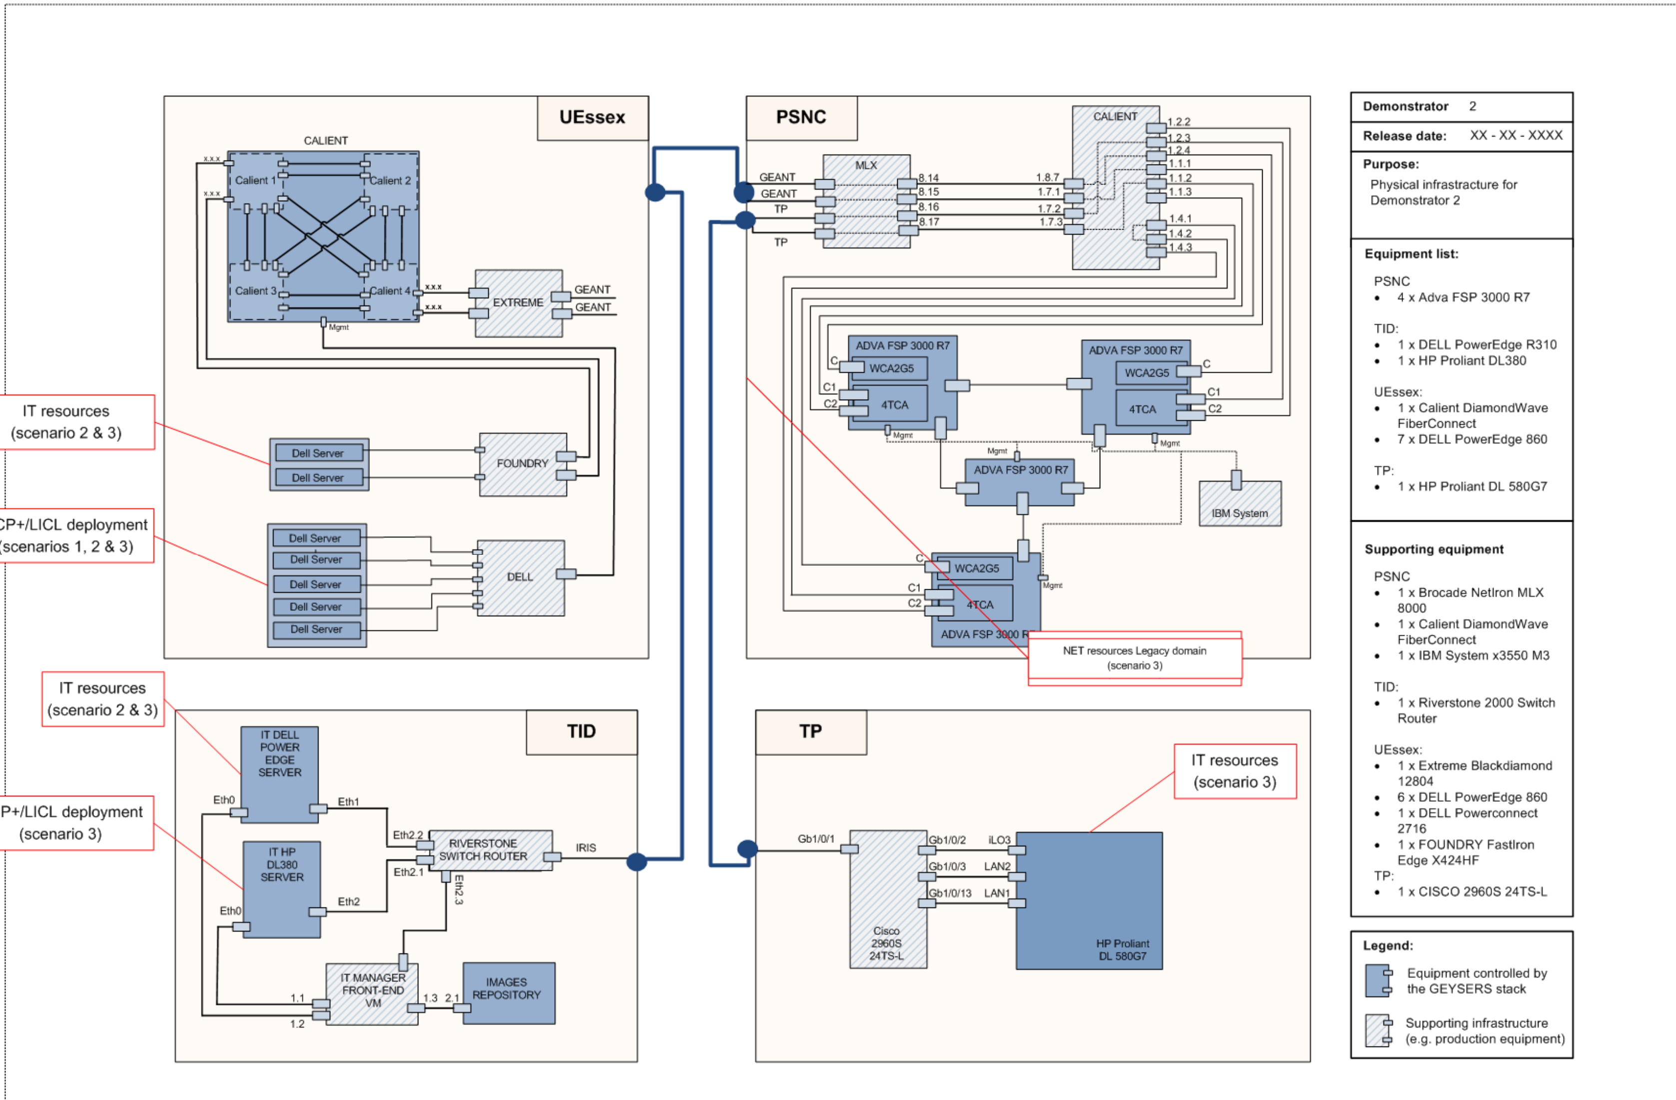

Figure 33: Overall test-bed for Demonstrator 2

4.3.4 Virtualized infrastructure of this Demonstrator

The Virtual Infrastructures considered for the three selected scenarios are shown in Figure 34, Figure 35, Figure 36, and Figure 37.

The first scenario (Figure 34) is characterised by a single-domain Virtual Infrastructure including four nodes and two IT end-points with server and storage resources. This is just a logical representation of the infrastructure topology since, as explained in the previous section, no real physical infrastructure resources will be used in this scenario but only a set of NCP+ controllers operating with LICL stubs configured to emulated LSC or FSC network nodes.

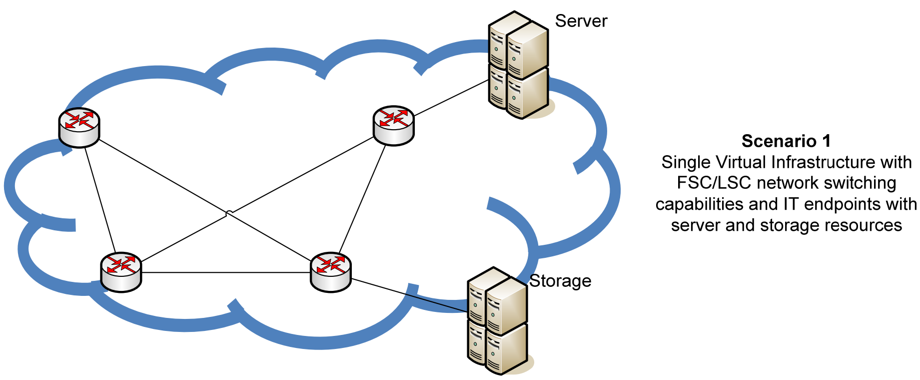

Figure 34: Demonstrator 2, scenario 1 – Virtual Infrastructure schema

:

The second scenario (Figure 35) includes two isolated Virtual Infrastructures “A” and “B” created over a common physical infrastructure composed of FSC network resources from UEssex test-bed and IT resources from the TID test-bed. The topology of the Virtual Infrastructure is the same as in the previous scenario.

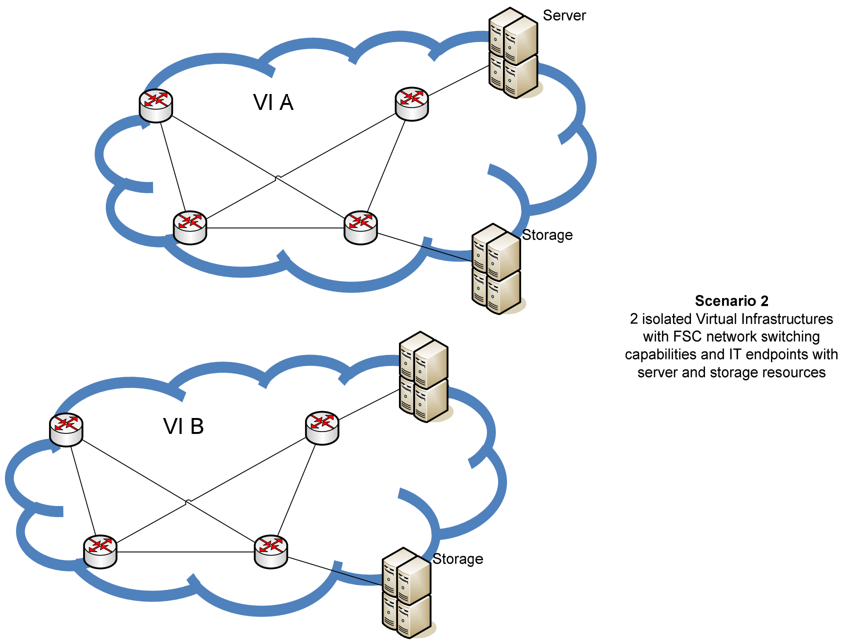

Figure 35: Demonstrator 2, scenario 2 – Virtual Infrastructure schema

The third scenario (Figure 36) is based on a single Virtual Infrastructure organised in two different routing domains "A" and "Z", the former based on FSC switching technology and the latter on LSC switching technology, each of them connecting IT end-points with server or storage resources. The FSC part of the Virtual Infrastructure is created over the physical network resources from the UEssex test-bed, while the LSC part over the network resources from the PSNC test-bed (ADVA FSP 3000R7, using 4TCA cards). TID and TP test-bed are used for IT resources.

A more complete version of the third scenario (Figure 37) includes also an additional transit domain, representing a legacy domain. This domain is not part of the Virtual Infrastructure controlled by the GEYSERS system, but is a physical network domain, with LSC switching technology, that is deployed over the ADVA FSP 3000R7 (using WCA2G5 cards) installed in the PSNC test-bed.

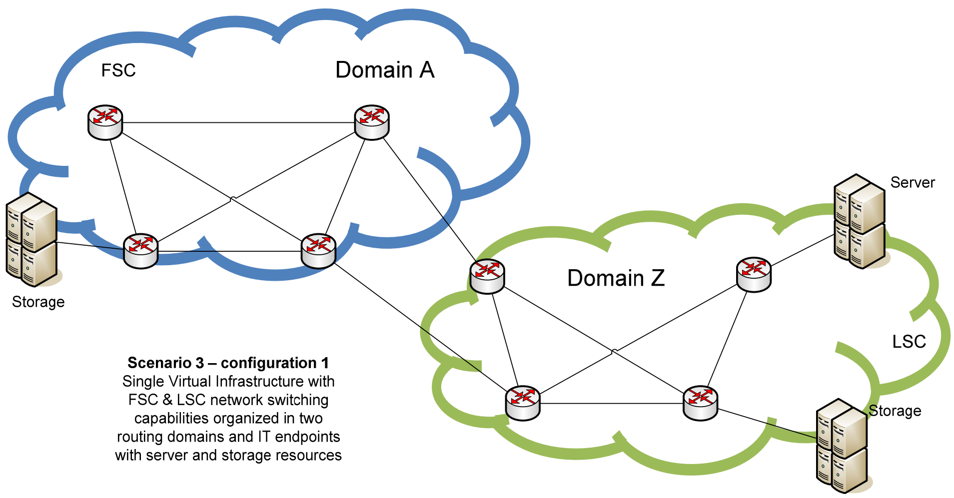

Figure 36: Demonstrator 2, scenario 3 – Virtual Infrastructure schema (1)

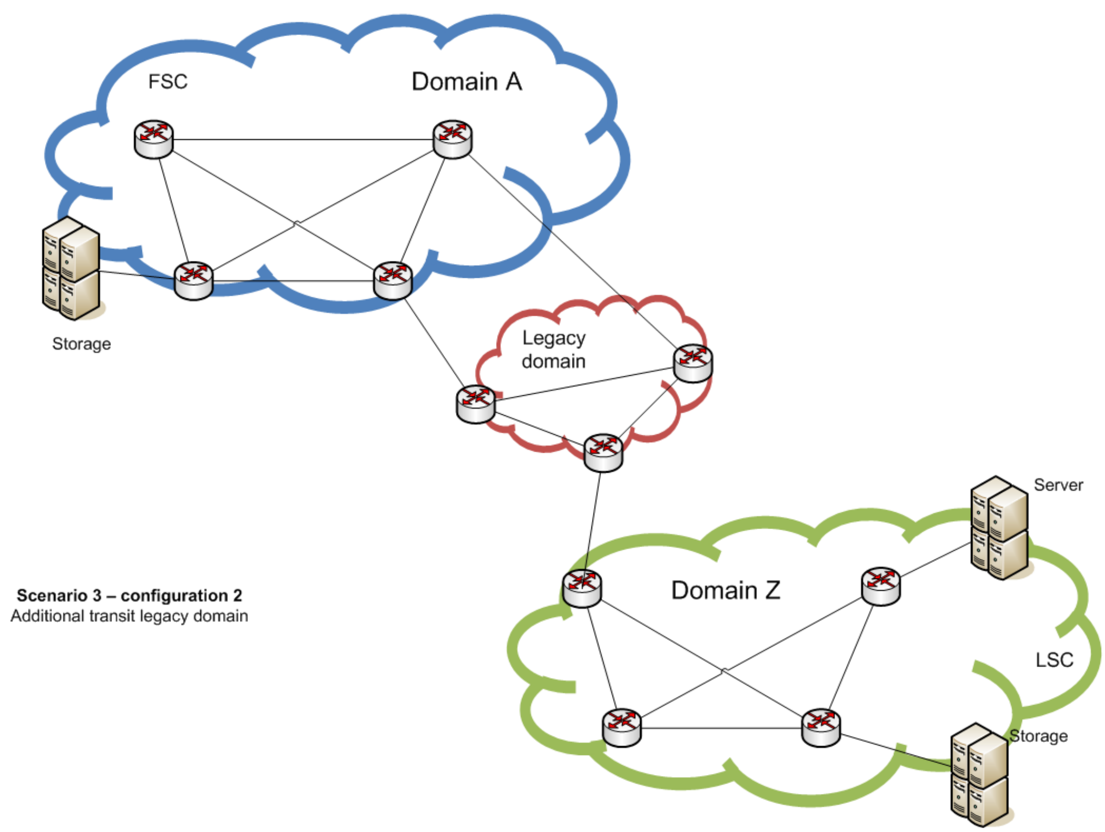

Figure 37: Demonstrator 2, scenario 3 – Virtual Infrastructure schema (2)

4.4 DEMONSTRATOR 3

4.4.1 Introduction and description

The objective of this Demonstrator is to dynamically adjust the available compute, storage and network resources for an Enterprise Information System (EIS) based on the demand from the EIS users. The intention is to motivate methods of synchronising the scaling of IT infrastructure and network capacity dynamically. The number of EIS users can vary up and down, as well as the frequency and size of their requests. SAP leads this Demonstrator and has implemented a workload and payload generator to emulate changing user and transaction loads on a virtual infrastructure. Two alternatives are to be compared: firstly, allowing the infrastructure to perform scaling on behalf of the EIS, synchronising the scaling of the virtual infrastructure (compute and network) against the user load. The second case is having the EIS perform and coordinate the scaling by using the SML's Request API. In both cases, the benefits for cost and overall "satisfaction" of the EIS provider's operational objectives, without disrupting the application users' service level objectives and experience is assessed.

An Enterprise Information System (EIS) is a query-intensive application-server and database system that serves various concurrent users. In a dynamic EIS, the number of users and frequency of queries changes, such that the network and resource demand changes. The capacity of the network and computational power of the servers involved should then reflect this demand without being over-provisioned. Since the objective is to cause the virtual resources in a virtual infrastructure serving an EIS to be scaled up and down based on the user's load, the input data required is for the generation of load on IT resources as well as on the network. Moreover, it is necessary to state what the expectations of users are in order to validate the response of the architecture to the changing user loads. This data can be obtained from empirical analysis of an existing EIS or based on existing benchmarks for Enterprise information systems including analytics and business information warehouses. There are also existing commercial and open-source load generation tools that can be used to simulate the changing of user loads and the generation of payloads to pass over the network. This can be randomised data or test data from a real-world system. The latter case might be difficult from the perspective of sample data sensitivity from real systems. The experiment should show scaling under different types of load conditions. Load can be created from the client of the system under study or from the business of the servers dealing with other requests from different clients in parallel.

The EIS application consists of multiple components, each of which is an OSGi bundle. They can be deployed in one container or distributed over multiple containers. In this report, a completely distributed approach where each component runs in its own container is assumed. Each component offers multiple services that are transparently discovered by the underlying OSGi framework. Furthermore, a WSDL is generated for each component so that external applications can communicate with the components as well. The workflow for the EIS scaling experiments is shown below.

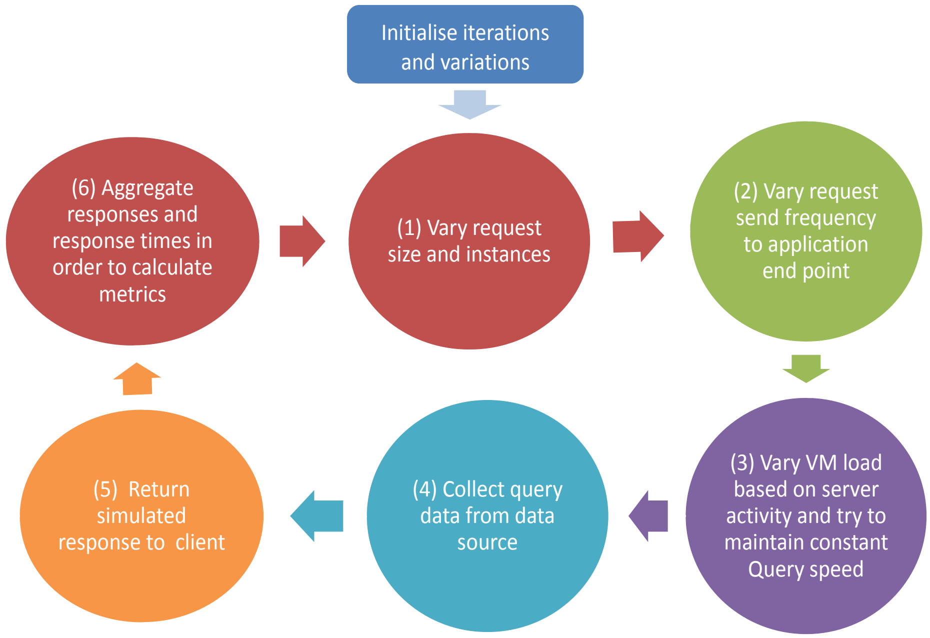

Figure 38: Experimental Workflow for Demonstrator 3

The first stage is to initialise and setup the virtual infrastructure required for the EIS scaling. This includes setting up the necessary monitoring probes and controllers in the infrastructure required to implement dynamic scaling. The experiment then continues as follows in a series of iterations:

(1) Starting with an initial request size and simulation of concurrent instances, the first stage is the creation of the EIS requests from the client machines to the VM with the EIS compute instance.

(2) The request send-frequency varies over time, as well as the size and nature of payload. The EIS server is implemented to support these different requests and seeks to handle the requests with a constant response time

(3) The locally-generated loads at the VM1 are varied in order to emulate shared infrastructure conditions at the server.

(4) Each request includes a query to be submitted to the EIS database. These are expected to be executed within a response time threshold although the size and frequency of requests is varied. A response is generated to show that the request has been handled and the time at which it was received.

(5) The response generated at the server is returned to the client machine. The response should contain its send time, receive time at the server, processing time and return time when stored at the client.

(6) The client aggregates and stores all responses so that they can be traced to the original requests.

The experiment is repeated for a number of iterations in three different scenarios, using infrastructures that are increasingly complex:

- **Scenario 1**: reliance on best-effort network and static configuration of infrastructure where there is no variation with load. It might also be possible to implement an application-layer load balancer that starts up new instances of the EIS server on new VMs when there is a detected load.

- **Scenario 2**: over- or median- provisioned resources based on a calculation of the maximum set of resources that would be required. The load should hence never exceed the available resources, although there is excess available during periods of low load.

- **Scenario 3**: dynamic scaling of the network infrastructure according to the load on the EIS. This is the main scenario of the Demonstrator, showing the capabilities of the GEYSERS architecture.

It is predicted that there will be advantages for Scenario 3 over Scenarios 1 and 2 based on cost and effectiveness. However, Scenario 3 is the most difficult configuration to implement of the three, since this is based on the novel mechanisms to be developed within GEYSERS. The local, physical test-beds required do not vary from the descriptions in Demonstrator 1 and 2.

4.4.2 Physical test-bed provided for this Demonstrator

The same physical test-bed will be used for all scenarios in the Demonstrator, since the aim is to compare the results across scenarios. The physical test-bed uses resources from three IT resource providers (Lyatiss, UvA and TP) and one network resource provider (TID). The three IT resource providers are used to emulate a multi-cloud, where multiple providers are used to make physical hosts available for virtual workloads. In the following, the details about the physical resources allocated for this Demonstrator are presented:

Lyatiss plays the role of a primary cloud provider in the Demonstrator, enabling the creation of VMs for EIS instances. Their CloudWeaver appliance is used behind the Lower-LICL to manage the creation of the initial VMs, as well as instances required dynamically in Scenario 3 of the Demonstrator.

UvA is also a cloud provider in the Demonstrator, but referred to as a “secondary provider”, whereby their IT resources are engaged when there is a need for load balancing or redirection. UvA uses OpenNebula as the resource manager behind the Lower-LICL.

TP is used as the composite or multi-cloud provider in this Demonstrator, where they play the role of VIP. Their resources are hence used to install the Upper-LICL, which requires connectivity to the Lower-LICLs at Lyatiss and UvA. TP will be used as the site where the NCP+ and SML are deployed for the scenario, as this seems to be a reasonable assumption that the same organisation will act as VIP and VIO.

TID is the network provider used to connect all three sites. This enables control over both the scaling of IT servers and the network resources used in the Demonstrator.

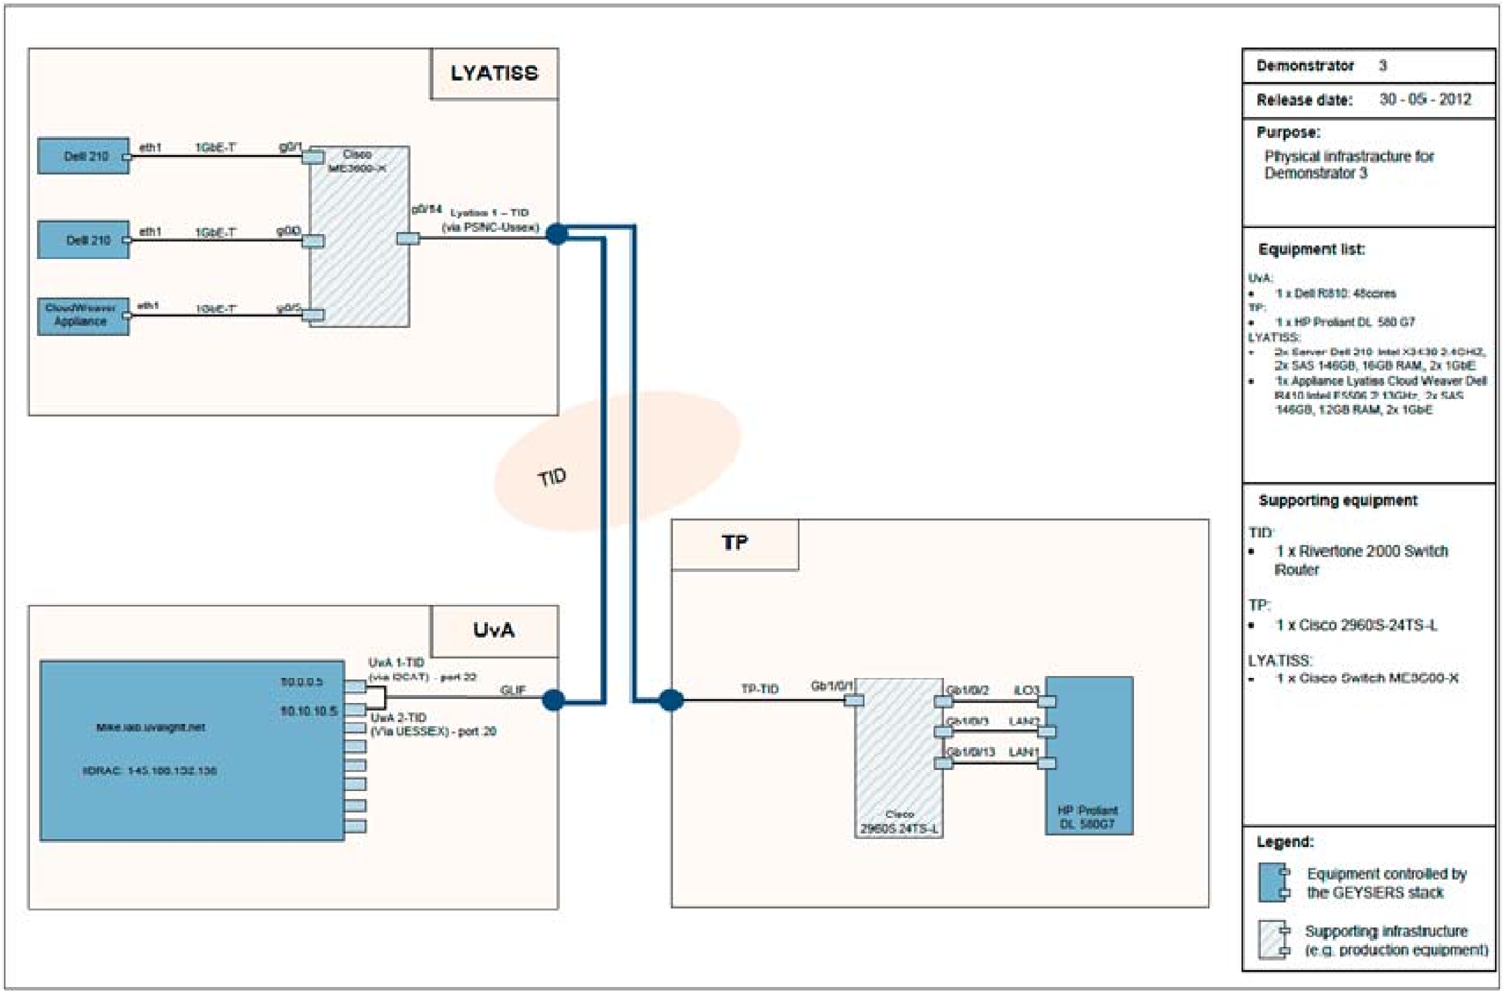

Figure 39: Physical Infrastructure for Demonstrator 3

4.4.3 Physical resources planning and virtualization

Table 12 lists the usage of the physical infrastructure for the 3 scenarios, while Figure 40 illustrates the topology and usage of the GEYSERS components in the scenarios.

	Scenario 1 "Best Effort"	Scenario 2 "Median Provisioning"	Scenario 3 "Dynamic Scaling"
NET PHY Infrastructure	**TID** provides access to their optical infrastructure consisting of 4 ADVA FSP 3000 F7 switches. These are just configured with VLAN tagging.	**TID** still used but the bandwidth is set to a determined median limit for the VLAN.	**TID's** configures switches to support dynamic scaling scenario.
IT PHY Infrastructure and Lower LICL Controllers	**Lyatiss:** 2 x Dell 210 servers. One used for Load-balancer and the other for the EIS application instance (s). Also assume availability of probes for CPU, RAM and power usage. **TP:** traffic generator (1 Agilent N2X router tester module)	**Lyatiss:** same resources a scenario 1 but more EIS instances started initially **TP**: same traffic load generated	**Lyatiss:** same resources as scenario 1 and 2 **UvA:** DAS 3 Blade. Note that UvA used as secondary resource provider **TP** – In addition to traffic generator will use the test-bed server: HP Proliant DL580 G7
Machines to deploy NCP+ controllers	N.A.	N.A.	**TP:** current plan is to use TP's server for the NCP+ as it will implement the VIP. If this proves inefficient during trials, we will use UvA's servers as an alternative.
Machines to deploy Upper LICL	**TP** used as primary site for deploying the Upper-LICL. **UvA** used in case planned capacity is unavailable.	(Same)	(same)
Machines to deploy SML	**TP** used as primary site for deploying the Upper-LICL. **UvA** used in case planned capacity is unavailable.	(Same)	(same)

Table 12: Demonstrator 3 – overview of physical resources

The distribution and topology of resources described in Table 12 is shown in Figure 40. SAP is also providing a VM for representing the application provider and end user perspective on the Virtual Infrastructure.

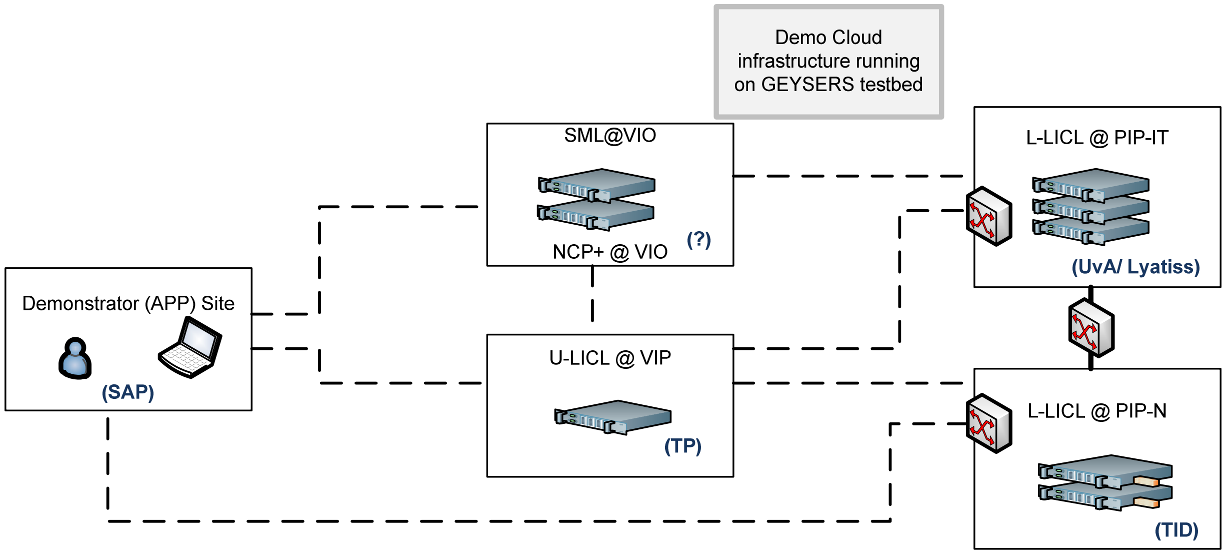

Figure 40: Topology of Infrastructure for Demonstrator 3

4.4.4 Virtualized infrastructure of this Demonstrator

The Virtual Infrastructure should appear as a single domain to the application provider (Figure 41). The IT resources for the EIS and load generators are deployed on VMs in different physical domains but connected by the same network provider (TID).

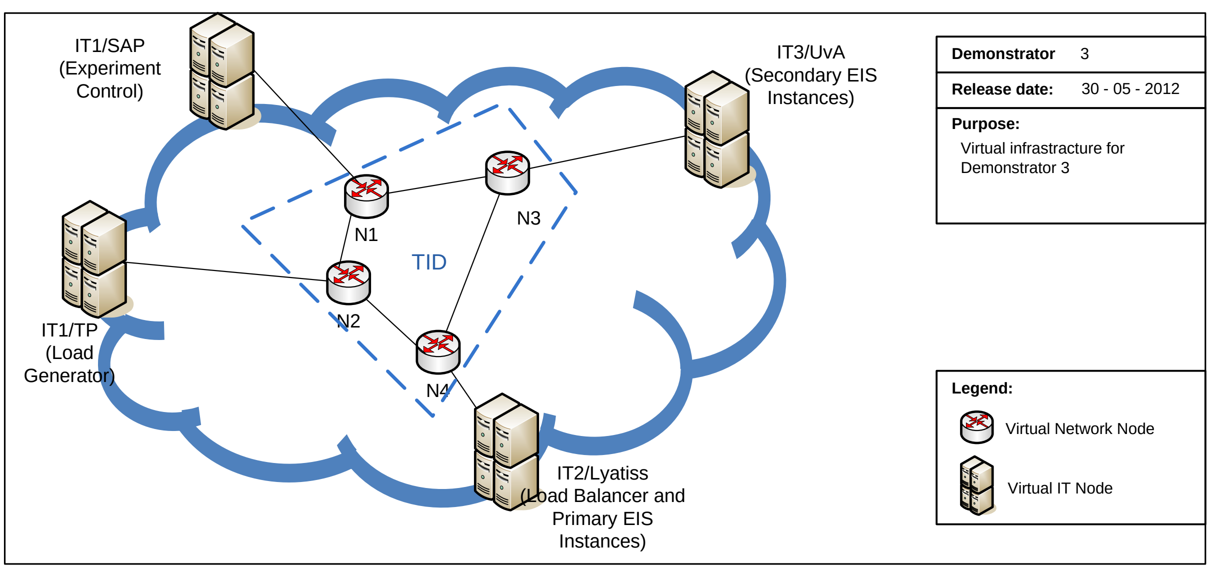

Figure 41: Experimental Workflow for Demonstrator 3

4.5 DEMONSTRATOR 4

4.5.1 Introduction and description

Demonstrator 4 is aimed at the demonstration of advanced network and IT management functionalities, with a specific focus on network infrastructure re-planning. The virtual infrastructure re-planning service is offered by the Virtual Infrastructure Provider (VIP) to the Virtual Infrastructure Operator (VIO), allowing the VIO to request the modification, up-scaling or down-scaling (e.g. upgrade of link capabilities, modification of network topology) of the leased virtual infrastructure in order to optimise the network resource utilisation. The Demonstrator will evaluate and compare manual and dynamic methods for the VI re-planning, by providing usage examples for both of them.

Two main GEYSERS architecture components: LICL and NCP+ take part within this Demonstrator. LICL is providing virtual infrastructure re-planning capabilities whereas NCP+ is responsible for automatic or manual triggering re-planning actions and adapting for dynamical infrastructure changes.

The physical infrastructure provided for Demonstrator 4 should allow for testing and demonstration of all re-planning capabilities defined for this Demonstrator:

- Link bandwidth re-planning - VIO operator is increasing/decreasing virtual link bandwidth,
- Connectivity re-planning - VIO operator is adding virtual connectivity between any two nodes,
- Link re-planning - VIO operator is removing a virtual link,
- Node re-planning - VIO operator is removing a virtual node.

4.5.2 Physical test-bed provided for this Demonstrator

Demonstrator 4 will be deployed over a subset of the overall GEYSERS test-bed, composed of five different sites (PSNC, Lyatiss, IRT/ALU-I, UvA, TP). Most of the sites are interconnected via GÉANT and local NRENs. In this section, the local-test-beds are briefly described, and the overall physical infrastructure for GEYSERS Demonstrator 4 is provided.

PSNC local test-bed resources – PSNC local test-bed (Figure 42) is composed of three ADVA FSP 3000 RE2 devices and one Calient DiamondWave FiberConnect switch which are managed by GEYSERS software. All tributary ports of ADVA DWDM system are connected to Calient switch which perform both supporting (five ports are configured manually) and operational actions (four ports are configured by GEYSERS software). Additionally, Calient switch connects PSNC optical equipment to another supporting Brocade NetIron MLX 8000 switch responsible for establishing of VLAN Q-in-Q interconnections with Lyatiss and UvA and TP test-beds. The last supporting equipment within PSNC test-bed is IBM System x3550 M3 server which will allow deploying Lower-LICL and NCP+ VMs.

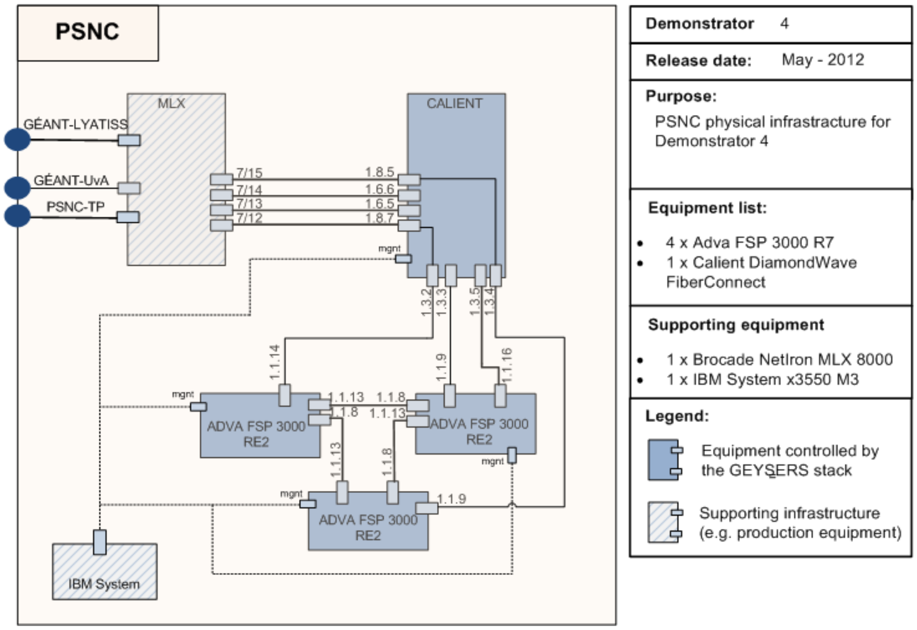

Figure 42: PSNC test-bed for Demonstrator 4

Lyatiss local test-bed resources – Lyatiss deploys (Figure 43) an IT-only local test-bed, composed by two Dell R210 servers installed in IN2P3 (Lyon) Datacenter. The hypervisor adopted is KVM; VMs and GEYSERS stack software can be deployed as needed. Additionally, a Cisco ME 3600-X is used for establishing VLAN dot1Q interconnections with the IRT/ALU-I test-bed, and VLAN Q-in-Q interconnections with the PSNC test-bed.

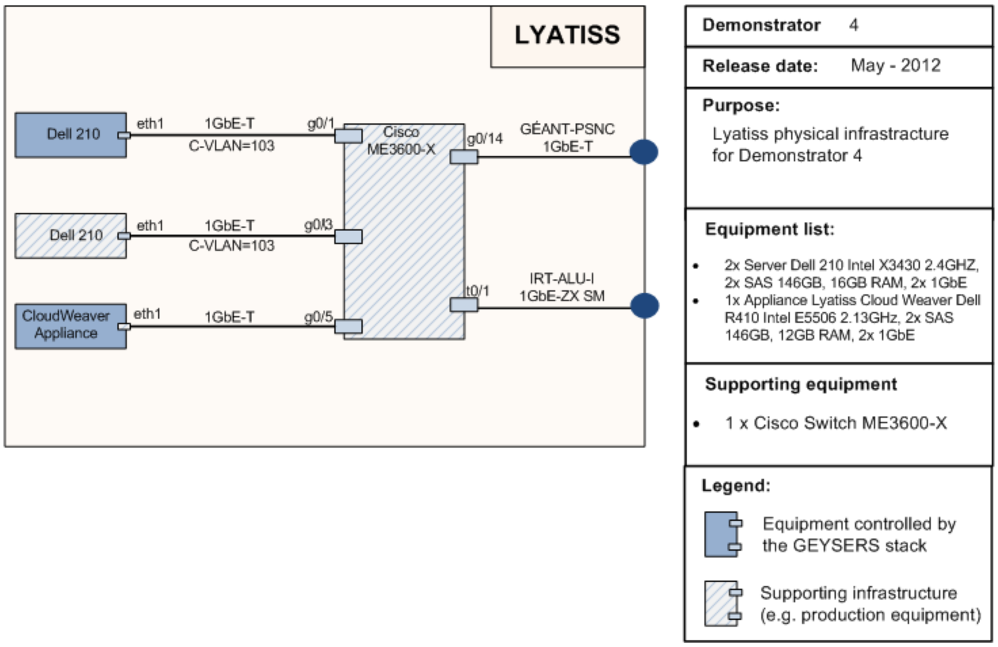

Figure 43: Lyatiss test-bed for Demonstrator 4

IRT/ALU-I local test-bed resources – IRT and ALU-I jointly deploy (Figure 44) a local test-bed for Demonstrator 4, composed of an optical network node and an IT server. IRT basically provides connectivity towards the Lyatiss test-bed in Lyon over its optical network, and it offers public IP Internet access and co-location services for ALU-I equipment. In fact, all the equipment is installed at IRT's premises in Milan Caldera, Building C. The Cisco 1841 router acts as a CPE and is deployed as an IRT standard solution for managed public IP Internet access. The other equipment is provided by ALU-I. The optical network node is an Alcatel-Lucent 1850 TSS-160, a Packet-Optical Transport switch of the Alcatel-Lucent TSS product family. The IT server is an HP Z600 Workstation, which will run several VMs, plus a lover LICL instance. A Cisco 3600 router is also deployed by ALU-I for the remote management of their equipment.

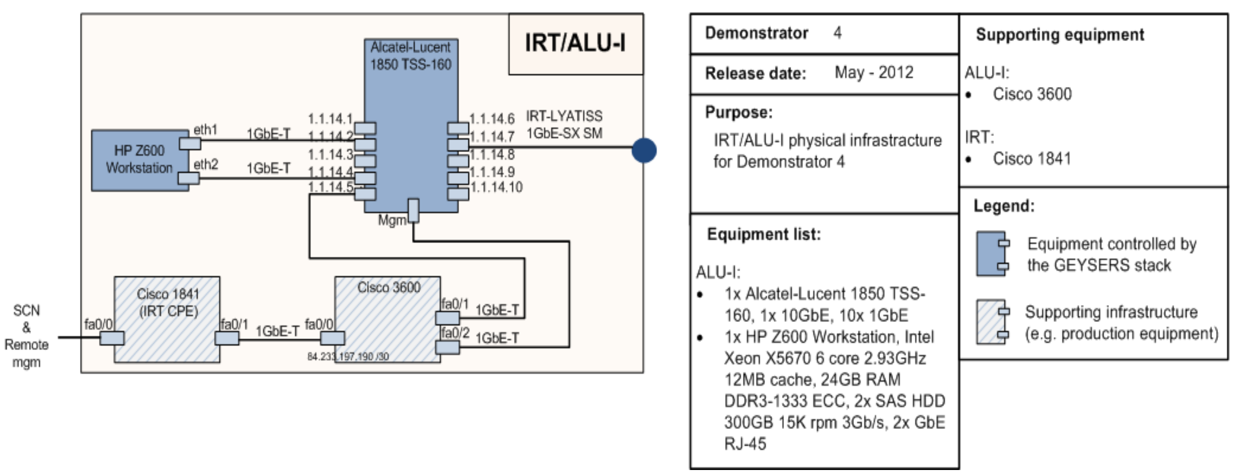

Figure 44: IRT/ALU-I test-bed for Demonstrator 4

UvA local test-bed resources – The UvA test-bed (Figure 22) provides a cloud-machine (Dell R810: 48cores, 128GB-RAM, 600GB storage) for on-demand computation and storage resources. These resources are provided with the aid of OpenNebula installed and the necessary software to deploy the Lower-LICL. This test-bed provides access via two GLIF lightpaths, to i2CAT and to UEssex.

TP local test-bed resources – The TP local test-bed (Figure 45) includes an HP Proliant DL580 G7 that will be used to provide a set of physical IT resources. These resources will be considered in combination with the physical network resources within the PSNC test-bed as a physical domain owned by a single PIP. This means that the overall set of resources located into the PSNC and TP test-beds will be managed by a single instance of a Lower-LICL (instantiated on the PSNC servers).

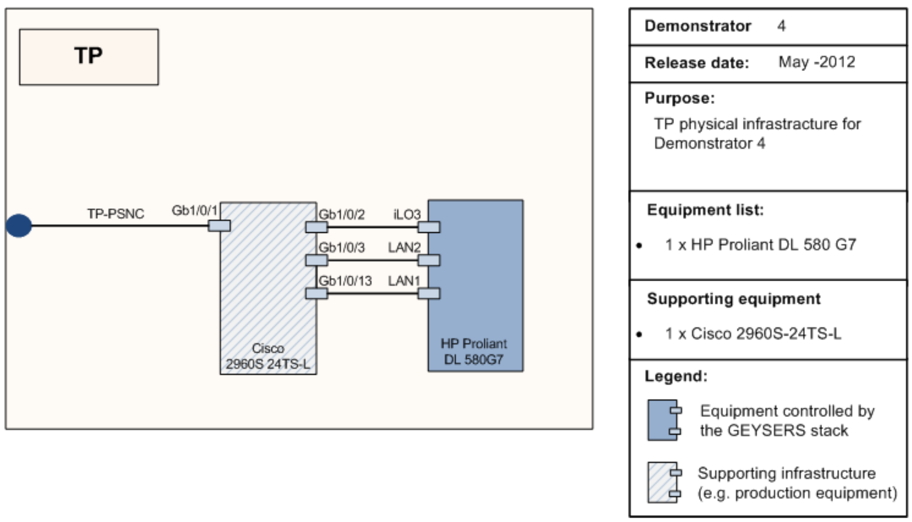

Figure 45: TP test-bed for Demonstrator 4

The overall test-bed for the GEYSERS Demonstrator 4 is depicted in Figure 46. In this figure, interconnections between local test-beds are presented with values of VLAN tag(s). The single VLAN values in the figure imply the usage of 802.1Q standards with a given tag value. The double VLAN values in the figure (i.e. S-VLAN, C-VLAN) imply the usage of the 802.1QinQ standard with given s-tag and c-tag values. The PSNC-UvA data link is established using UEssex as an intermediate hop. Similarly, PSNC-IRT data links are established via a statically configured hop in Lyatiss.

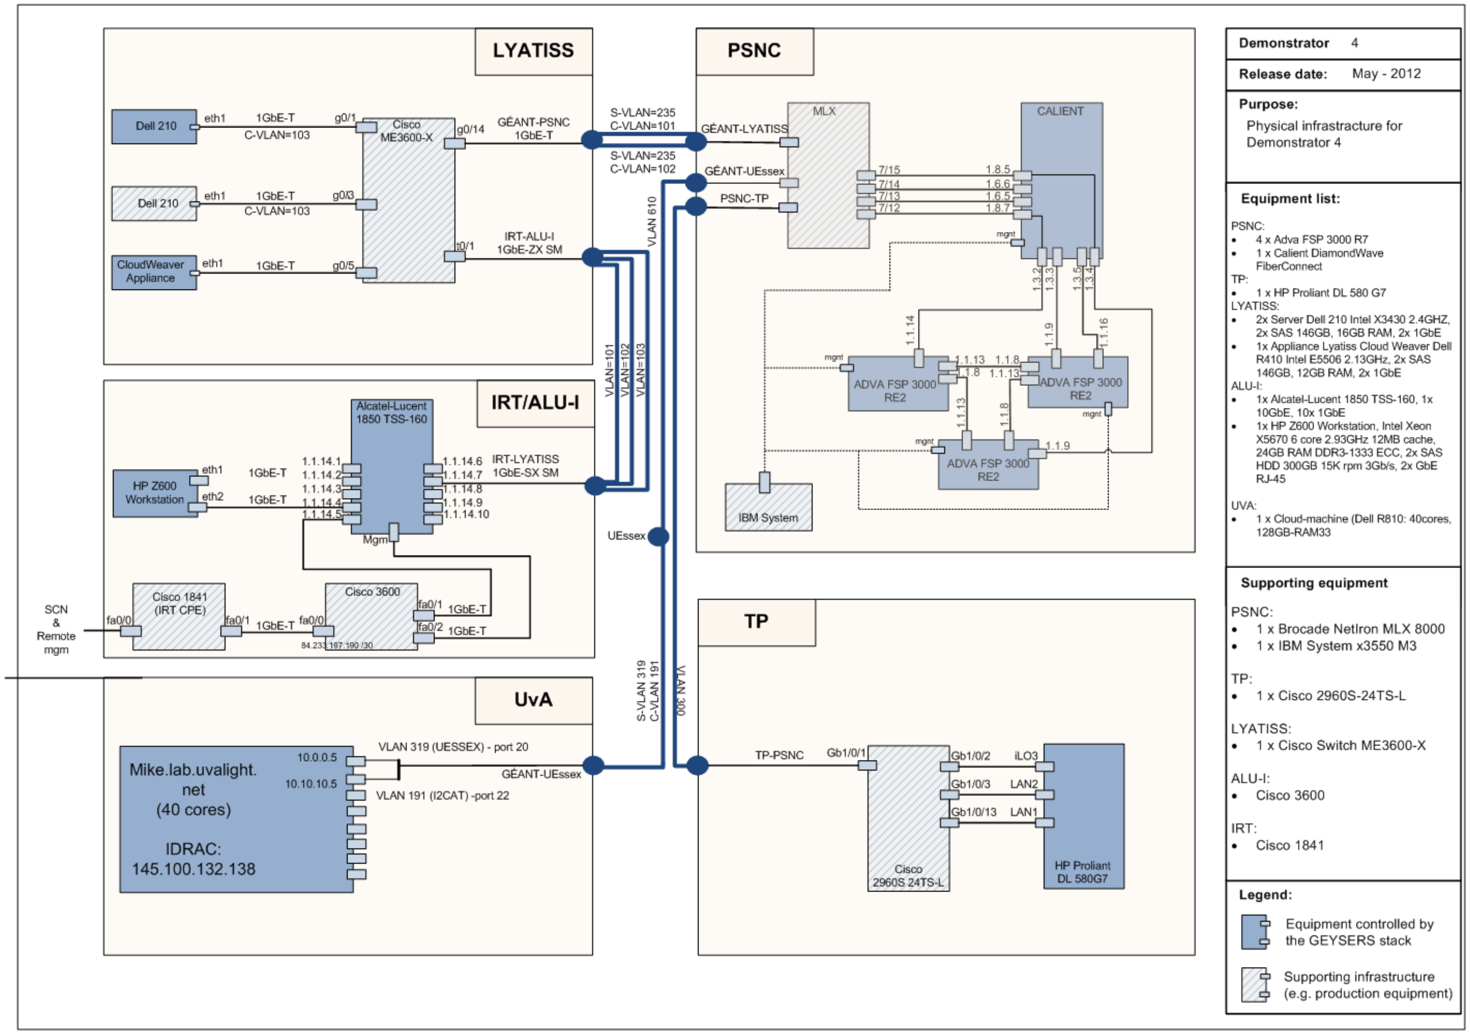

Figure 46: Overall test-bed for Demonstrator 4

4.5.3 Physical resources planning and virtualization

Physical resources provided for the Demonstrator 4 are provided by five independent organisations: PSNC, Lyatiss, IRT/ALU-I, TP and UvA. For this reason, five Physical Infrastructure Providers (PIPs) have been defined, which deploy Lower-LICL software. PSNC and IRT/ALU-I are managing the network infrastructure as PIP-N. Lyatiss, TP and UvA are managing the IT infrastructure as PIP-IT. Lower-LICL software offers part of the managed physical infrastructure to the Virtual Infrastructure Provider (VIP). In this Demonstrator, TP is performing the role of the VIP and deploys Upper-LICL software, responsible for creating a single virtual infrastructure composed of logical resources offered by all five PIPs. Finally, PSNC is using a given virtual infrastructure as VIO and deploys enhanced Network Control Plane (NCP+). It is important to note that during these scenarios two actors (PSNC and TP) play two roles at the same time: PSNC is a PIP-N and VIO, TP acts as a PIP-IT and VIP.

Figure 47 presents the organisational structure of the GEYSERS-defined roles performed by organisations participating in the Demonstrator 4. Additionally, GEYSERS software elements are installed within each of organisations and physical infrastructure (hardware equipment and data links) controlled by the GEYSERS software. Physical resources are shown in a logical form: this means that the supporting equipment (such as switches) is not shown here. It is, however, used to interconnect IT resources (servers), but does not play an active role in the GEYSERS environment. Therefore, this equipment can be treated as a part of a physical link – from the point of view of the user (VIP in this case, as it is the entity that will actually create a VI, and VIO as this an organisation that will actually use it) and the actual physical connection method is not important since it will not be visible to the users.

IT resource virtualization refers primarily to segmenting and isolating portions of the GEYSERS CPU, memory and storage on a physical server, such that a different guest Operating System can be run.

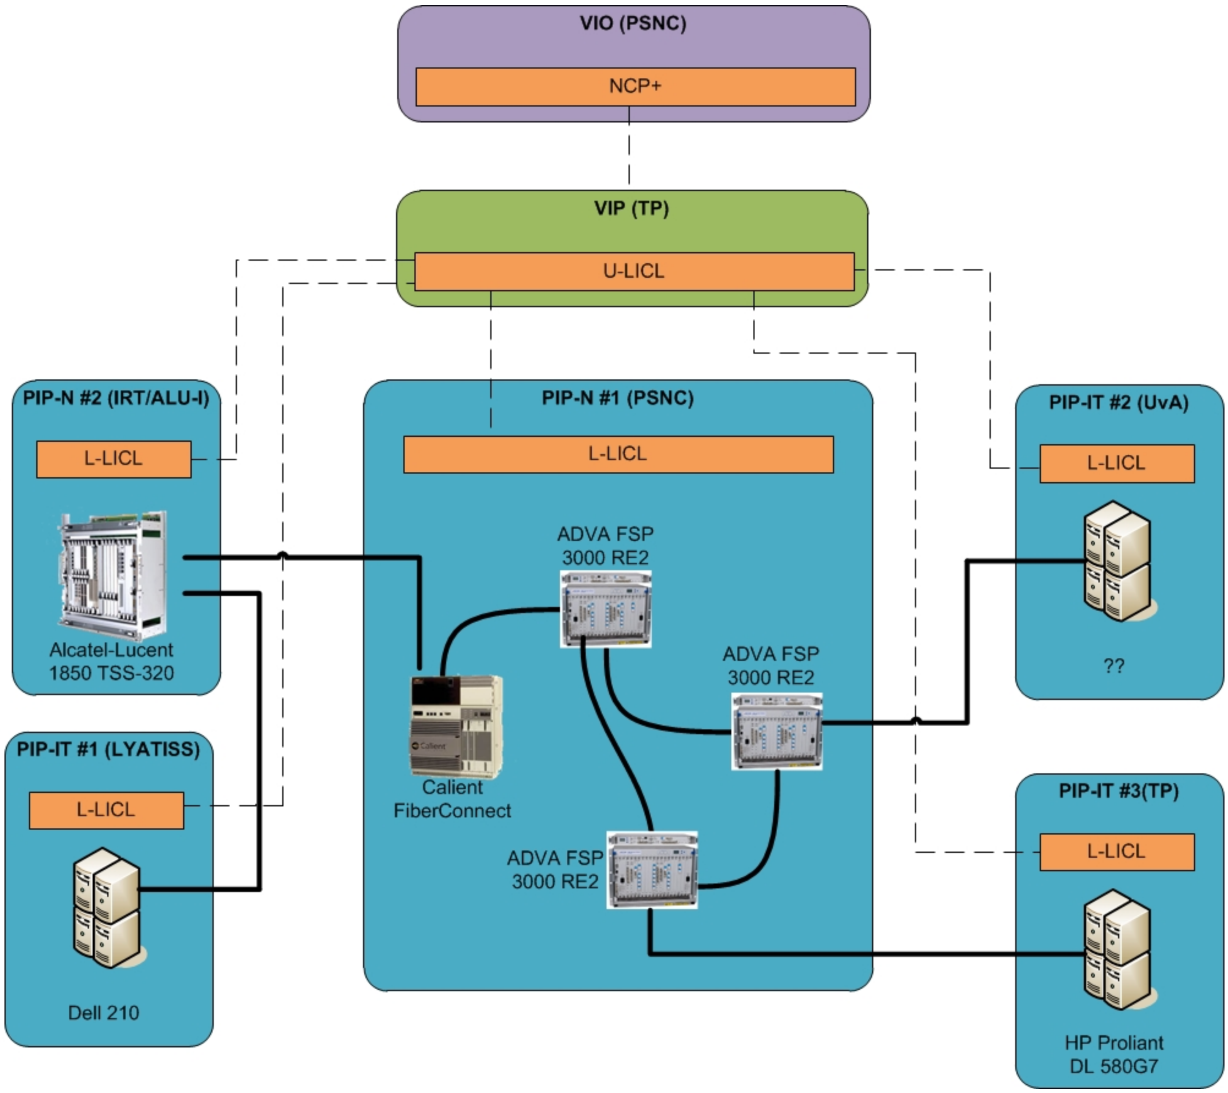

Figure 47: Demonstrator 4 – Business role assignment and GEYSERS components deployment

Each organisation, taking part in the Demonstrator 4 as PIP-N, PIP-IT, VIP or VIO role, has to deploy some components of the GEYSERS software.

Table 13 presents supporting infrastructure requirements for the software deployment. Each software component is provided in form of a VM which has to be installed using a physical server located in a particular organisation. Each physical infrastructure is based in a different geographical location.

GEYSERS component	Deployment requirements
NCP+ controllers	• GMPLS+ controllers (5 VM) • NIPS Server + AAI server (1 VM) • Child PCE+ server (1 VM) • AAI server (1VM) • VRM server (1VM) Total: 1 server required (PSNC)
LICL controllers	• Lower LICL (1 VM) Total: 1 server required (PSNC) • Lower LICL (1 VM) Total: 1 server required (IRT/ALU-I) • Lower LICL (1 VM) • CW appliance (1 VM) Total: 2 servers required (Lyatiss) • Upper LICL (1 VM) • Lower LICL (1 VM) Total: 1 server required (TP) • Lower LICL (1 VM) • CW appliance (1 VM) Total: 2 servers required (UvA)

Table 13: Demonstrator 4 – GEYSERS LICL and NCP+ components deployment requirements

For the sake of independence, VMs with different physical connections (Ethernet ports) should be used for each VM instead of VLANs.

4.5.4 Virtualized infrastructure of this Demonstrator

In the Demonstrator 4, only one virtual infrastructure is created by the LICL and it is controlled by NCP+ as a single administrative and operational domain. It is composed of five virtual network nodes, three virtual IT nodes and eleven virtual links interconnecting all nodes. A characteristic issue of the virtual infrastructure within this Demonstrator is that the virtual infrastructure is changing its topology depending on VIO requests. The virtual infrastructure presented in Figure 48 represents the maximum sized VI which is limited by the available physical infrastructure deployed for the Demonstrator 4.

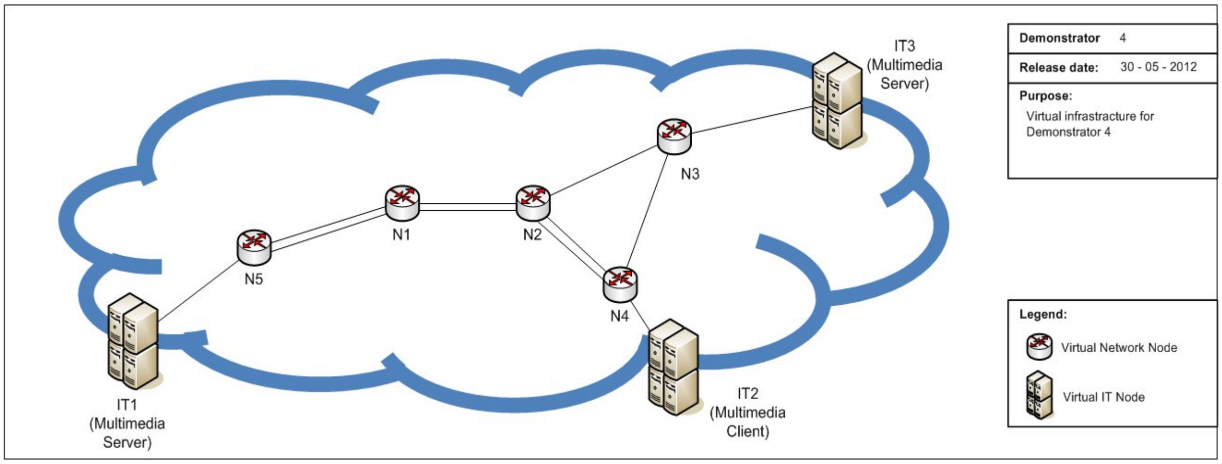

Figure 48: Virtual Infrastructure for Demonstrator 4

Table 14: Demonstrator 4 virtual nodes, presents the details of each virtual node like related physical equipment, its location and virtual node capabilities. Virtual network nodes N1 and N5 have port switching capability which is related to the full network abstraction provided by the LICL layer. Virtual network nodes N2-4 have lambda switching capability and in the case of these nodes, the NCP+ must recognise their lambda continuity and directionality constrains in order to calculate a physically possible path. Virtual IT nodes 1 and 3 are providing multimedia services such as video streaming whereas Virtual IT node 2 is a consumer of multimedia streams.

Virtual node	Physical equipment based on	Located in	Capabilities
Network node 1	Calient DiamondWave FiberConnect	PSNC	Port switching
Network node 2, 3, 4	ADVA FSP 3000 RE-II	PSNC	Lambda switching (lambda continuity and directionality constrains)
Network node 5	Alcatel-Lucent 1850 TSS-160	IRT	Port switching
IT node 1	Dell 210 Intel X3430	Lyatiss	Multimedia server
IT node 2	?	UvA	Multimedia client
IT node 3	HP Proliant DL 580 G7	TP	Multimedia server

Table 14: Demonstrator 4 virtual nodes

5 CONCLUSIONS AND NEXT STEPS

This document is an update of the test-bed high-level specification provided in [D5.1]. This deliverable D5.1.1 offers a detailed description of the local test-beds deployed by each involved partner both for data and control planes and the interconnections between them. The specification of the local test-bed capabilities enables the precise identification of the devices, the capabilities and roles that each partner offers in each of the Demonstrators by their partial or their full local test-bed infrastructure. These local test-beds constitute the basis for the definition and deployment of the Demonstrators.

The provisioning of IT and network resources constitutes one of the key outcomes shown in this document. Physical resource providers (PIPs) integrate in their local test-beds GEYSERS modules and virtualization tools. The virtualization planning is carried out by segmenting the physical resources, resulting in a better utilisation of these resources based on the GEYSERS modules and tools. These virtualized infrastructures are then ready to be managed by VIPs. Enabling the virtualization technology openness with the approach of physical resource adapters is one of the emergent capabilities of the GEYSERS framework. This permits the lower levels of the stack that implement virtualization functions to be selected from a wide range of technologies that exist in the marketplace and as open source today, extending further the value chain market. On the other hand, the virtualization of optical nodes can be achieved by partitioning or aggregating, while optical links (i.e. network resources) are virtualized taking into account the bandwidth granularity or switching technology to be used.

Demonstrators constitute the focus of WP5. A key outcome of the document is the exhaustive definition of each of the Demonstrators matching the use cases and scenarios described in [GEYSERS-D1.5] which show the added value of the GEYSERS capabilities. Each Demonstrator is focused on different technical items showing the potential applicability of the GEYSERS architecture and technology:

- Exploitation of physical devices through virtualization mechanisms taking into account scalability and optimum resources utilisation.
- The benefits of the LICL and NCP+ for managing/operating the virtualized resources and infrastructures in the network.
- SML and NCP+ interactions to provide the VIOs with enhanced mechanisms to offer dynamic services (unicast, restricted anycast, etc) specifically tailored to application requirements upon on demand requests.

For each Demonstrator the physical resources corresponding to the local test-beds as well as the virtualization planning and the resulting virtualized infrastructures have been defined. By means of the Demonstrator scenarios, the document shows both the physical and virtualized infrastructure and the management processes carried out by PIPs, VIPs and VIOs.

Future steps in WP5 will focus on the lead partners' finalisation of the deployment of the local test-beds and inter-connecting them to other partners' local test-beds in order to validate the final development of the Demonstrators. The story-line to show the GEYSERS architecture functionalities and performances within each Demonstrator is currently being refined. This will guide the process to show all of GEYSERS' capabilities, technical solutions and modules developed in the context of the project.

In parallel, the software and hardware integration activities will continue. The results of the integration activities will feed the GEYSERS Demonstrators with the integrated software stack, and enable experiments in a real, networking environment. At least two face-to-face meetings are planned in the forthcoming period to finalise the integration work between the software components being developed in WP3 and WP4, taking into account the specification of cross-layer interfaces detailed in WP2.

The impact of further steps in WP5 will reflect the achievements done in WP2, WP3 and WP4, and the feedback based on validation results will be provided to the technical WPs to refine both architecture design and prototype implementation.

In the context of the cross WP5-WP6 interaction, new dissemination opportunities will be studied, related to the development carried out in WP5. In the forthcoming months, a closer collaboration is expected with industrial partners involved in both work packages to identify and detail the exploitation plan and study the impact of potential use cases implementation in local test-beds of selected industrial GEYSERS stakeholders.

The final version of the integrated GEYSERS software stack will be prepared in WP5 and shared with external users for download and dissemination. WP5 has been supporting the development work packages, i.e. WP3 and WP4, in making a decision on the software licensing scheme to be used in the project. The shared stack placed on the public GEYSERS web site in collaboration with WP6 will reflect the actual decision on the software licensing scheme made by the software developers of each module of the GEYSERS project.

6 REFERENCES

[REF 1] GEYSERS-D1.5. Collected Use Cases and Scenarios. Available at: http://wiki.geysers.eu/images/8/8b/GEYSERS_D1.5_final.pdf

[REF 2] GEYSERS-D2.2 – Update. GEYSERS overall architecture & interfaces specification and service provisioning workflow, May 2011. Available at: http://wiki.geysers.eu/images/3/38/WP2_GEYSERS_D2.2_update.doc

[REF 3] GEYSERS-D2.6. Refined GEYSERS architecture, interface specification and service provisioning workflow. Available at: http://wiki.geysers.eu/images/7/7f/D2.6-final.pdf

[REF 4] GEYSERS-D3.2. Preliminary LICL Software release, January 2012. Available at: http://wiki.geysers.eu/index.php/D3.2

[REF 5] GEYSERS-D3.3. LICL Sub-systems release, April 2012. Available at: http://wiki.geysers.eu/images/b/b2/GEYSERS_D3.3_v1.0.docx

[REF 6] GEYSERS-D4.1 GEYSERS Deliverable D4.1, GMPLS+/PCE+ Control Plane architecture, November 2010. Available at: http://wiki.geysers.eu/images/2/29/GEYSERS_WP4_D4.1_v1.0.doc

[REF 7] GEYSERS-D5.1. GEYSERS test-bed Implementation, July 2011. Available at http://wiki.geysers.eu/images/b/b4/GEYSERS-D5.1_v1.6-final.pdf

[REF 8] OpenNebula: http://OpenNebula.org/

[REF 9] CloudWeaver: http://www.lyatiss.com/resources/

[REF 10] Libvirt: http://libvirt.org/

[REF 11] KVM: http://www.linux-kvm.org/page/Main_Page

[REF 12] XEN: http://xen.org/

[REF 13] VMW: http://www.vmware.com/

[REF 14] 802.1q IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks. Available at: http://standards.ieee.org/findstds/standard/802.1Q-2011.html