GN3plus Open Call: Technical Annex B – GÉANT OpenFlow Facility

Issue Date: 01-04-2013

THE SEVENTH FRAMEWORK PROGRAMME

The Seventh Framework Programme focuses on Community activities in the field of research, technological development and demonstration (RTD) for the period 2007 to 2013

Table of Contents

Table of Figures

Table of Tables

0 Executive Summary

This technical annex is provided in support of the GN3plus Open Call. It describes the GÉANT OpenFlow Facility available for use by GÉANT Open Calls.

- Theme 1: Innovative Uses of GÉANT Network Facilities.
    - Topic 2: Software defined networking

The aim of this document is to provide sufficient technical detail to allow proposers to put together a technical proposal for this topic.

The GÉANT OpenFlow testbed is a facility designed to support software defined networking experiments and prototyping and is located in five GÉANT Points of Presence (PoPs) in the following cities:

- London
- Frankfurt
- Vienna
- Zagreb
- Amsterdam

1 Using OpenFlow to Realise a Shared Testbed

1.1 Why OpenFlow?

OpenFlow is a technology that allows the creation of virtual switching devices. For the implementation of Network slices, OpenFlow switches can be interconnected using Layer 2 connectivity virtualisation technologies.

The OpenFlow specification defines an open protocol to program the flow table in switches and routers. A network administrator can partition traffic into production and research flows. Researchers can control their own flows by choosing the routes their packets follow and the processing they receive. In this way, researchers can try new routing protocols, security models, addressing schemes, and even alternatives to IP. Therefore, in principle, OpenFlow could be used as a mechanism to separate the production traffic from the experimental traffic, with the former processed as in standard routing, switching equipment. On the top of OpenFlow devices, a controller-computing element is expected in order to control uniformly and collectively the network elements of an OpenFlow infrastructure. In particular, a controller adds and removes flow-entries from the flow table on behalf of experiments. For example, a static controller might be a simple application running on a PC to statically establish flows to interconnect a set of test computers for the duration of an experiment.

1.2 Supporting Network Experimentation

One of the objectives of the GN3plus project is to provide a next-generation pan-European network and related services that meet the communications needs of research communities in all fields. Such needs include both a transport facility for production data and a network environment where experiments can be conducted.

To prevent the best-effort production traffic of commodity services from being disrupted by high-bandwidth applications and experiments, it makes sense to separate them. This also enables researchers to modify the behaviour of infrastructure elements, such as traffic routing, which could not be realised on the production infrastructure.

The GÉANT network meets these requirements by implementing a GÉANT OpenFlow Facility complete with associated support services. This facility aims to deliver a platform that is capable of creating multiple (virtual) networks on top of the GÉANT production environment.

1.3 Use Cases

The GÉANT OpenFlow facility aims to support applications and user communities requiring SDN functionality within the GÉANT Testbeds as a Service (TaaS).

Two major categories of use cases have been identified:

1. Using OpenFlow as a traffic engineering mechanism to manage programmatically the backbone capacity and paths to serve specialised applications and protocols at the end systems. Typical examples of such use cases include virtual machine live migration across remotely located data centres over a WAN and efficient big data transfer utilising OpenFlow-enabled traffic engineering [MPTCP].
2. Using OpenFlow to deliver "vanilla" Layer 2 slices for Layer 2 (and above) research and experimentation on the actual network data and control plane technologies to be carried out. Research on new protocols or capabilities such as that of [EthOAMOF] fall into this category.

2 GÉANT OpenFlow Facility

2.1 Background

The GÉANT OpenFlow Facility has been developed to add OpenFlow capabilities in the GÉANT Testbeds as a Service (TaaS). The facility is deployed on top of the GÉANT backbone production environment and is so far a single-domain environment.

GÉANT OpenFlow Facility supports:

- Announcement of the facility's network resources (Layer 2 links and OpenFlow-enabled switches) that can be used for the OpenFlow-controlled slice topologies interconnecting computing resources.
- Announcement of the available computing resources per Point of Presence (PoP) of the facility that is attachable to the slice topologies.
- A reservation mechanism for the network and computing resources.
- Management and control plane functionalities for the reserved resources.

2.2 Facility Design

Computing resources are offered as Virtual Machines (VMs) upon dedicated physical servers using Xen[1] [Xen] hypervisor-based virtualisation. Network resources are offered utilising software-based OpenFlow switches based on Open vSwitch (OvS) [Open vSwitch] and network links that interconnect the OpenFlow software switches.

Two general-purpose servers are installed in each of the five GÉANT PoPs hosting the GÉANT OpenFlow Facility PoPs (for details of the GÉANT OpenFlow Facility general server specification, see 2.5Appendix B). Each server at a PoP (see Figure 2.1) is either:

- The host of a software-based OpenFlow switch (Open vSwitch) on top of a native Linux Debian distribution, or

[1] The Xen hypervisor provides computing resource virtualisation. It was chosen due to its compatibility with the OFELIA Control Framework (OCF), realising the management plane of the GÉANT OpenFlow Facility.

- The host of a Xen hypervisor for the instantiation of multiple VMs that can be allocated to user slices.

The data plane topology of the GÉANT OpenFlow Facility is configured as a full mesh graph, so that every OpenFlow switch has direct connectivity with all of the other OpenFlow switches. This is achieved by:

- A back-to-back connection of each OpenFlow switch to the local production MX router of GÉANT, through four 1 Gigabit Ethernet ports.
- Interconnecting the GÉANT OpenFlow Facility PoPs over the GÉANT backbone using pseudowires (L2MPLS VPNs) configured on the MX routers. This setup eliminates the need for VLAN switching on GÉANT gear and its associated limitations to the OpenFlow Facility's slicing capabilities.

The network slicing technique used by the GÉANT OpenFlow Facility dynamically allocates one or a range of VLAN IDs to each user slice. In this way, the network is partitioned per Open vSwitch interface (physical or logical) and user traffic is distinguished by the VLAN ID. VLANs can be involved in experimentation within the slice when a set of VLAN IDs is allocated to a slice. Thus the control logic of a slice is defined by the slice controller permitting the involvement of the reserved VLAN IDs to control logic decisions.

A simple example is that of an experimenter who can use his own range of VLAN IDs for routing purposes on top of the OpenFlow topology, forcing the OpenFlow switch to forward Ethernet frames based on non-standard algorithms, rather than having them handled by a legacy broadcast domain.

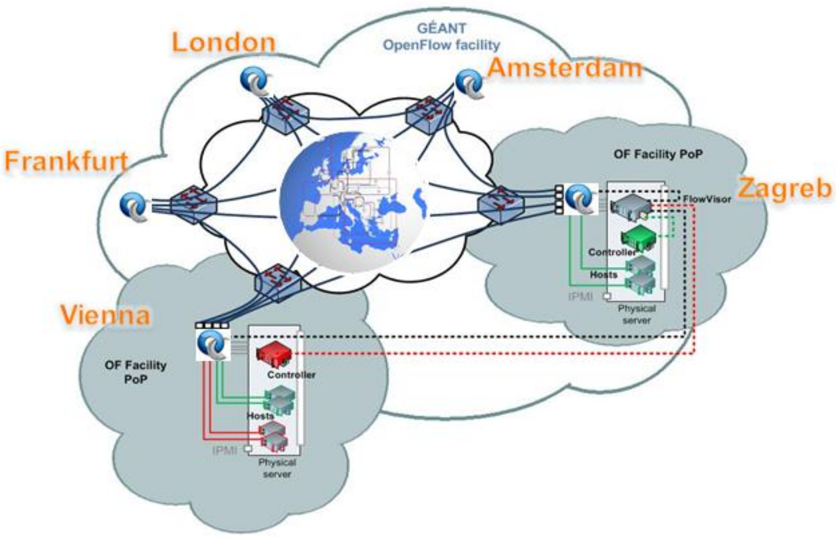

Figure 2.1: A generic view of the GÉANT OpenFlow Facility

An overview of the GÉANT OpenFlow Facility is shown in Figure 1. The GÉANT OpenFlow Facility is co-located in five GÉANT PoPs, in Vienna, Zagreb, London, Amsterdam and Frankfurt.

The facility's management and control plane elements/software are hosted in the Frankfurt PoP. This includes the GÉANT OpenFlow Control Framework (GOCF) and the FlowVisor software. The GÉANT OpenFlow Facility's control and management plane is shown in Figure 2.2.

End users can connect to the GÉANT OpenFlow Facility in order to transmit/receive traffic with three options[2]:

(a) Using Virtual Machine (VMs), hosted on servers offered by the facility.

(b) Installing their own equipment within the OpenFlow Facility PoPs.

(c) Connecting their remote labs to the OpenFlow Facility through an (emulated or native) Ethernet circuit provided by their local NREN or an ISP.

In option (a) above, the traffic that the VMs exchange with the facility has to comply with Ethernet standards. This means that the VMs can be used as traffic producers/consumers but they cannot be used to extend the OpenFlow facility. For such functionality, the users have to make use of options (b) or (c).

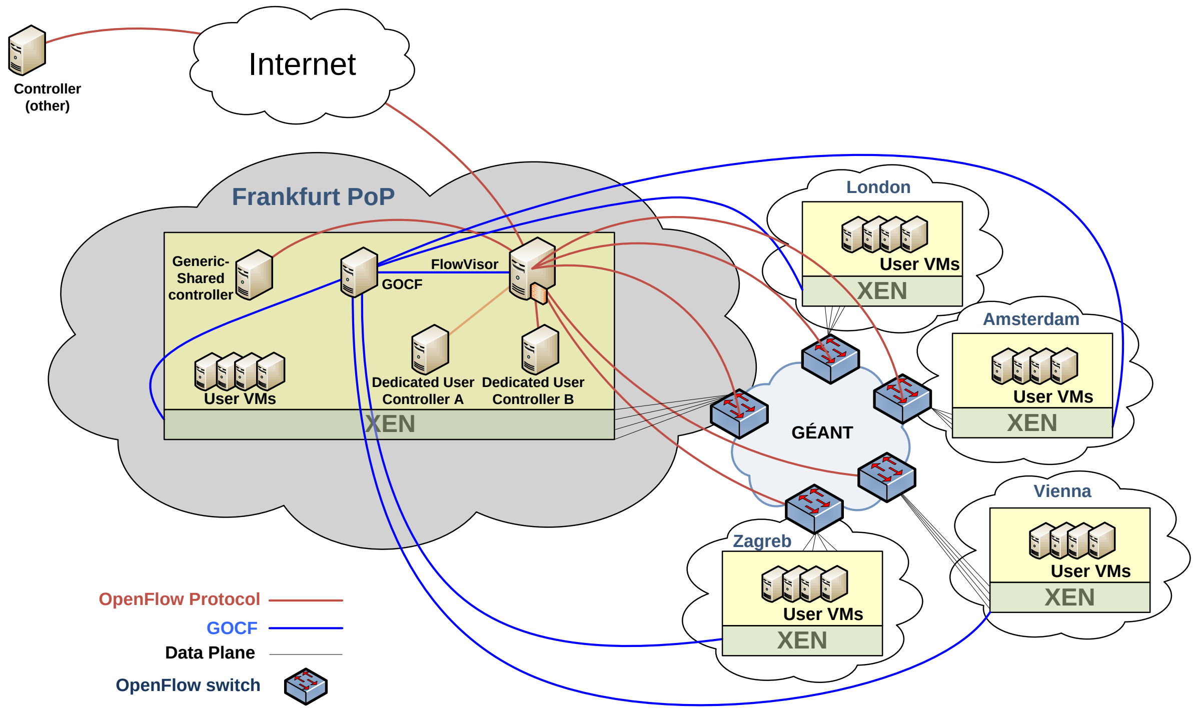

Figure 2.2: GÉANT OpenFlow Facility control and management plane

2.3 GÉANT OpenFlow Control Framework

The orchestration software deployed on top of the GÉANT OpenFlow Facility to implement management plane functionality is the OFELIA Control Framework (OCF) [OCFRepos]. OCF is developed within the OFELIA FP7 project [OFELIA], and it has been adapted to fulfil the GÉANT OpenFlow Facility requirements.

The main functionalities of GOCF are:

2 Not all of the options are available in the early phases of the facility's lifetime.

- Resource allocation and instantiation: The GOCF supports resource allocation, instantiation and de-allocation for any type of resource (e.g. an OpenFlow slice or a virtual machine out of a server farm).
- Experiment-based resource allocation: The resource allocation/de-allocation is performed per slice, with a slice being the smallest indivisible entity composed by the resources necessary to carry out an experiment. Slices are isolated from each other, even though they might share the same infrastructure substrate.
- Authentication and Authorisation (AA) and policy framework: GOCF supports the necessary mechanisms for authentication and authorisation (in several “scopes” or layers).
- Usability: End users/experimenters have access to comprehensive and easy-to-use user interface(s). In this sense, the main focus of the development is towards a web-based user interface.

An OpenFlow-aware network control plane is built on top of the facility’s OpenFlow proxy controllers so as to slice the network flowspace of the entire data plane topology. FlowVisor [FlowVisor] is the selected proxy controller supporting the facility’s network slicing and decoupling of forwarding (switches) and control (controller) elements. FlowVisor implements the network flowspace sharing and allocation per user controller logic. A single instance of a FlowVisor is deployed at the facility, controlling the OpenFlow switches and acting as a proxy between them and the experimenters’ OpenFlow controllers. A second FlowVisor may be added in the future for redundancy.

Each experimenter’s OpenFlow controller determines the control logic of a certain flowspace or slice. A controller can be either dedicated to a user/slice or shared (see Figure 2.2).

The controllers, FlowVisor and OpenFlow switches communicate utilising the public Internet over the management ports on the servers hosting them. Thus the control plane of the facility is implemented “out-of-band” at the IP layer, over the GÉANT network.

Authentication and Authorisation is handled by the Expedient tool [Expedient]. Expedient is a pluggable web user interface (UI) and clearinghouse originally developed by Stanford University, and adapted and extended to be a part of the OCF.

The Expedient clearinghouse functionality addresses user management, authentication and authorisation together with the projects’/experiments’ administration. It can be integrated with an LDAP server in order to support authentication against a users’ directory.

For the purposes of the GÉANT OpenFlow Facility, the GOCF Expedient instance has been integrated with Active Directory[3] , so that users in the GÉANT community can be approved and gain access to the facility using their credentials. Authorisation requires the definition of the appropriate user groups in the Active Directory. The GÉANT OpenFlow Facility distinguishes between an admin/management user and end users/researchers.

- Admin/Management user: This user is able to create projects and accept slice requests, and handles the general operation and management of the infrastructure.

[3] Provided by DANTE Ops.

- End User/Researcher: End users/researchers can submit through the UI projects, create slices and perform experiments on top of the OpenFlow facility. They have to belong to a project, which can be shared by multiple researchers.

2.4 Facility Behaviour

The current data plane setup of the facility and the underlying substrate result in an environment with 30 data plane links, five software switches controlled by the OpenFlow protocol and the specialised handling of VLAN tagged packets.

Due to the geographically dispersed nature of the OpenFlow enabled devices, flow establishment delays should be expected. Proactive or reactive flow recording and correlation with flow number constraints are monitored as part of operation of the facility. Long distances between OpenFlow-enabled devices and the OpenFlow controller in wide area environments introduce delay to the flow establishment process. Flows can be established using either proactively distributed or reactively distributed flow entries.

2.5 Operational Environment of the Facility

The facility is supported by several GÉANT services and therefore can be considered as a customer of the GÉANT network. A set of access points to the GÉANT services form the demarcation points between the facility and underlying GÉANT network.

On the layer above, facility slices administered by several end users/experimenters can operate concurrently. An administration interface between the facility's operator and the experimenters is defined. End users/experimenters are able to manipulate their own functional subset of the facility's resources taking advantage of the provided OpenFlow facility services.

A 3-tier architecture for operations (see Figure 2.3) is the result of using the OpenFlow facility on top of GÉANT to deliver TaaS. The lowest level (Tier-1) is the view of the GÉANT backbone production environment operators. On this level, operators are responsible for the physical infrastructure (bare-metal servers of the facility, physical connections towards the GÉANT backbone routers, management LANs at the PoPs) as well as GÉANT connectivity and other services supporting the facility (L2MPLS VPNs across the facility PoPs, IP connectivity across the facility's control plane elements (controllers), perimeter firewalling). On the second level (Tier-2), the OpenFlow facility operators are responsible for the operations of the services provided to the end users/experiments (such as monitoring, internal firewalling) including the required resources like server hypervisors, virtual machines, the OpenFlow proxy controller, control and managements capabilities of the OpenFlow slices. On the third level (Tier-3) the experimenter himself is operating a subset of OpenFlow resources, such as the capabilities of his allocated virtual machines and a set of network flows through his own OpenFlow controller, to fulfil the concepts of user-controlled slices and TaaS.

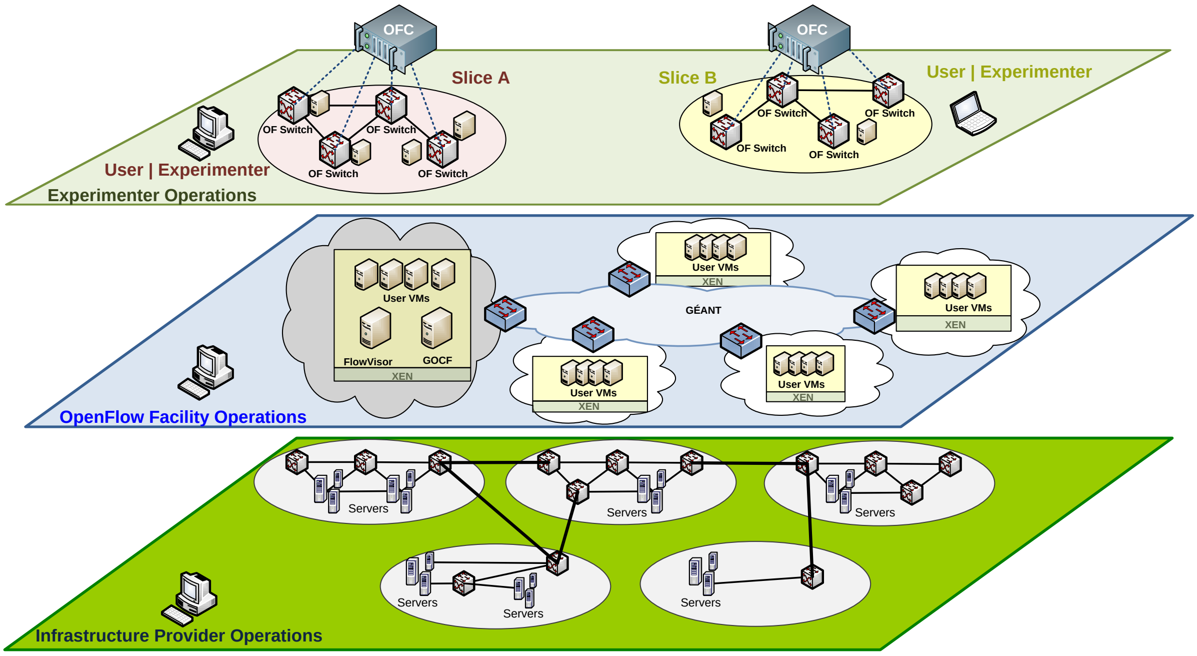

Figure 2.3: A three-tier architecture for operations

Based on this architecture, a list of services at different layers, operational tasks and corresponding responsibilities is provided in Table 2.1.

In order to further illustrate the concepts behind the 3-tier architecture, a closer look at part of the services delivered from Tier-1 to Tier-2 and associated demarcation points is shown in Figure 2.4. Built-in interfaces, labelled with “BltIn/1-eth0”, are used by the control and management plane of the OpenFlow facility. The service delivered to these interfaces by Tier-1 (the substrate) is IPv4/IPv6 Internet connectivity over the GÉANT backbone. The interfaces acting as OpenFlow data plane ports (highlighted with circles) participate in the facility’s data plane. For these, point-to-point L2 Ethernet connectivity is required, delivered by the L2 MPLS pseudowires over GÉANT depicted in Figure 2.4.

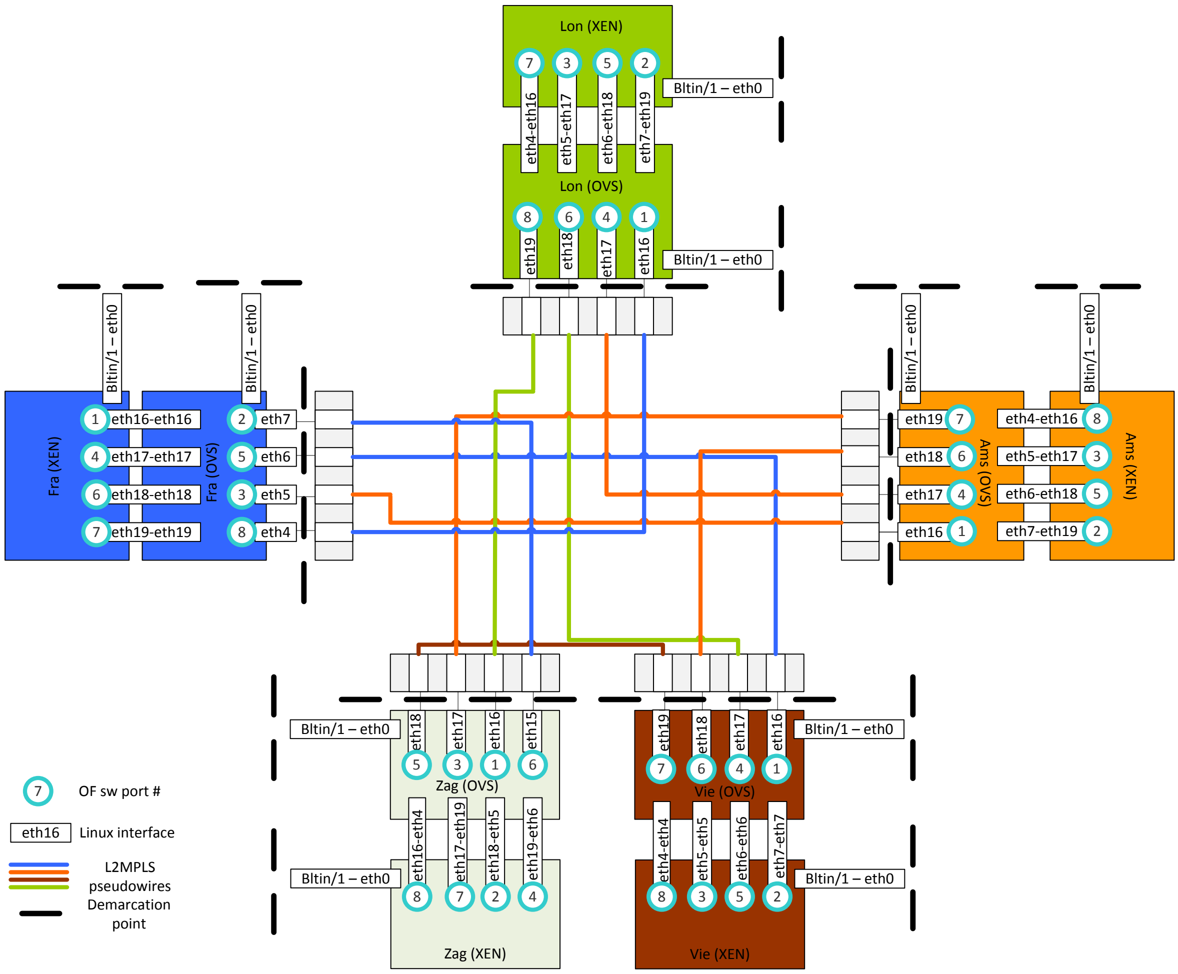

Figure 2.4: Demarcation points for the network connectivity services

Services running outside the scope of Tier-1 (the substrate), such as the GOCF, are closely related to the experimental scope of the facility and introduce non-standard operations for Tier-2.

Table 2.1 presents the operational tasks at all tiers and the corresponding assignments to the different operations entities. Generally speaking, production GÉANT services that are required by the OpenFlow facility are operated by the GÉANT NOC at Tier-1, while facility-specific components, exclusively developed for the facility operations, are operated by Subject Matter Experts (SMEs) from GN3plus Joint Research Activity 2 Technology Testing for Specific Service Applications Task 1 OpenFlow/SDN for Specialised Applications (JRA2 T1) at the moment of writing[4]. The Table also highlights slice operations at Tier-3 exercised by slice users.

Operational Tasks	Service Level/Provider	Operations Entity
Cabling	Substrate Provider	GÉANT NOC
Power Supply	Substrate Provider	GÉANT NOC

[4] In the long term, Tier-2 operations will be devolved to the operations entity defined for the facility.

Operational Tasks	Service Level/Provider	Operations Entity
Air Conditioning	Substrate Provider	GÉANT NOC
Physical Installation/Maintenance	Substrate Provider	GÉANT NOC
IPv4/IPv6 Network Connectivity (for the control & management plane)	Substrate Provider	GÉANT NOC
IPv4/IPv6 Address Space Allocation	Substrate Provider	GÉANT NOC
Domain Name System	Substrate Provider or OpenFlow Facility provider	GÉANT NOC
Perimeter Firewalling	Substrate Provider	GÉANT NOC
L2 data plane connectivity	Substrate Provider	GÉANT NOC
Operating System Administration (Xen & OvS servers)	OpenFlow Facility provider	JRA2 T1
Open vSwitch Administration	OpenFlow Facility provider	JRA2 T1
OpenFlow Proxy Controller (FlowVisor) Administration	OpenFlow Facility provider	JRA2 T1
Xen Virtual Machine Administration	OpenFlow Facility provider	JRA2 T1
GOCF Software	OpenFlow Facility provider	JRA2 T1
Experimenters Support	OpenFlow Facility provider	JRA2 T1
Experimenters Registration	OpenFlow Facility provider	JRA2 T1
Virtual Machine operations (create, delete, start, shutdown)	User/Experimenter	User/Experimenter
OpenFlow controller operations	User/Experimenter	User/Experimenter
Slice topology (create, edit, delete)	User/Experimenter	User/Experimenter

Table 2.1: Operational tasks across the Tiers and assignments to operations entities

As part of the future work on the facility and foreseen service offerings, a thorough documentation of the operational processes at all three Tiers, vertical communication across operational entities from the different Tiers (e.g. Tier-1 to Tier-2 interactions), ticketing, request fulfilment and troubleshooting are essential.

Appendix A Definition of Terms

Terminology	Description
OpenFlow	A standardised communication interface acting between the respective control and data-forwarding planes of an SDN architecture [ONFSDN]
OpenFlow Facility	A physical infrastructure of interconnected data forwarding and control plane elements adhering to and communicating in accordance with the OpenFlow specification
OpenFlow Slice	A functional subset of an OpenFlow Facility's resources allocated for exclusive use by a specific End User/Experimenter
OpenFlow Switch	A software or hardware implementation of an OpenFlow-specification-compliant switch
OpenFlow Controller	An OpenFlow-specification-compliant agent implementing control plane functionality for one or more OpenFlow Switches by managing their flow tables
OpenFlow Proxy Controller	A special-purpose controller, permitting the OpenFlow Facility slicing to different End Users by handling OpenFlow control messages from End User controllers to the OpenFlow Facility switches and vice versa, acting as a proxy and policy-enforcing entity
End Users/Experimenters	OpenFlow Facility users provided with access to the facility resources and allocated with a facility slice for use/experimentation. NRENs as well as academic and research institutions and research projects or groups/consortia connected to/through NRENs are potential End Users of the OpenFlow Facility
OpenFlow Facility substrate	The set of physical resources upon which OpenFlow Facility slices are implemented, including carrier network-provided backbone links, switches, servers

Appendix B GÉANT OpenFlow Facility: General Server Specification

Terminology	Description
Server	Ten (10)
Rack Units	≤ 2
Number of CPUs	≥ 2
Number of cores per CPU	≥ 4
CPU Cache Size	≥ 6 Mbyte
CPU Frequency	≥ 2.60 GHz
Memory Size	≥ 16 Gbyte
RAID Controller	RAID 1/5 with SAS HDDs & ≥ 4 HDDs support
HDDs	≥ 2 x SAS 146 GB
Network Interfaces	≥ 12 x Gigabit
Out-of-Band Server Management Card	1
Redundant Power Supply	1
Optical Drive	1
4-Hour Mission Critical Support	1 year

Appendix C GÉANT OpenFlow Facility PoP Addresses

PoP	Address
Frankfurt	InterXion Weismüllerstrasse 21-23 BUILDING FRA3 1st Floor Rm 2-1C Frankfurt 60314 Germany
Vienna	InerXtion Louis-Häfliger-Gasse 10, Vienna 1210 Austria
Amsterdam	VANCIS B.V. RmS145 1st Floor Science Park 121 Amsterdam 1098 XG Netherlands
London	Telecity 8-9 Harbour Exchange London E14 9GE England
Paris	InterXion 45 Ave Victor Hugo Batiment 260 Aubervilliers Paris 93534 France

References

[Expedient]	http://www.openflow.org/wk/index.php/Expedient
[EthOAMOF]	R. van der Pol, S. Boele, F. Dijkstra, "OpenFlow Demo, IEEE 802.1ag Ethernet OAM", Nov 12-18, 2011, SCInet Research Sandbox (SC11, Seattle, USA)
[FlowVisor]	https://github.com/OPENNETWORKINGLAB/flowvisor/wiki
[MPTCP]	R. van der Pol, S. Boele, F. Dijkstra, A. Barczyk, G. van Malenstein, J.H. Chen, J. Mambretti, "Multipathing with MPTCP and OpenFlow", Supercomputing 2012 SCInet Research Sandbox Proposal, June 2012
[OCFRepos]	OFELIA Control Framework GitHub repository http://fp7-ofelia.github.com/ocf/
[OFELIA]	OFELIA FP7 project http://www.fp7-ofelia.eu/
[ONFSDN]	Open Networking Foundation, "Software-Defined Networking: The New Norm for Networks", ONF White Paper, 13 April 2012 https://www.opennetworking.org/images/stories/downloads/white-papers/wp-sdn-newnorm.pdf
[Open vSwitch]	http:// http://openvswitch.org/
[Xen]	http://www.xen.org

Glossary

AA	Authentication and Authorisation
API	Application Programming Interface
AS	Autonomous System
GE	Gigabit Ethernet
GOCF	GÉANT OpenFlow Control Framework
ID	Identifier
IP	Internet Protocol
ISP	Internet Service Provider
JRA	Joint Research Activity
JRA2	Gn3plus JRA2 Technology Testing for Specific Service Applications
JRA2 T1	JRA2 Task 1 OpenFlow/SDN for Specialised Applications
L2	Layer 2
MPLS	Multi-Protocol Label Switching
NDDI	Network Development and Deployment Initiative
NOC	Network Operations Centre
NREN	National Research and Education Network
OCF	OFELIA Control Framework
OvS	Open vSwitch
OFELIA	OpenFlow in Europe: Linking Infrastructure and Applications
PoP	Point of Presence
REN	Research and Education Network
SDN	Software-Defined Networking
SME	Subject Matter Expert
TaaS	Testbeds as a Service
UI	User Interface
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPN	Virtual Private Network