28-04-2016

Deliverable D13.3
Proposed Network Architectures – White Paper

Deliverable D13.3

Contractual Date:	30-04-2015
Actual Date:	28-04-2015
Grant Agreement No.:	691567
Work Package/Activity:	13/JRA1
Task Item:	Task 3
Nature of Deliverable:	R (Report)
Dissemination Level:	PU
Lead Partner:	CARNet
Document Code:	GN4-1-16-37A1D
Authors:	Damir Regvart (CARNET), Tasos Karaliotas (GRNET), Daniel Arbel (IUCC), Sonja Filiposka (MARNET), Migiel de Vos (SURFNET), Yuri Demchenko (UvA), Kurt Baumann (SWITCH), Jan Radil (CESNET), Marcin Garstka (PSCN), Robert Pekal (PSCN), Reggie Cushing (UvA)

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 691567 (GN4-1).

Abstract
This deliverable presents the research of JRA T3 and defines a blueprint for a cross-domain and cross-technology information and services exchange platform for the NREN/GÉANT community. It proposes a Zero Touch NaaS concept to deliver an agile programmable network deployed using different traditional and innovative technologies such as Alien Wavelengths, spectrum sharing, the wire-free concept, etc.

Table of Contents

Table of Figures

Executive Summary

Traditional networks are gradually becoming obsolete, as they are unable to handle the increasing strain placed on their available resources by the various "heavy-weight" requirements of emerging applications. Network data flows in this highly dynamic environment, where the computing focus has shifted from local servers to the cloud, have vastly increased in volume, while at the same time being less tolerant to delays and jitter due to their time-sensitive nature.

As real-time stream processing of big data on high-performance computing resources in the cloud continues to grow, it is of fundamental importance that the next generation networks implemented by GÉANT and the NRENs are capable of agilely handling and adapting to this burden. This means that a future network architecture to support the emerging demands of the R&E community must rely on an agile backbone and aggregation links that are part of a highly automated, programmable network, capable of responding to individual user requirements by exhibiting application and cloud awareness. The future GÉANT network must maintain standards of efficiency and resiliency while providing a flexible and scalable service delivery infrastructure providing dynamic end-to-end connections on-demand that will offer a higher quality of experience (QoE) to its end users.

A high QoE can only be delivered if the underlying network is dynamic and flexible enough to be responsive and able to adapt to an environment in a rapidly changing environment in which complex services are constantly evolving. The network must not only provide automatic connections for devices, but even more importantly the automatic provisioning of network resources must be implemented on all layers based on application and service requirements. This will enable a zero-effort, highly efficient utilisation of network resources for end users, who will experience the agile network that provides connectivity to their needed services as Zero Touch (ZT). Using the Zero Touch paradigm, the underlying network will be able to adapt to the requirements of the requested service and control all flows, with the aim of providing a high level of QoE for the user no matter which device is used to connect to the network.

This convergence between the services offered and the delivery network infrastructure has many benefits for both end users and network administrators. The Zero Touch network-as-a service-concept offers a higher value service to the end users while simultaneously significantly lessening the burden on network administrators for provisioning and managing the network devices. The concept of automation offered by the Zero Touch paradigm also involves end users and developers as active participants in the control of their traffic flows. Enabling the end connection points to have an active role in traffic management ensures that the requested requirements and constraints are taken into consideration.

However, in order to fully exploit the Zero Touch paradigm, there must be agile responsiveness to end-user requirements along the whole end-to-end connection, which in most cases falls across multiple network domains, which requires seamless cooperation between the GÉANT network and the NRENs. This can be achieved by developing an orchestration layer that will practically collapse the network domain boundaries to provide user-controlled configuration and provisioning of a full end-to-end path in the multi-domain environment of the R&E community. In this way, a complete Zero Touch network as a service (NaaS) can be offered to the end user. Based on this model the user can set up and tear down multi-domain connections between any two (or more) points that can be reached in the merged GÉANT/NREN networks.

Given the above, user QoE could be further enhanced by establishing a tight, federation-based cooperation between GÉANT and the NRENs as network infrastructure providers and the service providers (i.e. commercial cloud providers, or specialized research oriented providers). It should be noted here that NRENs maintain a high level of cooperation, where commercial networks do not. Therefore, technologies/architectures requiring a high level of cooperation between networks (as in an inter-domain multipath environment) are unlikely to be developed by commercial network service providers, which places NRENs in a good position to launch this new networking paradigm.

With such a federation-based cooperation in place, a complete Zero Touch service can be offered as part of a service catalogue, and once requested by a user can be provided via an end-to-end connection established on-demand using the underlying ZT NaaS. This integration between the network and the services offered on top of it opens up possibilities for establishing a marketplace where the service providers can offer their services to the end users.

Joint Research Activity 1 – Task 3 has researched and defined a blueprint for the cross-domain and cross-technology information/services exchange platform for the NREN/GÉANT community. The task also aimed to test and demonstrate current, emerging and potential new technologies to provide a proof of concept for the implementation of innovative intelligent network services within GÉANT. The demo environment developed could be extended to create a testbed environment for emerging network technologies in the future.

The research presented in this white paper is a natural extension of the scope of connection-oriented services such as Bandwidth-on-Demand and Multi Domain VPN to an end-to-end service. The ZT NaaS concept presented can be employed to deliver an agile programmable network deployed using different traditional and innovative technologies such as Alien Wavelengths, spectrum sharing, the wire-free concept, etc. The key results from the research presented here point towards an innovative framework for delegating local networking resources to remote users in a controlled way.

The research especially focused on investigating the feasibility of cross-domain network solutions that can support various (complex) services benefiting from the open exchange philosophy. By bringing the NRENs, GÉANT and commercial service providers together, capitalising on the federated use of network resources, openness, and integration, a service architecture based on an application- and cloud-aware intelligent programmable network is proposed. The architecture directly supports the network requirements and demands defined within the user applications, in particular scientific workflows using big data, by including network topologies and network QoS requirements in the applications' description and design, in particular by extending the OASIS TOSCA standard accepted by the industry.

Open Cloud Exchange (OCX) was adopted as a use case on top of which the Zero Touch Network as a Service model was applied to demonstrate the benefits of the multi-domain Zero Touch paradigm. A demonstration presented at the Super Computing conference in Austin in November 2015 showed the value of the approach that has been adopted, which brings the benefits of the on-demand user-controlled agile network infrastructure closer to the public, exhibiting outstanding QoE and flexibility.

1 Introduction

The last decade has witnessed a great expansion in the field of data acquisition and processing, led by the proliferation of services offered to end-users. Traditional networks are gradually becoming obsolete, as they are unable to handle the increasing strain placed on their available resources by the various "heavy-weight" requirements of emerging new applications.

To investigate possible solutions to the needs identified in the area of future network design, the main goal of Joint Research Activity 1 – Task 3 was to research and define the blueprint for the cross-domain and cross-technology information and services exchange platform for the NREN/GÉANT community. The task also aimed to test and demonstrate current, emerging and potential new technologies to provide a proof of concept for the implementation of innovative intelligent network services within GÉANT.

The document first introduces the Zero Touch concept, based on a bottom-up approach starting from network device provisioning. The current implementations of Zero Touch provisioning are also investigated as well as the possibility of extending the concept to a multi-domain agile network that can provide self–servicing. The section on ZT OCX extends the Zero Touch paradigm to the OCX architecture. The different complex multi-cloud service provisioning scenarios using Zero Touch combined with OCX are discussed, further focusing on the OCX architecture and the components of the proposed automated network-as-a-service (NaaS) provisioning.

This is followed by a discussion of the integration of the NaaS provisioning with the development process for new cloud-based applications that can benefit from user-controlled end-to-end connection. As a further step towards the creation of composite cloud- and application-aware networks, the possibility of setting up a marketplace for publishing and offering services to end-users on top of the ZT NaaS model is also examined. Finally, the conclusion summarises the research findings. The technical details of the Super Computing demonstration are included in Appendix A.

2 Zero Touch Networking

In today's constantly evolving scenario, the network design and optimisation process must take into account the continuous addition of new services, especially those requiring specific QoS levels, such as minimum bandwidth and latency for real-time multimedia streaming. This moreover places a greater operational burden on network administrators as it makes the day-to-day network management and optimisation increasingly complex and time consuming [NMS]. A flexible and accessible Zero Touch network, therefore, should not only respond to user requirements seamlessly and automatically, but also be easily managed and implemented so as not to overburden network managers.

Substantial effort has been invested in the field of network management automation [NMA] where the Zero Touch concept emerged. The main goals of the proposed solutions are to satisfy clients' demands for efficient yet seamless networking. In order to achieve this, network engineers are looking at employing and developing solutions that will automate at least some daily administration tasks, such as collecting and analysing network-related data and automatically adjusting network configuration parameters to achieve the desired performance levels and offer the end user a stable networking experience. This concept of automation through integration of network planning, configuration, and optimisation into a single, mostly automated, process requiring minimal manual intervention is not new. The main objectives when implementing these automated solutions are operational and capital expenditure reductions by diminishing human involvement in different operational tasks, as well as the optimisation of network capacity, coverage, and service quality. Over time, these efforts have led to the Zero Touch concept where human intervention in the process of network provisioning, configuration or management is reduced to the bare minimum ("zero" human intervention). However, as discussed later on, this high level of automation of the network configuration process is not only beneficial during the network provisioning phase, but can also be used to give end users partial control of the network and allow them to assert their requirements for their own traffic flows. Thus the complete Zero Touch networking experience involves a highly automated network configuration and monitoring process that is overseen by both network administrators and end users.

The existing Zero Touch solutions are being promoted as versatile and flexible, and able to fulfil a plethora of specific and service provisioning needs. These solutions could be deployed in many NRENs to suit a variety of aims such as network optimisation, achieving a higher quality of experience (QoE), and better operational efficiencies.

This section explores the trends that could leverage the Zero Touch potential to enhance the services GÉANT and the NRENs offer to end clients (seamless wireless connections as well as more advanced services) and the possibilities for developing new services that can be efficiently provisioned using the Zero Touch approach.

2.1 Zero Touch Provisioning Concept

As discussed previously, at its core Zero Touch Provisioning (ZTP) is an automation solution that is designed to reduce errors and save time when network administrators need to bring new infrastructure online. In other words, it is a kind of plug-and-play behaviour for network devices, or, more precisely, the ability of Zero Touch-enabled products to facilitate remote device provisioning where even a non-technical person can simply take a unit out of the box, connect it to the network and power it up, without the need for them to do anything else. This is one example of the potentials of Zero Touch; however, the full range of features it enables extends well beyond remote network device provisioning.

The heart of the ZTP process is actually a reinvention of the automated server deployment procedure, which is a typical tool that has been used in IT ever since the introduction of the first Linux servers [ZTPNW]. Rather than using command-line interfaces (CLI) to configure systems one at a time, server administrators can use automation tools to roll out the operating system software, patches and packages on new servers automatically. Advanced scripting capabilities also allow administrators to tailor the boot configuration of these systems with profiles for specific applications. As all of this is particularly useful in data centres, but can also be reapplied in the network devices environment.

Network devices have traditionally been managed via the CLI. For example, switches are typically coupled with pre-loaded proprietary network operating systems. Network technicians use CLI or the manufacturers' own tools to provision the device, a process that can be broken down into the following basic groups of steps (see Figure 2.1):

1. **Offline**: A newly acquired network device comes with a pre-installed OS. When unpacked for the first time, the device is kept offline while the administrator checks the operating system version and performs all of the necessary updates: OS patches, bug fixes, or any new feature updates as necessary. When the updates are complete, the administrator connects directly to the devices and sets up the basic configuration using the CLI. The initial configuration is set up to establish basic network connectivity. This includes parameters such as administrator and user authentication information, the management IP address and default gateway, and basic network services (DHCP, NTP, etc). The processes for enabling the chosen L2 and L3 network protocols are also examples of the bootstrap process.
2. **Online**: Once the initial OS and configuration has been verified, the device can be installed into the environment (racked and cabled), where the configuration can be further customised (either locally via the console or using a remote access protocol). These final configurations are specific to the application and location within the network.

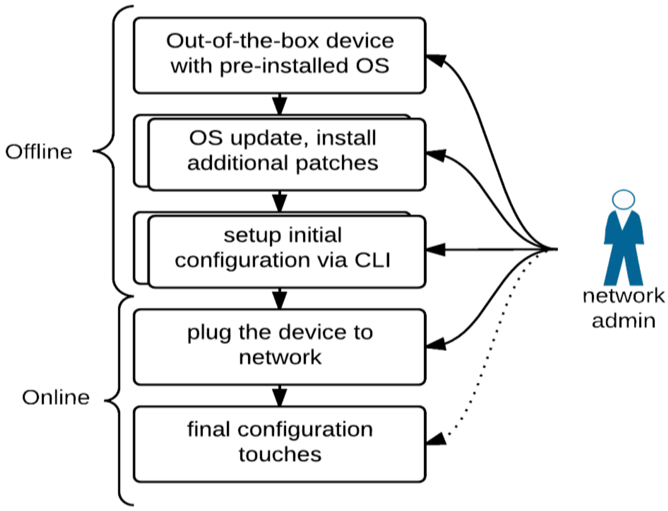

Figure 2.1: Illustration of the steps involved in the traditional network device provisioning process

A completely different scenario arises if, instead of the traditional provisioning process, Zero Touch Provisioning is used. With ZTP, all new devices connected to the ZTP environment are able to function without any manual CLI or GUI intervention. When using ZTP, first the device needs to be physically installed (rack and cable). Once the device is powered on, it will use standard network protocols to fetch everything it needs for provisioning. The process usually includes the following steps (with small variations depending on the vendor) (see Figure 2.2):

1. Power on the device and connect it to the network immediately.
2. The device will send a DHCP query and get an assigned IP address that will enable network connectivity and management.
3. The DHCP server will also provide the device with information on how to contact the ZT engine (one or several, usually TFTP, servers) in order to obtain the necessary operating system image and receive and activate the right configuration file based on the application profile.

Zero Touch Provisioning of a new device in a network occurs in two phases. The first is the basic ZTP process that includes general provisioning, such as downloading of the proper software and configuration files to the ZTP device. This is accomplished by using the additional options of the DHCP protocol so that the DHCP server not only responds with an IP address, but also sends the image name, config file name, and the location of these files to the device.

The device will then download the image and configuration files from the indicated location, compare them with the running versions, and, if these are newer, install the downloaded software and reboot. The download and installation completes the basic ZTP process. In the second phase, the device automatically installs a device-specific configuration based on its location. The instruction that triggers the device to download and install a device-specific script is contained within the standard configuration file from phase one. The script will effectively enable the device to download a second configuration file and merge the existing configuration with the newly downloaded one.

Note that there is no need to identify the MAC address/Serial number of the device; the device obtains a specific configuration based on where it plugs into the network. With ZTP, a new device has its port configuration and its IP address automatically provisioned based on the requirements of its location.

Additionally, when an inoperable device is swapped out, the replacement will automatically be configured correctly.

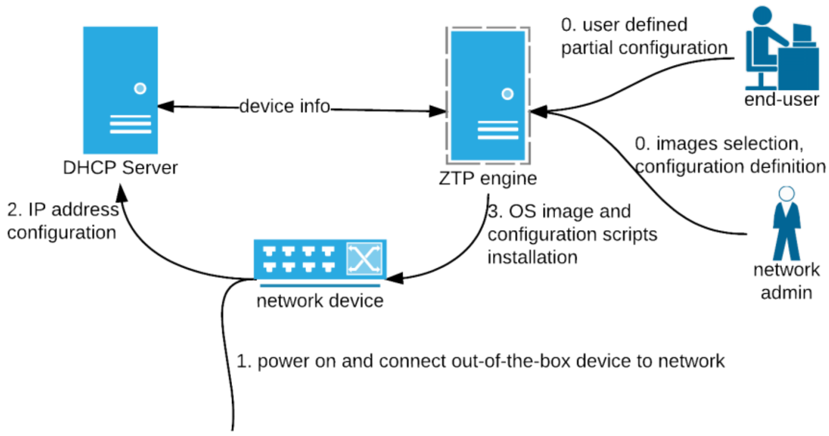

Figure 2.2: Illustration of the process of Zero Touch network device provisioning

Obviously, for the ZTP procedure to be successful, it is presumed that all of the necessary information already resides in the ZT engine. This means that the proper device operating system image and specific configurations need to be decided and defined at the heart of the ZT control (i.e. the ZT engine) prior to connecting the new device to the network. This part of the ZTP process still requires human intervention, however it is completely decoupled from the setting up process for the device itself.

In this model, once the network administrator sets up the IP address scheme via the DHCP server, and the OS and configuration files on the TFTP server, ZTP can be used to effectively roll out tens, hundreds, or thousands of networking devices, all fully customisable. In this way, ZTP can not only save time, money, and resources, but also dramatically reduce the NRENs' time-to service.

Figure 2.2 also shows that another option exists for defining the configuration of the provisioned device. This is in fact its most interesting aspect since it not only greatly influences the increased ZT capability for interconnectivity and process automation on a large scale, but is also responsible for enabling the desired flexibility of the network. Mainly this involves supplying the full or partial device configuration to the ZT engine via an intuitive web-based interface, or, even more effectively, using web services. This opens up the possibility of allowing at least partial end-user control of network devices. This feature has the potential to singlehandedly enable end-users to exercise control over the network infrastructure when they require access to a specific service. It would however be considerably more beneficial if this type of remote configuration were to be automatically triggered along with the service request, thus effectively hiding the unnecessary details from the end-user. This idea will be explored in more detail later on in this document.

It is likely therefore that all NRENs could benefit from implementing the ZTP approach, thus reducing the time for provisioning new equipment, but also for updating and adjusting their existing devices and the network services running on top of them. The greatest benefits are apparent for the network operations centre (NOC) during their daily routine business, as well as in the possibility of an unmanned network setup and configuration for different NREN clients, e.g. remote setup of new devices on the premises of schools or research facilities using Zero Touch.

2.2 Existing Zero Touch Provisioning Implementations

When discussing the Zero Touch concept, it should be stressed that all major networking vendors are moving in the direction of ZTP by providing solutions that enable newly connected network devices to be fully configured automatically in plug-and-play fashion. ZTP is now part of the services portfolio of a large number of network equipment vendors, such as Juniper [JZTP], Cisco [CZTP] and Ciena [ZTESP] to name just a few. In light of this trend, it is even more important to implement ZTP and network automation practices within NRENs.

2.2.1 Cisco

Different vendors nonetheless have slightly different implementations of the general approach described above. For example, the Cisco Networking Services Zero Touch feature [CZTP] provides a Zero Touch deployment solution where the device contacts a Cisco Networking Services configuration engine to retrieve its full configuration automatically. This capability is made possible through a single generic bootstrap common configuration file across all service providers' end customers subscribing to the services. Within the Cisco Networking Services framework, customers can create this generic bootstrap configuration without device-specific or network-specific information. The Cisco Dynamic Host Control Protocol (DHCP) Zero Touch feature enables a device to retrieve configuration files from the remote DHCP server during initial deployment with no end-user intervention.

2.2.2 Juniper

With ZTP in Juniper Networks [JCZTP], a new device has its port configuration and its IP address automatically provisioned based on the requirements of its location. When a switch is connected to the network and powered up with its factory default configuration, the ZTP process on the switch downloads the appropriate software as well as the configuration file for the device. The basic ZTP process provides a standard configuration file based on the type of device. The second phase is device-specific, taking advantage of Juniper Networks built-in automation capabilities. During the second phase, the device automatically installs a device-specific configuration based on its location. Should the device be removed and another be plugged at that location, the new device will have the same specific configuration installed.

2.2.3 Ciena

Similarly, Ciena offers Zero Touch Ethernet services provisioning [ZTESP] that provides automated service configuration, which compared to the traditional turn-up of new Ethernet services is measured in minutes instead of months. In this scenario, the NOC preconfigures the service profile and uploads it to a server. The customer's IT technician only needs to install the equipment in-rack, connect fibre and power up. The equipment will self-configure after which auto-triggered service activation testing will start. Once the equipment has been successfully tested, the client is online and the NOC monitors services in real-time. With the Zero Touch capability, upgrading from 1 GigE to 10 GigE can be achieved in 10 minutes, with self-configuration at the push of a button. With the real-time monitoring feature offered to customers, so that they can view their real-time service performance using a cloud-based portal, the customer churn is reduced by 2-3%, rapid-fire deployment revenues are accelerated and costs are lowered by as much as 75%.

2.2.4 Brocade

Using the Brocade VCS Fabric Technology [BRCD] a Zero Touch scale-out of the data centre can be achieved. The scale-out approach provides several benefits, such as network expansion over time, horizontal scaling of spine switches as the number of leaf switches increases, and the creation of resilient network fabrics, eliminating “single point of failure” and potential downtime.

This fabric technology provides automation through zero-touch provisioning, self-forming trunks, and Logical Chassis, enabling the addition, movement, and deletion of capacity automatically without configuration changes to the existing network. There is integrated resiliency by providing multipath at L1 to L3. Brocade VCS Fabric technology is designed from the ground up to facilitate and optimise scale-out architectures. Zero-touch provisioning enables simple rapid deployment. As one of the many Brocade fabric automation capabilities, the switches are preconfigured so that when a new switch is deployed only a power and network connection is required for the switch to become a member of the fabric. Inter-switch links are automatically formed between new members and all of the switches in the fabric. As a foundation to simplifying scale-out architecture, this method of installation eliminates the use of a manual process and reduces complexity.

Brocade claims that this Fabric technology enables 50% lower operational expense (OPEX) and doubled network utilisation.

2.2.5 Overture

ZTP in Overture products allows a technician to simply take a unit out of the box, install it, connect it to the network power it up and check various LEDs, without the need for manual configuration. This feature is currently available on the entire Ethernet over Copper product line and a growing number of other products [OVT].

With Zero Touch, the Central Office (CO) functions as a central management hub that automatically detects all subtended Customer Premise Equipment (CPEs) and sends the proper configuration and any needed software updates, without the need for manual intervention. The capabilities included with ZTP are:

- The CO device automatically discovers CPE devices connected to it.
- The CO device keeps a copy of all the configuration data for each CPE. When the CPE first connects to the CO, the CO overwrites all the configuration data on the CPE with the data stored on the CO.
- Because the CO device keeps a copy of all the configuration data for each CPE, the CPE can be provisioned even if it is not present.
- The CO device and subtending CPE are managed as single system, simplifying management and reducing IP consumption.

- The SNMP Agent on the CO device encompasses all the CPEs in its management domain and allocates interface index values (ifIndices) for all the CPE Interfaces.
- The CO device collects all the alarms on the CPE devices it manages, and raises and clears the alarms as if they were originally raised by the CO device itself.
- When the CPE first connects to the CO device, the CO ensures the CPE is running the correct software revision. If it is not, the CO pushes a new software image to the CPE. The service provider only needs to update the software on the CO, not the CPEs attached to the CO.
- Auto-bonding automatically creates the bonded group of TDM circuits or copper interfaces that deliver service.

With all of these features ZTP saves time, money, and resources by dramatically reducing a service provider’s time-to service and the number of truck rolls needed to establish and maintain the Carrier Ethernet services. By automating device configuration and management, Zero Touch also removes the risk of misconfigurations and human error during deployment and eliminates the need to maintain the latest software version on inventory and spares.

2.2.6 Allied Telesis

The Zero Touch Configurator (ZTC) from Allied Telesis [AT] is a framework especially designed for WAN networks that enables the network components to be remotely configured. ZTC consists of two parts: Client and Server. The ZTC Client is embedded in the CPE firmware. The ZTC Server is based on Java and manages the configuration data for the network devices. It basically supports all devices that have a ZTC Embedded Client implementation available. Its main purposes are device configuration, firmware release update and events logging.

The client periodically asks the ZTC Embedded Server for a newer configuration (configuration pull). When a new configuration for the device is available, the client downloads and executes it. All the connections made by the device are outbound. This implies that no listening service has to be started on the device. This minimises security issues related to accepting inbound connections. All users who access the configuration database must be authenticated. The ZTC Server manages profiles for users with different degrees of visibility and permissions.

The ZTC Server architecture is structured in functional modules that communicate via a standard interface (TFTP, HTTP and SOAP).

2.2.7 Aruba

Large wireless equipment vendors also offer a form of wireless ZTP. For example, Aruba [Aruba] offers a ZTP service where remote access points can receive their configurations automatically without any need to configure them from the controller itself. There are also controller-less solutions that are based on one dynamically elected access point that can automatically distribute the network configuration to the rest of the access points in the wireless network. Aruba Instant uses innovative Virtual Controller technology to deliver enterprise-grade WLAN capabilities such as robust security, performance, and scalability. It can be set up in minutes with minimal IT assistance and is managed centrally through Aruba AirWave.

The Aruba Instant WLAN is comprised of multiple 802.11n APs and Virtual Controller technology. It can be deployed as an overlay to an existing wired LAN thus eliminating the need for redesign or modifications of the wired infrastructure. All that is required is to configure one Aruba Instant AP over the air designed as a wizard-driven process. Since over-the-air provisioning is being used, there is no need to modify an IP address to configure Aruba Instant. After the AP is powered up and connected to the LAN, the Aruba Instant user interface login page can be automatically accessed through a browser to set up the configuration. To configure additional Aruba Instant APs, they just need to be connected and powered up. The first configured AP automatically becomes a primary Aruba Instant Virtual Controller and configures all the other APs. Aruba Instant is a fully distributed architecture. In the event of failure of a primary Virtual Controller, another Aruba Instant AP automatically takes on the role with no disruption. The primary Virtual Controller operates like any other Aruba Instant AP with full WLAN functionality. Multiple Aruba Instant APs can be configured per Layer 2 subnet, or Virtual Controller group, and enterprises can have as many subnets or Virtual Controller systems as needed on a campus or in a building.

The Instant Virtual Controller technology provides security, consistently high performance, scalability, and other enterprise-class network access services without requiring a dedicated controller. Utilising an adaptive, self-organising wireless grouping, the Virtual Controller technology supports multiple APs across wired LANs and over the air through the mesh, enabling easy scaling.

The idea of automatically uploading device information from one network device to another via a cloud-based service is patented by Aruba in the Zero Touch provisioning US patent no. 20140122674 A1. The patent describes the Zero Touch Provisioning process in a generalised environment for configuration of multiple interconnected devices in a cloud. To start the Zero Touch Provisioning process the configuration of the first device is uploaded to the first device. The rest of the process runs from the first device, which may need to download additional information about configuration of other devices and their interconnection, if this is not obtained during the first phase. The configurations or VM images for the subsequent devices are obtained or replicated from the first device Individual devices are configured either by uploading the necessary information to the devices themselves or are controlled by the first device. The general idea covered by the patent can be used in many different network scenarios, but cannot be extended to a multi-domain and multi-vendor environment. It also does not support an operational phase for the provisioned infrastructure.

Other features are available in addition to the provisioning to deliver enhanced Zero Touch capabilities, such as the Adaptive Radio Management (ARM) technology that automatically manages the WLAN's 2.4-GHz and 5-GHz radio bands to optimise Wi-Fi client performance and mitigate RF interference. This ensures that each AP uses the optimal channel and transmit power for its RF environment. The management technology additionally offers priority traffic handling, channel load balancing, band steering, airtime fairness and other quality-of-service (QoS) controls to ensure that the available Wi-Fi bandwidth is fairly distributed to all mobile devices on the WLAN.

2.2.8 Optical Zero Touch

Zero Touch networking and/or provisioning concepts can be extended to all layers of the ISO/OSI reference model. It should also be noted that different vendors use different names for the same concepts, for example touchless, elastic, programmable, or intelligent services or simply services. Although the focus is mostly on Ethernet and virtual local area networks (VLANs stitching), lower layers are also of interest to many Internet Service Providers, including in the academic and research community worldwide, especially for new applications such as accurate time transfer or ultra-stable frequency transfer. The stringent requirements for such new applications can be satisfied only with optical fibre and appropriate optical networking equipment.

Optical networking has evolved in the past few years from being rather static to being dynamic and flexible. This has been achieved thanks to new kinds of equipment, most notably reconfigurable optical add/drop multiplexers (ROADMs). The details of ROADMs functions are out of scope for this document, but directionless, colorless and contentionless ROADMs have to be supported to provide full Zero Touch functionalities. Software implementations of Wavelength Switched Optical Networks and Generalized Multi-Protocol Label Switching protocol suites (WSON/GMPLS) are also needed. Applying the concept of Internet Protocol over Dense Wavelength Division Multiplexed networks (IPoDWDM), Ethernet switches and routers may easily be integrated and provide continual integration of the packet and optical layers.

2.2.8.1 *Cisco Systems*

The Cisco Network Convergence System (NCS) [CNCS] can deliver such touchless capabilities on a very large scale and for ultra-long-haul (ULH) distances. Advanced ROADMs can respond to new bandwidth requirements, route traffic in case of network failures and also change network topologies dynamically in highly meshed networks. The total number of degrees of ROADM nodes is 32, with flex spectrum capabilities. Flex spectrum (or flexigrid) can be used to optimise the total system capacity and optical reach using different modulation formats. The NCS offers WSON and GMPLS capabilities to support dynamic service provisioning and restoration over touchless DWDM networks, and also cooperates with routers and switches to automate service provisioning, eliminate human error, and support quick failure recovery.

2.2.8.2 *ECI Telecom*

ECI launched its ELASTIC networks strategy [ECI] to ‘enable customers to seamlessly and cost-effectively adapt in tandem with the demands of the continuously evolving telecom market’. The key component of ELASTIC is SmartLight™, which includes four interconnected functional blocks: a transport hardware layer, a control layer, a Software Defined Networking (SDN) applications layer, and a security layer. Open application programming interfaces (APIs) are used for full functionality in multi-vendor environments, Network Function Virtualization (NFV) functionality on key platforms and software-upgradable clients.

2.2.8.3 *Infinera*

Infinera [INF] is well known for pioneering ideas with photonic integration. DWDM transport and Optical Transport Networks (OTN) switching technologies address new challenges in the Terabit Era. Infinera has introduced a New Simplified Model of Layer C and Layer T. Layer T is an intelligent Transport layer that combines Ethernet services, MPLS, OTN switching and even optical DWDM transmission ROADM functionalities. Layer C is a Cloud Services layer, with integration of services such as security, voice and others. This is made possible with the help of NFV, which enables services running on different hardware devices to be turned into software services that can run on a single hardware platform, and of SDN, which provides abstraction of the services and the functions of the Layer T and presents standardised application APIs so that the control can be offloaded to Layer C.

2.3 Bare Metal Zero Touch Approach (and SDN/NFV)

The existing Zero Touch provisioning and seamless configuration solutions offered by some of the biggest networking vendors briefly outlined above are indicative of the ever-increasing focus on Zero Touch networks. However, these only provide a first set of the benefits that future intelligent, Zero Touch, agile networks can offer. Nevertheless, it is important to keep in mind that the ultimate ZT goal for GÉANT and the NRENs can be seamlessly achieved only by employing ZTP using a vendor-agnostic hybrid approach. Another requirement is the hybrid approach that includes centralised automation and distributed control of network elements configuration. This is needed first of all because of the intrinsic nature of the networks in question (having centralised automation on the level of GÉANT can help develop seamless workflows for any collaborative efforts, while the controls must always stay distributed to each individual NREN independently). It also brings us closer to the ultimate holistic vision of a self-configurable user-centric dynamic network where control is not only held inside the NOC, but is partially delegated to the end user upon request.

In the complex NREN and GÉANT environment, where vendor lock-in is to be avoided at all costs, a vendor-agnostic integrated Zero Touch approach must be considered. A "one solution fits all" model in this respect can be achieved only by decoupling hardware and software [HSD]. This decoupling is part of the concept of network function virtualization [NFV], software defined networking [SDN], bare metal switches and white box switches [BMWBS], as well as of other similar solutions that contribute to a fast-paced and dynamic network environment.

For example, bare metal switches[1] are built using off-the-shelf hardware that may provide fewer features than proprietary chips, but enable a lower-cost and more flexible switching alternative. White box switches on the other hand present a more interesting option in that they are poised to be open enough to allow network vendors and users to choose their network operating system based on their organisation's needs. These switches do not come with a preloaded network operating system, therefore a third-party network operating system (such as Big Switch Light OS, Cumulus Linux OS or Pica8 PicOS) must be procured separately or built in-house. Several vendors support these ideas and have fast become the major players in the field of dynamic networks. Cumulus Networks [CN] created the first Linux operating system for network hardware and also partnered with Dell to offer its Linux network OS on Dell's networking hardware. Big Switch [BS] has made the move to bare-metal switches in combination with SDN.

The Open Network Install Environment [ONIE], which started as an Open Compute Project [OCP] as pioneered by Facebook, is the main initiative that defines an open "install environment" for bare metal networking, aimed at facilitating the installation (or removal) of any network operating system on any ONIE-based switch. Today, there are approximately 20 ONIE-based platforms in flexible 40G, 10G and 1G configurations available with PowerPC or x86 CPUs providing the key building blocks for a modern network. Administrators can thus select network hardware and software platforms independently, making the best choices for each scenario. Open Network Linux [ONL] is a Linux distribution for bare metal switches that uses ONIE to install Linux onto on-board flash memory. It is a part of the OCP and is a component in a growing collection of open source and commercial projects of its kind. ONL is in

[1] It should be noted that in the literature different definitions for bare metal and white box switches are used interchangeably and produce confusion in the networking community. For example, according to Forrester's taxonomy bare metal switches are switches from original design manufacturers with no operating system, while white box switches are made with off-the-shelf components, i.e. a commodity-based bare metal switch with preloaded network operating system.

essence similar to a hybrid PC BIOS and Grub/LILO/Sysimage and comprises a mini-Linux installation of around 4MB, with the option of a traditional control plane (e.g. Cimulus Linux or PicOS) or SDN control plane (e.g. Switch Light OS).

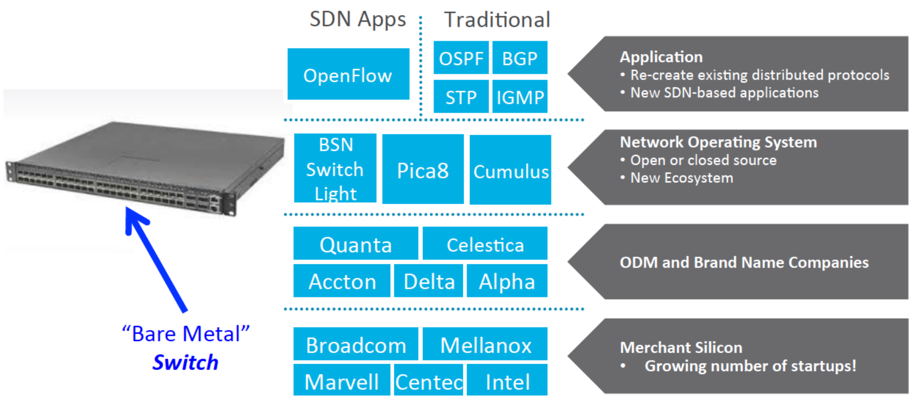

Figure 2.3: Open networking architecture using bare metal devices [BMWBS]

Using the bare metal paradigm, networking devices can be thought of as having a similar lifecycle to servers, from the initial loading of the OS, to provisioning the device for a specific application, and finally decommissioning the device (see Figure 2.3). All of this can be achieved by leveraging the experience and wealth of tools from Linux and the generic server industry, with part of the lifecycle involving ZTP for full deployment without human intervention. ZTP features can therefore be seen to be a natural extension of the bare metal model and to further increase the flexible, dynamic nature of the network. The ZTP process for these devices works in a similar way to that for the processes described earlier: when a device is physically connected to the network it is booted with a default configuration, after which it attempts to upgrade the OS software automatically and to auto-install a configuration file from the network using information provided by the DHCP server. Some of the existing ZTP applications for bare metal switching provided by PicOS [Pic8] are:

- **Using templates with Puppet/Chef to configure switches in a compute cluster:** Extending the model of cluster server provisioning, these tools create recipes or templates to quickly deploy additional racks of servers in the cluster. In addition to the server templates, the manager also adds a template for the Top of Rack (ToR) switches. As racks of servers are added to the cluster, both network and server platforms are configured and integrated into the existing cluster.
- **Using scripting to install Linux extensions:** A service provider writes a shell script that is downloaded to a switch with ZTP at boot time. The shell script downloads, installs and configures an OpenVPN RPM. Finally, it sends a message to the NOC with its logs and configuration information. After the switch reboots, managers can securely access the switch from the remote NOC for provisioning.

The scripting and templates approach not only alleviates the process of provisioning new devices in the network, but also enables a high level of network flexibility and agility when combined into the continuous network management cycle.

3 Enabling Zero Touch Operation and Management with Network-Wide Orchestration

The initial bootstrapping and configuration of network devices that enable zero-touch provisioning of new additions in the network is just the first step in the Zero Touch concept.

Beyond this is the automatic deployment of network management that allows device discovery, SNMP management, alarm management and event logging, network-wide upgrades and rollback, etc. When enabling automatic network management, the result is a centralised representation of the distributed network state that provides a single point of integration and network-wide visibility and analytics. All of the network management-related features can be added via partial customisation of the initial configuration files during the Zero Touch Provisioning stage, after which the newly added devices will be automatically added to the list of managed devices within the network management system (NMS) and all relevant management information will be pushed/pulled.

Leveraging the possibility for a centralised representation of the distributed network, the Zero Touch concept can be further expanded to include network operation and management by enabling network-wide orchestration. This opens up possibilities for deployment of advanced services across the entire network domain incorporating seamless mobility, load balancing and quality of service.

Zero touch orchestration can in fact be seen as an enabler for self-organising networks (SON) [C-SON]. In this way the inherent complexity of deploying and managing a feature-rich network (WAN in the case of many NRENs) can be simplified by automating the underlying configuration and monitoring tasks. To do this, Zero Touch orchestration must include the ability of deciding which is the best network-wide configuration, based on traffic flow and client requests from the network. The role of vendor-agnostic centralised SON (C-SON) is critical towards achieving that goal, as opposed to that of the distributed D-SON which is generally vendor-centric and proprietary. Consequently, the need to rationalise, coordinate, and manage the existing D-SON islands with C-SON is increasing. To allow a broader overview of network elements and a better coordination across a wide geographic area, C-SON typically sits at the same level as the Operation Support System (OSS). Ultimately, networks will likely be managed using a mix of centralised and distributed SON or a tactical combination of elements of each in a hybrid solution, known as hybrid SON.

This Zero Touch framework is expected to further improve the OSS/BSS (Operations Support System and Business Services System) responsibilities [MON] by providing rapid deployment, while enabling high availability and advanced features such as performance-based routing. In other words, it provides the means for achieving the ultimate QoE for all users within the community. By adopting the Zero Touch orchestration approach, the NRENs would turn their networks into agile dynamic environments that can instantly respond to traffic pattern changes and user demands. However, this would also have a negative impact in terms of the additional stress on operations, as constantly changing resources and services would overwhelm current operational models. A simple solution to this operations problem would be integration with the ZTP platform.

In the integrated environment, the ZTP platform notifies the service assurance platform when provisioning new services and resources. Once notified, operations can monitor and manage the change without the need for manual intervention. Automation ensures that the system stays up to date in real-time. This prevents the typical confusion caused by change management or provisioning in real-time. The technical aspects of an example integration (Monolith's Assure Now [MAN]) are shown in Figure 3.1 below. A unified approach to service assurance is provided by ensuring that the ZTP platform sends standard open API calls (SOAP/XML) to provision the device, KPIs, services, and topology necessary for monitoring the change.

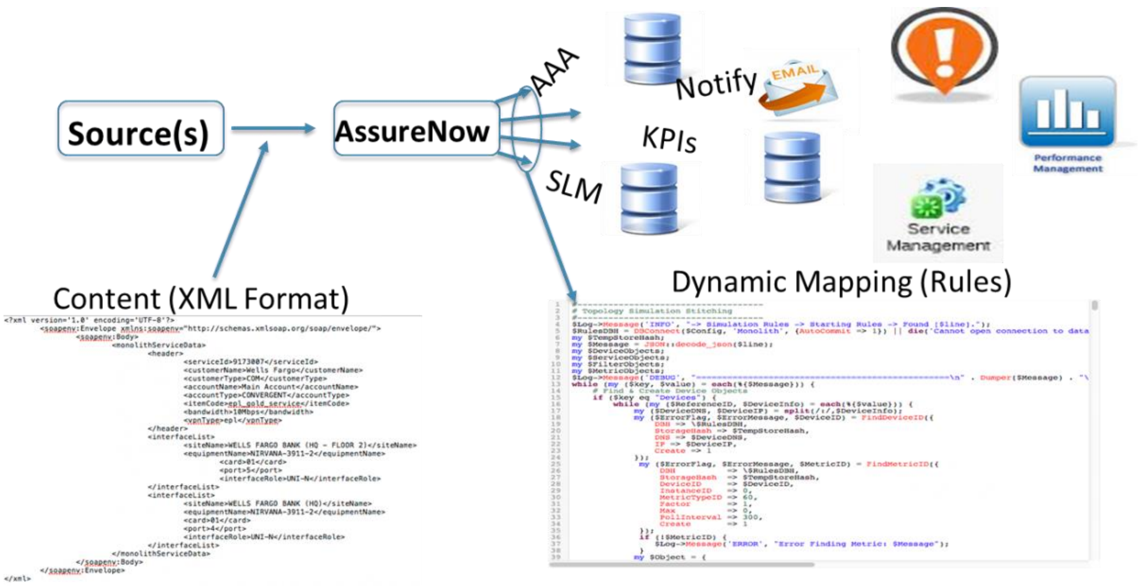

Figure 3.1: ZTP platform integration using AssureNow [MAN]

By investigating vendors for this type of use case an interesting solution was identified in Big Switch networks that leverage bare metal SDN fabric with an SDN software layer. These provide ready-to-deploy hyperscale network solutions [BCF], mainly targeted at next-generation data centres. Using Open whitebox switches and SDN controller technology, the solution offers application agility through automation, massive operational simplification through use of SDN and dramatic cost reduction thanks to HW/SW disaggregation. The Big Cloud Fabric provides L2 switching, L3 routing and, L4-7 service insertion and chaining with no single point of failure and support of headless mode operations. The SDN fabric architecture is based on the separation of the network's data and control planes, followed by centralisation of the control plane functionality. The network's policy plane, management plane and much of the control plane are externalised from the hardware device itself.

Using an SDN controller (cluster of VMs), the network state is centralised but hierarchically implemented with a high degree of programmability and automation. Since the configuration, automation and most of the troubleshooting are handled via the controller, the number of management consoles involved in provisioning new physical capacity or new logical apps is only one (the controller console). This results in great time savings, reduced error rates and simpler automation designs. If the controller console further exposes a web-based GUI, a traditional networking-style CLI and REST API, it can become a very powerful management tool. In addition to this, configuration in the CLI, GUI or REST API can be based on the concept of logical end users. Each end user can have administrative control over a logical L2/L3/policy design that connects the edge ports under his control. The controller then must provide the intelligence to translate the logical design into optimised entries in the forwarding tables of the network devices that are en route between the edge ports.

Another relevant example is the Glue Intelligent Engine with central policy-based orchestration which is part of the Cisco supported Gluware SDN-based software stack [CIWAN]. Gluware, designed by Glue Networks, is an SDN solution that simplifies the inherent complexity of deploying and managing a feature-rich wide area network (WAN) where the underlying configuration details, orchestration and ongoing monitoring tasks are handled automatically. This is achieved by using a proprietary, rules-based expert system that incorporates deep engineering knowledge from best-of-class approaches and industry best practices. It goes beyond typical existing network management solutions by adding an intelligent architecture layer on top of traditional deployment tools, and provides validated, production-proven, best-practice architectures for automated deployment and management using best-practice configuration templates together with extensive policy verification before applying configuration changes to the network. In addition, complex networking features can be safely enabled without the need for expert low-level networking knowledge, allowing clients to confidently manage feature-rich networks.

Using this type of intelligent orchestrating layer, the network can be turned from static and inflexible to agile and dynamic, able to adjust to the demands that are placed on it. Moreover, this type of solution can enable users to define WAN or other types of policies. However, in order to protect the network from incidental errors that can lead to hard-to-diagnose problems, any policy change needs to be verified by intelligent engine to validate that a requested policy would actually work given specific hardware and OS versions, network status, etc. The engine will orchestrate the changes throughout the network only after this check is completed.

3.1 Zero Touch Network as a Service

Network virtualisation has a great impact on the future development of the NREN networks and services they provide. With the expanding usage of bandwidth-hungry, low-latency cloud-based applications, the clients' key demands steadily turn towards on-demand, assured, user-centric network services. In order to answer this growing need, NRENs need to strive for changes towards virtual network management that will enable agile delivery of new, dynamic, on-demand services with rapid new service and technology operationalisation that give performance and security guarantees.

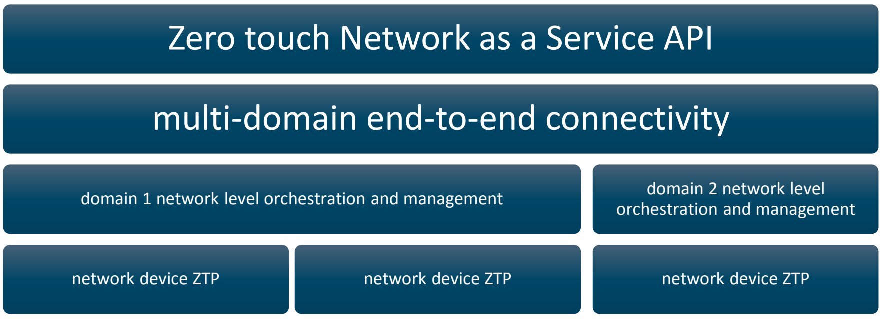

Figure 3.2: Zero touch network level orchestration and management model

Since the service providers can potentially be accessed from within a different NREN than the one where the client resides, the question arises of orchestrated delivery of automated connectivity as a service across multiple network domains and multiple internal multi-technology networks [ONC].

The goal is to empower the end user by offering a connection-on-demand service. The end user can self-choose the type of connectivity and define the higher-level service requirement from the network. Each connection should allow for service elasticity, i.e. the possible change of service characteristics during its lifetime. This could be considered, in its broadest sense, as a Zero Touch Network as a Service model (ZT NaaS), Figure 3.2.

The ultimate challenge for the Zero Touch approach in GÉANT is the implementation of a service/network configuration solution where service continuity is unaffected by service creation/modification/removal, which can be achieved by automating the network configuration changes initiated by network administrators or operational policy change [NetCloud]. It should be possible to set up services in a matter of minutes, and continuous delivery status and performance monitoring should be available to the end user. The SLAs provided to the end user should be monitored proactively to verify that there are consistent performance guarantees across all domains involved. The NREN (or GÉANT) who acts as a connectivity provider should expose on-demand self-service ordering with a fully orchestrated operations solution and integration between internal partners and external end points.

In order to provide automatic ordering, design, testing and activation of network services, this scenario requires an implementation of API integration together with a seamless integration between the higher-layer service access points and the controllers (possibly SDN based), and the network function virtualization (NFV)-based orchestration. The programmable NFV infrastructure is needed to provide rapid instantiation of new services. The solution also needs to ensure the quality of service and satisfaction of the requirements for end-to-end network connectivity set out in the agreed SLAs. The solution requires secure APIs for cross-domain exchange of performance information together with service assurance applications, as well as service integrity checks.

One possible direction towards the development of the multi-domain Zero Touch connectivity-as–a-service solution is the one taken by the Zero-time Orchestration, Operations and Management (ZOOM) project pushed by TMF and other companies [ZOOM], which enjoys wide support from the MEF community [MEF]. The project is inspired by the MEF’s “Third network vision” for an agile delivery of new, dynamic, on-demand services. The “Third Network vision” (see Figure 3.3), ensures new service and technology operationalisation and provides assured performance and security guarantees through the orchestrated delivery of automated NaaS across multiple network operators’ domains.

The MEF community strongly advocates the use of NFV to automate the process. Rapid instantiation of new services can be achieved using the programmable NFV infrastructure and the seamless integration of the SDN controllers can be orchestrated on the NFV level and beyond with fully orchestrated BSS/OSS systems. The fundamental idea behind this model is to design service-aware networks that reflect the highly dynamic virtual environment. The virtual environment makes it both possible and essential to radically rethink operations management, and open up previously unimaginable expectations for service personalisation, speed, flexibility, automation, and user centricity. In this environment, Zero Touch translates into to self-service and adaptive automation where services can be fully personalised by the end users themselves to fit their preferences and requirements.

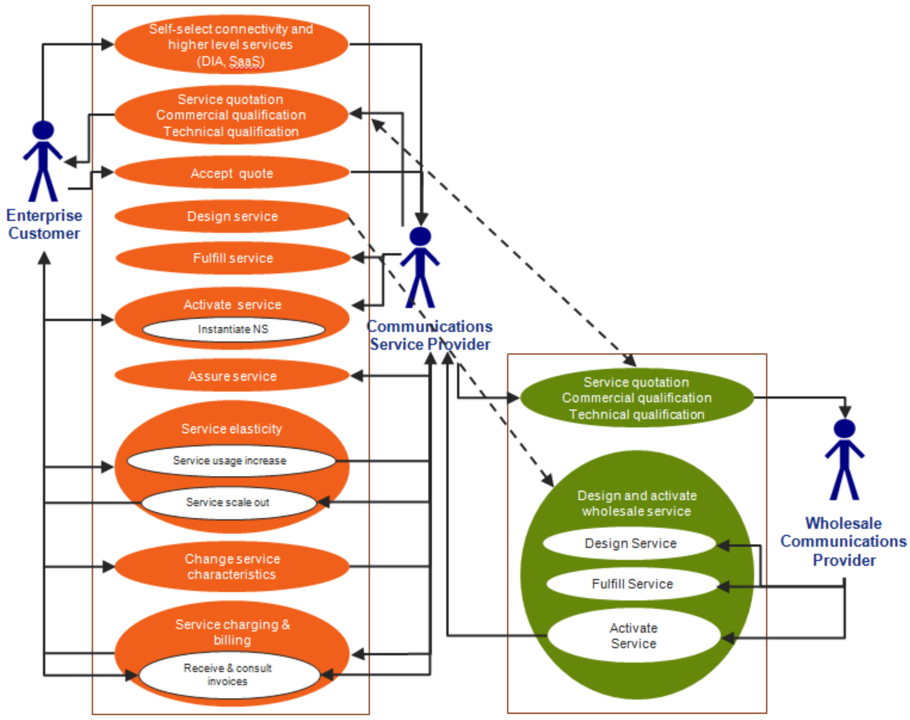

Figure 3.3: Agile, assured and orchestrated with NFV use case diagram [MEF]

To achieve this goal, the knowledge and approaches utilised for cloud services provisioning and SDN design must be applied to networks and services management. Emerging approaches such as NetOps, ServOps, and DevOps can accommodate the high rate of change and desired flexibility required of virtualised NREN networks. Given the current interest in network automation, it is imperative that the correct tools are chosen for the right tasks. As the understanding of network automation toolsets matures, so will the capabilities and styles of tools available. So far, network automation has been approached from a simplified perspective, due more to toolset limitations rather than to a lack of desire to automate more complex scenarios. Using NetOps, administrators can create templates and insert variable names or placeholders, then use a configuration build tool to replace those placeholders with actual variable names to complete configuration.

DevOps are drawing more attention due to the increased similarities in programmability and software-based configuration between network devices and servers. As they stand, DevOps tools such as Ansible and Puppet are intended to prepare servers for the software that the development team requires to be installed on specific systems. The tools themselves do not cover end-to-end logic and validation. Templating the configuration of more complex network scenarios only partly resolves this. Real automation requires the ability to store persistent objects containing JSON or XML (returned network device data), the ability to query a database (network resources), ability to communicate via a southbound API (NETCONF, REST, SNMP, CLI), and the ability to perform the equivalent of unit tests based on test-driven development methodologies.

While Ansible, Puppet, Chef, Salt and other similar tools are able to perform some simple network tasks, they need to be further developed to enable control to be delegated to a user/application to request a set of network resources. A user web portal can also be used for these purposes. The combined power of the enhanced "Ops" should provide the full set of desired features.

All the above must be taken into consideration when defining a ZT NaaS architecture to provide personalisation for the end user by dynamically assembling the necessary physical and/or virtual network components. This architecture must reside on open dynamic APIs that will provide transparent end-to-end management. The automated management needs to be event-driven and policy-based with integrated analytics from an expert-based system. Service requests must be executed rapidly with zero human intervention on a self-service basis with a high degree of end-user customisation.

By adopting all the best approaches in terms of network elements, network domain and multi-networks integration, the future ZT NaaS network architecture stands to fulfil the requirements of future services and end users. By building an agile, dynamic, programmable network that is highly automated and self-configurable, GÉANT and the NRENs will ensure that they are in a position to fully take on the challenges posed by virtualised environments and offer their users the quality of experience that they themselves select.

4 Zero Touch Provisioning Cloud Infrastructure

4.1 Motivation for ZTP Cloud Infrastructure

The need for Zero Touch Provisioning is not limited to network provisioning, whether at the level of single devices, the network, or a complex multi-domain network infrastructure, but should be extended to the services the network offers. This is especially true in the case of cloud services. To ensure that the user can use the full potential of the cloud-network cooperation, the notion of Zero Touch self-service must be applied to the complete service lifecycle that is empowered using a cloud service delivery infrastructure (CSDI).

The surge in cloud services imposes additional requirements on network operators and service and product providers, who require innovative network solutions to set up reliable and highly performing cloud-ready networks to tackle the increasing demands of this changing landscape efficiently. However, network configuration and installation require highly skilled personnel adept at configuring a variety of network elements. In addition, the operational costs involved in provisioning and managing large multi-vendor networks covering multiple technologies have been increasing in recent years. Coupled with increasing scarcity of human resources and increasing costs of real estate, this is a “perfect storm” for service providers that is leading to renewed interest in solutions that can unify network management and provisioning across multiple domains.

In response to this growing need, GÉANT and the NRENs need to strive for changes in network management that will enable agile delivery of new, dynamic, on-demand services, with rapid new service and technology operationalisation that provides performance and security guarantees. This also involves the need to provide orchestrated delivery of automated connectivity as a service across multiple network domains and multiple internal multi-technology networks [RANC].

Modern research is increasingly data intensive, relying on both specialist scientific instruments and cloud-based facilities and services for data storage and processing. The inherently distributed yet collaborative nature of research is pushing the development of cloud-based application infrastructures that include services and resources from multiple cloud service providers [ICCE]. This ongoing trend has placed great stress on the network infrastructure, which needs to provide high-performing connectivity between the researchers as clients and the multitude of cloud service providers that offer the needed set of resources. In these scenarios, the quality of the cloud services as perceived by the end user is crucial for the continuous usage and operation of cloud-based applications.

It is therefore very important that the underlying network infrastructure can provide an elastic and dynamic cloud-based experience for the end users working with big data sets who have high computing demands. This is even more complex to achieve when combined with the demand for a dynamically configured network infrastructure provisioned on-demand used to offer these services. The majority of the resources that are currently available and in use are cloud based and supported by corresponding automated cloud deployment tools. However, these tools do not cover external interconnections to other cloud networks.

To address this issue, the GÉANT project proposed the use of the Open Cloud eXchange [OCX] as a component and central hub of the GÉANT-based cloud-aware backbone infrastructure. A marketplace model is used to provide a single point access and delivery of services offered by multiple cloud providers to NREN users and the European research community.

A number of projects, such as GEYSERS [GEYSERS], CYCLONE [CYCLONE], and GN4-1 [GN4-1] have been addressing multi-cloud and inter-cloud network infrastructure delivery. Among the aims of these projects are the development of tools and solutions to enable connectivity between cloud providers and NRENs or campus networks over the GÉANT network infrastructure. Examples include SlipStream [Slip] for multi-cloud services deployment, OpenNaaS [ONaaS] for inter-cloud network provisioning, and Open Cloud eXchange [OCX] for smooth multi-provider cloud services delivery over the GÉANT and NREN infrastructure. However, these solutions and components are not all integrated into a single infrastructure service provisioning system or platform.

This section explores the possibility of leveraging the Zero Touch Provisioning concept and the ZT NaaS model discussed previously to implement a flexible, virtualised and software-defined OCX. A specific use case scenario, ZT OCX, is used below to investigate whether the ZTP concept could be used as a common framework to build an integrated infrastructure services provisioning environment for complex (multi-)cloud based applications.

4.2 ZTP Cloud Infrastructure Components

The results of a recent study on cloud computing performance [HPC] have shown that networking performance has a significant impact on the quality of cloud services, and in many cases data communications cause bottlenecks that affect the ability of clouds to support high-performance applications. Therefore, networks with QoS capabilities are an indispensable element of high-performance cloud computing. The significant role that networking plays in cloud computing calls for a holistic vision involving both computing and networking resources. Such a vision requires the underlying network infrastructure to be opened up and exposed to upper layer applications, thus enabling combined management, control, and optimisation of computing and networking resources for cloud service provisioning.

This leads to a convergence of networking and cloud computing systems to form a composite network-cloud service provisioning system. Due to the complexity of networking technologies and protocols, exposure of network functionalities in a cloud environment is only feasible with appropriate abstraction and virtualisation of networking resources.

Implementing and operating ZT OCX services in a distributed multi-provider and multi-domain environment will require dedicated infrastructure that supports distributed control over large-scale services. Additionally, resource deployment will be needed, in many cases requiring complex workflow execution that defines both the sequence of resource deployment, activation and adjustment to local environments (depending on the current state of the provider resources), and resource interaction.

When developing a consistent model for ZT-enabled complex multi-cloud/inter-cloud applications over the GÉANT and NREN network infrastructure, the cloud experience gained in deploying large-scale services should be reused, and extended with the ZT-enabled network services control as shown in Figure 4.1. Increasing proliferation of SDN and NFV at all levels of the network infrastructure will facilitate the use of cloud automation tools.

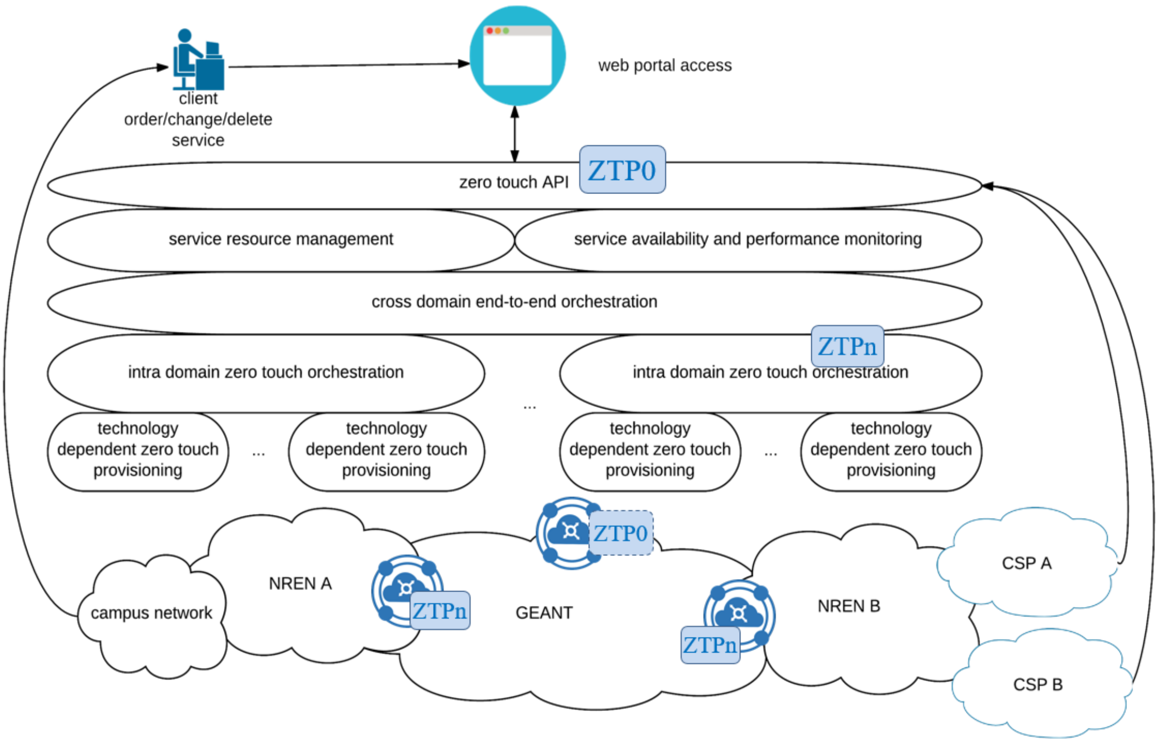

Figure 4.1: Zero touch Network as a Service model (OCX use case scenario)

Network connectivity provisioning in a multi-cloud environment will require dedicated infrastructure components, including intra-domain and inter-domain ZTP configuration servers that will enable automated service provisioning over existing multi-domain/multi-NREN network infrastructure. These servers can be co-located or integrated with the OCX [OCXIC] being developed both by the GN4-1 project and as a component of the Intercloud Architecture Framework [ArchFW] and Intercloud Federation Framework [FACHIE]. OCX's role is to provide connectivity and a network information exchange point to deliver cloud services from cloud data centres to user locations in NRENs and campuses (see Figure 4.1). In this case the OCX Marketplace will simplify API exchange between user applications and the multiple CSPs and network providers involved.

4.3 Open Cloud Exchange (OCX)

4.3.1 Open Cloud Exchange Architecture (and Design Principles)

The grid has always been seen as the traditional computing environment for e-Science applications [CCS]. Although several tools and frameworks have been developed for grid platforms to manage and handle big data, these tools are not widely adopted by scientists owing to the complexity of their installation and configuration. Lately, many big data scientists have been turning to cloud computing, attracted by its benefits [ICG], especially the uniform interface to a dynamically scalable underlying resource provided by virtualisation, which conceals the heterogeneity of physical resources and geographic location, as well as faults.

However, the cloud computing environment also suffers from a number of problems [GCC], including direct support being limited to single user or single organisation access, and low performance of computing resource and data integration using network transfers from outside the cloud. Cloud-aware networks are increasingly needed if the full benefits of collaborative big data e-Science are to be exploited. These power users demand increased flexibility of services at the network level as well as enhanced capabilities and stability. This means use of an ordinary Internet connection to access a single or composite cloud service is unacceptable where high requirements are included in the cloud service SLAs.

One current strategy to deal with this issue is to create direct dedicated connections from the power users to the cloud service data centres. However, in this case the Cloud Service Providers (CSPs) only offer a dedicated port at one of their data centres and a dedicated link must be provided by the customer. An example is the AmazonWeb Services (AWS) Direct Connect service that can provide dedicated port speed from 50Mbps to 10 Gbps [ADC], while Microsoft Azure offers a similar service called ExpressRoute [MAER]. In both cases dedicated connectivity can be provided only at certain specific CSP data centres and it is not available everywhere across Europe. Furthermore, dedicated links provisioning remains the responsibility of the customer side and the links are dedicated to only one CSP. GN3plus therefore proposed an Open Cloud Exchange (OCX) [NACS] solution to enable an on-demand direct link to be established from any power user to any CSP(s) by leveraging the existing GÉANT and NREN network. GÉANT's implementation of the OCX concept (gOCX) has been tested in several demo environments [gOCX] and has been proven to provide high-performing network connections on top of which composite multi- and inter-cloud services involving multiple cooperating user groups can be run, including for distributed real-time video streaming applications.

gOCX brings together cloud service users from the R&E community with different CSPs by establishing direct L0-L2 on-demand connections between the users and the providers. Designed with power cloud users in mind, the main goal of the gOCX architecture is to provide a dedicated infrastructure for efficient, fast, reliable and cost-effective connectivity–as-a-service, to facilitate inter-cloud computing federations.

gOCX is defined as:

1. a dedicated network infrastructure, implemented on top of the GÉANT/NRENs infrastructure, used exclusively for the provisioning of cloud services, providing performance and quality that are difficult to achieve with standard Internet access;

2. a marketplace that offers:
   a. a service directory, in which CSPs publish their services and users can discover and subscribe to services,
   b. “connectivity as a service” via automatic on-demand link provisioning,
   c. an SLA Repository and Clearing house; and
3. a Trusted Third Party that can facilitate dynamic federation service agreements and trust establishment.

The OCX instances are interconnected via backbone links engineered and dimensioned in such a way that certain performance metrics can be guaranteed. Downstream, OCX instances connect users with CSPs wishing to lease and offer cloud services. Finally, upon a user’s request, OCX provides connectivity between any two or more OCX access points in a secure and isolated manner. OCX access ports can multiplex various services on one port using VLANs for logical traffic separation.

The main benefits provided by the OCX infrastructure have been demonstrated in a demo scenario [gOCX]. These include the use of dedicated links towards the CSPs that provide substantially improved data transfer performances between the clients and the CSPs, while completely bypassing the best-effort Internet. A major argument in favor of gOCX is that all previously established L2 connections can be reused for future cloud service delivery so that after initial setup clients can use the cloud service transparently. Furthermore, since the connections to the cloud providers take place on L2 (or lower), the available cloud services can easily be expanded with more advanced features (e.g. extending the customer network into the cloud). The use of multi-domain L2 services additionally ensures traffic isolation.

The connectivity services offered to power users by gOCX should provide the same elasticity and dynamicity that are the hallmarks of cloud services. To this end, a networking setup standardisation should be imposed for all gOCX participating parties with the aim of providing an almost automated provisioning procedure via a well-organized web portal, or standardized API, that would make instant connectivity between the CSP and the user available over gOCX.

The gOCX service offers virtual, isolated and secure extensions to the direct connections to users’ networks and their requested cloud services, providing a connection between service users and providers over the GÉANT backbone. The demo scenarios have highlighted the benefits of direct connectivity, but also shown that the task of setting up all the direct connection links needed to interconnect the parties involved is time consuming and fragmented [NACS]. Additionally, when setting up such connections manually there is a high chance of human error, delays due to poor synchronisation and misunderstandings. This means that the true benefits of the gOCX architecture will be available, usable and transparent to end users only if the connection setup process is performed automatically, preferably by virtualising all OCX instances and using a web service-based platform that will enable Zero Touch Provisioning

Considering the complexity of provisioning the Cloud Services Delivery Infrastructure (CSDI) the following challenges need to be addressed to make the OCX concept a viable solution:

- The process of configuration of the direct link across multiple networks that are under different administrative domains needs to be automated providing on-demand link establishment in a matter of minutes.

- End users should be provided with an easy-to-use light web interface that will enable them to choose from the service offers and order a connection to one or more CSPs with a given SLA.
- The demand for the direct link establishment within the higher-level service that is being used should be seamlessly integrated in the Network as a Service (NaaS) scenario.

The requirement for minimum configuration and lots of automation is particularly critical. To address this, Zero Touch Provisioning could provide a solution for the easy, on-demand, fast and flexible integration of direct connections to CSPs with high-level composite services.

4.3.2 OCX Requirements

In the use case scenario examined, the ZT NaaS offers the client on-demand connectivity-as-a-service by allowing the end user to choose the type of connectivity based on the required higher-level services, and providing the ability for elastic connection where the connection's characteristics can be modified during its lifetime.

The ultimate challenge for the Zero Touch approach is the implementation of a service and network configuration solution that provides service continuity that is not adversely impacted by service creation, modification and removal through automating configuration changes.

Figure 4.1 illustrates the Zero Touch NaaS model, where end users typically demand a seamless experience for all connectivity towards one or more cloud service providers (CSPs), including both those within their domain and those that can be reached across other domains. Service setup should be possible in a matter of minutes, and continuous delivery status and performance monitoring should be available to the client. Using the performance monitoring data provided by the underlying network management system the client will be able to verify that the minimum SLA levels are achieved and that consistent performance guarantees are enforced across all domains involved.

The connectivity provider also must expose on-demand self-service ordering for all available CSPs using a fully orchestrated operations solution and integration between internal partners and the CSPs.

In order to provide automatic ordering, design, testing and activation of network services, this scenario requires implementation of API integration as well as seamless integration between the higher layer service access points. The solution also needs to provide assurance for the QoS and requirements for the end-to-end network connectivity set out in the agreed SLAs. APIs should additionally be secured for the cross-domain exchange of performance information, service assurance applications and service integrity checks.

4.3.3 OCX Functional Components and Design

The proposed gOCX architecture has been proven to successfully provide the required network performances for high-performing cloud computing. In the next stage, the best implementation of the gOCX instances to enable the needed network-cloud composite system should be determined. At the same time, the system should rely on the Zero Touch NaaS model as discussed to provide ease of use for the end user.

In the gOCX context, a Zero Touch NaaS should provide:

- From the client perspective: on-demand fast provisioning of cloud services that can be requested via a web interface or well-defined integrated API;
- From the CSP perspective: automatic provisioning of requested services together with service portfolio offerings;
- From the GÉANT network management perspective: automatic yet dynamic setup and configuration of direct (short-lived) links, which will ultimately be combined to provide an end-to-end direct path with the requested QoS parameters.

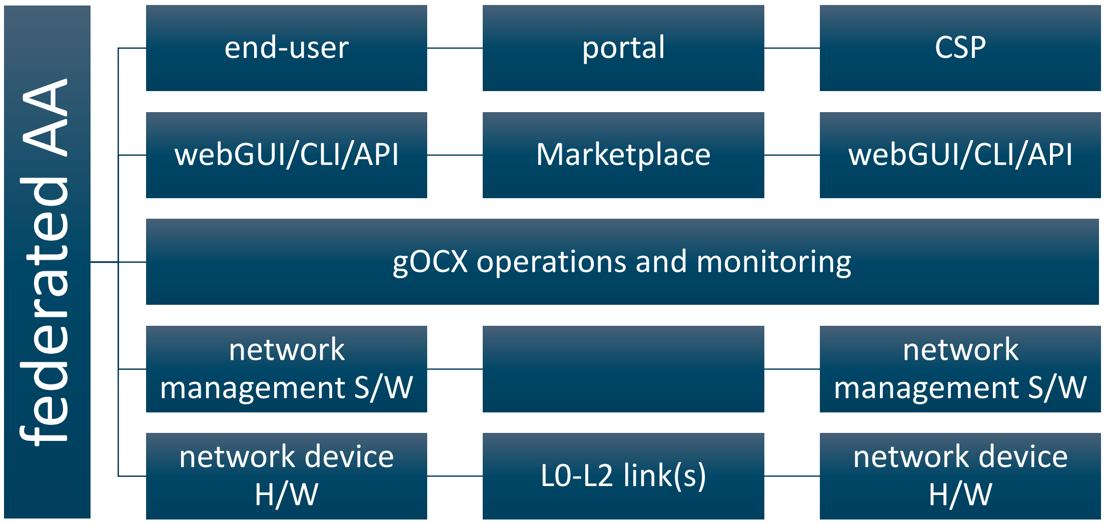

Figure 4.2: gOCX building blocks

Several gOCX building blocks were identified for the implementation of these requirements at the network level:

1. **Network H/W** – the device that accepts the physical link. It may be an optical cross-connect, WDM device, Ethernet switch, router or a group of such devices. Because of OCX's hierarchical design, the gOCXs deployed will probably comprise a group of dedicated devices to provide high performance for multiple requests at the GÉANT network level, and a smaller part of a larger device (e.g. one physical port) at the NREN level. This block will be provisioned using the technology-dependent ZTP at the device level.
2. **Network Management** – The S/W that controls the H/W operation. Its task is to poll for status and performance data, intercept traps for critical status changes and send configuration commands. This block is controlled at the level of the domain where it resides. Intra-domain Zero Touch orchestration procedures should be used for this purpose.
3. **Marketplace** – this is a S/W block that serves as the front end of the gOCX. CSPs use it to offer services, and clients can browse, search and demand available services. Each demand is accompanied by automatically defined or user-specific network QoS parameters that should provide the necessary characteristics of the direct link that will be established. The setup of the direct link on-demand functionality is implemented at the level of cross-domain end-to-end orchestration according to the ZTP NaaS model.

It is anticipated that additional blocks, such as federated authentication and authorisation, status, and accounting will be needed for full gOCX implementation.

4.3.4 Enabling Technologies

4.3.4.1 *Connection On-demand Setup*

To enable Zero Touch self-service high-performance connections to CSPs for end users, OCX needs to orchestrate the process of setting up the user-CSP(s) connection(s) across multiple domains. This includes finding a path between the user and the CSP over one or more OCX instances and setting up BoD services on the full set of multi-domain segments along this path. Considering the two-layered hierarchical topology of OCX instances (at the GÉANT and NREN levels) the path-finding feature must also allow for load balancing over multiple connections to a same CSP, as one CSP will ideally have multiple points of connection to one or more OCX instances.

Another important aspect is that the path setup phase should be executed in multiple steps:

1. Verify whether an acceptable (feasible) path exists under the given constraints (these may include service type, CSP, set of QoS parameters, etc.).
2. If a path exists, this should be reserved for the user who requested it.
3. Path setup should be monitored in real time to ensure that the defined constraints are all guaranteed.
4. Should a problem arise, the connection must be automatically re-established over a different path if available resources exist.
5. When a user no longer desires the connection, the path should be released so that the resources are free to be reserved by other users.

All of these features define the southbound interface requirements for the uppermost ZT OCX layer. Using this API, the layer will be able to interact with a multi-domain BoD service that will implement all of the lower layers of the OCX design, thus effectively implementing the OCX overlay infrastructure on top of the GÉANT network.

4.3.4.2 *SDN-Based Design*

The complete solution must be modular so that it can be scaled and extended where necessary. Components should be loosely coupled rather than "locked into place". In this way, the proposed solution can run in a multi-vendor environment and rely on integration to build the whole system. This also calls for the use of Open API between each component. The Northbound API between Management and Orchestration and Applications, for example, can use REST API, NetConf, SOAP or even XML RPC with a standard data format such as XML and JSON. The Southbound API from Management and Orchestration towards network resources can use NetConf, OpenFlow, SNMP, Telnet/SSH or even device-specific open API.

Studying the different possibilities in terms of concepts and models that can be used to implement the defined gOCX building block functionalities makes evident the necessity for maximum flexibility and transparency. The task considers that gOCX should be based on the Network Virtualization Functions [NFV] architecture that sharply decouples the S/W from the H/W components. In addition to their virtualisation, NFV also provides for the connection of network functions through a process known as "Service Function Chaining" (SFC), to derive benefits such as time reduction for service provisioning, and dynamic and elastic scaling for network services.

Furthermore, the network programmability that is made possible by the Software Defined Network [SDN] concept allows seamless communication at all levels, from hardware to software, and finally to end users (network operators). Programmability makes applications aware of the network and the network aware of applications. This enables greatly improved use of resources and opens up the potential for new applications with the additional prospect of revenue generation (e.g., flow metering), in which cost plans can be defined based on level of service provision. So far, SDN capabilities have been proven in a single domain environment, so the main challenge lies in implementing cross-domain end-to-end orchestration to bring everything together in a highly dynamic virtual environment.

Using SDN to separate the control and data planes and NFV to enable network functions and software to run on any open standards-based hardware, however, are not alone sufficient to provide a new service offering unless they are overseen by a Services Orchestration that enables automation, provisioning and interworking of physical and virtual resources. The combination of these capabilities offers a path to agile and dynamic networks – the confluence of SDN, NFV, and orchestration provides service automation so that GÉANT and the NRENs can address new possibilities and dynamically modify existing services to best benefit their end users.

The Zero Touch services provisioning model can be further empowered by using the DevOps [DO] technology that allows combining services development, deployment and operation management into one production cycle supported by management and orchestration tools. For cloud-based services and applications, the DevOps platforms and tools are built on top of popular cloud server deployment automation tools such as Puppet, Chef and Ansible [Review], however their network configuration capabilities only allow intra-cloud network configuration and must be extended to be successfully deployed in a multi-domain environment.

A good example of fusion between cloud-originated technologies and SDN is the recent development of the Network Automation and Programmability Abstraction Layer with Multivendor support system [NAPALM] that implements a common set of functions to interact with different network Operating Systems using a unified API. NAPALM supports several methods to connect to devices, manipulate configuration or retrieve data, and uses Ansible [Ansible] to configure network devices as programmable devices. Ansible gives certain benefits compared to other tools for network deployment and management, as it does not require a node agent and runs all processes over SSH. This simplifies its use for configuring network devices from multiple vendors.

For the final, front end to user component, the use of REST [RF] web services is advocated. These provide a richer yet straightforward and easy user experience that is customer-centric and as transparent as possible, while providing end users with the needed level of control of the virtual network over which the data will travel [SONV].

4.3.4.3 *OCX Marketplace*

The northbound interface of the ZT OCX can be implemented with a set of RESTful API calls that should essentially enable easy integration of ZT OCX into a more user-friendly environment. The only users that it is foreseen will directly interact with the ZT OCX API are the cloud application developers who can use this API to ensure that the underlying network will provide the necessary performances for the applications developed. This means that the API needs to be developed in such a way that will allow simple integration with different DevOps.

Another possibility is to provide a web portal that will enable end users to order the necessary network services from OCX. This web portal could be conceived as a meeting point where CSPs can offer their services to the users, who can in turn select the characteristics they require for the network connecting to the CSPs that provide their chosen services. This web portal could act as a Cloud Marketplace offering a catalogue of services and enabling the transparent interaction of users with ZT OCX, and providing a truly Zero Touch connection setup process.

Today cloud-powered service development is focused on the distributed complex applications that are multi-cloud based [CNCSP]. This means the final application or service offered to the end user is actually composed of number of interlinked cloud services that may be offered by different CSPs. This is one area in which the OCX marketplace can play a crucial role to alleviate the burden on multi-cloud application developers. A detailed service portfolio can help developers locate providers for needed resources and decide on the final set of development tools that are compatible with their CSP's choice.

The entire service can be developed applying DevOps by re-using existing collections of recipes or cookbooks where system configurations are expressed in a declarative language, adjusting them to satisfy the specific service requirements, and finally provision them using cloud automation and Zero Touch tools. The cloud-based services and applications development and provisioning workflow can be enriched by using TOSCA (Topology and Orchestration Specification for Cloud Applications) [TOSCA], which defines a flexible format for the description of the application topology and interrelationship between the components included in the provisioning workflow.

In order to provide the expected QoS for the uninterrupted high performance of this type of complex applications that use federated inter-cloud services, a holistic integrated approach to cloud application development and network infrastructure must be implemented. This means that the underlying OCX-based network infrastructure must be exposed to the higher application layer, as only in this way can the necessary optimisation be achieved to develop a cloud-aware network. It is therefore proposed that the OCX marketplace should also include a Zero Touch NaaS API so that the appropriate (long and short-term) virtual circuits are automatically set up according to each communication pair's demands, based on the multi-cloud workflow that defines the complete data flow between the services. In this way, the data flow from one service to another will be transported under the conditions as requested by the controlling application.

Figure 4.3: GÉANT's Cloud Marketplace

The provisioning of the seamless, Zero Touch-based implementation of this convergence of networking and cloud services is far from trivial. For high-performing multi-cloud applications this task would include the dynamic setup and tearing of a number of virtual connections, wherein the optimal path between a given set of gOCX instances needs to be computed according to the current complete OCX topology and network status. The complexity of the problem increases exponentially when the gOCX instances are placed in different network domains, considering that setting up one dynamic virtual link translates into a multi-domain network-wide orchestration process.

Given the complexity of the task, the solution is only feasible when applying complete network function abstraction and network resource virtualisation as provided by NFV and SDN. These are necessary tools to enable the agility of the underlying network infrastructure through programmability and to provide combined control and management when implementing the ZTP concept. The exposed ZT NaaS API from the OCX marketplace will enable invoking of the desired setup for the underlying network infrastructure using native cloud deployment tools. In this way, the newly developed complex multi-cloud applications (which are usually HPC or Big Data-oriented collaborative applications in the R&E environment) will provide composite network-cloud services solutions that enable flexible, dynamic network service composition across heterogeneous networking systems.

4.4 OCX Demonstration

With the aim of investigating the suitability of NFV and SDN as the basis for providing the OCX ZT NaaS features, a multi-cloud application for real-time UHD video editing was developed and demonstrated at the SC15 conference in Austin, Texas [SC15]. The demonstration focused on the possibilities of using NFV combined with Service Function Chaining [SFC] through the OCX programmable network (SDN). NFV was used to move the functions traditionally performed by specialised equipment to the virtualised cloud environment, which provides high flexibility and scalability. This makes it possible to manage network traffic by creating a chain of network functions. Adding network functions to active traffic with full transparency and in real-time requires a programmable network, hence the reliance on an SDN-based network. Functions can be added and removed in a programmable network at any time, in any order.

In the demo (see Appendix A), using a web-based interface, the end-user could choose different video editing functions, effectively defining the SFC workflow. The SFC generated API calls to the gOCX programmable instance (implemented as an OpenDayLight controller in Amsterdam) to direct the traffic flow between the various CSPs, since in each one a different type of video editing function is implemented. The generated traffic flow was 3.2Gbit/s. Using mostly 100G links, with no video functions activated. The traffic came straight back to Austin from Amsterdam with a delay of around 240ms.

The SFC demonstration was implemented using VLAN stitching and MAC rewriting so that different paths and functions were kept separate. This proved to be quite complex and not as flexible as intended, but the possibilities for SFC are still in development with a new protocol for identifying service function paths in the network called Network Services Headers (NSH).

The application presented at SC15 has successfully demonstrated the integration of network and cloud services in a federated and distributed infrastructure, and will be used as the basis to define an OCX ZT NaaS API. The setup is planned to evolve into a testbed for live testing of ZT NaaS combined with different multi-cloud-based applications. The next step is to analyse the complete network monitoring and management data available from the demo and decide accordingly on the way SDN will be integrated in the gOCX instances.

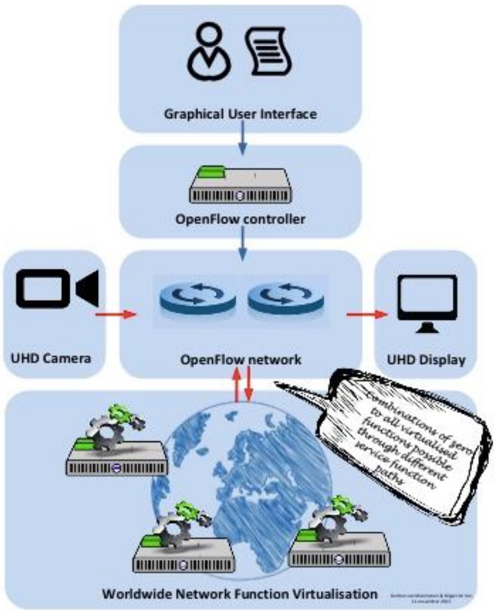

Figure 4.4: SC15 demonstration scenario

The first order testbed setup also provides an example of the possible future application functions solutions that could become part of the OCX marketplace.

Based on the NFV paradigm, a large number of additional services can be developed and offered on the marketplace, creating a pool of virtualised functions that end users can request in innovative combinations based on their requirements, using methodologies such as SFC to create custom multi-cloud solutions that will enrich the R&E community. In this sense, the Marketplace is envisioned to become the focal point of standard and integrated cloud services that dynamically interact with the agile programmable network infrastructure for high-performance delivery of targeted data flow.

5 Towards a Cloud-Aware Network Architecture

Collaborative video editing that requires complex interconnected infrastructure provisioning with multi-path data transfer and network connection with guaranteed QoS (e.g. latency and jitter) is just one use case among an increasing number of e-Science and Big Data applications. All such applications require on-demand access to large amounts of computational and storage resources and to transfer large amount of data at high speed. The proposed ZT OCX is the starting point for developing a full cloud-aware network architecture overlay with the aim of addressing the requirements of future network services.

To effectively provision complex cloud infrastructures on-demand, a consistent description of all infrastructure components or functional layers is required. This should include a general description of scientific applications and workflow, cloud resources, and intra-cloud and inter-cloud network infrastructure. All components of this environment need to be mutually aware of each other's functionality. In particular the (inter)cloud network infrastructure needs to be aware of the requirements and functionality of the applications and cloud resources/infrastructure, and vice versa.

Making this kind of services or functionality available to researchers and NREN users will bring numerous benefits such as increased research effectiveness and highly efficient network infrastructure usage. The GÉANT network is well positioned to implement such an environment, which will bring all the needed elements together in an overall cloud-aware layer.

For example, the Helix Nebula project identified that cloud providers want to optimise both the intra-cloud network infrastructure and the inter-cloud infrastructure to achieve optimal performance of scientific applications and effective cloud resources usage. At the same time, researchers expect to be able to control some of the cloud infrastructure components so that they can optimise their scientific workflow applications by managing the network that interconnects them. This use case was discussed within the framework of a GÉANT cooperation during the high-level design stages of OCX.

A Cloud-Aware Network Architecture (CANA) provides functionality to support user applications and scientific workflows and can fully integrate with both NaaS and SDN as network layer technologies. All infrastructure components and individual developments should be integrated into a common architecture and design environment to enable direct support for the network requirements/demand from the user applications, in particular for scientific workflows, and the inclusion of network topology descriptions and network QoS requirements in application design.

The envisioned future cloud-aware network architecture should make use of and combine the various technology components that are available and being developed within the research and GÉANT community, including:

1. Inter-domain network connectivity/infrastructure control and management services such as BoD, Autobanh [AB] and PerfSONAR [PS].
2. Network Service Interface [NSI] and Network Markup/Description Language (NML/NDL) [NML] developed by OGF and implemented in the GÉANT network.
3. Scientific and business workflow description languages and tools such as Kepler [KEP], Business Process Definition Metamodel [BPDM], etc.
4. The recently developed OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) standard, which immediately gained acceptance by industry, and which has developed a metamodel for defining IT services (see Figure 5.1).

The first set of components listed above enable dynamic, flexible, on-demand, high QoS network connections across multiple domains. AutoBahn is an example of a BoD provisioning tool with NSI support, while perfSONAR is a widely deployed test and measurement tool for monitoring network performance and verifying network setup. perfSONAR is designed to provide federated coverage of paths, and help establish end-to-end usage expectations, all using a uniform interface that allows for the scheduling of measurements, storage of data in uniform JSON based formats, and scalable methods to retrieve data and generate visualisations. This extensible system can be modified to support new metrics, and provides different options for data presentation.

High-performance networks offer advanced network services to end users with differing requirements, and a second set of components would allow users, applications or middleware to request network services from one or more service providers through a network service interface (NSI). The network service setup then requires configuration, monitoring and orchestration of network resources under particular agreements and policies. NSI is built as a generic network service interface in order to provide interoperability in a heterogeneous multi-domain environment. It defines information exchange, and the required messages, protocols and operational environment. Network Markup Language can be used as a basis to tackle naming and data modelling of network objects, topology description and sharing, and discovery of services and network capabilities. This is an extensible specification that is intended to be network infrastructure-agnostic, scalable, multi-layered and multi-domain.

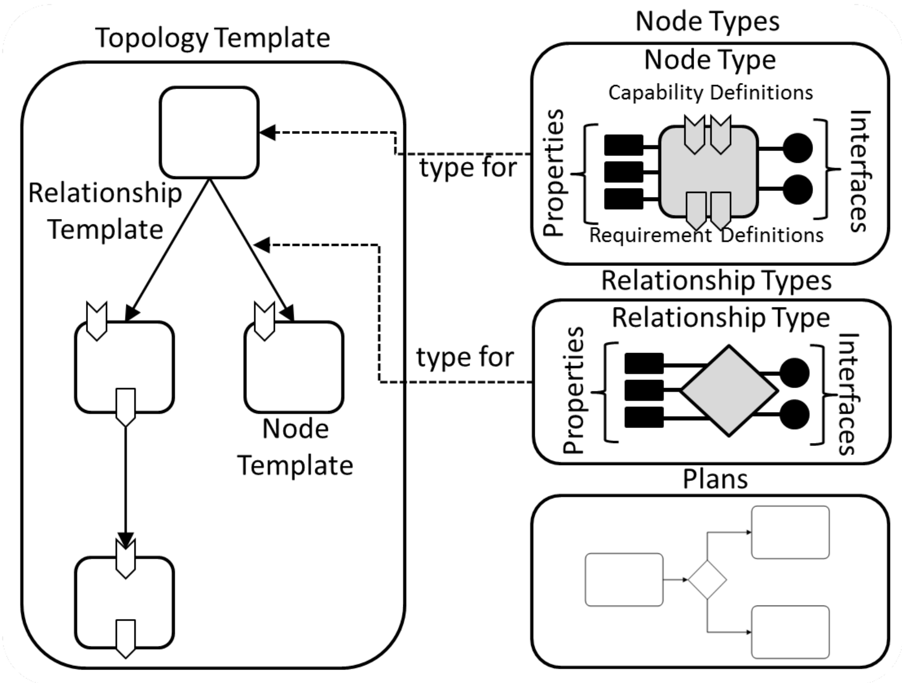

Figure 5.1: Main elements defining a service as depicted by TOSCA

The consistent deployment and operation of complex multi-provider, multi-cloud services and applications will require a precise actionable description of the underlying network infrastructure that can be adapted to possible variations in the available local connectivity services. This functionality is not available in current application description formats such as TOSCA or used by DevOps and cloud automation tools. On the other hand, the existing L0-3 network description languages and network management tools need to be extended to support scientific workflow integration and standards. They can be then used for seamless integration with the cloud-based applications so that the application can request the required network connections to be set up on-demand. This means that the languages and tools of the workflow description need to be able interface with the network using NSI so that they can send requests to create, modify, or tear down connections.

The workflow definition also needs to be adapted to the cloud environment where the modelling of data flow from one step to another is more complex when compared to the traditional grid environment where the workflow is a simple directed graph. However, using semantic annotation of the workflow that is based on a well defined ontology, several advanced features can be added.

The BPDM on the other hand offers the ability to integrate process models for workflow management processes, automated processes, and collaborations between units. It supports the description of the collaboration between participating entities, which is essential for multi-cloud applications. It also has the ability to exchange process specifications between modeling tools, and between tools and execution environments, using XMI.

The core TOSCA specification provides a language to describe service components and their relationships using a service topology. It can be used for describing the management procedures that create or modify services using orchestration processes. The combination of topology and orchestration in a Service Template describes what must be preserved across deployments in different environments to enable interoperable deployment of cloud services and their management throughout the complete lifecycle (e.g. scaling, patching, monitoring, etc.) when the applications are ported over alternative cloud environments.

As already discussed, TOSCA and the workflow descriptions can be extended using network description and management standards such as NML/NDL and NSI (standards developed and widely used by the GÉANT community). However, future research must also take into account the progress made by Standards Developing Organizations (SDOs) such as ONF and ETSI, which have strong affiliations to Network Function Virtualization (NFV). The requirement for infrastructure portability and replication, on-demand provisioning and cloudbursting when additional resources are needed must also be taken into account.

CANA is not just about adding network infrastructure control from application and workflow tools. Its ultimate purpose is Zero Touch control , hence to enable the simplified use and access of advanced network services for QoS-critical applications, extending the range of user services for community members.

6 Conclusions

Networking systems are facing the challenge of rapidly developing and deploying new functions and services to support the diverse requirements of various applications. In order to live up to users' expectations for fast, but also highly elastic, dynamic and scalable networking, fundamental changes are required in the way heterogeneous networks cooperate to support the wide spectrum of applications they offer. A promising approach within the networking research community in addressing these challenges is the decoupling of hardware and software in networking components, leading to the virtualisation of all networking resources. Service-oriented network virtualisation can enable the implementation of a network-as-a-service (NaaS) paradigm that allows the network infrastructure to be exposed and accessed as network services.

In this way the ultimate aim of Zero Touch, i.e. network flexibility and agility, is achieved by implementing the automation of network provisioning. The Zero Touch paradigm enables fast provisioning and configuration changes of networking devices, executed remotely and automatically, triggered by outside API calls or via a centralised cloud-based portal. This effectively transforms the underlying static network into a vibrant user-centric programmable network that can respond to requirements in a matter of minutes. The final stage of Zero Touch networking is to offer the user control over the network (more precisely, the user's traffic on the network), leading to a fully application-aware network that can offer NaaS services to the upper layer application. The application can use the NaaS services to request special treatment of its traffic as it flows through the network.

This white paper has investigated the current status of Zero Touch networking and beyond, exploring the possibilities of building a ZT NaaS architecture that will support future application-aware programmable agile networks. Possible implementation scenarios for ZT NaaS are proposed and discussed, with particular focus on NFV, SDN, and bare metal switching in combination with DevOps as technologies to enable the Zero Touch concept in a multi-domain complex networking environment.

Special attention is given to enabling the Zero Touch concept for cloud services since this area is likely to continue to expand and grow in the years to come. Users consuming cloud services offered through various providers are pushing to blend these service offerings into a richer experience. A dedicated cloud service delivery infrastructure is needed to enable network and computing service providers, content providers, and end users to offer and consume collaborative services. The gOCX concept is proposed as a solution to provide self-ordered, on-demand connectivity with guaranteed QoS between users and CSPs across the multi-domain R&E GÉANT and NREN network. To ensure that end users will truly benefit from gOCX services, user requests must be processed automatically using Zero Touch provisioning. By integrating the ZT NaaS model in gOCX the necessary level of automation and flexibility for on-demand connectivity as a service could be achieved. This vision requires the underlying networking infrastructure to be opened and exposed to upper-layer applications. This enables combined control, management, and optimisation of computing and networking resources for cloud service provisioning, leading to the creation of a converged composite network–cloud service provisioning system. The necessary exposure of network functionalities in a cloud environment can only be achieved through abstraction and virtualisation of networking resources, mainly using NFV and supporting technologies such as SDN and bare metal switching.

The proposed ZT OCX architecture has the potential to become a focal point for seamless (multi)cloud services. Through the Cloud Marketplace, a power user or a cloud application developer could access all of the cloud-related tasks they need to perform to create and use complex cloud-based services. The Marketplace portfolio could also include on-demand connection requests that are sent to and processed by the underlying intelligent cloud-aware network. This would comprise a stepping stone towards achieving a holistic vision involving the optimal use of the networking and computing resources available to the GÉANT community as a whole.

Appendix A Open Cloud Exchange and SDN-Enabled CSDI Demonstration

A.1 Technical Description of the Supercomputing 2015 Demo

One of the main goals of the defined scenario is to demonstrate the value of ZT OCX by using network hardware that can be operated by a separate controller (limited to the NetherLightOCX in this demo), while exploiting NFV in the cloud, and by chaining the network and service functions together. The complete demo scenario aims to demonstrate the power of ZT OCX in enabling an efficient usage of the network infrastructure for supporting multi-cloud applications that allow users to activate different processing tasks on different cloud service providers, while also controlling and directing the big data stream that flows in and out of each CSP.

The demonstration also enables a representation of the NFV features in a real-time environment, allowing the end user to visualise the possibilities offered by the programmable network. This is accomplished by allowing the demo users to alter the network path in real time. Traffic follows the user-selected path through zero or more network functions. The process of defining/altering the network path in real time is also known as Service Function Chaining (SFC).

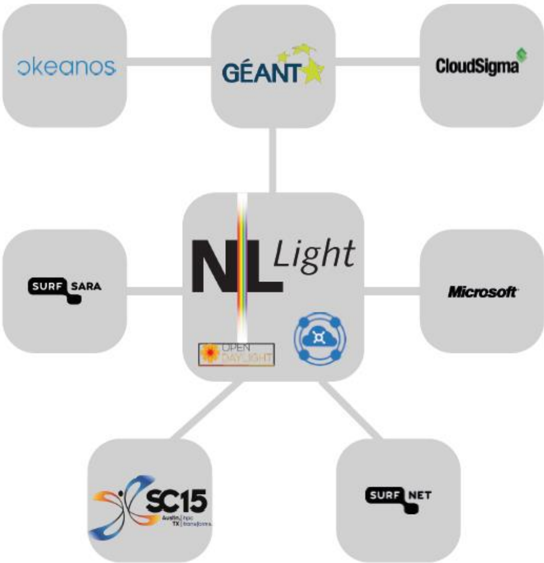

Figure A.1: Parties involved in the demo

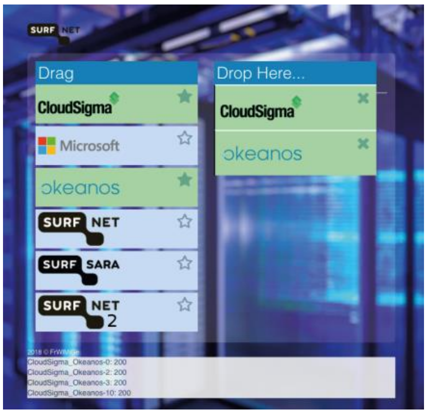

Figure A.2: The graphical user interface used by the clients in the demo

The SC15 [SC15] booth in Austin, Texas had a point-to-point connection to the NetherLight SDN platform in Amsterdam. This SDN platform was controlled by the OpenDayLight controller instance. Since NetherLight is an Open Cloud eXchange (OCX) it also provided a direct connection to all cloud service providers involved in the demo. The GÉANT OCX (gOCX) infrastructure model is distributed over the GÉANT core network and several NRENs. This allows cloud providers to physically connect at one of the OCX locations and have access to the entire GÉANT/NREN user community. Participating service providers for this demonstration are SURFsara (Netherlands), Okeanos (Greece), Microsoft (Netherlands), CloudSigma (Switzerland) and the SURFnet testbed (Netherlands), see Figure A.1 for all interconnected parties.

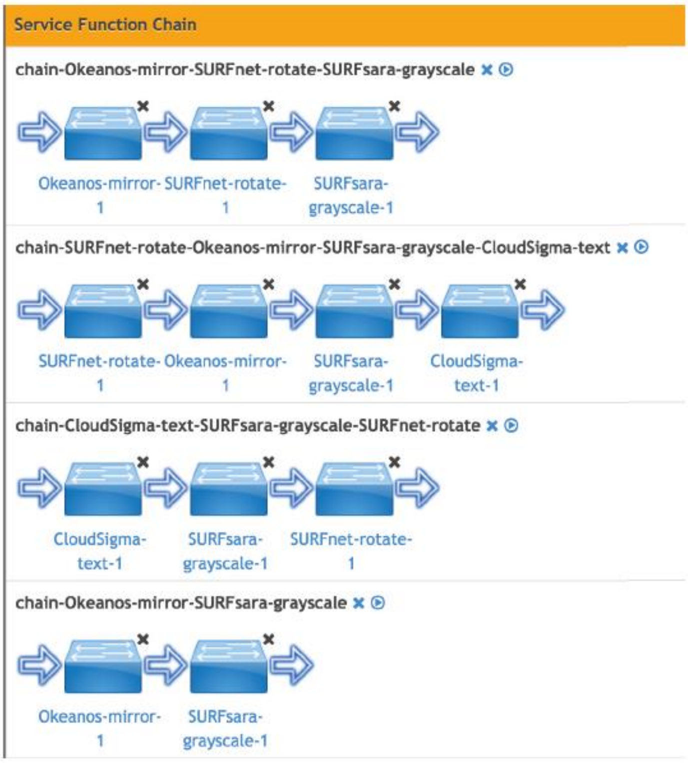

Figure A.3: Example video processing paths that can be defined using SFC

While a UHD camera at the SURF booth in Texas captures video frames of the booth visitors, an NFV-capable setup in the Netherlands receives the video stream and applies various video effects on the fly. Each video effect is created by processing the video frames at a different cloud service provider, while the user has complete control of the order in which the available video effects are applied. The manipulated video stream is then returned to the SURF booth, where it is displayed on a UHD screen. The user can choose several video transformations in a custom-made control GUI (see Figure A.2), redirecting the stream through different network functions in Europe. For the purposes of the demo, the following video transformations have been implemented: insert text, insert logo, grayscale, mirror and flip. Figure A.3 depicts different paths the video stream can take depending on the choices available to the end user.

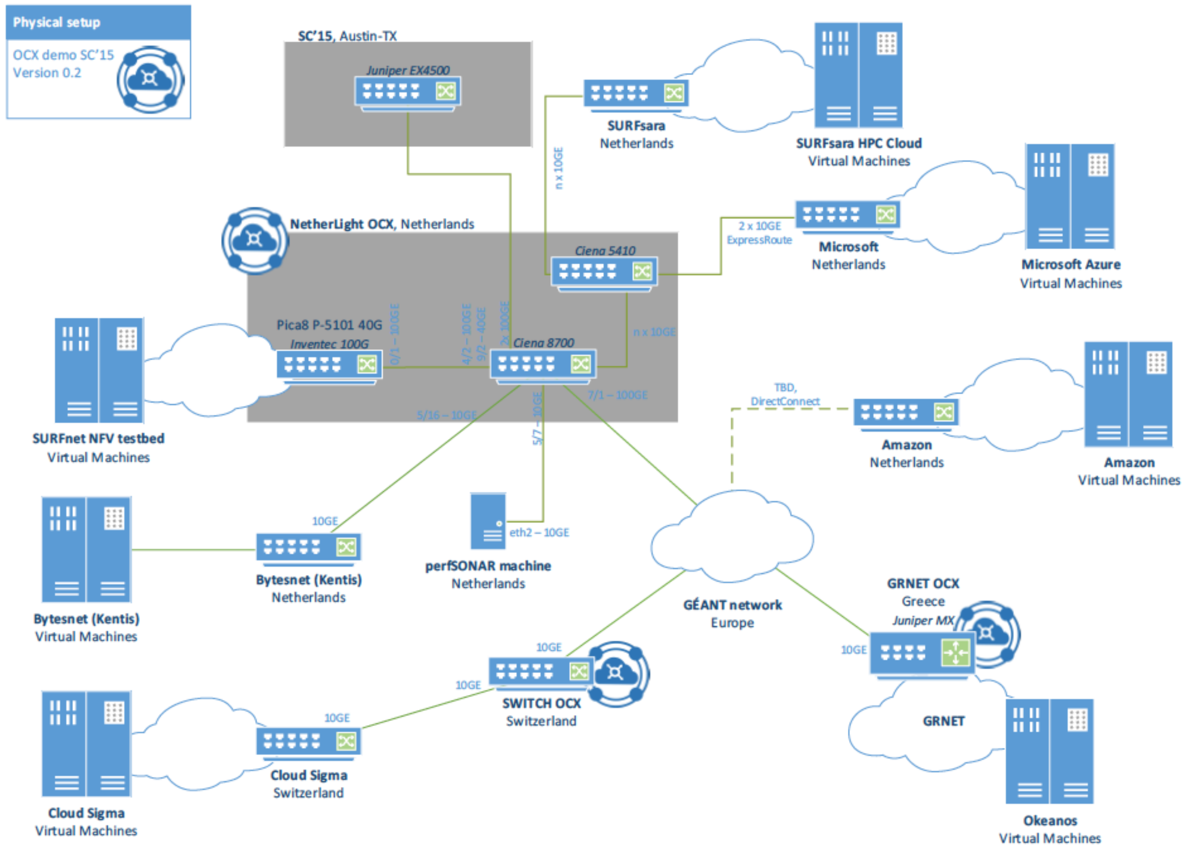

Figure A.4: Network setup of the demo scenario

In Figure A.4 the complete network setup for the demo scenario is presented. As mentioned, the NetherLightOCX is used as a Software Defined eXchange(SDX) in this scenario. This means that for this demo the "one switch, one controller" principle is being applied (this is one of the points to be upgraded for future demo versions of ZT OCX). While the other paths are currently static, in future attempts the envisioned multi-domain-controlled environment will be deployed.

The software used for the controller is OpenDayLight (ODL) – Lithium-SR1, chosen in view of its having the largest and most diverse active developers' community, since support is a key issue in experimental endeavours. During the demo setup, the team concluded that ODL is lagging slightly behind on ONOS, but should be able to catch up in time. The controller was used to control the Inventec/Pica8 bare metal switch, which was running PicOS OpenFlow firmware that was upgraded using an ODL application/plugin. With this choice of hardware, the demo also enabled more in-depth research into the possibilities of bare metal switches and their interconnection with SDN- and NFV-based networks.

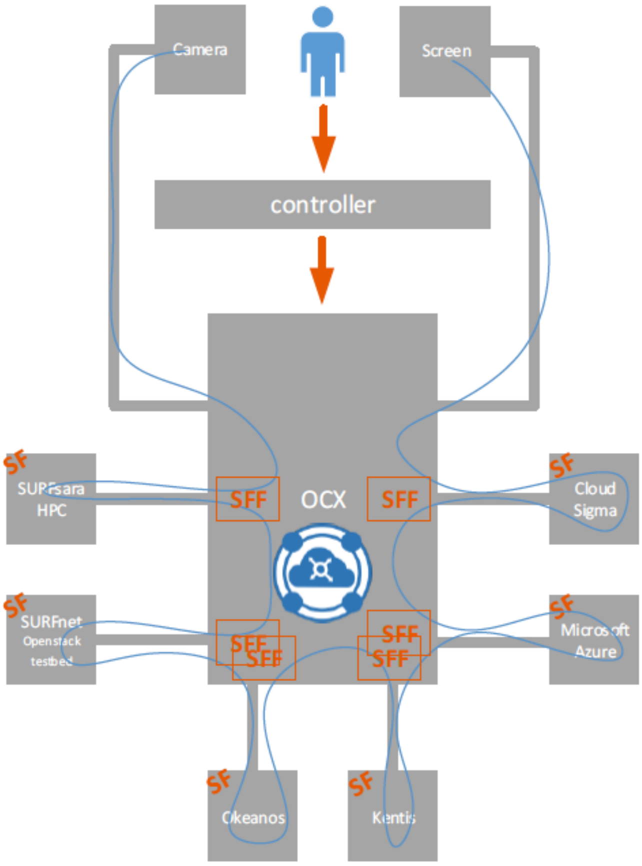

Figure A.5: SFFs and SFs setup in the ZT OCX demo scenario

The main actions taken in the demo involve active control over the video traffic flow using the purpose-built GUI. This possibility is mainly based on the Service Function Chaining feature. This provides the ability to define an ordered list of a network services. These services are then "stitched" together in the network to create a service chain. SFC is conceptually related to Policy Based Routing in physical networks but is typically thought of as a Software Defined Networking technology. Fundamentally, SFC causes network packet flows to route through a network via a path other than the one that would be chosen by routing table lookups on the packet's destination IP address. It is most commonly used in conjunction with Network Function Virtualization when recreating a series of network functions in a virtual environment that would have traditionally been implemented as a collection of physical network devices connected in series by cables.

The video sender (master in the SC15 demo scenario) is a JVC GY-HMQ10 camera with a Design DeckLink4K Extreme capture card, equipped with a 10GE NIC and running Ubuntu 14.04 LTS. The video receiver (slave in the SC15 demo scenario) is a Panasonic TX 55CX700E with GeForceGTX 970 running Ubuntu 14.04 LTS, also equipped with a 10GE NIC.

The Ultragrid application [UG], especially designed for low-latency, high-quality video network transmissions, is installed on both sender and receiver and enables sending and receiving of a 4K uncompressed video transferred using Jumbo frames. On the VMs running on the cloud providers, the HD-rum-transcode software is used for the transcoding actions applied on the video stream. This was made possible thanks to the support provided by the UltraGrid developers from CESNET [CESNET].

Appendix B OCX Functional Components and Design

Based on the high level ZT OCX design discussion given in Section 4 of this white paper, this Appendix sets out the essential functionalities of the ZT OCX main building blocks that are needed to achieve the desired goals of Zero Touch provisioning, connectivity on-demand and ease of selection and use of CSP services in a complex multi-cloud application environment.

In this sense, the Cloud Marketplace acts as a repository where CSPs offer different services or subscriptions, and users can browse and choose to subscribe to services, with or without customised SLAs, which can be provided by activating the underlying ZT OCX infrastructure. The main information kept in the repository is the catalogue of services with their type and description, the available APIs that can be used to access them, and the available SLAs that can be applied when connecting to the given CSP. Each CSP that is directly connected to one or more OCX instances will have a set of services recorded in this repository.

The minimum set of functions needed in order to use this repository are:

1. Subscribe to (for users) and publish (for CSPs) service.
2. Modify service (CSPs) / subscription (users).
3. Delete service (CSPs) / Unsubscribe (users).
4. Search/filter catalogue (users).

Access to this repository needs to be provided in multiple forms so that it can be easily used standalone or by integrating its features into a custom-built solution. Hence, the marketplace repository should be accessed via open API, CLI or RESTfull web services.

The visualisation component of the Marketplace is envisioned to be a web portal that should become the extension of today's GÉANT's SA7 cloud catalogue, which so far is a static directory of CSPs. The users of the web portal are the CSPs for publishing information and cloud application developers (who need to be familiarised with what is available for use before starting to develop applications), as well as citizen scientists (big data power users) who may wish to browse, subscribe to and use services directly from the portal itself.

The Cloud Marketplace is positioned on top of the ZT OCX stack and needs to interact with the top ZT OCX layer to provide on-demand dynamic and flexible high-performance direct connections from the end users to the CSP(s) via the OCX instances. As discussed previously, the ZT OCX should provide a northbound open API for accepting the requests from the marketplace or another third party. The minimum set of functionalities that this API should provide are:

1. tryconnection(at least two ServiceTerminationPorts (from CSP into OCX input stream port, from user to OCX, outport from OCX), QoS arguments, timeFrame), the output of this call is whether the connection demand is feasible or not
   - the call should be overridden with possibilities for multipoint connections
2. if connection is feasible then setupconnection() can be used. The output of this call is the status of the setup: success/failure
   - if connection is successfully setup, it needs to have its own service/flow/connection id(s)
3. modify connection(modified parameters). The output of this call is the status of the setup: success/failure
4. destroy connection() ex.: the CSP can call it if the user stops paying. If the user does not call it upon finishing the work, the connection will be destroyed automatically (based on the timeout of connection, inactivity).

Using the API the user can reserve connections for use in a future timeslot, just as it is implemented in the one-domain BoD services.

An additional set of functions are needed to provide real-time monitoring of the established connections. This monitoring will allow both parties to verify the agreed SLA parameters by polling info from the lower level components (domain and device level). At the same time, the monitoring data will enable connection assurance.

Evidently, in order to provide seamless integration, the complete design must support an existing AA framework that will not only take care of the access rights, but will extend the users' profiles into the CSPs environment.

B.1.1 OCX Use Scenario

In order to demonstrate a typical usage of the complete marketplace and ZT OCX solution, a scenario is considered where the user is a multi-cloud application developer who needs to use a set of cloud services and ensure high-performance connections to the CSPs that offer the chosen services.

To achieve the task of developing a multi-cloud application, the developer will need to go through a given set of steps that involve using the proposed solution (the rest of the development steps that do not include the proposed solution are omitted):

1. Build an application abstract model
   a. This includes answering the question "what I want to accomplish"
   b. As well as the question "what I need in order to do it"
      i. The answer to this will define the types of cloud and network resources that are needed such as storage, VMs, network parameters, and software components
2. After defining the needed resources, the marketplace can be used in order to browse the catalogue for the required components: VMs, app, data, CSPs
   a. The results of this step is a list of candidate CSPs that can be used
3. Upon deciding on the CSPs that are going to be used, the next step is to decide on the tools for development such as Chef, Puppet, Ansible, Slipstream, etc.

   a. The choice of tool will depend on the chosen CSPs’ support and on the features and familiarity of the developer with the tools themselves
4. Using the chosen tool, the physical infrastructure and resources needed for the application need to be described
   a. This is accomplished by writing so-called recipes or a cookbook wherein the workflow of the application is hardcoded
   b. The application infrastructure (topology) also needs to be described
      i. This can be accomplished using a tool such as OASIS TOSCA
      ii. The infrastructure will define the necessary connections that need to be set up using calls to ZT OCX
   c. An application layer diagram that describes the interaction of the multiple different components is to be defined
   d. Based on the previous information, the data workflow is to be defined (ex. the output from VM1 and VM2 on CSP A needs to be mixed with the output from VM1 and VM2 on CSP B inside a VM on CSP C in order to obtain the result for the end user)
5. Based on the application infrastructure topology and the devised workflow, the multi-cloud application is to be developed as a composite service
   a. service function chaining is defined according to the workflow
   b. the workflow and SFC will provide the final information on the order and duration of the establishment of the on-demand direct connections between the entities in the OCX infrastructure.

Appendix C Dissemination of Work

The work presented in this white paper has been published in a number of papers and posters, and presented at several conferences:

[1] Demchenko, Yuri, Sonja Filiposka, Raimundas Tuminauskas, Anastas Mishev, Damir Regvart, Kurt Baumann, and Tony Breach. "Enabling Automated Network Services Provisioning for Cloud Based Applications Using Zero Touch Provisioning." NetCloud workshop, part of the Utility and Cloud Computing conference, Cyprus (2015)

[2] Demchenko, Yuri, C. Dumitru, S. Filiposka, T. Matselyukh, D. Regvart, M. de Vos, T. Karaliotas, K.Baumann, D. Arbel, C. de Laat. “Open Cloud eXchange (OCX): A Pivot for Intercloud Services Federation in Multi-provider Cloud Market Environment.” Proc. IEEE International Conference on Cloud Engineering (IC2E) (2015)

[3] Suerink, Tristan, Tasos Karaliotas, Damir Regvart, Yuri Demchenko, Kurt Baumann, François Kooman, Wladimir Mufty, Hans Trompert, Ronald van der Pol, Gerben van Malenstein, Migiel de Vos. “Real-Time Worldwide Service Function Chaining Using A Programmable Network With Network Function Virtualisation”, SuperComputing conference (2015)

[4] Filiposka, Sonja, Yuri Demchenko, Daniel Arbel, Anastas Mishev, Migiel de Vos, Tasos Karaliotas and Damir Regvart. “Enabling High Performance Cloud Computing Using Zero Touch Provisioning”, Telfor (2015)

[5] Demchenko, Yuri, Jeroen van der Ham, Cees de Laat, Migiel de Vos, Damir Regvart, Sonja Filiposka, Kurt Baumann, Taras Matselyukh, Tasos Karaliotas, Daniel Arbel, Eduard Escalona, Alex Marvin, Tony Breach. “Open Cloud eXchange (OCX): Bringing Cloud Services to NRENs and Universities.” Terena TNC (2014)

[6] Filiposka, Sonja, Yuri Demchenko, Migiel de Vos, Tasos Karaliotas, Damir Regvart. “Distributed cloud services based on programmable agile networks.” Terena TNC (2016) accepted, full paper pending

[7] Khasnabish, B., J. Chu, S. Ma, N. So, P. Unbehagen, M. Morrow, M. Hasan, Y. Demchenko, Y. Meng. “Cloud Reference Framework.” IETF Internet draft, ver. 8.0 (2015)

References

[AB]	Multi-domain Bandwidth on Demand service pilot (2011)
[ADC]	Amazon Direct Connect
[Ansible]	Ansible IT automation tool
[AnsibleCM]	Hall, D. “Ansible Configuration Management.” Packt Publishing Ltd (2015)
[ArchFW]	Demchenko, Y., M. Makkes, R.Strijkers, C.Ngo, C. de Laat. “Intercloud Architecture Framework for Heterogeneous Multi-Provider Cloud based Infrastructure Services Provisioning.” The International Journal of Next-Generation Computing (IJNGC), 4-2 (2013)
[Aruba]	Aruba Networks. “Aruba Instant: Ultra-Rapid Enterprise Wi-Fi Deployment with Zero Touch Provisioning.” Product Overview (2013)
[AT]	Allied Telesis. “ZTC Zero Touch Configurator.” Datasheet (2007)
[BCF]	Big Switch. “The Big Cloud Fabric – Hyperscale Networking Design Principles.” Whitepaper (2014)
[BMWBS]	Kemp, Glen. “What separates white-box software from bare-metal switches?” TechTarget (2014)
[BPDM]	Conrad Bock. “Introduction to the Business Process Definition Metamodel.” NIST (2008)
[BRCD]	Brocade. “Brocade data center technology leadership.” (2015)
[BS]	Sherwood, Rob. “Tutorial: White box / bare metal switches.” Big Switch Networks, Open Network User’s Groups (2014)
[C-SON]	Téral, S. “The Logical Rise of C-SON. Why C-SON is About to Rule the World,” Infonetics Research White Paper (2015)
[CCS]	Rimal, B. P., A. Jukan, D. Katsaros, Y. Goeleve. “Architectural Requirements for Cloud Computing Systems: An Enterprise Cloud Approach.” Journal of Grid Computing 9 (2011): 3–26.
[CESNET]	https://www.cesnet.cz/?lang=en
[CIWAN]	Cisco. “Cisco Intelligent WAN: Orchestrate Change from a Central Location.” White Paper (2014)
[CN]	Cumulus Networks. “Bare Metal Switches — Is There a Cost Benefit?” BusinessBrief (2015)

[CNCS] Cisco. “Cisco Optical Transport Solutions: The Cisco NCS 2000” (2015)

[CNCSP] Huanga, Jun, Guoquan Liub, Qiang Duan. “On Modeling and Optimization for Composite Network–Cloud Service Provisioning.” Journal of Network and Computer Applications. 45 (2014)

[CompServ] Fehling, Christoph, and Ralph Mietzner. "Composite as a service: Cloud application structures, provisioning, and management." it-Information Technology Methoden und innovative Anwendungen der Informatik und Informationstechnik 53.4 (2011): 188-194.

[CZTP] Cisco Networking Services. “Zero Touch Cisco Networking Services Configuration Guide.” Cisco IOS Release 15M&T (2012)

[CYCLONE] CYCLONE Project

[DO] Httermann, Michael. DevOps for developers. Apress, 2012.

[ECI] “ECI TO DEBUT ITS ELASTIC NETWORK STRATEGY” (2015)

[FACHIE] Demchenko, Yuri, C. Lee, C.Ngo, C. de Laat. “Federated Access Control in Heterogeneous Intercloud Environment: Basic Models and Architecture Patterns.” IEEE Third International Workshop on Cloud Computing Interclouds, Multiclouds, Federations, and Interoperability (Intercloud), Proc IEEE International Conference on Cloud Engineering (IC2E) (2014)

[FNA] Wessing, H., K. Bozorgebrahimi, A. Tzanakaki, B. Belter, S. Naegele-Jackson, A. Metz, P. Skoda, J. Vojtech, V. Olifer. “Future Network Architectures.” GN3 Deliverable D12.3 (DJ1.1.1) (2015)

[GCC] Rings, T., G. Caryer, J. Gallop, J. Grabowski, T. Kovacikova, S. Schulz, I. Stokes-Rees. “Grid and Cloud Computing: Opportunities for Integration with the Next Generation Network.” Journal of Grid Computing 7(2009):375–393.

[GEYSERS] GEYSERS Project

[GN4-1] http://www.geant.org/Projects/GEANT_Project_GN4-1/Pages/Home.aspx

[gOCX] Demchenko, Yuri, at al. “GÉANT Open Cloud eXchange (gOCX): Architecture, Components, and Demo Scenario.” Supercomputing conference (2014)

[HSD] Malekzadeh, R. “The Advent of Decoupling of Hardware and Software,” CIO Review – Networking (2014)

[HPC] Jackson, K. R., K. Muriki, S. Canon, S. Cholia, and J. Shalf. “Performance analysis of high performance computing applications on the Amazon Web services Cloud.” Proc. IEEE Int. Conf. on Cloud Computing Technology and Science (2010):159–168.

[ICCE] Toosi, Adel Nadjaran, Rodrigo N. Calheiros, and Rajkumar Buyya. “Interconnected Cloud Computing Environments: Challenges, Taxonomy, and Survey.” ACM Comput. Surv. 47-1-7 (2014)

[ICG] Andronico, G., R. Barbera, A. Fornaia, S. Monforte. “An attempt to integrate Clouds in Grids.” Proc. of the International Symposium on Grids and Clouds and the Open Grid Forum Academia Sinica (2011)

[INF] Infinera. “Infinera Intelligent Transport Network, Evolve to the Terabit Era” (2015)

[JCZTP] Juniper Networks. “Configuring Zero Touch Provisioning in Branch Networks.” (2015)

[JZTP] Juniper Networks. “Understanding Zero Touch Provisioning.” Juniper TechLibrary (2014)

[KEP] Taylor, I.J., Deelman, E., Gannon, D.B., Shields, M. (Eds.). “Workflows for e-Science: Scientific Workflows for Grids.” 530 p. Springer (2007

[MAER] Microsoft Azure Express Route

[MAN] Monolith Software. “Unifies Service Assurance.” AssureNow Overview (2014)

[MEF] MEF. “The Third Network: Lifecycle Service Orchestration Vision.” Whitepaper (2015)

[MON] Ennis, Shawn. “Meeting the Agility that Zero-Touch Provisioning Demands of Operations.” Monolith Software (2014)

[NACS] Regvart, D. at al., “Network Architectures for Cloud Services White Paper: gOCX.” MS101 (MJ1.2.1), GN3+ (2014)

[NAPALM] Network Automation and Programmability Abstraction Layer with Multivendor support (NAPALM)

[NGN] Pathan, Al-Sakib Khan, Muhammad Mostafa Monowar, Zubair Md. Fadlullah. “Building Next-Generation Converged Networks: Theory and Practice.” CRC Press (2013)

[NetCloud] Demchenko, Yuri, Sonja Filiposka, Raimundas Tuminauskas, Anastas Mishev, Damir Regvart, Kurt Baumann, and Tony Breach. "Enabling Automated Network Services Provisioning for Cloud Based Applications Using Zero Touch Provisioning." NetCloud workshop, part of the Utility and Cloud Computing conference, Cyprus (2015).

[NFV] ETSI GS NFV 002. “Network Functions Virtualisation (NFV) Architectural Framework.” V1.2.1 (2014-12)

[NMA] Kim, Hyojoon, and Nick Feamster. "Improving network management with software defined networking." Communications Magazine, IEEE 51.2 (2013): 114-119.

[NML] Jeroen van der Ham, Freek Dijkstra, Roman Łapacz , Aaron Brown. “The Network Markup Language (NML) A Standardized Network Topology Abstraction for Inter-domain and Cross-layer Network Applications.” TERENA TNC (2013)

[NMS] Rao, U. H. “Challenges of Implementing Network Management Solution.” Int. Journal of Distributed and Parallel Systems (IJDPS) 2.5 (2011):67-76

[NSI]	Eduardo Escalona**. "**OGF Network Service Interface (NSI)." TERENA ISoD Workshop (2009)
[OCP]	Open Compute Project, Networking Specs and Designs (2015)
[OCX]	Demchenko, Y., J. Van Der Ham, C. Ngo, C. De Laat, T. Matselyukh, E. Escalona, and S. Filiposka. "Open cloud exchange (OCX): Architecture and functional components." Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th Int Conf on, 2 (2013): 81-87.
[OCXIC]	Demchenko, Yuri, C. Dumitru, S. Filiposka, T. Matselyukh, D. Regvart, M. de Vos, T. Karaliotas, K.Baumann, D. Arbel, C. de Laat. "Open Cloud eXchange (OCX): A Pivot for Intercloud Services Federation in Multi-provider Cloud Market Environment." Proc. IEEE International Conference on Cloud Engineering (IC2E) (2015)
[ODSFC]	Kumbhare, Abhijit, and Vinayak Joshi. "Opendaylight Service Function Chaining Use-cases." Ericsson (2014).
[ONaaS]	OpenNaaS Network provisioning platform
[ONC]	Vicat-Blanc, P., Soudan, S., Figuerola, S., et al. "Bringing Optical Networks to the Cloud: an Architecture for a Sustainable future Internet." Springer Lecture Notes in Computer Science, 6656 (2011): 307-322.
[ONF]	ONF. "L4-L7 Service Function Chaining Solution Architecture." ver. 1.0, (2015)
[ONIE]	Brune, Curt. "Open Network Install Environment." Open Compute Summit (2014)
[ONL]	Open Network Linux, Documentation (2015)
[OVT]	Overture. "Zero Touch Provisioning." Product Features (2012)
[PIC8]	Pica8 open networking. "Bare Metal Networking Leveraging "White Box'' Thinking." Whitepaper (2015)
[PS]	PerfSonar
[QDQE]	Klaus Diepold. "The Quest for a Definition of Quality of Experience". QualinetNewslet (2012)
[RANC]	Sharkh, M. A. , M. Jammal, A. Shami, A. Ouda. "Resource Allocation in a Network-Based Cloud Computing Environment: Design Challenges." Cloud Networking and Communications, IEEE Communications Magazine(2013):46-52.
[RF]	Richardson, Leonard, and Sam Ruby. RESTful web services. O'Reilly Media, Inc., 2008.
[Review]	Venezia, Paul. "Review: Puppet vs. Chef vs. Ansible vs. Salt" (2013)
[SC15]	Suerink, Tristan, Tasos Karaliotas, Damir Regvart, Yuri Demchenko, Kurt Baumann, François Kooman, Wladimir Mufty, Hans Trompert, Ronald van der Pol, Gerben van Malenstein, Migiel de Vos. "Real-Time Worldwide Service Function Chaining Using A Programmable Network With Network Function Virtualisation", SuperComputing conference (2015)

[SDN] Sezer, S., S. Scott-Hayward, P. K. Chouhan, B. Fraser, D. Lake, J. Finnegan, N. Viljoen, M. Miller, N. Rao. “Are We Ready for SDN? Implementation Challenges for Software-Defined Networks, Future Carrier Networks.” IEEE Communications Magazine (2013).

[SFC] Halpern, J., and C. Pignataro. Service Function Chaining (SFC) Architecture. No. RFC 7665. 2015.

[Slip] SlipStream Application Deployment Platform

[SONV] Duan, Q., Y. Yan, A. V. Vasilakos. “A Survey on Service-Oriented Network Virtualization Toward Convergence of Networking and Cloud Computing.” IEEE Transactions On Network and Service Management 9-4(2012):373-392

[Telfor] Filiposka, Sonja, Yuri Demchenko, Daniel Arbel, Anastas Mishev, Migiel de Vos, Tasos Karaliotas and Damir Regvart. “Enabling High Performance Cloud Computing Using Zero Touch Provisioning”, Telfor (2015)

[TNC14] Demchenko, Yuri, Jeroen van der Ham, Cees de Laat, Migiel de Vos, Damir Regvart, Sonja Filiposka, Kurt Baumann, Taras Matselyukh, Tasos Karaliotas, Daniel Arbel, Eduard Escalona, Alex Marvin, Tony Breach. “Open Cloud eXchange (OCX): Bringing Cloud Services to NRENs and Universities.” Terena TNC (2014)

[TNC16] Filiposka, Sonja, Yuri Demchenko, Migiel de Vos, Tasos Karaliotas, Damir Regvart. “Distributed cloud services based on programmable agile networks.” Terena TNC (2016) submitted, under review

[TOSCA] “Topology and Orchestration Specification for Cloud Applications.” ver. 1.0. Candidate OASIS Standard. (2013)

[UG] UltraGrid Open-source software for high bandwidth, low latency network transmissions, CESNET (2015)

[ZOOM] Zero-time Orchestration, Operations and Management (ZOOM), TeleManagement Forum (2015)

[ZOOM2] TMF. “Zero Touch Network-as-a-Service: Agile, Assured and Orchestrated with NFV.” TMForum (2015)

[ZTESP] Gowan, B. “Infographic: Zero-Touch Ethernet Services Provisioning.” Ciena Insights Blog, (2013)

[ZTPNW] Chai, C. “Zero Touch Provisioning Can Help the Network World Catch up to Server Advances.” Network World (2015)

Glossary

AP	Access point
API	Application Programmable Interface
ARM	Adaptive Radio Management
BOD	Bandwidth on Demand
BPDM	Business Process Definition Metamodel
BSS	Business Support System
CANA	Cloud Aware Network Architecture
CLI	Command Line Interface
CSDI	Cloud Services Delivery Infrastructure
CSP	Cloud Service Provider
DHCP	Dynamic Host Configuration Protocol
DWDM	Dense Wavelength Division Multiplexing
ETSI	European Telecommunications Standards Institute
GMPLS	Generalized Multiprotocol Label Switching
GUI	Graphical User Interface
H/W	Hardware
IPoDWDM	IP over DWDM
JSON	JavaScript Object Notation
MAC	Media Access Control
NaaS	Network as a Service
NAPALM	Network Automation and Programmability Abstraction Layer with Multivendor Support
NCS	Network Convergence System
NDL	Network Description Language
NFV	Network Function Virtualization
NGN	Next Generation Networks
NIC	Network Interface Card
NML	Network Markup Language
NOC	Network Operations Center
NREN	National Research & Education Network
NSI	Network Service Interface
NTP	Network Time Protocol
OASIS	Organization for the Advancement of Structured Information Standards
OCP	Open Compute Project
OCX	Open Cloud eXchange
ODL	OpenDayLight
OGF	Open Grid Forum

ONF	Open Networking Forum
ONL	Open Network Linux
ONIE	Open Network Install Environment
OPEX	Operational Expense
OS	Operating System
OSS	Operations Support System
OTN	Optical Transport Networks
QoS	Quality of Service
QoE	Quality of Experience
R&E	Research & Education
REST	Representational State Transfer
RF	Radio Frequency
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RPC	Remote Procedure Call
SDN	Software Defined Networks
SDO	Standards Developing Organization
SDX	Software Defined Exchange
SF	Service Function
SFC	Service Function Chaining
SFF	Service Function Forwarder
SNMP	Simple Network Management Protocol
SOAP	Service Oriented Architecture Protocol
S/W	Software
TMF	TeleManagement Forum
TOSCA	Topology and Orchestration Specification for Cloud Applications
ULH	ultra-long-haul
VM	Virtual machine
WAN	Wide Area Network
WLAN	Wireless Local Area Network
WSON	Wavelength Switched Optical Network
XMI	XML Metadata Interchange
XML	Extended Markup Language
ZOOM	Zero-Touch Orchestration Operations and Management
ZT	Zero touch
ZTN	Zero Touch Networking
ZTP	Zero Touch Provisioning