SDN orchestration architectures and their integration with Cloud Computing Applications

Arturo Mayoral[a,*], Ricard Vilalta[a], Raul Muñoz[a], Ramon Casellas[a], Ricardo Martínez[a]

[a]*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC)*
Parc Mediterrani de la Tecnologia, Av. Carl Friedrich Gauss 7
08860 Castelldefels (Barcelona), Spain.

Abstract

Emerging cloud-based applications, running in geographically distributed data centers (DCs), generate new dynamic traffic patterns which claim for a more efficient management of the traffic flows. Geographically distributed DCs interconnection requires automatic and more dynamic provisioning and deletion of end-to-end (E2E) connectivity services, through heterogeneous network domains. Each network domain may use a different data transport technology but also a different control/management system. The fast development of Software Defined Networking (SDN) and the interworking with current control plane technologies such as Generalized Multi-protocol Label Switching (GMPLS), demand orchestration over the heterogeneous control instances to provide seamless E2E connectivity services to external applications (i.e. cloud computing applications).

In this work, we present different orchestration architectures based on the SDN principles which use the Path Computation Element (PCE) as a fundamental component. In particular, a single SDN controller orchestration approach is compared with an orchestration architecture based on the Application Based Network Operations (ABNO) defined within the International Engineering Task Force (IETF), in order to find the potential benefits and drawbacks of both architectures. Finally, the SDN IT and Network Orchestration (SINO) platform which integrates the management of Cloud Computing infrastructure with the network orchestration, it is used to validate both architectures by evaluating their performance providing two inter-DC connectivity services: E2E connectivity and Virtual Machine (VM) migration.

Keywords:
SDN, ABNO, Orchestration, AS-PCE, Cloud Computing, VM migration

1. Introduction

Most emerging internet applications relay upon cloud-based services geographically distributed among different data centers (DCs). DCs are generating an increasing amount of dynamic, variable horizontal traffic [1] (i.e. databases synchronization or virtual instances migration), and vertical traffic between customers and DCs. The augment on the number of dynamic application-driven requests of network resources' reservation, is moving network operators to find new network architectures which can provide automatic and efficient management of the setup and release of connectivity services. These new architectures must provide end-to-end (E2E) connectivity across different network domains. Intra-DC networks require a very dynamic, packet-based traffic control, while long-haul optical transport networks, which transport the inter-DC traffic, have carrier-grade, multi-domain control requirements.

Network management and control systems are suffering an unprecedented evolution since Software Defined Networking (SDN) has emerged as the future networking paradigm. SDN proposes separating the
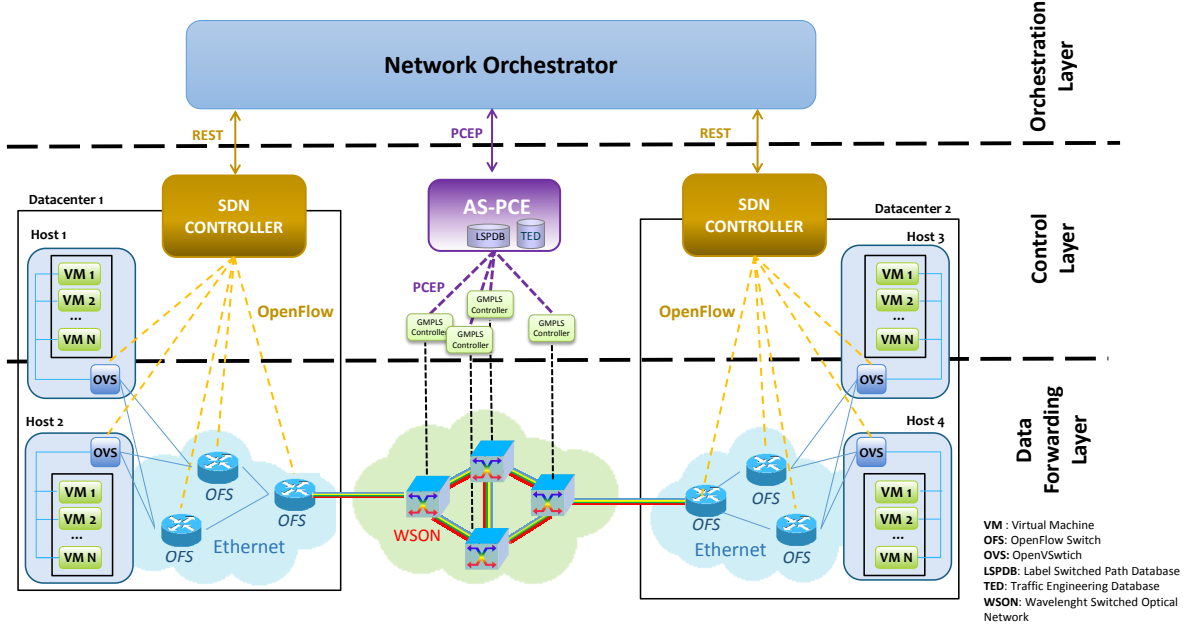

Figure 1: Network orchestration architecture.

control logic from the switching infrastructure by removing the 'intelligence' from the forwarding elements and concentrating it into a logically centralized SDN controller. OpenFlow (OF) [2] is a standard protocol developed within the Open Networking Foundation (ONF), which allows externally define the forwarding behavior of the network infrastructure by characterizing the traffic as a combination of flow rules based on the packet headers. It has become the preferred SDN interface between control and data planes. SDN allows cutting costs from network infrastructure due to dedicated hardware can be replaced by software-based switches installed on cheaper Commercial Off the Self (COTS) servers. This is one of the main reasons why SDN is attracting so much interest from a wide spectrum of the networking industry (especially on DC operators segment).

On the other hand, Generalized Multi-Protocol Label Switching (GMPLS) in combination with the Path Computation Element (PCE) is a mature technology with more than ten years of standardization progress, which offers a carrier-grade control solution for automatic circuit provisioning in Wavelength Switched Optical Networks (WSON). The Active Stateful PCE (AS-PCE) [3] has been demonstrated as a robust and effective management solution for the dynamic establishment and release of optical circuits or Label Switched Paths (LSPs) in GMPLS-controlled optical networks [4].

Operators which have already deployed GMPLS-based control solutions need to assure the return of investment of their current deployments, thus any network upgrade needs to account on existing control technologies. In this context, it arises the need of coordinating or orchestrating multiple, heterogeneous control planes. Inter-working between different control planes requires a higher, master entity (referred here as Network Orchestrator - NO) which automatically coordinates the processes to establish and release E2E connections through different network domains controlled by different control instances. A graphical description of this network scenario can be viewed in Figure 1.

Recently, the ABNO architecture [5] has been designed within the IETF, based on standard protocols and components to efficiently provide a network orchestration solution for multi-layer and multi-domain networks. In this article, we are presenting a full-defined ABNO implementation, with a modular, plugin-based, architecture to orchestrate multiple southbound controllers. Its northbound Application Programmable Interface (API) has been designed following the Representational State Transfer (REST)-ful principles to allow external IT applications [6] (i.e. cloud computing management systems) directly request E2E connectivity services into the network. In addition, a previously presented orchestration approach [7] based on a single

SDN controller is compared with the ABNO orchestration architecture.

The paper is structured as follows: in section 2 an overall description of the orchestration process and different orchestration approaches is presented, also Cloud Computing and the need of integrating network resources into a jointly orchestration is examined and introduced. In section 3, two different SDN orchestration architectures: ABNO architecture and Single SDN Controller Architecture (SC-Arch); are thoughtfully described and compared. To conclude the section, the SDN IT and Network Orchestration (SINO) application and its integration within the two previously described architectures, is presented. And finally, section 4 presents the experimental validation of both network orchestration architectures in the SDN/NFV Cloud Computing platform and Transport Network of the ADRENALINE Tesbed, by measuring and comparing the setup delay of two typical inter-DC communication uses cases: E2E channel provisioning and Virtual Machine (VM) migration.

2. State of the Art

This section reviews the current state of the art on SDN orchestration and Cloud Computing. It is organized in two subsections: first one goes into the description of the orchestration process and the global perspective of the current orchestration approach; and second one, it makes a brief introduction to Cloud Computing and the need of integrating network resources into a jointly orchestration to provide E2E network transport services for the interconnection of the Cloud Computing infrastructure.

2.1. SDN Orchestration

Firstly, we define here network orchestration as the coordination and automation of the establishment and release of multiple independent network connections (usually performed by different control instances) for the provisioning of E2E connectivity services through heterogeneous network domains (which might be composed of different network technologies).

The SDN orchestration approach is based on an hierarchical architecture where a software-based, logically centralized entity provides E2E communication through different transport networks (Ethernet/DWDM) and/or control technologies (SDN/OF, GMPLS/PCE). This hierarchical approach has been proposed before in [8] where multiple OpenFlow-enabled PCEs are orchestrated by a parent-PCE to provision E2E connections. In [9], the ABNO architecture is used in a international Testbed to orchestrate E2E connectivity services across several optical network transport technologies (Optical Packet Switching - OPS, Elastic Optical Networks - EON [10] and WSON) controlled by different SDN controller distributions.

Differently, in this work we explore the orchestration of a network scenario composed by SDN/OF network domains interconnected by a GMPLS/PCE controlled transport network, with a logically centralized Network Orchestrator (NO). Network domains are interconnected by border links shared between two nodes which belong to different domains. This heterogeneity, among the control technologies considered, introduces the need of implementing multiple southbound interfaces in the NO and also introduce new elements of discussion on the design of the orchestration architecture (i.e. level of topology abstraction and multi-layer path computation). Now, we propose to outline the set of requirements that should be fulfilled by the orchestration layer in the previously presented scenario.

1. Translation of the external application connectivity service requests to the configuration of the different control plane instances. Definition of a standard and extensible northbound API to support customer service requests and offering network control abstraction to customer applications.
2. Discovery and inventory of the control instances and the network topology. Full physical network topology information is not strictly required in the orchestration layer, but at least, the inter-domain connectivity and an abstracted view of the network domains are required.
3. E2E path calculation across the different network domains. Domain selection or full path computation depending on the level of topology abstraction.
4. Provisioning and restoration of the E2E connectivity. Programmability of the different network controllers through specific provisioning interfaces depending on the underlying control technology.

5. Event handling, notification support of changes in the network (failures, topological changes...).

In the following paragraphs we present in detail different protocol/interface alternatives for the design and implementation of a SDN orchestration architecture which fulfill the previous outlined requirements.

Topology discovery.
The NO may compose its network topology by the cooperation of the underlying network controllers which can advertise its intra-domain topologies using different protocol or interfaces. Most of the SDN controllers implement custom RESTful APIs to offer network topology to external applications, but also other possible interfaces are attracting a lot of interest, this is the case of the NETCONF protocol [11] and its RESTful based version RESTCONF [12], both based on the YANG modeling language [13].

The topology discovery in the orchestration layer can be done in a reactive or proactive manner. In the proactive approach, the NO recovers the network topology to the control plane instance every time it need to refresh its working copy to perform a new path computation. RESTful interfaces are connection-less interfaces and they do not inherently support asynchronous notifications. This feature may constrain some orchestration implementations which employ RESTful as a topology discovery interface, to operate only in a proactive manner.

Another alternative to expose the network topology by a control instance is the northbound distribution of the Link State and TE information using the Border Gateway Protocol (BGP-LS). The BGP-LS UPDATE message has been recently extended [14] to allow multi-protocol network layer reachability information (NLRI) and can be used to advertise the link-state topology gathered by the control instance. This approach was experimentally validated in [15] for the advertisement of intra and inter-domain TE Links by child-PCEs (c-PCEs) to a higher parent-PCE (p-PCE) responsible of the inter-domain path calculation. BGP-LS is a session-oriented protocol which can provide to the Orchestration Layer asynchronous notifications about changes into the network topology. A BGP-LS speaker instance is required in both the NC component and into the NO. This approach would represent a possible implementation of the reactive approach of topology discovery.

Regarding the control layer, the preferred solution for most SDN controllers is the combination of the Link Layer Discovery Protocol (LLDP) and OF. This technique consist on sending periodic OFPT_PACKET_OUT messages encapsulating an LLDP packet to the network elements (NE) within the network domain. These packets are forwarded to a specific port according with the forwarding rule included in the OF message. When another OFS receives a LLDP packet and there is no specific flow action for that packet, the OFS executes the OFPT_PACKET_IN action and encapsulating the LLDP message and re-send it to the controller. This way the SDN controller can determine the network topology.

Similarly, in the GMPLS control plane an Internal Gateway Protocol (IGP), i.e. Open Shortest Path First (OSPF) protocol, can be used to share the topological information between GMPLS nodes. The AS-PCE can sniff as well those IGP's packets and building the TED based on the gathered information.

Path Computation.
The E2E path computation can be performed in a different manner depending on the level of topology abstraction and the number of transport layers. Depending on the level of abstraction, the NO may just have access to an abstracted view of the network consisting in the domain connectivity with inter-domain links an abstract representation of the network domains [16]. In this case, the intra-domain path computation is delegated to the lower, per-domain controllers and the NO only performs the domain selection of the controllers involved in the E2E path calculation.

Another alternative is the discovery of the complete physical network topology by the NO. In this case the NO is responsible of calculating the full path across the network, which in the proposed scenario comprises different layers. In the multi-layer scenario, a separate path computation instance (i.e. a PCE) for each layer topology can be employed, or a single path computation instance, with network visibility of all the transport layer topologies, can use multi-layer aware algorithms to calculate the routes.

Connectivity provisioning.

The E2E connectivity provisioning involves the orchestration of different control plane and forwarding technologies. The NO is responsible for the implementation of the interfaces or protocols exposed by the control plane to forward the orders from the orchestration layer to the control layer. Also, the generalization of different transport technology connections into a flexible and common data structure is a key requirement to be able to offer abstracted information to upper layers through the NBI.

Most SDN controllers offers a custom provisioning REST API which input parameters are aligned with those used on the OFPT_FLOW_MOD messages to insert flows into the OF-enabled switching devices. A flow is defined by one or more matching rules which range from Layer 2 to Layer 4 packet headers, and an action (forward to a specific port, drop packet, forwarding to the controller, etc...) inserted into the forwarding device's OF table.

The AS-PCE can instantiate or remove LSPs into the network using the PCEP initiate request message (PCInitiate) [17]. This message includes the endpoints and the computed explicit route object (ERO), defining the route and resources to be traversed and allocated by the LSP. After the connection is successfully established, a PCEP Report Message (PCRpt) is generated to notify to the AS-PCE the successful LSP establishment and its management (e.g., deletion, modifying attributes, etc). The PCEP protocol can be use as a NBI of the AS-PCE to expose a programming interface to the orchestration entity.

2.2. Cloud Computing and Datacenter interconnection

Cloud computing services are becoming an essential part of any enterprise IT infrastructure. Cloud has introduced a new application paradigm where storage and server infrastructures are hosted and shared. These new paradigm has allowed cost reduction and innovation in services and applications. These new innovations are based on sliced servers, introducing a pay-as-you-go model. In this context, the concept of Platform as a Service (PaaS) is introduced as a new service model for networking the future Internet. Thus, it is one of the faster emerging business for Internet Server Providers (ISP).

A data center (DC) refers to any large, dedicated cluster of computers that is owned and operated by a single organization. DCs interconnection is one of the biggest problems that service providers have to face. DCs have been spread geographically to reduce services' latency to the end user, and that has led into an exponential growth on the inter DC traffic [18]. Moreover, the DC interconnection cannot rely upon static, coarse granular and expensive connection pipes, but need to be adjusted to the traffic demand for both Cloud and Network providers can take full advantage of the existing network infrastructure [19].

The integration of the network control and management systems with cloud-based applications relays upon the promotion of open, standard interfaces between network control/orchestration systems and upper applications, which allow gathering the relevant information from the network and directly programming the network behavior. Standard interfaces allow IT and Network applications being developed independently hiding the internal implementation details and offering abstracted services, such as the virtual machine creation/deletion or the E2E connectivity provisioning.

3. Proposed Network orchestration architectures

In this section two different SDN orchestration approaches are described and compared in terms of: their features and characteristics of each approach; the implementation requirements; and control overhead they introduce in the network orchestration process.

3.1. Orchestration based on a single SDN controller (SC-Arch)

Firstly, we start describing an orchestration approach based on a single SDN controller architecture (SC-Arch) (Figure 2). This architecture handles the network orchestration by a single, logically centralized SDN controller. It directly controls the OF network domains and delegates the control of the optical transport network to the AS-PCE.

On top of the SDN controller architecture a NBI interface is responsible of offering the services to external application which must be general and not constrained to control specifics. A SDN controller can implement
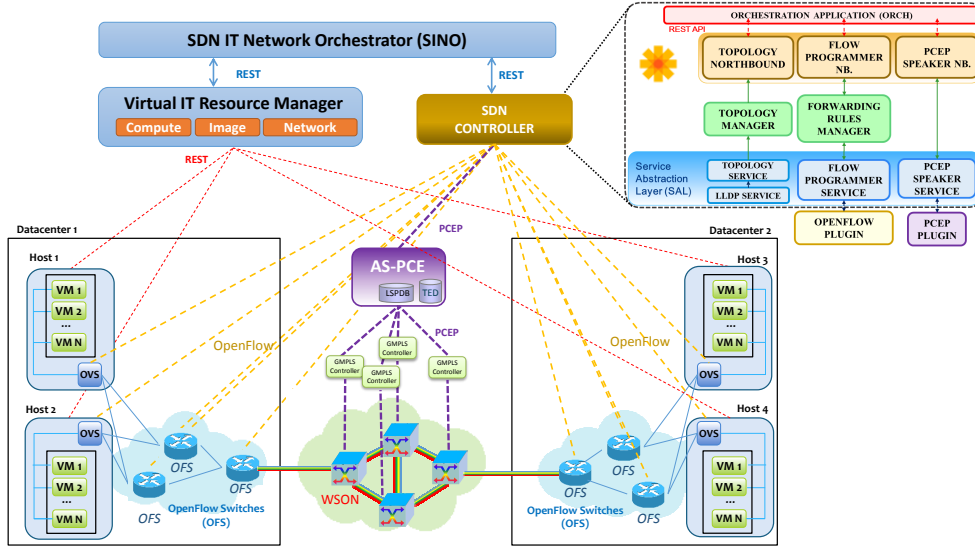
Figure 2: Architecture based on a single SDN controller.

different protocol interfaces as southbound plugins. The SDN controller handles the control technology abstraction internally by means of a Service Abstraction Layer (SAL) which exposes device services to applications hiding specific networking protocol plugins characteristics. The SAL determines how to fulfill the requested service irrespective of the underlying protocol used between the controller and the network devices.

Among the different open-source SDN controllers available, we have chosen the OpenDayLight (ODL) project [20] for implementing the SDN controller. OpenDaylight SDN controller already includes the PCEP protocol as a southbound plugin, however a PCEP-Speaker service has been created to manage the dynamic establishment of LSPs into the GMPLS/PCE control plane. The main building blocks of the internal architecture of the proposed modified ODL controller is shown in the top-right part of Figure 2.

In our implementation, the SDN controller does not recover the whole L0 optical network topology (BGP-LS could be used to recover this information), the optical domain topology is abstracted as a single-node, whose ports, represent the border nodes of the optical transport network. The inter-domain connectivity is dynamically loaded from a external configuration file when the SDN controller starts. Consequently, the path computation, inside the optical domain, is delegated to the AS-PCE.

On top of the SC-Arch, an internal orchestration application (ORCH) has been created in order to handle the required mechanisms to solve the E2E provisioning in the multi-layer network. The E2E provisioning workflow diagram can be seen in Figure 3.

The ORCH recovers the network topology from the Topology Manager Service (TMS). It is also responsible of the path calculation between the E2E connectivity request's endpoints. If the ORCH fails to obtain a feasible route in the network, it explores the inter-domain connectivity information to detect if the E2E service request failed due to the lack of connectivity through the optical domain. If so, a new virtual L0 connection is established the ORCH sends a Path Computation Request (PCRequest) to the AS-PCE to obtain a feasible route between the border OFS. After a successfully response from the AS-PCE, a LSP establishment request with the pre-calculated route encoded into an Explicit Route Object (ERO) is sent to the AS-PCE through the PCEP-Speaker service. After the AS-PCE notifies the PCEP-Speacker Service the effective LSP creation, the new links is discovered by the Topology Manager between the disconnected OFSs through the LLDP mechanism described in the sub-section 2.1.

Finally, the SDN controllers provisioning the E2E connection by calculating the route into the OF network topology managed by the TMS between the E2E request's endpoints and it configures the OFS forwarding tables through OFPT_FLOW_MOD messages containing the match rules (source and destination hosts MAC addresses) and the action (input port/output port).
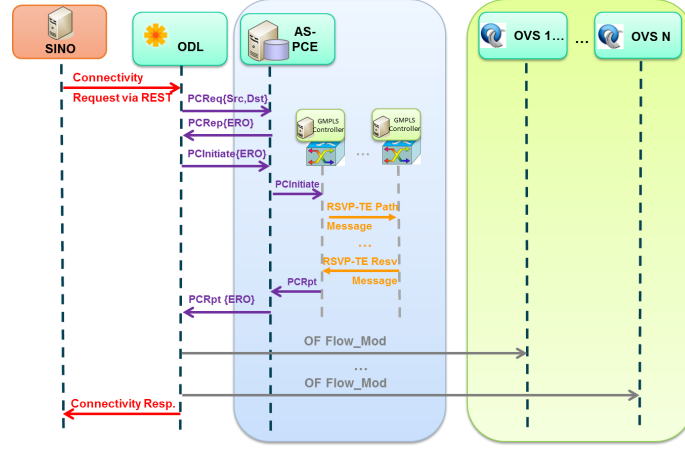
6

Figure 3: E2E provisioining workflow with Single SDN controller orchestrtion achitecture.

3.2. ABNO architecture for multi-domain, multi-controller orchestration

In this section, we proposed a network orchestration approach based on the ABNO architecture orchestrating a dedicated control instance for each transport network: a SDN controller and an AS-PCE (Figure 4). Firstly, the different modules that conforms the ABNO architecture are presented, describing how this approach fulfills the set of requirements enumerated in the sub-section 2.1.
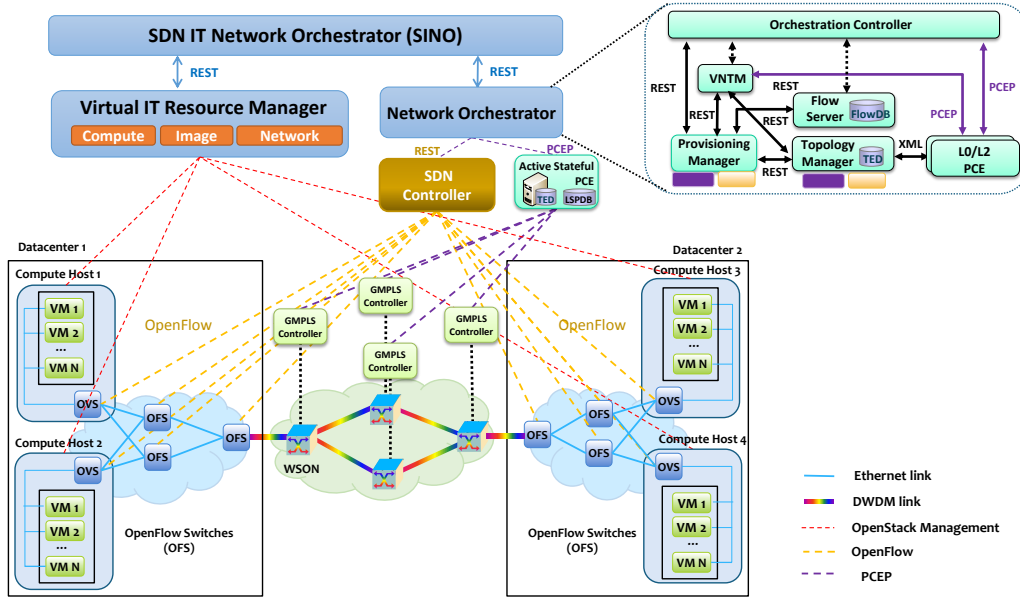

Figure 4: ABNO architecture for multi-domain, multi-controller orchestration

On top of ABNO architecture, the Orchestration Controller (OC) handles the internal workflow between different ABNO modules and it processes the incoming requests from the NBI. The NBI offers, to SINO an other external applications, the CRUD operations for the E2E connectivity service. It has been implemented by a RESTful API.

The Topology Manager (TM) is the component responsible of gathering the network topology from each control domain and building the TED. The TED includes all the relevant information about network links and nodes, and it is used by the Path Computation Element (PCE) for calculating routes across the network. In our implementation, the TM recovers the physical network topology of each network domain and the inter-domain connectivity from Summary. From this information, the TM builds a complete multi-domain topology and a separated topology, built by filtering the whole topology based on the TE information of the links, for each transport layer technology. The TM implements dedicated plugins for: (1) the custom SDN controller RESTful interface; (2) a proprietary interface of the AS-PCE based on a raw socket TCP and XML/JSON encoding. The topology discovery is performed in a proactive manner.

The Virtual Network Topology Manager (VNTM) is the responsible of the multi-layer management. In the proposed scenario, the VNTM is in charge of the set-up of Layer-0 (L0)/DWDM optical connections to satisfy upper layer's connectivity demands. For instance, if a Layer-2 (L2)/Ethernet connectivity service request cannot be served because there is not an available route through the L2 network topology, the VNTM manages the set-up of a L0 connectivity request between the network domains to which belongs the L2 endpoints. To do this, the VNTM requests the inter-domain topology to the TM, to find the L0 border node pair between whom request the connection. After the successfully establishment of the L0 connection, the VNTM notifies the TM the creation of a new virtual link and the related virtual ports on the border nodes into the L2 network topology. The mapping of the L0 connections and L2 virtual links is responsibility of the VNTM too.

The Provisioning Manager is the module which translates the connectivity requests, processed by the ABNO controller and the VNTM, into the corresponding provisioning request messages of the underlying network controllers. The Provisioning Manager implements a provisioning plugin for each different network controller connected to it. In the proposed architecture it implements the custom SDN controller's provisioning REST API and the PCEP with Stateful and PCE-initiated LSP Setup extensions, for the communication with the AS-PCE. All the established connections (both L0 and L2) are stored in the Flow server by the provisioning manager. The Connection is the data structure exchanged between the different ABNO components in our implementation, and it consists on the following parameters:

- **Endpoints**. Source and destination nodes, described as: {Router_ID, Interface_ID}

- **Path**. List of hops traversed by the connection, each one described as: {Router_ID, Interface_ID}

- **Transport_layer**. (L0, L2)

- **Forwarding_rules**. Matching_Rules and Action similar to the OpenFlow equivalents [2]

- **Connection_type**. (Unidirectional, Bidirectional)

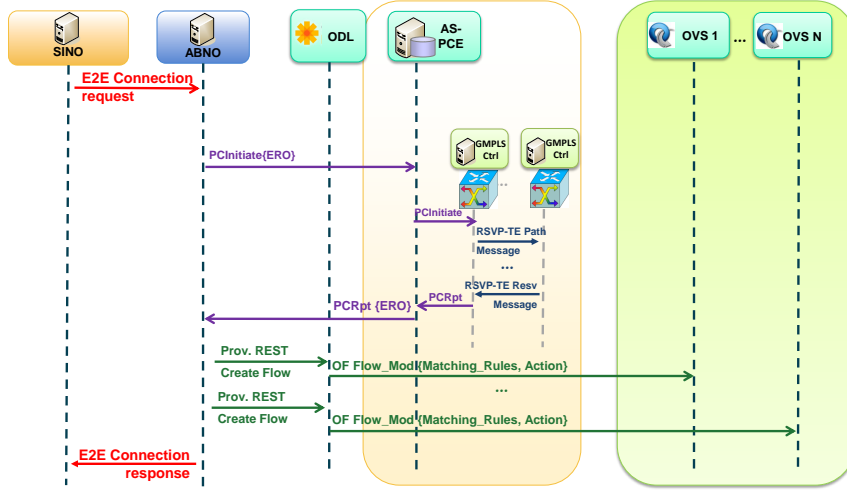Figure 5: E2E provisioining workflow with ABNO orchestrtion achitecture.

The provisioning workflow of an E2E service, involving the inter-DC network, can be seen in Figure 5. The internal ABNO workflow of the E2E service involves the creation of a virtual link by the VNTM through the PCEP plugin. A PCInitiate message is sent to the AS-PCE to setting up the creation of a new LSP to interconnect the OF network domains. This request includes a pre-calculated path through the optical network and the border node's ports connected to the border links which interconnects the OF domains with the optical transport network. These Endpoints are represented as Unnumbered Interfaces composed by the router-ID and the Interface-ID encoded in 64 bits. After the creation of the virtual links the OC request the path computation into the L2/ Ethernet layer and sends the Provisioning Requests to the SDN controller through the Provisioning Manager interface. These requests contains the MAC addresses of the source and destination E2E service Endpoints as matching rules for the OFSs configuration.

3.3. Orchestration approaches comparison

After describing the two orchestration approaches we present in this work, we summarize and evaluate the differences between them in Table 1:

9

	ABNO	Single SDN controller
Features		
Topology discovery	Full topology discovery, proactive mode	Abstract optical domain topology, reactive mode
Provisioning Connections	Proactive	Proactive
Path Computation	Dedicated PCEs for each layer, full path computation	Domain selection, intra-domain path computation delegated to control instances
Interfaces		
Provisioning: Topology discovery:	REST, PCEP REST, BGP-LS	OpenFlow, PCEP LLDP
Network control overhead	Internal PCEP communication, REST communication between ABNO components	

Table 1: Network orchestration approaches summary

In the first place, the two proposed orchestration solutions follow a similar hierarchical approach. The main architectural difference between them, is the complete separation between control and orchestration layers in the ABNO approach. The ABNO architecture allows delegating some of the control tasks to underlying control instances, and this feature let us to think that it will present less scalability problems in large network deployments, in comparison with the SC-Arch architecture. Regarding the SC-Arch approach, the control and orchestration are made by the same SDN controller combining the direct control of network equipment with the integration of other control instances without disclosing the internal domain network details.

Regarding the implementation of these solutions, the E2E service provisioning in a proactive manner in both cases. Topology discovery is made proactively within the ABNO-based architecture by implementing custom REST interfaces between the TM and the per-domain controllers and reactivelly in the SC-Arch approach. While the ABNO architecture recovers the complete, non-abstracted, physical topology, the SC-Arch approach recovers only an abstracted view of the optical domain and delegates the intra-domain path computation to the AS-PCE. Finally, let us to mention that the ABNO-based orchestration approach introduces some control overhead by implementing each ABNO architecture component as independent entities and using PCEP and REST interfaces for the communication between them.

3.4. Cloud Computing integration with SDN orchestration architectures

In this section we are going to describe the integration of the Cloud Computing platform with the already presented SDN orchestration architectures. The purpose is to provide integrated IT and Network resources orchestration.

The Cloud Controller is named Virtual IT Resource Manager (VIRM) platform, which is responsible of the management the creation/ migration/deletion of VM instances (computing service), disk images storage (image service), and the management of the VM network interfaces (networking service). The VIRM handles multiple compute servers geographically distributed in the DC locations.

The computing service manages the VM into the compute hosts (Hosts 1-4 in Fig.2 and Fig.4). A compute service agent is running in each host and controls the computing hypervisor (e.g., KVM) used for the creation/deletion of the VMs. The image service handles the disk images which are used as templates

for VM file systems. The connectivity between VMs and virtual switches within the hosts is managed by the networking service. The networking service is responsible for the handling of all network elements, such as switches, firewalls, etc. We propose to modify the networking service, in order to limit the networking services to the creation of the virtual interfaces, the attachment of the virtual interfaces to the virtual switches and finally the offer of a DHCP service for the VMs to get the assigned IP address. The NO will be the responsible for the control and management of the connectivity services between the different network domains. The SINO is introduced in order to coordinate the VIRM and the NO.

The SINO controls the VIRM through a RESTful API, used to both trigger the VIRM actions and get the necessary information about the running VM instances.

The SINO has been previously presented in [21], for the orchestration of IT and Network resources. The SINO exposes a variety of services, among which are: VM Create, Read, Update, Delete (CRUD) operations; E2E CRUD connectivity; and VM migration. It exposes all these services through a NBI implemented through a RESTful API. In section 4 we use the SINO for the selected use cases to evaluate the performance of the two NO architectures described in the following sections

4. Experimental validation

The experimental validation of the previously described architectures has been carried out by the implementation of two typical uses cases demanded by Cloud Computing applications: E2E connectivity provisioning and seamless virtual machine migration.

The SDN/NFV cloud computing platform and transport network of the ADRENALINE Tesbed, which consists on the SINO, the VIRM and the NO, is the reference scenario used for obtaining all the experimental results that will be shown below. The overall architecture was described in section 3.4 and shown in Figures 2 and 4.

The cloud computing platform has been implemented using the Openstack Havanna release, specifically, the VIRM corresponds to the Cloud Controller node in the OpenStack nomenclature. The VIRM controls four Compute Hosts located in two different Data Center (DC) locations, which are responsible of storing the VMs created in the cloud. All cloud components components has been deployed into physical servers with 2 x Intel Xeon E5-2420 and 32GB RAM each.

The proposed network scenario consists of two intra-DC network domains based on Ethernet transport technology and an inter-DC optical transport network. Each intra-DC network consists of four OpenFlow Switches (OFS) deployed using standard Custom Off The Shelf (COTS) hardware and running OpenVSwitch (OVS), which are all controlled by a single SDN/OpenFlow controller (OpenDaylight in both architectures). Each DC border switch has been implemented using COTS hardware, a 10 Gb/s XFP tunable transponder and OVS technology. The inter-DC network is a GMPLS-controlled Wavelength Switched Optical Network (WSON) which consists on: 2 ROADMs and 2 OXCs providing re-configurable (in space and in frequency) E2E lightpaths, deploying a total of 610 km of G.652 and G.655 optical fiber, with six DWDM wavelengths per optical link.

4.1. E2E connectivity service

For this use case, four different VMs has been deployed on each of the compute hosts distributed in the two DC locations. The experiment will measure the setup delay introduced, by both orchestration architectures, by providing an E2E channel between two of these VMs. Each service request received in the NO demands a L2 connectivity service between the OFS' ports connected to two randomly chosen VMs. Each connection created in the network is removed before the next service request arrives, thus the blocking probability is theoretically zero and all requests have uniform network conditions.

In Figure 6, the setup delay distribution of 150 requests experiment are shown for both architectures. The results show two different groups of setup delays in both architectures: a set of L2 service requests between co-located VMs (same DC) provisioned in less that 1.5 seconds; and another group involving the setup of an optical connection into the GMPLS domain (inter-DC communication), the setup delay in this case varies in the range of 2.5 to 5 seconds.

11

Figure 6: E2E workflow: a)Setup delay distribution under the ABNO-Arch, b)Setup delay distribution under the SC-Arch

Architecture	ABNO		Single SDN controller	
Request group	Intra-DC	Inter-DC	Intra-DC	Inter-DC
Mean	0.945s	2,867s	0,681s	3,730s
Std.Dev*	0,091s	0,142s	0,178s	0,589s
CV**	0,111	0,053	0,291	0,161

* Standard Deviation
** Coefficient of variation = (Std.dev) / (Mean)

Table 2: Results Statistics

If we analyze the Intra-DC requests results (Table 2), we can see better performance achieved by the SC-Arch (lower mean value). This result is motivated by the fact that ABNO introduces more orchestration overhead mainly by the usage of the PCE in the path computation.

The inter-DC group of requests shows more differences. Firstly, mean value of setup delay under ABNO architecture is almost one second less, however, it is remarkable the big dispersion obtained in the SC-Arch results (Std.Dev = 0,589 and CV = 0.16). The dispersion of the setup dalay in the SC-Arch is caused by the responsiveness of the SDN controller to the effective connectivity through the inter-DC connection. In the other side, when ABNO NO receives the response from the GMPLS control plane of the effective establishment of the LSP starts the Virutal Link creation and the provisioning of the L2 connection. Therefore it shows a more deterministic and predictable behavior and globally better performance than the SC-Arch.

4.2. Seamless Virtual Machine Migration

The second use case proposed to evaluate the overall network orchestration architectures is the seamless migration of a VM. In the proposed scenario, the VM1 is connected to the VM2 (both initially running into the DC-1), the experiment has been designed to migrate VM1 to DC-2. The migration comprises the use of the inter-DC network and it will be evaluated by measuring the disruption time of the connectivity between VM1 and VM2. The process involves six steps, detailed in Figure 7.a:
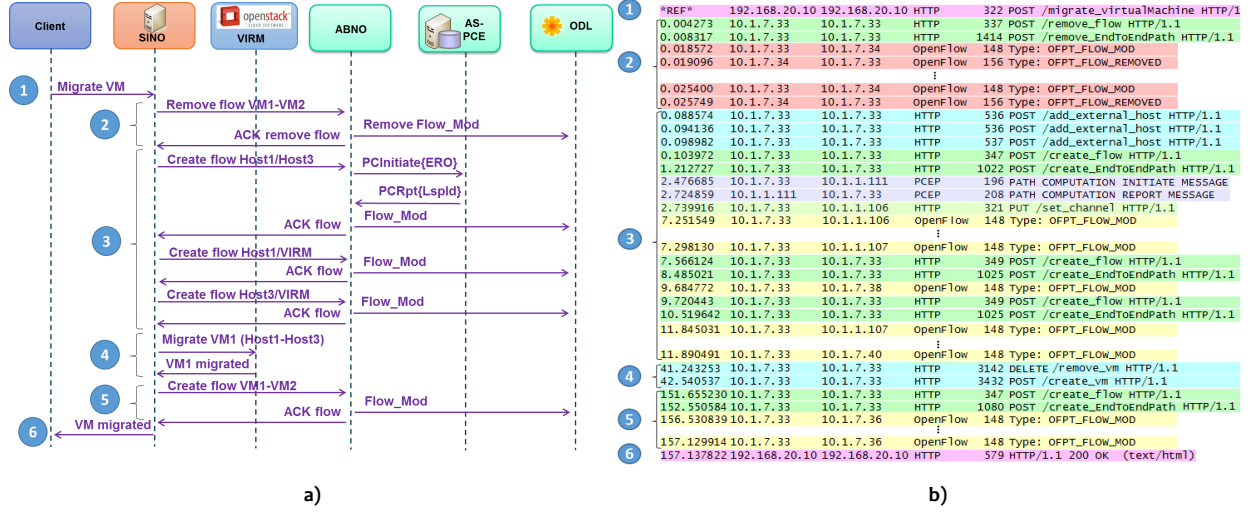
12

Figure 7: a) VM migration flow diagram; b) Wireshark capture of SINO commands.

1. *Virtual machine migration request trigger*: SINO receives a VM1 migration request through HTTP request.

2. *Deletion of VM connections*: all the flows established in the network in which the target VM is involved must be removed before start the migration. The SINO requests to NO the deletion of all flows involving VM1. The NO deletes all of them by sending the orders to the SDN controller which at once will send the OFPT_FLOW_REMOVED messages to the OFSs.

3. *Inter-DC connectivity*: the VIRM, the source and destination hosts must be interconnected during all the migration process in order to exchange the VM images data using the VIRM image service. The SINO adds each of the three hosts (VIRM, Host1 and Host3) as external hosts into the NO, and request an end-to-end path between each pair. The inter-DC connectivity between hosts involves the establishment of a Label Switched Path (LSP) into the WSON by sending a Path Computation Initiate Message to the AS-PCE.

4. *IT migration*: First, VM1 is paused and snapshot in the source host and it is saved into the VIRM. Once the VM1 snapshot is ready, VM1 is removed and a new instance based on VM1 snapshot is deployed into destination server.

5. *Restore VM connections*: the SINO restores the connections between the migrated VM and all the rest of VMs to which it was connected. The SINO requests a new flow between VM1 and VM2 to NO. The NO's PCE computes a path, but since the previous optical LSP between DCs is reused, it is only necessary to create the flows of the OFSs.

6. *Acknowledgement of virtual machine migration*: The SINO notifies the client of VM1 successful migration

Figure 7.b shows the network traces captured at the SINO's server. It includes the HTTP communication through the REST interfaces between SINO, VIRM and NO. These modules are running in the same server and that is why the IP address is the same in the messages exchanged between those entities. For the network orchestration, it can be seen how the NO request the AS-PCE the creation of a LSP through the optical network for the inter-DC communication and finally, the OFPT_FLOW_MOD messages are send first, for the deletion of flows and the for creating them again at the end of the instance migration.

Figure 8: a) VM migration traffic (packets/s) received in DC2 over time; b) ICMP traffic capture between VM1 and VM2 during VM1 seamless migration.

This use case has been evaluated only with the ABNO architecture. The most significant part of the disruption time of the migration process is introduced in the IT orchestration (VM snapshotting and creation) and not in the network orchestration this is the reason only the ABNO experimental validation results have been shown. As it has been shown in the previous subsection, the differences between E2E connectivity provisioning of both architectures are in the order of hundred of miliseconds. These differences are not relevant, in the scope of the whole disrupting time measured in the migration experiment (Figure 8), to make a comparison in these terms for this use case. However, except for the network orchestration workflow and the migration workflow would be the same in both cases.

In Figure 8.a is shown the incoming traffic to DC-2 sampling each 5s the packets received in the border OFS of DC-2. A significant increase of the number of packets occurs after the trigger of creation of a VM1 into Host3 due to the download of the VM1 image from VIRM. Finally, a traffic capture of the dialog between VM2 and VM1 is shown in Figure 8.b which provides the total disruption time occurred due to the seamless migration process. We can observe that after a downtime of 158s the connectivity between VM1 and VM2 is recovered. It can be observed that the MAC address of VM1 has changed due to the block migration.

5. Conclusions and Future Work

This work has present an in-depth discussion about the network orchestration problem in multi-domain networks comprising different transport technologies and control planes. The integration between network orchestration and potential client applications (i.e cloud computing) has also been considered by presenting IT and SDN orchestration architecture (SINO-VIRM-NO) and validated through all the experiments carried out in the study.

Two possible network orchestration approaches: ABNO and SC-Arch, have been discussed and compared through the definition of their architectures and implementations. They have been experimentally validated by measuring the setup delay introduced in the orchestration process of the provisioning of E2E connectivity services in the reference scenario.

ABNO architecture has presented more deterministic and predictable results in the setup delay distribution measured in the E2E connectivity use case, than the SC-Arch approach. In terms of absolute performance the results has not shown significant differences. ABNO architecture also provides a more scalable solution for the orchestration problem because completely separates orchestration and control layers. This

14

separation can really matters in network scenarios where each domain belongs to different administrative entities which may not want to disclose their internal information with other management systems.

To continue the work presented in this article, the integration of resilience mechanisms in the orchestration architecture will be a necessary step to provide a more close to reality solution for the problem assessed here, the reactive approach on topology discovering is a key-element for this integration. Another important consideration is the heterogeneously of the interfaces to communicate the Orchestration Layer and the Control instances, a standardized solution for a common interface is needed when different SDN controllers (provided by different providers) will be considered for the control of multiple SDN domains.

Acknowledgment

References

[1] T. Benson, A. Akella, and D. Maltz,"Network traffic characteristics of data centers in the wild" in Proceedings of the 10th annual conference on Internet measurement. ACM, 2010, pp. 267?280

[2] Open Networking Foundation, "Open Networking Forum: OpenFlow Switch Specification", ONF v.1.3.0, June, 2012.

[3] E. Crabbe, I. Minei, J. Medved, R. Varga, "PCEP Extensions for Stateful PCE", *draft-ietf-pce-stateful-pce-09*, IETF, 2014.

[4] R. Casellas, R. Martinez, R. Muñoz, R. Vilalta, L. Liu, T. Tsuritani, and I. Morita, "Control and Management of Flexi-Grid Optical Networks with an Integrated Stateful PCE and OpenFlow controller [Invited]", in *Journal of Optical Communications and Networking* Vol. 5, No. 10, pp. A57-A65, 2013.

[5] D. King, and A. Farrel, "A PCE-based Architecture for Application-based Network Operations", IETF draft *draft-farrkingel-pce-abno-architecture-14*, 2014.

[6] Fielding, R.T. (2000), "Architectural Styles and the Design of Network-based Software Architectures", Ph.D. Thesis. University of California, Irvine, EEUU.

[7] A. Mayoral, R. Vilalta, R. Muñoz, R. Casellas, R. Martinez, "Experimental validation of automatic lightpath establishment integrating OpenDaylight SDN controller and Active Stateful PCE within the ADRENALINE Testbed", in *International Conference on Transparent Optical Networks (ICTON)* 6-10 July 2014, Graz (Austria).

[8] R. Casellas, R. Muñoz, R. Martínez, R. Vilalta, L. Liu, T. Tsuritani, I. Morita, V. López, O. G. de Dios, and J. P. F. Palacios, "Sdn based provisioning orchestration of openflow/gmpls flexi-grid networks with a stateful hierarchical pce", in Proc. of Optical Fiber Communication Conference and Exposition (OFC), San Francisco, California, USA, March 2014.

[9] Y.Yoshida, A. Maruta, K. Kitayama, M. Nishihara, T. Tanaka, T. Takahara, J. C. Rasmussen, N. Yoshikane, T. Tsuritani, I. Morita, S. Yan, Y. Shu, M. Channegowda, Y. Yan, B.R. Rofoee, E. Hugues-Salas, G. Saridis, G. Zervas, R. Nejabati, D. Simeonidou, R. Vilalta, R. Muñoz, R. Casellas, R. Martínez, M. Svaluto, J. M. Fabrega, A. Aguado, V. López, J. Marhuenda, O. González de Dios, J. P. Fernández-Palacios, First international SDN-based Network Orchestration of Variable capacity OPS over Programmable Flexi-grid EON , in Proceedings of PDP Optical Fiber Conference (OFC), 14 March 2014.

[10] International Telecommunications Union, "Spectral grids for WDM applications: DWDM frequency grid", ITU-T Recommendation G.694.1, February 2012.

[11] R. Enns et al., "Network Configuration Protocol (NETCONF)", IETF RFC 6241, Jun 2011.

[12] A. Bierman , M. Bjorklund, K. Watsen, and R. Fernando, "RESTCONF Protocol", draft-ietf-netconf-restconf-04, February 2014.

[13] M. Bjorklund, "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", IETF RFC 6020, Oct 2010.

[14] Gredler, H., Medved, J., Previdi, S., Farrel, A., and Ray, S., "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution, January 2015.

[15] M. Cuaresma, F. Muñoz, S. Martinez, A. Mayoral, O. Gonzalez de Dios, V. López, J.P. Fernández-Palacios, "Experimental Demonstration of H-PCE with BGP-LS in elastic optical networks", in European Conference on Optical Communication (ECOC), We.4.E.3, Sep 2013.

[16] R. Casellas, R. Martinez, R. Muñnoz, L. Liu, T. Tsuritani, I. Morita, and M. Tsurusawa,"Dynamic virtual link mesh topology aggregation in multi-domain translucent wson with hierarchical-pce", OSA Optics Express Journal, vol. 19, no. 26, December 2011.

[17] E. Crabbe, I. Minei, S. Sivabalan, R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", draft-ietf-pce-pce-initiated-lsp-03, IETF, 2014.

[18] "Cisco Global Cloud Index: Forecast and Methodology 2013-2018". `http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html`

[19] R. Nejabati, S. Peng, D. Simeonidou, "Role of Optical Network Infrastructure Virtualization in Data Center Connectivity and Cloud Computing", in Proc. of Optical Fiber Communication Conference (OFC), OM3E.1 2013,

[20] OpenDaylight, A. "Linux Foundation Collaborative Project."

[21] A. Mayoral, R. Vilalta, R. Muñoz, R. Casellas, R. Martinez, J. Vílchez, "Integrated IT and Network Orchestration Using OpenStack, OpenDaylight and Active Stateful PCE for Intra and Inter Data Center Connectivity", in *European Conference on Optical Communications*, We.2.6.6, Cannes, 2014.