Interoperability of Lightpath Provisioning Systems in a Multi-domain Testbed

Alfred Wan, Paola Grosso, and Cees de Laat

University of Amsterdam – System and Networking Engineering research group
Science Park 107, 1098 XG Amsterdam, The Netherlands
{wan,p.grosso,delaat}@uva.nl

Abstract. On-demand services are a key feature of Future Internet architectures. Already today research networks around the world provide dedicated optical circuits (lightpaths) to scientists, to offer a higher quality of network service. Tools to request and instantiate these lightpaths within a single domain exist, but interoperability of the various provisioning systems in a multi-domain scenario is still in its infancy. We present here our work for end-to-end multi-domain circuits reservation and setup. We worked in a large-scale research testbed composed by different provisioning systems, spanning multiple network domains. We developed translation modules at the provisioning system boundaries that allow for seamless path reservation and setup. We conclude identifying the main elements for interoperability and provide some guidelines for future integration of research network provisioning systems.

Keywords: network provisioning systems, service and control planes, federated testbeds.

1 Introduction

Future Internet architectures will certainly provide on-demand network services, away from the one-size-fits-all feature of the current Internet, and toward a tailored approach to users and applications requirements.

Network researchers are focusing on this issue, backed up by several funding efforts. The European Commission promotes the Future Internet Research and Experimentation (FIRE) initiative [1]. The National Science Foundation in the USA does the same with the Global Environment for Network Innovation (GENI) project [2].

Research and Education Networks (NRENs) around the world have tackled the problem in the last years. Their answer has been to offer separated network circuits to users, called *lightpaths*. They have built hybrid network architectures, with coexisting IP and lightpath services in the same physical infrastructure. The SURFnet6 network in the Netherlands is an example: it has offered use of lightpaths since 2006 to the Dutch research community and academic centers.

T. Magedanz et al. (Eds.): TridentCom 2010, LNICST 46, pp. 412–427, 2011.

e-Science and Grid applications are the driving force behind this network model. These applications rely on guaranteed bandwidths and fixed latencies; concurrent traffic disturbs or, worse, completely disrupts their functioning. Traditional routed IP services are not sufficient, and scientists clearly profit from having a dedicated network path. For example, the high-energy physics experiments at the Large Hadron Collider at CERN [3] distribute their data to the computing centers via lightpaths; in very long interferometry radio astronomy experiments [4], telescopes also send the recorded signals to a correlation center via lightpaths.

We foresee that many of the solutions that satisfy scientists in the somehow closed world of research networks will be adopted in other communities, for large and small businesses and ultimately single users.

Requests for circuits within a single NREN are nowadays relatively easy to satisfy; each domain provides a reservation system to its end users, while ad-hoc software tools instantiate and configure the paths. The challenges arise whenever the circuits need to span multiple network domains, for the lack of interoperability between the various systems.

We present in this article the software tools we develop to facilitate interoperability along a multi-domain lightpath chain; we present the results obtained in several experiments; based on these we outline necessary future work to further ease the use of multi-domain lightpaths.

2 Provisioning Systems

Dedicated circuits in hybrid networks can be implemented at various network layers. Layer1 circuits are carried over a DWDM infrastructure, where different services run on different wavelengths [5]. Layer2 circuits are Ethernet paths extending into the WAN. A multi-domain circuit may comprise several segments at different layers, where (de)adaptation functions (de)encapsulate the data between layers [6][7].

Lightpaths can be static and persist for very long periods of time, to form an optical private network. They can also be dynamic and short-lived, being automatically unconfigured once the reservation time has expired.

A lightpath provisioning system always performs two major tasks, independently from the characteristics of a circuit:

- It provides the interfaces for path reservations to the end users. The user can make requests without being aware of the underlying technologies used e.g. SONET or Ethernet. This component constitutes the Network Service Plane (NSP);
- It configures the network equipment. These are the internal components that are vendor and equipment specific. We refer to them as the Control Plane (CP).

2.1 Network Service Planes

We identify four main components in a NSP, as shown in Fig. 1.

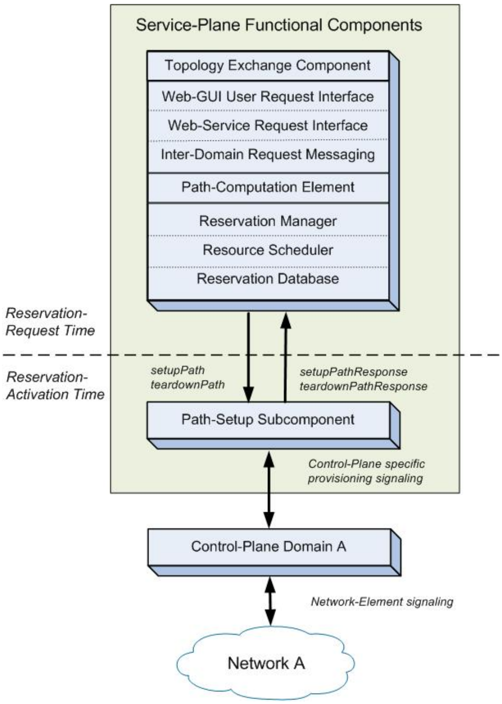

Fig. 1. The four components of the Network Service Plane. From top to bottom: the Topology Exchange Component, the Reservation Request Interface, the Path Computation Element and the Reservation Management System.

This model is based on observations of the design of existing provisioning systems. The User-Controlled LightPath (UCLP) provisioning system was one of the first tools that provided on-demand dynamic lightpaths. Later versions of UCLP [8] and other independent efforts rendered designs with common functionalities though with slightly different characteristics.

We distinguish between components that operate at reservation request time and components that operate at reservation activation time.

At reservation request time we have:

- A *Topology Exchange Component.* This exposes the local topology and facilitates topology exchange with other domains, a step necessary to setup multi-domain paths;
- A *Reservation Request Interface*. This allows users and neighboring domains to make reservation requests for paths in the local domain, and for the local domain

torequest reservations in other domains. In the most common scenario, a multi-domain path request originates from a user that utilizes a Web-GUI (*Web-GUI User Request Interface*). A user can also use a Web-Service client to do this (*Web-Service Request Interface*). If a domain receives a request that involves another domain, the NSP has to forward it to that domain and check its status using an *Inter-Domain Request Messaging* component;
- A *Path Computation Element.* This calculates a path between two endpoints specified in the request given the topologies of the domains involved;
- A *Reservation Management System*. This checks the availability of resources when it receives a request, schedule resources when a request can be accommodated and stores it in its local database.

The component that bridges the NSP and CP is the *Path-Setup Subcomponent* (PSS). It operates at reservation activation time. The PSS translates path-setup requests that carry the general request parameters, e.g., bandwidth, duration, VLANs, etc., to CP specific messages, e.g., path-refresh (RSVP-TE) messages that it then sends at regular intervals to the CP until the reservation ends.

3 The Combined Harmony-IDC Testbed

NRENs have developed independently their provisioning systems. These efforts resulted in software suites that cannot immediately interoperate. Although the purpose of each NREN NSP is the same, the functionalities it offers to the user differ; the communication between the NSP and the CP, for example the way to specify time and bandwidth parameters, or the way the communication is secured, vary in all implementations.

Several initiatives are underway to specify how different NREN NSPs should interact to successfully create inter-domain connections. One example is the Network Service Interface (NSI) workgroup in OGF [9]; one of its goals is to provide the interface between network domains for interoperability in a heterogeneous multi-domain environment. However, these initiatives attempt to formulate specifications from a top-down perspective, while the NSP designs of the various NRENs are often idiosyncratic and it is therefore challenging to unify their functions.

We have opted instead for a bottom-up approach and choose a specific usecase to determine what it takes to create multi-domain NSP interoperability and data-path setup.

We performed our research and development in a large-scale testbed, consisting of network resources from several Phosphorus [10] partners, UvA, SURFnet, i2CAT and VIOLA, and of a DICE partner [11], Internet2. Fig.2 depicts the testbed.

Our goal was to create a working lightpath starting in the Internet2 network and ending up in one of the Phosphorus domains, and vice versa.

In our testbed the i2CAT and VIOLA domains are controlled by the Harmony provisioning system. The UvA/SURFnet and the Internet2 domain are under control of the IDC provisioning system.

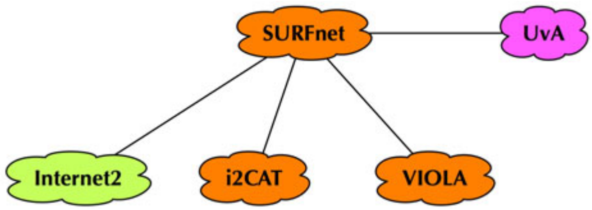

Fig. 2. The combined Phosphorus-DCN testbed where we conducted our work. We show a high level overview of the interconnected domains. Concerning the Network Service Plane, Internet2, SURFnet and UvA are under control of an IDC whereas i2CAT and VIOLA are under control of Harmony. SURFnet runs DRAC, i2CAT runs ARGIA and VIOLA runs ARGON.

3.1 Harmony

We participated in the development of Harmony, the NSP of the EU-funded project Phosphorus. The project has ended, but the network infrastructure is still available for experimentation.

The Phosphorus project aimed to make applications aware of the available computational and networking resources, and to provide dynamic, adaptive and optimized use of heterogeneous network infrastructures connecting these resources. In the project network domains connected to each other via inter-domain lightpaths, and each one ran its own NSP. Three NSPs are in use within the Phosphorus community:

- ARGIA, a commercial system used in the i2CAT network in Spain. It provides time slices or lightpaths to users or organizations, by virtualization of the network resources. An end user can create circuits on-demand based on the applications needs;
- ARGON, an MPLS/GMPLS enabled system used in the VIOLA testbed in Germany. ARGON stands for Allocation and Reservations in Grid-enabled Optical Networks; it provides interfaces for lightpath reservations;
- DRAC, a Nortel commercial product used in the SURFnet network in the Netherlands. DRAC is the Dynamic Resource Allocation Controller that also fulfills the purpose of lightpath reservation.

Harmony provides the APIs to control these three different NSPs.

3.2 IDC

We also contributed to the Inter-Domain Controller (IDC), the provisioning system of the DICE partners. The IDC is an architecture that facilitates both intra-domain path provisioning as well as the creation of circuits across domains. The IDC protocol defines formats and exchanges for:

- Network resource reservation requests messages;
- Intra-domain topology descriptions;

- Inter-domain topology exchange messages;
- Intra-domain network-element signaling and inter-domain request forwarding.

All functional components processing these requests and descriptions can be accessed through Web-Service (WS) interfaces. The WS interfaces can be categorized into:

- User to Network Interface (UNI) that provides users means to interact with the reservation system OSCARS;
- Internal Network to Network Interface (INNI) through which the IDC interacts with local resources, such as network elements;
- External Network to Network Interface (ENNI) that allows other IDCs to interact with the local IDC facilitating multi-domain path reservation and provisioning [12]. The possibility to interact with non-DCN Interdomain Brokers (IDBs) was not present when we started our work. We developed the ENNIs for this purpose as we will describe in the upcoming section.

  Furthermore, IDC interfaces to the DRAGON CP [13], that configures the network equipment to create paths.

4 Service-Plane to Service-Plane Translation Modules

In general, NSPs interpret advance reservation requests for lightpaths in such a way that they are stored in the reservation database unambiguously and consistently. Unambiguity regards the meaning of the parameters pertaining to the reservation, such as endpoint designations, reservation duration and bandwidth specification; consistency avoids conflicts such as overlapping reservations and requests exceeding the available bandwidth.

This also holds for reservation modification and reservation status inquiry requests, albeit in a somewhat lesser degree because while accommodating new reservation requests the NSPs already interpreted, disambiguated and consistently stored them in the reservation database.

When combining heterogeneous networks and NSPs to facilitate multi-domain lightpath reservation and provisioning, the syntax and semantics of the requests may have to be translated in order for the parameters in the requests to conform to the required format of the receiver, and for the requests to have the same operational semantics. For instance, endpoint-designations in the topologies of two domains may be different and have to be translated into the proper format of the receiving domains. In the case of a reservation modification, one domain may allow modifications of a reservation that has already been activated, while in another domain this may not be allowed: this causes differences in the effect of the same operation in different domains. For instance, in one domain a reservation may be cancelled while it is active, whereas in another domain this is not allowed.

We identified two general types of translations necessary at the boundaries between two heterogeneous NSPs, i.e., 1) Web-Service operation translations, and 2) Endpoint translations. For both types we produced code to perform the translations.

4.1 Web-Service Operation Translations

Operation translation requires that input used in a Web Services call in a certain NSP is mapped onto the correct input components in another.

Our first effort was to create translation modules to interface Harmony to IDC. Table 1 shows the operations we implemented.

Table 1. List of supported Web-Service operations. The first column indicates the name of the Web-Service operation. The second column shows if the operation is supported by IDC; the third column indicates support from Harmony. The last and fourth column shows whether we have implemented a translator service for the specific operation.

Operations	IDC	Harmony	Request translator
createReservation	✓	✓	✓
cancelReservation	✓	✓	✓
queryReservation (IDC) getStatus (NSP)	✓	✓	✓
isAvailable (NSP)		✓	
modifyReservation	✓	✓	
listReservation (IDC) getReservations (NSP)	✓	✓	✓
createPath (IDC) activate (NSP)	✓	✓	
refreshPath (IDC)	✓		
tearDownPath (IDC)	✓		

The translation may be:

- A parameter conversion whenever the two NSPs use different data formats to describe the same parameter. For example, start and end time of reservations are expresses in seconds from Epoch in one system and as time-day-month-year in the other;
- Be required to overcome that the results of a Web Service operation produce inconsistent or different results when invoked in the two systems.
- Necessary due to different security mechanisms deployed in the two systems. The IDC requires Web Services operation to be authenticated. IDC identifies requestors with X.509 certificates. Harmony on the other hand allows all users in its network domain to make requests. Our translation module will use a single X.509 certificate for all IDC requests coming from Harmony users.

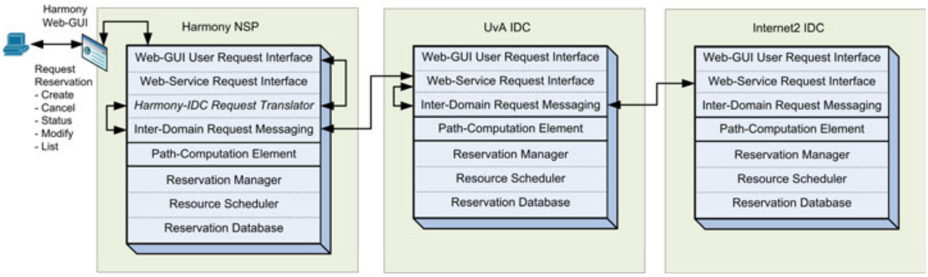

Fig. 3. The NSP components involved in multi-domain reservation requests. In the Harmony NSP (on the left) we insert our translator module; it translates a user request from Harmony format to the IDC format. Conversely, the responses sent to Harmony-NSP are translated from the IDC format to the format required by the Harmony NSP.

Fig.3 shows the insertion of our Harmony-IDC Request Translator among the components present in the Harmony NSP (on the left). The GUI User Request Interface couples to the Request Translator and the Request Translator interfaces internally to the Inter-Domain Request Messaging component.

4.2 Endpoint Translations

To provide interoperability between the NSPs in our testbed we also needed to identify the network endpoints in the proper format. Harmony and IDC express topologies in different manners.

Harmony identifies endpoints with Transport Network Addresses (TNAs). It uses for this a dotted decimal notation. IDC uses instead a URN identifiers for endpoint based on the topology schema developed by the OGF Network-Monitoring Working Group (NM-WG) [14], with specification of domain:node:port:link. Listing 1 shows an example of endpoint described according the IDC format.

We manually created topology descriptions for all the endpoints in the testbed in the proper formats. Listing 1 shows the XML structure used by IDC for the identification of endpoints. Using this schema an IDC endpoint would be described as:

link-id = urn:ogf:network:domain=uvalight.net:node=SURFnetOME4T_I2:port=s1p1:link10.9.2.10

Listing 1 – The XML schema representing a port in the IDC syntax.

```
<domain id="domain-id">
  <node id="node-id">
     <port id = "port-id">
        <link id="link-id">
            <remoteLinkId>remoteId</remoteLinkID>
        </link>
      </port>
   </node>
</domain>
```

Fig. 4 is a detailed representation of our topology. The goal is to connect a Phosphorus endpoint to an Internet2 endpoint. The Phosphorus endpoint is at the top left of the figure; the Internet2 endpoint is at the bottom left. In this use case both Web Services operation and endpoint translations take place.

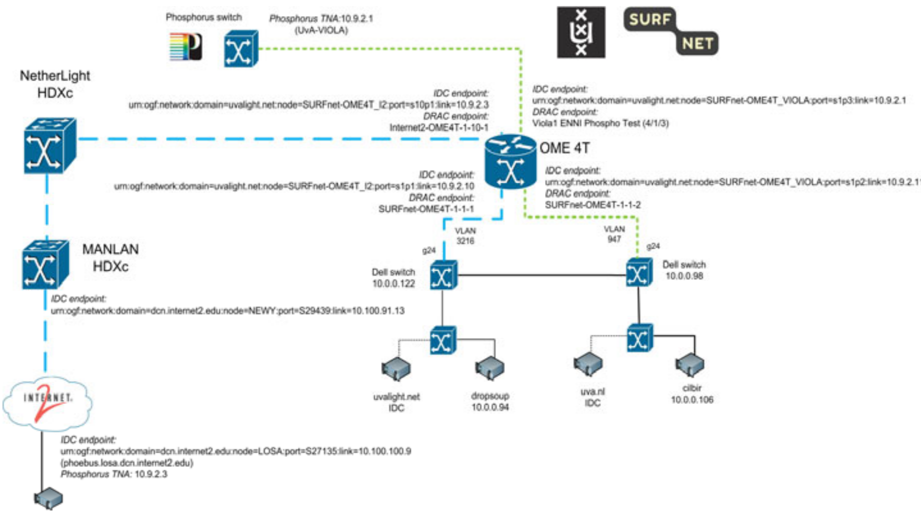

Fig. 4. Path setup scenario between a Phosphorus and an Internet2 host. Different line-patterns represent different VLAN tagging of the traffic running over the data-connections. The OME 4T is a Nortel OME 6500 and provisions Layer1 cross-connects between its ports. The Netherlight and MANLAN HDXc-s provide carrier services over TDM.

Harmony identifies the Phosphorus endpoint with TNA 10.9.2.1 and the Internet2 node with TNA 10.9.2.3. So the reservation request in Harmony contains TNA pair (10.9.2.1, 10.9.2.3) as the (*src_TNA, dst_TNA)* parameters for the *createReservation* Web Service operation. But the devices along the path between the source and the destination are under control of IDC and not Harmony. Harmony makes use of our Request Translator and performs:

- An endpoint translation, from the (*src_TNA, dst_TNA*) to an IDC (*src_URN, dst_URN*);
- A Web Service request translation, where Harmony uses our request translator to create a corresponding IDC Web Service *createReservation* request.
- A *createReservation* Web Service call to the IDC NSP.

Harmony in this manner delegates the actual setup to the IDC provisioning system, while offering to the user a single interface.

5 Service-Plane to Control-Plane Translation Module

The biggest challenge for Harmony/IDC interoperability stems from a fundamental incongruence in their respective designs.

The Harmony NSP is designed to interoperate with other single-domain NSPs that are built on top of CPs tailored to their individual domains, e.g., ARGON in the VIOLA domain. For this reason, CP functionalities such as path-setup and path-teardown components are absent in the Harmony NSP.

Conversely, the IDC design does not facilitate interoperability with any other NSP than another IDC. Moreover, CP functionality is an integral part of the IDC design and is tailored to interface with the DRAGON CP. The only possibility the IDC offers to interface with CPs other than DRAGON is to modify the Path-Setup Subcomponent (PSS) that signals path-setup, path-refresh and path-teardown events to the underlying CP.

The SURFnet/UvA network depicted on the right of Fig.4 is under control of IDC. Still IDC cannot configure paths because the DRAGON control plane cannot configure the OME testbed and Layer2 switches present in this portion of the network. The OME devices, a Layer1 switch from Nortel, can only be configured with DRAC; the Layer2 Dell switches can be controlled by using SNMP. In the following section we describe the high-level design of the PSS modifications we made to allow configuration.

5.1 Path-Setup Subcomponent

Fig. 5 depicts the components participating in the path setup through the two domains controlled by the IDC, the Internet2 domain and the UvA/SURFnet domains we saw in Fig.4.

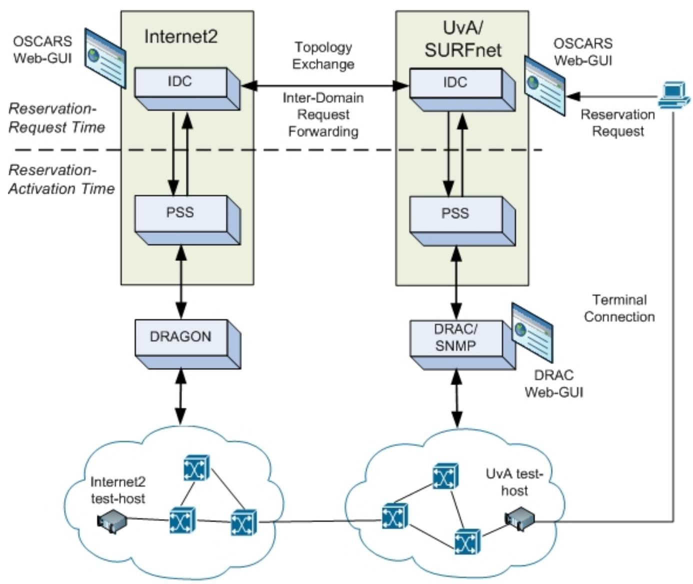

Fig. 5. The Internet2 domain and the UvA/SURFnet domain communicate in our testbed. At reservation activation time the Path Setup Subcomponent (PSS) communicate with the control plane to configure the network. In the Internet2 domain the PSS communicates with DRAGON; in the UvA/SURFnet domain the PSS triggers DRAC and SNMP calls to our equipment.

At reservation-activation time the IDC signals the CPs through the PSS. In the UvA/SURFnet domain we modified the PSS to signal the Layer1 network elements (OME 6500s) through interfacing with DRAC, and the Layer2 network elements (Dell Powerconnect 5324s) by interfacing with the SNMP agents running on these switches.

The parameters that are passed to the PSS are a Global Reservation ID (GRI), the requested bandwidth, a VLAN number, and a path consisting of a list of intra-domain hops calculated by the Path Calculation Element at reservation-request time.

The PSS has two main tasks: 1) to analyze the path and discover if there are ingress/egress hops (e.g., from/to the Internet2 domain) linked to intra-domain hops; these hop-pairs will be translated into DRAC request to set up a cross-connects on the OME 4T, and 2) to set up the Layer2 switches in order to retag traffic running between the Harmony and Internet2 domains, in case of VLAN mismatches.

The path contains an ingress/egress hop from the Internet2 domain to an internal UvA/SURFnet node, the PSS would perform the following translation:

- IDC hop urn:ogf:network:domain=uvalight.net:node=SURFnet-OME4T_I2:port=s10p1:link=10.9.2.3 translates into DRAC endpoint Internet2-OME4T-1-10-1;
- IDC hop urn:ogf:network:domain=uvalight.net:node=SURFnet-OME4T_I2:port=s1p1:link=10.9.2.10 translates into DRAC endpoint: SURFnet-OME4T-1-1-1

After mapping the ingress/egress hop to intra-domain hop pairs, the PSS calls DRAC with the corresponding DRAC-endpoint pairs to set up the Layer1 cross-connects on the OME 4T. In the case of the scenario depicted in Fig. 5 the only necessary cross-connect to be provisioned on the OME 4T is between *Internet2-OME4T-1-10-1* and *SURFnet-OME4T-1-1-1*. Subsequently, the port of the switch connected to the OME 4T will be configured by the PSS to accept traffic with the VLAN number included in the arguments of the PSS call, and to tag outgoing traffic with that VLAN number as well.

6 Workflow

Once the translation modules were ready we could prove the feasibility of an inter-domain path in our IDC-Harmony combined testbed. In this section we use an example end-to-end path to illustrate all the steps in the provisioning workflow. We use an example request for a circuit starting in the VIOLA domain and ending in the IDC domain.

6.1 Reservation Time

The user interacts with the NSP through the Reservation Request Interface. The UNI has two variants: a WS-UNI or a web-browser based UNI (WBUI). In both cases the user provides the reservation system with the parameters of the request: the source and destination IP addresses or transport-network addresses (TNAs), bandwidth, start time and duration. In the case of the IDC it's also possible to provide a VLAN number that will be used for the end-to-end Layer2 path.

Submitting requests is restricted to authorized users, however, the user authentication and authorization methods between Harmony and IDC differs considerably as we described before. In Harmony, requests from both the WBUI as well as from the WS-UNI are allowed from machines within the Harmony VPN. Both the IDC and DRAC WBUI require username/password authentication. The IDC WS-UNI requires messages to be signed by the user's private key, and DRAC authenticates WS requests through an included username/password token.

In Fig.6 we show the workflow associated to the setup of a path from the VIOLA domain to the Internet2 domain.

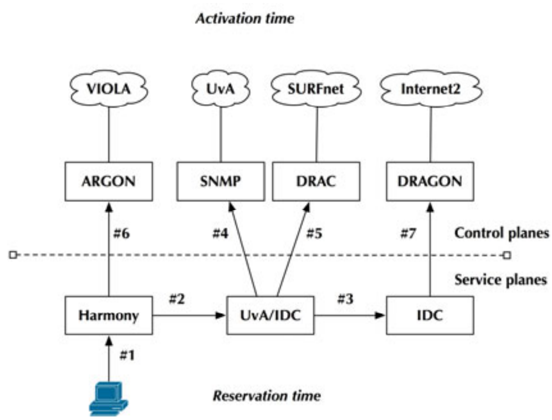

Fig. 6. Schematics of the workflow for a lightpath request starting in the VIOLA domain and ending in the Internet2 domain. The NSP operate at reservation time and the control planes operate at activation time.

At reservation time the user sends a requests for a path from VIOLA to Internet2 through the Harmony Web-GUI reservation form (*step #1*). The Harmony request is translated into an IDC request that is subsequently sent to the UvA IDC (*step#2*). This makes use of our Web Services and end-point translation modules.

The UvA IDC determines that the request comprises a multi-domain hop to the Internet2 domain, which causes the request to be forwarded to the Internet2 IDC (*step#3*). If this request can be accommodated a reservation is subsequently made in the Internet2 IDC and the UvA IDC.

6.2 Activation Time

At the start time of the reservation, both Harmony and the IDC signal the underlying control planes to set up the respective paths. The signaling consists of several distinct and heterogeneous processes:

- *DRAC-OME signaling* (*step #4*). The UvA IDC signals DRAC through its Path Setup Subsystem (PSS). DRAC processes this request as an immediate reservation and sets up the requested path through a Nortel OME testbed. This process sets up the data-paths in the OME testbed and uses Nortel/SURFnet proprietary code and algorithms.

- *UvA-IDC-SNMP signaling* (*step #5*). The VLANs used in the Internet2 and Phosphorus domains differ. The switches in the UvA domain have to be configured to convert VLAN tags between the ones used by Phosphorus and Internet2. The signaling consists of sending SNMP commands through control connections from the IDC to the Dell switches. Here we use our PSS translation code.
- *Harmony NSP-ARGIA signaling* (*step #6*). This step is optional, because the reservation Harmony creates in ARGIA can both be of type automatic or signal-on-activation. Harmony-ARGIA (and other CPs in the Phosphorus testbed) signaling happens through modules dubbed 'adapters' that enable ARGIA and the Harmony NSP to exchange information about e.g., reservations and topology.
- *IDC-DRAGON signaling* (*step #7*). As in the case of IDC-DRAC signaling, the PSS signals the DRAGON CP to set up the reserved path at activation time. This happens through RSVP-TE messages.

7 Demonstrations and Results

We demonstrated path setup as described in the previous section between Internet2, i2CAT and VIOLA at Supercomputing 2008 (SC08; held in Austin Texas, USA) and the Internet2 Spring Member Meeting in April 2009 (Washington DC, USA). The latest demonstration has been at the Oct. 2009 GLIF workshop in Korea.

At SC08 and at the Internet2 Member meeting the focus was on the NSP-to-NSP translator (section 4). The demonstration in Korea focused on the service-plane to control-plane translation module (section 5.1). Fig.7 is the graphical representation of a path created with our software: one endpoint is in Amsterdam and the other is in Los Angeles.

Fig. 7. The graphical overview of an end to end path realized with our software. In the specific case we created a dedicated circuit between Amsterdam and Los Angeles. This makes use of our service-plane to control-plane software.

8 Challenges and Future Work

Currently in our testbed the request translation only flows from the Harmony NSP to the IDC and not the other way around (see Fig.3). To achieve full interoperability, we intend to develop a translator for the IDC that will be plugged in among the standard IDC components.

We also intend to enhance the IDC-DRAC integration. Currently, the IDC is designed to completely control the domain it resides in, and acts as both an SP as well as a CP. This means that the IDC is the only authoritative entity in one single domain and disregards the SP/CP that may already exists in the domain. In case of the IDC-DRAC interaction this means the IDC only sends path setup and path teardown requests to DRAC, while DRAC itself has a very sophisticated reservation system. Possible inconsistencies can arise because of this when the IDC tries to set up paths in the domain DRAC controls and there is alreday an existing reservation made by a DRAC user in the local domain.

A more general question concerning integration of provisioning system is whether interoperability will be achieved by following a centralized model or a federated model. Fig.8 exemplifies the two architectures.

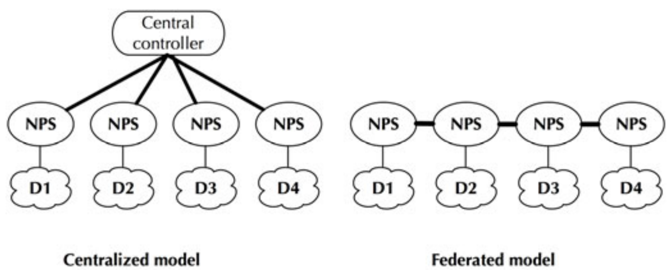

Fig. 8. Two possible architecture models for interoperability: a centralized model (left) and a federated model (right)

Centralized architectures see a vertical communication pattern (called north-south communication), where different domains relinquish control to a central management plane. This is the model implemented in Harmony. Federated architectures maintain a horizontal communication pattern (often called east-west communication). This chain model is the one realized by IDC.

We cannot predict which model will prevail. Our work provides tools to bridge one to the other by allowing a centralized model (Harmony) to co-exist and communicate with a federated model (IDC).

Furthermore, our work lead us to identify four components necessary for interoperability of provisioning systems that will need to be addressed in future work:

- *Terminology consistency*. Lightpaths is an all-encompassing term: at time lightpaths are configured purely at Layer1, at times they are Layer2 Ethernet paths. Consistent naming, clear technology information is a prerequisite for configuration across domains;
- *Topology descriptions consistency*. Network domains require some knowledge of the connected domain topology to set up on-demand paths beyond network borders.

For this reason domains must exchange network topologies. But what constitutes the essential and necessary information is matter of study. We expect standard topologies will come from efforts in the Open Grid Forum community, such as the one in the Network Markup Language working group [15];
- *Consistency of information on the status of resources*. Currently Harmony and IDC operate independently from each other at activation time: this means, a path that had been guaranteed during reservation might fail to be provisioned in one of the domains. Harmony or IDC will not be notified of this by the other NSP. As a result, the user will not have an end-to-end path available, while the NSP has guaranteed the existence of one. Notifications of resource availability and status in domains along a path can be a way to maintain consistency.
- *AAA features consistency*. This implies that higher or lower control on the resources must be harmonized across domains. As we mentioned in Sec. 4.1 Harmony and IDC have very different authentication models: with IDC authentication requiring X.509 certificates individually assigned to users. User authentication is a necessary condition to assign different permission levels and control of resources [16]. All NSPs should in the future implement this model.

9 Conclusions

Several challenges are present when trying to perform inter-domain circuit provisioning. Systems that work perfectly in one domain need to be interfaced to interoperate. We showed our solution to this problem in the case of a combined Harmony-IDC testbed. We developed translation modules at the service-plane-to-service-plane level, and at the service-plane-to-control-plane level. With these modules we could show the setup of intercontinental multi-domain lightpaths.

We identified many fundamental inconsistencies when providing tools for interoperability between Harmony and IDC. The creation of standardized interfaces between network domains, as pursued for instance in the NSI group at OGF, will provide the answer. Our work helps in that effort as it contributes to the identification of the necessary interoperability components.

Acknowledgements. We are grateful to SURFnet that sponsored this research through its Research on Network program. We would like to thanks our collegues Alexander Willner, Christian de Waal, Jan Gassen at the University of Bonn; Joan Antoni Garcia Espin, Jordi Ferrrer Riera, Carlos Baez Ruiz at i2CAT; and John Vollbrecht and Andrew Lake at Internet2. We also thank our collegues at SURFnet, Gerben van Malenstein and Hans Trompert.

References

1. Future Internet Research & Experimentation (FIRE) Initiative website, http://cordis.europa.eu/fp7/ict/fire/
2. Global Environment for Network Innovations initiative website, http://www.geni.net
3. The Large Hadron Collier website, http://lhc.web.cern.ch/lhc/

4. The European VLBI website, `http://www.evlbi.org/`
5. de Laat, C., Radius, E., Wallace, S.: "The Rationale of the Current Optical Networking Initiatives". iGrid 2002 special issue, Future Generation Computer Systems 19(6) (2003)
6. Dijkstra, F., van der Ham, J., Grosso, P., de Laat, C.: A Path Finding Implementation for Multi-layer Networks. Future Generation Computer Systems 25(2), 142–146 (2009)
7. Kuipers, F., Dijkstra, F.: *Path Selection in Multi-Layer Networks*. Elsevier Computer Communications 32, 78–85 (2009)
8. Grasa, E., Junyent, G., Figuerola, S., Lopez, A., Savoie, M.: UCLPv2: a network virtualization framework built on web services. IEEE Communications Magazine 46(3), 126–134 (2008)
9. NSI-WG website, `http://forge.gridforum.org/sf/projects/nsi-wg`
10. Phosphorus project website, `http://www.ist-phosphorus.eu/`
11. DICE initiative website, `http://www.geant2.net/server/show/conWebDoc.1308`
12. Lehman, T., et al.: Control Plane Architecture and Design Considerations for Multi-Service, Multi-Layer, Multi-Domain Hybrid Networks". In: INFOCOM 2007 High-Speed Networks Workshop, 26th Annual IEEE Conference on Computer Communications, Anchorage, AK, pp. 67–71 (2007)
13. Lehman, T., Sobieski, J., Jabbari, B.: DRAGON: a framework for service provisioning in heterogeneous grid networks Export. IEEE Communications Magazine 44(3), 84–90 (2006)
14. NM-WG website, `http://nmwg.internet2.edu`
15. NML-WG website, `https://forge.gridforum.org/projects/nml-wg`
16. Gommans, L., Xu, L., Wan, F., Demchenko, Y., Cristea, M., Meijer, R., de Laat, C.: Multi-Domain Lightpath Authorization using Tokens. Future Generation Computing Systems 25(2), 153–160 (2008)