

Internet Engineering Task Force (IETF)
Request for Comments: 7491
Category: Informational
ISSN: 2070-1721

D. King
Old Dog Consulting
A. Farrel
Juniper Networks
March 2015

A PCE-Based Architecture for Application-Based Network Operations

Abstract

Services such as content distribution, distributed databases, or inter-data center connectivity place a set of new requirements on the operation of networks. They need on-demand and application-specific reservation of network connectivity, reliability, and resources (such as bandwidth) in a variety of network applications (such as point-to-point connectivity, network virtualization, or mobile back-haul) and in a range of network technologies from packet (IP/MPLS) down to optical. An environment that operates to meet these types of requirements is said to have Application-Based Network Operations (ABNO). ABNO brings together many existing technologies and may be seen as the use of a toolbox of existing components enhanced with a few new elements.

This document describes an architecture and framework for ABNO, showing how these components fit together. It provides a cookbook of existing technologies to satisfy the architecture and meet the needs of the applications.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7491>.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Scope	5
2. Application-Based Network Operations (ABNO)	6
2.1. Assumptions	6
2.2. Implementation of the Architecture	6
2.3. Generic ABNO Architecture	7
2.3.1. ABNO Components	8
2.3.2. Functional Interfaces	15
3. ABNO Use Cases	24
3.1. Inter-AS Connectivity	24
3.2. Multi-Layer Networking	30
3.2.1. Data Center Interconnection across Multi-Layer Networks	34
3.3. Make-before-Break	37
3.3.1. Make-before-Break for Reoptimization	37
3.3.2. Make-before-Break for Restoration	38
3.3.3. Make-before-Break for Path Test and Selection	40
3.4. Global Concurrent Optimization	42
3.4.1. Use Case: GCO with MPLS LSPs	43
3.5. Adaptive Network Management (ANM)	45
3.5.1. ANM Trigger	46
3.5.2. Processing Request and GCO Computation	46
3.5.3. Automated Provisioning Process	47
3.6. Pseudowire Operations and Management	48
3.6.1. Multi-Segment Pseudowires	48
3.6.2. Path-Diverse Pseudowires	50
3.6.3. Path-Diverse Multi-Segment Pseudowires	51
3.6.4. Pseudowire Segment Protection	52
3.6.5. Applicability of ABNO to Pseudowires	52
3.7. Cross-Stratum Optimization (CSO)	53
3.7.1. Data Center Network Operation	53
3.7.2. Application of the ABNO Architecture	56
3.8. ALTO Server	58
3.9. Other Potential Use Cases	61
3.9.1. Traffic Grooming and Regrooming	61
3.9.2. Bandwidth Scheduling	62
4. Survivability and Redundancy within the ABNO Architecture	62
5. Security Considerations	63
6. Manageability Considerations	63
7. Informative References	64
Appendix A. Undefined Interfaces	69
Acknowledgements	70
Contributors	71
Authors' Addresses	71

1. Introduction

Networks today integrate multiple technologies allowing network infrastructure to deliver a variety of services to support the different characteristics and demands of applications. There is an increasing demand to make the network responsive to service requests issued directly from the application layer. This differs from the established model where services in the network are delivered in response to management commands driven by a human user.

These application-driven requests and the services they establish place a set of new requirements on the operation of networks. They need on-demand and application-specific reservation of network connectivity, reliability, and resources (such as bandwidth) in a variety of network applications (such as point-to-point connectivity, network virtualization, or mobile back-haul) and in a range of network technologies from packet (IP/MPLS) down to optical. An environment that operates to meet this type of application-aware requirement is said to have Application-Based Network Operations (ABNO).

The Path Computation Element (PCE) [RFC4655] was developed to provide path computation services for GMPLS- and MPLS-controlled networks. The applicability of PCEs can be extended to provide path computation and policy enforcement capabilities for ABNO platforms and services.

ABNO can provide the following types of service to applications by coordinating the components that operate and manage the network:

- Optimization of traffic flows between applications to create an overlay network for communication in use cases such as file sharing, data caching or mirroring, media streaming, or real-time communications described as Application-Layer Traffic Optimization (ALTO) [RFC5693].
- Remote control of network components allowing coordinated programming of network resources through such techniques as Forwarding and Control Element Separation (ForCES) [RFC3746], OpenFlow [ONF], and the Interface to the Routing System (I2RS) [I2RS-Arch], or through the control plane coordinated through the PCE Communication Protocol (PCEP) [PCE-Init-LSP].
- Interconnection of Content Delivery Networks (CDNi) [RFC6707] through the establishment and resizing of connections between content distribution networks. Similarly, ABNO can coordinate inter-data center connections.

- Network resource coordination to automate provisioning, and to facilitate traffic grooming and regrooming, bandwidth scheduling, and Global Concurrent Optimization using PCEP [RFC5557].
- Virtual Private Network (VPN) planning in support of deployment of new VPN customers and to facilitate inter-data center connectivity.

This document outlines the architecture and use cases for ABNO, and shows how the ABNO architecture can be used for coordinating control system and application requests to compute paths, enforce policies, and manage network resources for the benefit of the applications that use the network. The examination of the use cases shows the ABNO architecture as a toolkit comprising many existing components and protocols, and so this document looks like a cookbook. ABNO is compatible with pre-existing Network Management System (NMS) and Operations Support System (OSS) deployments as well as with more recent developments in programmatic networks such as Software-Defined Networking (SDN).

1.1. Scope

This document describes a toolkit. It shows how existing functional components described in a large number of separate documents can be brought together within a single architecture to provide the function necessary for ABNO.

In many cases, existing protocols are known to be good enough or almost good enough to satisfy the requirements of interfaces between the components. In these cases, the protocols are called out as suitable candidates for use within an implementation of ABNO.

In other cases, it is clear that further work will be required, and in those cases a pointer to ongoing work that may be of use is provided. Where there is no current work that can be identified by the authors, a short description of the missing interface protocol is given in Appendix A.

Thus, this document may be seen as providing an applicability statement for existing protocols, and guidance for developers of new protocols or protocol extensions.

2. Application-Based Network Operations (ABNO)

2.1. Assumptions

The principal assumption underlying this document is that existing technologies should be used where they are adequate for the task. Furthermore, when an existing technology is almost sufficient, it is assumed to be preferable to make minor extensions rather than to invent a whole new technology.

Note that this document describes an architecture. Functional components are architectural concepts and have distinct and clear responsibilities. Pairs of functional components interact over functional interfaces that are, themselves, architectural concepts.

2.2. Implementation of the Architecture

It needs to be strongly emphasized that this document describes a functional architecture. It is not a software design. Thus, it is not intended that this architecture constrain implementations. However, the separation of the ABNO functions into separate functional components with clear interfaces between them enables implementations to choose which features to include and allows different functions to be distributed across distinct processes or even processors.

An implementation of this architecture may make several important decisions about the functional components:

- Multiple functional components may be grouped together into one software component such that all of the functions are bundled and only the external interfaces are exposed. This may have distinct advantages for fast paths within the software and can reduce interprocess communication overhead.

For example, an Active, Stateful PCE could be implemented as a single server combining the ABNO components of the PCE, the Traffic Engineering Database, the Label Switched Path Database, and the Provisioning Manager (see Section 2.3).

- The functional components could be distributed across separate processes, processors, or servers so that the interfaces are exposed as external protocols.

For example, the Operations, Administration, and Maintenance (OAM) Handler (see Section 2.3.1.6) could be presented on a dedicated server in the network that consumes all status reports from the network, aggregates them, correlates them, and then dispatches notifications to other servers that need to understand what has happened.

- There could be multiple instances of any or each of the components. That is, the function of a functional component could be partitioned across multiple software components with each responsible for handling a specific feature or a partition of the network.

For example, there may be multiple Traffic Engineering Databases (see Section 2.3.1.8) in an implementation, with each holding the topology information of a separate network domain (such as a network layer or an Autonomous System). Similarly, there could be multiple PCE instances, each processing a different Traffic Engineering Database, and potentially distributed on different servers under different management control. As a final example, there could be multiple ABNO Controllers, each with capability to support different classes of application or application service.

The purpose of the description of this architecture is to facilitate different implementations while offering interoperability between implementations of key components, and easy interaction with the applications and with the network devices.

2.3. Generic ABNO Architecture

Figure 1 illustrates the ABNO architecture. The components and functional interfaces are discussed in Sections 2.3.1 and 2.3.2, respectively. The use cases described in Section 3 show how different components are used selectively to provide different services. It is important to understand that the relationships and interfaces shown between components in this figure are illustrative of some of the common or likely interactions; however, this figure does not preclude other interfaces and relationships as necessary to realize specific functionality.

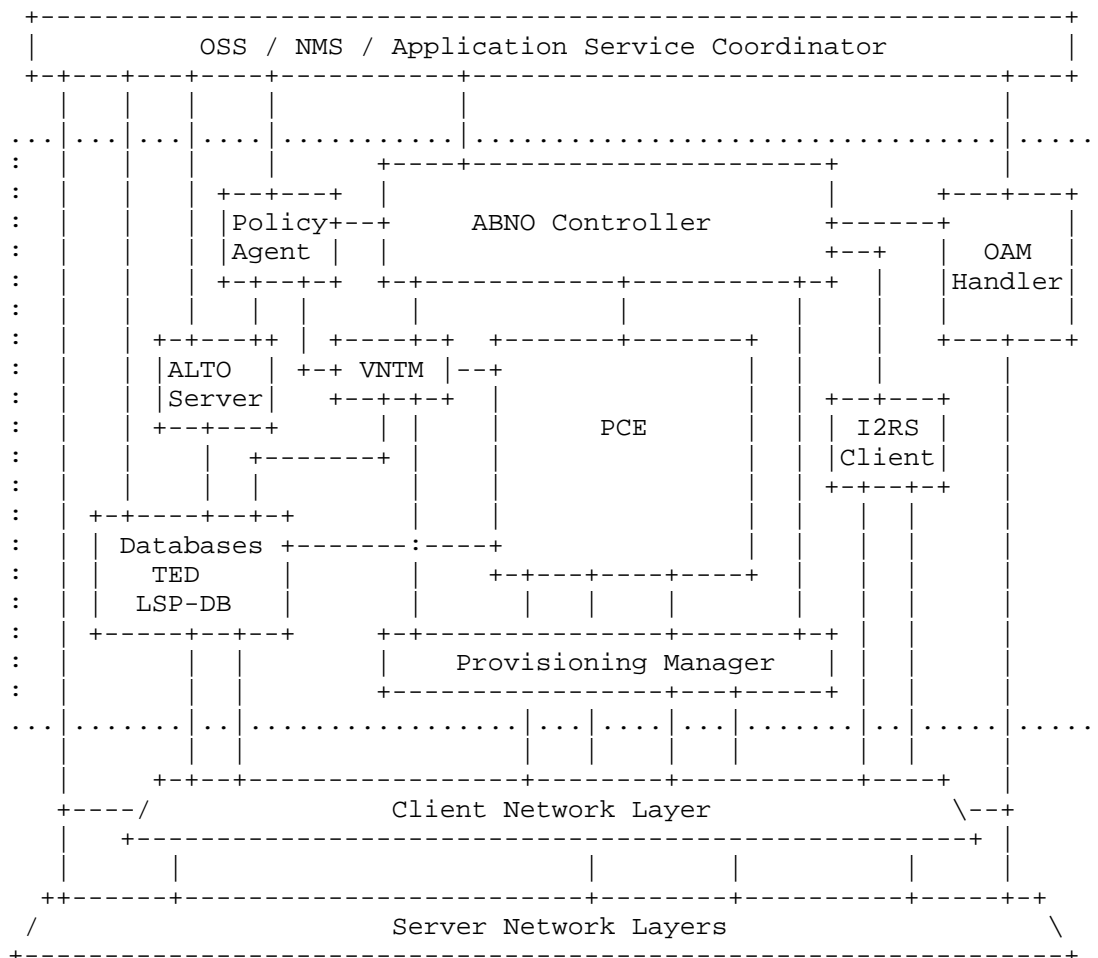


Figure 1: Generic ABNO Architecture

2.3.1. ABNO Components

This section describes the functional components shown as boxes in Figure 1. The interactions between those components, the functional interfaces, are described in Section 2.3.2.

2.3.1.1. NMS and OSS

A Network Management System (NMS) or an Operations Support System (OSS) can be used to control, operate, and manage a network. Within the ABNO architecture, an NMS or OSS may issue high-level service requests to the ABNO Controller. It may also establish policies for the activities of the components within the architecture.

The NMS and OSS can be consumers of network events reported through the OAM Handler and can act on these reports as well as displaying them to users and raising alarms. The NMS and OSS can also access the Traffic Engineering Database (TED) and Label Switched Path Database (LSP-DB) to show the users the current state of the network.

Lastly, the NMS and OSS may utilize a direct programmatic or configuration interface to interact with the network elements within the network.

2.3.1.2. Application Service Coordinator

In addition to the NMS and OSS, services in the ABNO architecture may be requested by or on behalf of applications. In this context, the term "application" is very broad. An application may be a program that runs on a host or server and that provides services to a user, such as a video conferencing application. Alternatively, an application may be a software tool that a user uses to make requests to the network to set up specific services such as end-to-end connections or scheduled bandwidth reservations. Finally, an application may be a sophisticated control system that is responsible for arranging the provision of a more complex network service such as a virtual private network.

For the sake of this architecture, all of these concepts of an application are grouped together and are shown as the Application Service Coordinator, since they are all in some way responsible for coordinating the activity of the network to provide services for use by applications. In practice, the function of the Application Service Coordinator may be distributed across multiple applications or servers.

The Application Service Coordinator communicates with the ABNO Controller to request operations on the network.

2.3.1.3. ABNO Controller

The ABNO Controller is the main gateway to the network for the NMS, OSS, and Application Service Coordinator for the provision of advanced network coordination and functions. The ABNO Controller governs the behavior of the network in response to changing network conditions and in accordance with application network requirements and policies. It is the point of attachment, and it invokes the right components in the right order.

The use cases in Section 3 provide a clearer picture of how the ABNO Controller interacts with the other components in the ABNO architecture.

2.3.1.4. Policy Agent

Policy plays a very important role in the control and management of the network. It is, therefore, significant in influencing how the key components of the ABNO architecture operate.

Figure 1 shows the Policy Agent as a component that is configured by the NMS/OSS with the policies that it applies. The Policy Agent is responsible for propagating those policies into the other components of the system.

Simplicity in the figure necessitates leaving out many of the policy interactions that will take place. Although the Policy Agent is only shown interacting with the ABNO Controller, the ALTO Server, and the Virtual Network Topology Manager (VNTM), it will also interact with a number of other components and the network elements themselves. For example, the Path Computation Element (PCE) will be a Policy Enforcement Point (PEP) [RFC2753] as described in [RFC5394], and the Interface to the Routing System (I2RS) Client will also be a PEP as noted in [I2RS-Arch].

2.3.1.5. Interface to the Routing System (I2RS) Client

The Interface to the Routing System (I2RS) is described in [I2RS-Arch]. The interface provides a programmatic way to access (for read and write) the routing state and policy information on routers in the network.

The I2RS Client is introduced in [I2RS-PS]. Its purpose is to manage information requests across a number of routers (each of which runs an I2RS Agent) and coordinate setting or gathering state to/from those routers.

2.3.1.6. OAM Handler

Operations, Administration, and Maintenance (OAM) plays a critical role in understanding how a network is operating, detecting faults, and taking the necessary action to react to problems in the network.

Within the ABNO architecture, the OAM Handler is responsible for receiving notifications (often called alerts) from the network about potential problems, for correlating them, and for triggering other components of the system to take action to preserve or recover the services that were established by the ABNO Controller. The OAM Handler also reports network problems and, in particular, service-affecting problems to the NMS, OSS, and Application Service Coordinator.

Additionally, the OAM Handler interacts with the devices in the network to initiate OAM actions within the data plane, such as monitoring and testing.

2.3.1.7. Path Computation Element (PCE)

PCE is introduced in [RFC4655]. It is a functional component that services requests to compute paths across a network graph. In particular, it can generate traffic-engineered routes for MPLS-TE and GMPLS Label Switched Paths (LSPs). The PCE may receive these requests from the ABNO Controller, from the Virtual Network Topology Manager, or from network elements themselves.

The PCE operates on a view of the network topology stored in the Traffic Engineering Database (TED). A more sophisticated computation may be provided by a Stateful PCE that enhances the TED with a database (the LSP-DB -- see Section 2.3.1.8.2) containing information about the LSPs that are provisioned and operational within the network as described in [RFC4655] and [Stateful-PCE].

Additional functionality in an Active PCE allows a functional component that includes a Stateful PCE to make provisioning requests to set up new services or to modify in-place services as described in [Stateful-PCE] and [PCE-Init-LSP]. This function may directly access the network elements or may be channeled through the Provisioning Manager.

Coordination between multiple PCEs operating on different TEDs can prove useful for performing path computation in multi-domain or multi-layer networks. A domain in this case might be an Autonomous System (AS), thus enabling inter-AS path computation.

Since the PCE is a key component of the ABNO architecture, a better view of its role can be gained by examining the use cases described in Section 3.

2.3.1.8. Databases

The ABNO architecture includes a number of databases that contain information stored for use by the system. The two main databases are the TED and the LSP Database (LSP-DB), but there may be a number of other databases used to contain information about topology (ALTO Server), policy (Policy Agent), services (ABNO Controller), etc.

In the text that follows, specific key components that are consumers of the databases are highlighted. It should be noted that the databases are available for inspection by any of the ABNO components. Updates to the databases should be handled with some care, since allowing multiple components to write to a database can be the cause of a number of contention and sequencing problems.

2.3.1.8.1. Traffic Engineering Database (TED)

The TED is a data store of topology information about a network that may be enhanced with capability data (such as metrics or bandwidth capacity) and active status information (such as up/down status or residual unreserved bandwidth).

The TED may be built from information supplied by the network or from data (such as inventory details) sourced through the NMS/OSS.

The principal use of the TED in the ABNO architecture is to provide the raw data on which the Path Computation Element operates. But the TED may also be inspected by users at the NMS/OSS to view the current status of the network and may provide information to application services such as Application-Layer Traffic Optimization (ALTO) [RFC5693].

2.3.1.8.2. LSP Database

The LSP-DB is a data store of information about LSPs that have been set up in the network or that could be established. The information stored includes the paths and resource usage of the LSPs.

The LSP-DB may be built from information generated locally. For example, when LSPs are provisioned, the LSP-DB can be updated. The database can also be constructed from information gathered from the network by polling or reading the state of LSPs that have already been set up.

The main use of the LSP-DB within the ABNO architecture is to enhance the planning and optimization of LSPs. New LSPs can be established to be path-disjoint from other LSPs in order to offer protected services; LSPs can be rerouted in order to put them on more optimal paths or to make network resources available for other LSPs; LSPs can be rapidly repaired when a network failure is reported; LSPs can be moved onto other paths in order to avoid resources that have planned maintenance outages. A Stateful PCE (see Section 2.3.1.7) is a primary consumer of the LSP-DB.

2.3.1.8.3. Shared Risk Link Group (SRLG) Databases

The TED may, itself, be supplemented by SRLG information that assigns to each network resource one or more identifiers that associate the resource with other resources in the same TED that share the same risk of failure.

While this information can be highly useful, it may be supplemented by additional detailed information maintained in a separate database and indexed using the SRLG identifier from the TED. Such a database can interpret SRLG information provided by other networks (such as server networks), can provide failure probabilities associated with each SRLG, can offer prioritization when SRLG-disjoint paths cannot be found, and can correlate SRLGs between different server networks or between different peer networks.

2.3.1.8.4. Other Databases

There may be other databases that are built within the ABNO system and that are referenced when operating the network. These databases might include information about, for example, traffic flows and demands, predicted or scheduled traffic demands, link and node failure and repair history, network resources such as packet labels and physical labels (i.e., MPLS and GMPLS labels), etc.

As mentioned in Section 2.3.1.8.1, the TED may be enhanced by inventory information. It is quite likely in many networks that such an inventory is held in a separate database (the Inventory Database) that includes details of the manufacturer, model, installation date, etc.

2.3.1.9. ALTO Server

The ALTO Server provides network information to the application layer based on abstract maps of a network region. This information provides a simplified view, but it is useful to steer application-layer traffic. ALTO services enable service providers to share information about network locations and the costs of paths between

them. The selection criteria to choose between two locations may depend on information such as maximum bandwidth, minimum cross-domain traffic, lower cost to the user, etc.

The ALTO Server generates ALTO views to share information with the Application Service Coordinator so that it can better select paths in the network to carry application-layer traffic. The ALTO views are computed based on information from the network databases, from policies configured by the Policy Agent, and through the algorithms used by the PCE.

Specifically, the base ALTO protocol [RFC7285] defines a single-node abstract view of a network to the Application Service Coordinator. Such a view consists of two maps: a network map and a cost map. A network map defines multiple Provider-defined Identifiers (PIDs), which represent entrance points to the network. Each node in the application layer is known as an End Point (EP), and each EP is assigned to a PID, because PIDs are the entry points of the application in the network. As defined in [RFC7285], a PID can denote a subnet, a set of subnets, a metropolitan area, a Point of Presence (PoP), etc. Each such network region can be a single domain or multiple networks; it is just the view that the ALTO Server is exposing to the application layer. A cost map provides costs between EPs and/or PIDs. The criteria that the Application Service Coordinator uses to choose application routes between two locations may depend on attributes such as maximum bandwidth, minimum cross-domain traffic, lower cost to the user, etc.

2.3.1.10. Virtual Network Topology Manager (VNTM)

A Virtual Network Topology (VNT) is defined in [RFC5212] as a set of one or more LSPs in one or more lower-layer networks that provides information for efficient path handling in an upper-layer network. For instance, a set of LSPs in a wavelength division multiplexed (WDM) network can provide connectivity as virtual links in a higher-layer packet-switched network.

The VNT enhances the physical/dedicated links that are available in the upper-layer network and is configured by setting up or tearing down the lower-layer LSPs and by advertising the changes into the higher-layer network. The VNT can be adapted to traffic demands so that capacity in the higher-layer network can be created or released as needed. Releasing unwanted VNT resources makes them available in the lower-layer network for other uses.

The creation of virtual topology for inclusion in a network is not a simple task. Decisions must be made about which nodes in the upper layer it is best to connect, in which lower-layer network to provision LSPs to provide the connectivity, and how to route the LSPs in the lower-layer network. Furthermore, some specific actions have to be taken to cause the lower-layer LSPs to be provisioned and the connectivity in the upper-layer network to be advertised.

[RFC5623] describes how the VNTM may instantiate connections in the server layer in support of connectivity in the client layer. Within the ABNO architecture, the creation of new connections may be delegated to the Provisioning Manager as discussed in Section 2.3.1.11.

All of these actions and decisions are heavily influenced by policy, so the VNTM component that coordinates them takes input from the Policy Agent. The VNTM is also closely associated with the PCE for the upper-layer network and each of the PCEs for the lower-layer networks.

2.3.1.11. Provisioning Manager

The Provisioning Manager is responsible for making or channeling requests for the establishment of LSPs. This may be instructions to the control plane running in the networks or may involve the programming of individual network devices. In the latter case, the Provisioning Manager may act as an OpenFlow Controller [ONF].

See Section 2.3.2.6 for more details of the interactions between the Provisioning Manager and the network.

2.3.1.12. Client and Server Network Layers

The client and server networks are shown in Figure 1 as illustrative examples of the fact that the ABNO architecture may be used to coordinate services across multiple networks where lower-layer networks provide connectivity in upper-layer networks.

Section 3.2 describes a set of use cases for multi-layer networking.

2.3.2. Functional Interfaces

This section describes the interfaces between functional components that might be externalized in an implementation allowing the components to be distributed across platforms. Where existing protocols might provide all or most of the necessary capabilities, they are noted. Appendix A notes the interfaces where more protocol specification may be needed.

As noted at the top of Section 2.3, it is important to understand that the relationships and interfaces shown between components in Figure 1 are illustrative of some of the common or likely interactions; however, this figure and the descriptions in the subsections below do not preclude other interfaces and relationships as necessary to realize specific functionality. Thus, some of the interfaces described below might not be visible as specific relationships in Figure 1, but they can nevertheless exist.

2.3.2.1. Configuration and Programmatic Interfaces

The network devices may be configured or programmed directly from the NMS/OSS. Many protocols already exist to perform these functions, including the following:

- SNMP [RFC3412]
- The Network Configuration Protocol (NETCONF) [RFC6241]
- RESTCONF [RESTCONF]
- The General Switch Management Protocol (GSMP) [RFC3292]
- ForCES [RFC5810]
- OpenFlow [ONF]
- PCEP [PCE-Init-LSP]

The TeleManagement Forum (TMF) Multi-Technology Operations Systems Interface (MTOSI) standard [TMF-MTOSI] was developed to facilitate application-to-application interworking and provides network-level management capabilities to discover, configure, and activate resources. Initially, the MTOSI information model was only capable of representing connection-oriented networks and resources. In later releases, support was added for connectionless networks. MTOSI is, from the NMS perspective, a north-bound interface and is based on SOAP web services.

From the ABNO perspective, network configuration is a pass-through function. It can be seen represented on the left-hand side of Figure 1.

2.3.2.2. TED Construction from the Networks

As described in Section 2.3.1.8, the TED provides details of the capabilities and state of the network for use by the ABNO system and the PCE in particular.

The TED can be constructed by participating in the IGP-TE protocols run by the networks (for example, OSPF-TE [RFC3630] and IS-IS TE [RFC5305]). Alternatively, the TED may be fed using link-state distribution extensions to BGP [BGP-LS].

The ABNO system may maintain a single TED unified across multiple networks or may retain a separate TED for each network.

Additionally, an ALTO Server [RFC5693] may provide an abstracted topology from a network to build an application-level TED that can be used by a PCE to compute paths between servers and application-layer entities for the provision of application services.

2.3.2.3. TED Enhancement

The TED may be enhanced by inventory information supplied from the NMS/OSS. This may supplement the data collected as described in Section 2.3.2.2 with information that is not normally distributed within the network, such as node types and capabilities, or the characteristics of optical links.

No protocol is currently identified for this interface, but the protocol developed or adopted to satisfy the requirements of the Interface to the Routing System (I2RS) [I2RS-Arch] may be a suitable candidate because it is required to be able to distribute bulk routing state information in a well-defined encoding language. Another candidate protocol may be NETCONF [RFC6241] passing data encoded using YANG [RFC6020].

Note that, in general, any combination of protocol and encoding that is suitable for presenting the TED as described in Section 2.3.2.4 will likely be suitable (or could be made suitable) for enabling write-access to the TED as described in this section.

2.3.2.4. TED Presentation

The TED may be presented north-bound from the ABNO system for use by an NMS/OSS or by the Application Service Coordinator. This allows users and applications to get a view of the network topology and the status of the network resources. It also allows planning and provisioning of application services.

There are several protocols available for exporting the TED north-bound:

- The ALTO protocol [RFC7285] is designed to distribute the abstracted topology used by an ALTO Server and may prove useful for exporting the TED. The ALTO Server provides the cost between EPs

or between PIDs, so the application layer can select which is the most appropriate connection for the information exchange between its application end points.

- The same protocol used to export topology information from the network can be used to export the topology from the TED [BGP-LS].
- The I2RS [I2RS-Arch] will require a protocol that is capable of handling bulk routing information exchanges that would be suitable for exporting the TED. In this case, it would make sense to have a standardized representation of the TED in a formal data modeling language such as YANG [RFC6020] so that an existing protocol such as NETCONF [RFC6241] or the Extensible Messaging and Presence Protocol (XMPP) [RFC6120] could be used.

Note that export from the TED can be a full dump of the content (expressed in a suitable abstraction language) as described above, or it could be an aggregated or filtered set of data based on policies or specific requirements. Thus, the relationships shown in Figure 1 may be a little simplistic in that the ABNO Controller may also be involved in preparing and presenting the TED information over a north-bound interface.

2.3.2.5. Path Computation Requests from the Network

As originally specified in the PCE architecture [RFC4655], network elements can make path computation requests to a PCE using PCEP [RFC5440]. This facilitates the network setting up LSPs in response to simple connectivity requests, and it allows the network to reoptimize or repair LSPs.

2.3.2.6. Provisioning Manager Control of Networks

As described in Section 2.3.1.11, the Provisioning Manager makes or channels requests to provision resources in the network. These operations can take place at two levels: there can be requests to program/configure specific resources in the data or forwarding planes, and there can be requests to trigger a set of actions to be programmed with the assistance of a control plane.

A number of protocols already exist to provision network resources, as follows:

- o Program/configure specific network resources
 - ForCES [RFC5810] defines a protocol for separation of the control element (the Provisioning Manager) from the forwarding elements in each node in the network.
 - The General Switch Management Protocol (GSMP) [RFC3292] is an asymmetric protocol that allows one or more external switch controllers (such as the Provisioning Manager) to establish and maintain the state of a label switch such as an MPLS switch.
 - OpenFlow [ONF] is a communications protocol that gives an OpenFlow Controller (such as the Provisioning Manager) access to the forwarding plane of a network switch or router in the network.
 - Historically, other configuration-based mechanisms have been used to set up the forwarding/switching state at individual nodes within networks. Such mechanisms have ranged from non-standard command line interfaces (CLIs) to various standards-based options such as Transaction Language 1 (TL1) [TL1] and SNMP [RFC3412]. These mechanisms are not designed for rapid operation of a network and are not easily programmatic. They are not proposed for use by the Provisioning Manager as part of the ABNO architecture.
 - NETCONF [RFC6241] provides a more active configuration protocol that may be suitable for bulk programming of network resources. Its use in this way is dependent on suitable YANG modules being defined for the necessary options. Early work in the IETF's NETMOD working group is focused on a higher level of routing function more comparable with the function discussed in Section 2.3.2.8; see [YANG-Rtg].
 - The [TMF-MTOSI] specification provides provisioning, activation, deactivation, and release of resources via the Service Activation Interface (SAI). The Common Communication Vehicle (CCV) is the middleware required to implement MTOSI. The CCV is then used to provide middleware abstraction in combination with the Web Services Description Language (WSDL) to allow MTOSIs to be bound to different middleware technologies as needed.

- o Trigger actions through the control plane
 - LSPs can be requested using a management system interface to the head end of the LSP using tools such as CLIs, TL1 [TL1], or SNMP [RFC3412]. Configuration at this granularity is not as time-critical as when individual network resources are programmed, because the main task of programming end-to-end connectivity is devolved to the control plane. Nevertheless, these mechanisms remain unsuitable for programmatic control of the network and are not proposed for use by the Provisioning Manager as part of the ABNO architecture.
 - As noted above, NETCONF [RFC6241] provides a more active configuration protocol. This may be particularly suitable for requesting the establishment of LSPs. Work would be needed to complete a suitable YANG module.
 - The PCE Communication Protocol (PCEP) [RFC5440] has been proposed as a suitable protocol for requesting the establishment of LSPs [PCE-Init-LSP]. This works well, because the protocol elements necessary are exactly the same as those used to respond to a path computation request.

The functional element that issues PCEP requests to establish LSPs is known as an "Active PCE"; however, it should be noted that the ABNO functional component responsible for requesting LSPs is the Provisioning Manager. Other controllers like the VNTM and the ABNO Controller use the services of the Provisioning Manager to isolate the twin functions of computing and requesting paths from the provisioning mechanisms in place with any given network.

Note that I2RS does not provide a mechanism for control of network resources at this level, as it is designed to provide control of routing state in routers, not forwarding state in the data plane.

2.3.2.7. Auditing the Network

Once resources have been provisioned or connections established in the network, it is important that the ABNO system can determine the state of the network. Similarly, when provisioned resources are modified or taken out of service, the changes in the network need to be understood by the ABNO system. This function falls into four categories:

- Updates to the TED are gathered as described in Section 2.3.2.2.
- Explicit notification of the successful establishment and the subsequent state of the LSP can be provided through extensions to PCEP as described in [Stateful-PCE] and [PCE-Init-LSP].
- OAM can be commissioned and the results inspected by the OAM Handler as described in Section 2.3.2.14.
- A number of ABNO components may make inquiries and inspect network state through a variety of techniques, including I2RS, NETCONF, or SNMP.

2.3.2.8. Controlling the Routing System

As discussed in Section 2.3.1.5, the Interface to the Routing System (I2RS) provides a programmatic way to access (for read and write) the routing state and policy information on routers in the network. The I2RS Client issues requests to routers in the network to establish or retrieve routing state. Those requests utilize the I2RS protocol, which will be based on a combination of NETCONF [RFC6241] and RESTCONF [RESTCONF] with some additional features.

2.3.2.9. ABNO Controller Interface to PCE

The ABNO Controller needs to be able to consult the PCE to determine what services can be provisioned in the network. There is no reason why this interface cannot be based on standard PCEP as defined in [RFC5440].

2.3.2.10. VNTM Interface to and from PCE

There are two interactions between the Virtual Network Topology Manager and the PCE:

The first interaction is used when VNTM wants to determine what LSPs can be set up in a network: in this case, it uses the standard PCEP interface [RFC5440] to make path computation requests.

The second interaction arises when a PCE determines that it cannot compute a requested path or notices that (according to some configured policy) a network is low on resources (for example, the capacity on some key link is nearly exhausted). In this case, the PCE may notify the VNTM, which may (again according to policy) act to construct more virtual topology. This second interface is not currently specified, although it may be that the protocol selected or designed to satisfy I2RS will provide suitable features (see Section 2.3.2.8); alternatively, an extension to the PCEP Notify message (PCNtf) [RFC5440] could be made.

2.3.2.11. ABNO Control Interfaces

The north-bound interface from the ABNO Controller is used by the NMS, OSS, and Application Service Coordinator to request services in the network in support of applications. The interface will also need to be able to report the asynchronous completion of service requests and convey changes in the status of services.

This interface will also need strong capabilities for security, authentication, and policy.

This interface is not currently specified. It needs to be a transactional interface that supports the specification of abstract services with adequate flexibility to facilitate easy extension and yet be concise and easily parsable.

It is possible that the protocol designed to satisfy I2RS will provide suitable features (see Section 2.3.2.8).

2.3.2.12. ABNO Provisioning Requests

Under some circumstances, the ABNO Controller may make requests directly to the Provisioning Manager. For example, if the Provisioning Manager is acting as an SDN Controller, then the ABNO Controller may use one of the APIs defined to allow requests to be made to the SDN Controller (such as the Floodlight REST API [Flood]). Alternatively, since the Provisioning Manager may also receive instructions from a Stateful PCE, the use of PCEP extensions might be appropriate in some cases [PCE-Init-LSP].

2.3.2.13. Policy Interfaces

As described in Section 2.3.1.4 and throughout this document, policy forms a critical component of the ABNO architecture. The role of policy will include enforcing the following rules and requirements:

- Adding resources on demand should be gated by the authorized capability.
- Client microflows should not trigger server-layer setup or allocation.
- Accounting capabilities should be supported.
- Security mechanisms for authorization of requests and capabilities are required.

Other policy-related functionality in the system might include the policy behavior of the routing and forwarding system, such as:

- ECMP behavior
- Classification of packets onto LSPs or QoS categories.

Various policy-capable architectures have been defined, including a framework for using policy with a PCE-enabled system [RFC5394]. However, the take-up of the IETF's Common Open Policy Service protocol (COPS) [RFC2748] has been poor.

New work will be needed to define all of the policy interfaces within the ABNO architecture. Work will also be needed to determine which are internal interfaces and which may be external and so in need of a protocol specification. There is some discussion that the I2RS protocol may support the configuration and manipulation of policies.

2.3.2.14. OAM and Reporting

The OAM Handler must interact with the network to perform several actions:

- Enabling OAM function within the network.
- Performing proactive OAM operations in the network.
- Receiving notifications of network events.

Any of the configuration and programmatic interfaces described in Section 2.3.2.1 may serve this purpose. NETCONF notifications are described in [RFC5277], and OpenFlow supports a number of asynchronous event notifications [ONF]. Additionally, Syslog [RFC5424] is a protocol for reporting events from the network, and IP Flow Information Export (IPFIX) [RFC7011] is designed to allow network statistics to be aggregated and reported.

The OAM Handler also correlates events reported from the network and reports them onward to the ABNO Controller (which can apply the information to the recovery of services that it has provisioned) and to the NMS, OSS, and Application Service Coordinator. The reporting mechanism used here can be essentially the same as the mechanism used when events are reported from the network; no new protocol is needed, although new data models may be required for technology-independent OAM reporting.

3. ABNO Use Cases

This section provides a number of examples of how the ABNO architecture can be applied to provide application-driven and NMS/OSS-driven network operations. The purpose of these examples is to give some concrete material to demonstrate the architecture so that it may be more easily comprehended, and to illustrate that the application of the architecture is achieved by "profiling" and by selecting only the relevant components and interfaces.

Similarly, it is not the intention that this section contain a complete list of all possible applications of ABNO. The examples are intended to broadly cover a number of applications that are commonly discussed, but this does not preclude other use cases.

The descriptions in this section are not fully detailed applicability statements for ABNO. It is anticipated that such applicability statements, for the use cases described and for other use cases, could be suitable material for separate documents.

3.1. Inter-AS Connectivity

The following use case describes how the ABNO framework can be used to set up an end-to-end MPLS service across multiple Autonomous Systems (ASes). Consider the simple network topology shown in Figure 2. The three ASes (ASa, ASb, and ASc) are connected at AS Border Routers (ASBRs) a1, a2, b1 through b4, c1, and c2. A source node (s) located in ASa is to be connected to a destination node (d) located in ASc. The optimal path for the LSP from s to d must be computed, and then the network must be triggered to set up the LSP.




Figure 2: Inter-AS Domain Topology with Hierarchical PCE (Parent PCE)

The following steps are performed to deliver the service within the ABNO architecture:

1. Request Management

As shown in Figure 3, the NMS/OSS issues a request to the ABNO Controller for a path between s and d. The ABNO Controller verifies that the NMS/OSS has sufficient rights to make the service request.




Figure 3: ABNO Request Management

2. Service Path Computation with Hierarchical PCE

The ABNO Controller needs to determine an end-to-end path for the LSP. Since the ASes will want to maintain a degree of confidentiality about their internal resources and topology, they will not share a TED and each will have its own PCE. In such a situation, the Hierarchical PCE (H-PCE) architecture described in [RFC6805] is applicable.

As shown in Figure 4, the ABNO Controller sends a request to the parent PCE for an end-to-end path. As described in [RFC6805], the parent PCE consults its TED, which shows the connectivity between

ASes. This helps it understand that the end-to-end path must cross each of ASa, ASb, and ASc, so it sends individual path computation requests to each of PCEs a, b, and c to determine the best options for crossing the ASes.

Each child PCE applies policy to the requests it receives to determine whether the request is to be allowed and to select the types of network resources that can be used in the computation result. For confidentiality reasons, each child PCE may supply its computation responses using a path key [RFC5520] to hide the details of the path segment it has computed.

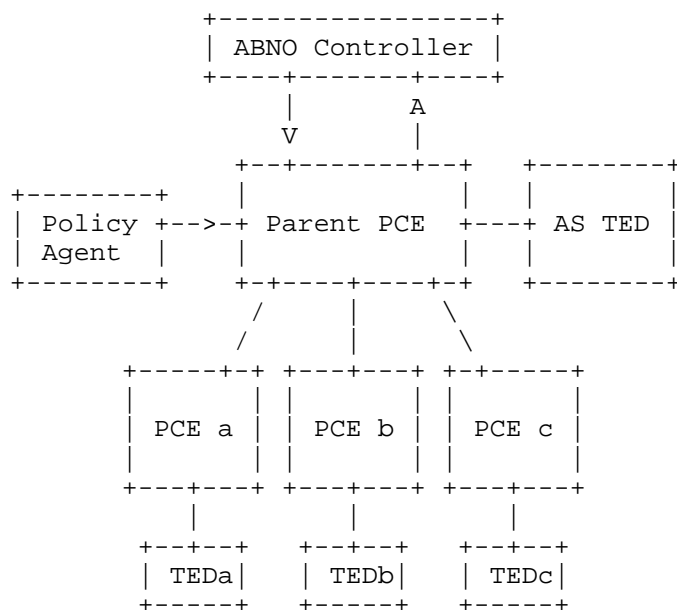


Figure 4: Path Computation Request with Hierarchical PCE

The parent PCE collates the responses from the children and applies its own policy to stitch them together into the best end-to-end path, which it returns as a response to the ABNO Controller.

3. Provisioning the End-to-End LSP

There are several options for how the end-to-end LSP gets provisioned in the ABNO architecture. Some of these are described below.

3a. Provisioning from the ABNO Controller with a Control Plane

Figure 5 shows how the ABNO Controller makes a request through the Provisioning Manager to establish the end-to-end LSP. As described in Section 2.3.2.6, these interactions can use the NETCONF protocol [RFC6241] or the extensions to PCEP described in [PCE-Init-LSP]. In either case, the provisioning request is sent to the head-end Label Switching Router (LSR), and that LSR signals in the control plane (using a protocol such as RSVP-TE [RFC3209]) to cause the LSP to be established.




Figure 5: Provisioning the End-to-End LSP

3b. Provisioning through Programming Network Resources

Another option is that the LSP is provisioned hop by hop from the Provisioning Manager using a mechanism such as ForCES [RFC5810] or OpenFlow [ONF] as described in Section 2.3.2.6. In this case, the picture is the same as that shown in Figure 5. The interaction between the ABNO Controller and the Provisioning Manager will be PCEP or NETCONF as described in option 3a, and the Provisioning Manager will be responsible for fanning out the requests to the individual network elements.

3c. Provisioning with an Active Parent PCE

The Active PCE is described in Section 2.3.1.7, based on the concepts expressed in [PCE-Init-LSP]. In this approach, the process described in option 3a is modified such that the PCE issues a direct PCEP command to the network, without a response being first returned to the ABNO Controller.

This situation is shown in Figure 6 and could be modified so that the Provisioning Manager still programs the individual network elements as described in option 3b.

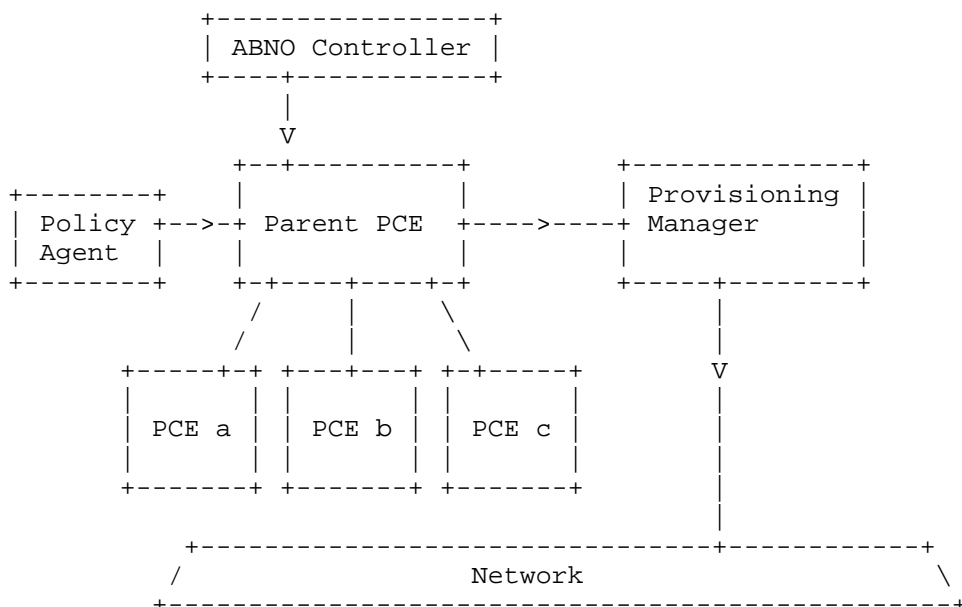


Figure 6: LSP Provisioning with an Active PCE

3d. Provisioning with Active Child PCEs and Segment Stitching

A mixture of the approaches described in options 3b and 3c can result in a combination of mechanisms to program the network to provide the end-to-end LSP. Figure 7 shows how each child PCE can be an Active PCE responsible for setting up an edge-to-edge LSP segment across one of the ASes. The ABNO Controller then uses the Provisioning Manager to program the inter-AS connections using ForCES or OpenFlow, and the LSP segments are stitched together following the ideas described in [RFC5150]. Philosophers may debate whether the parent PCE

in this model is active (instructing the children to provision LSP segments) or passive (requesting path segments that the children provision).

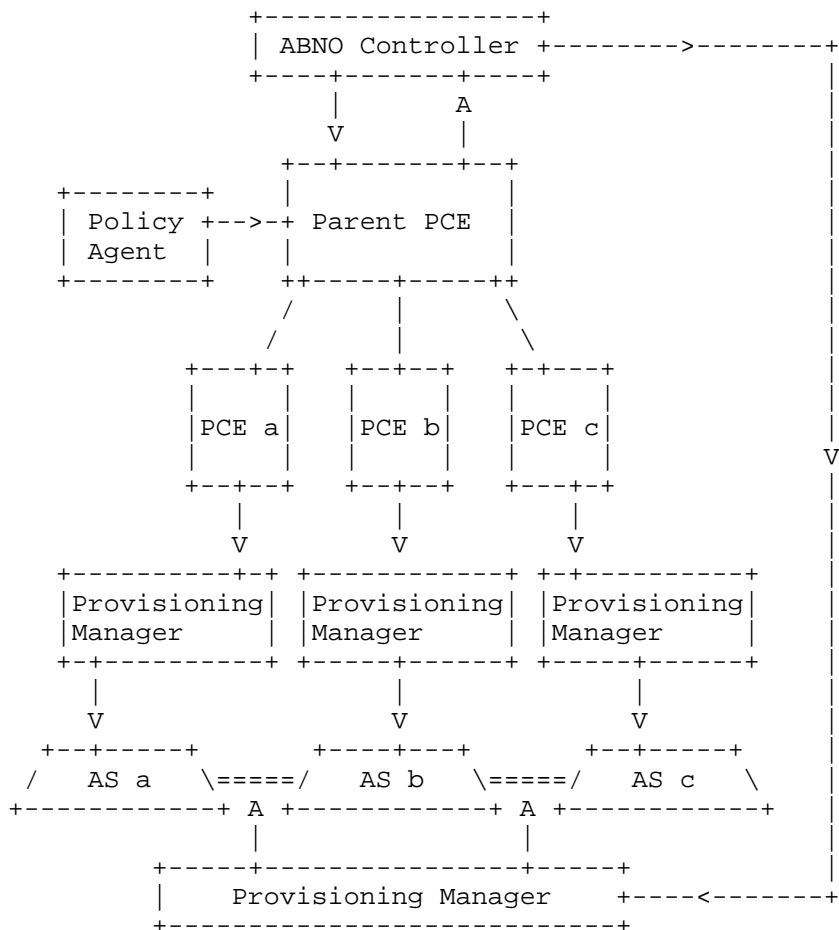


Figure 7: LSP Provisioning with Active Child PCEs and Stitching

4. Verification of Service

The ABNO Controller will need to ascertain that the end-to-end LSP has been set up as requested. In the case of a control plane being used to establish the LSP, the head-end LSR may send a notification (perhaps using PCEP) to report successful setup, but to be sure that the LSP is up, the ABNO Controller will request the OAM Handler to perform Continuity Check OAM in the data plane and report back that the LSP is ready to carry traffic.

5. Notification of Service Fulfillment

Finally, when the ABNO Controller is satisfied that the requested service is ready to carry traffic, it will notify the NMS/OSS. The delivery of the service may be further checked through auditing the network, as described in Section 2.3.2.7.

3.2. Multi-Layer Networking

Networks are typically constructed using multiple layers. These layers represent separations of administrative regions or of technologies and may also represent a distinction between client and server networking roles.

It is preferable to coordinate network resource control and utilization (i.e., consideration and control of multiple layers), rather than controlling and optimizing resources at each layer independently. This facilitates network efficiency and network automation and may be defined as inter-layer traffic engineering.

The PCE architecture supports inter-layer traffic engineering [RFC5623] and, in combination with the ABNO architecture, provides a suite of capabilities for network resource coordination across multiple layers.

The following use case demonstrates ABNO used to coordinate allocation of server-layer network resources to create virtual topology in a client-layer network in order to satisfy a request for end-to-end client-layer connectivity. Consider the simple multi-layer network in Figure 8.




Figure 8: Multi-Layer Network

There are six packet-layer routers (P1 through P6) and three optical-layer lambda switches (L1 through L3). There is connectivity in the packet layer between routers P1, P2, and P3, and also between routers P4, P5, and P6, but there is no packet-layer connectivity between these two islands of routers, perhaps because of a network failure or perhaps because all existing bandwidth between the islands has

already been used up. However, there is connectivity in the optical layer between switches L1, L2, and L3, and the optical network is connected out to routers P3 and P4 (they have optical line cards). In this example, a packet-layer connection (an MPLS LSP) is desired between P1 and P6.

In the ABNO architecture, the following steps are performed to deliver the service.

1. Request Management

As shown in Figure 9, the Application Service Coordinator issues a request for connectivity from P1 to P6 in the packet-layer network. That is, the Application Service Coordinator requests an MPLS LSP with a specific bandwidth to carry traffic for its application. The ABNO Controller verifies that the Application Service Coordinator has sufficient rights to make the service request.

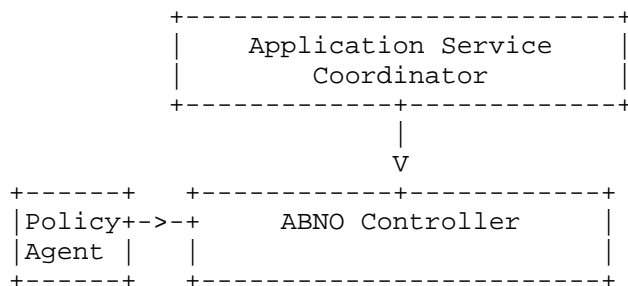


Figure 9: Application Service Coordinator Request Management

2. Service Path Computation in the Packet Layer

The ABNO Controller sends a path computation request to the packet-layer PCE to compute a suitable path for the requested LSP, as shown in Figure 10. The PCE uses the appropriate policy for the request and consults the TED for the packet layer. It determines that no path is immediately available.




Figure 10: Path Computation Request

3. Invocation of VNTM and Path Computation in the Optical Layer

After the path computation failure in step 2, instead of notifying the ABNO Controller of the failure, the PCE invokes the VNTM to see whether it can create the necessary link in the virtual network topology to bridge the gap.

As shown in Figure 11, the packet-layer PCE reports the connectivity problem to the VNTM, and the VNTM consults policy to determine what it is allowed to do. Assuming that the policy allows it, the VNTM asks the optical-layer PCE to find a path across the optical network that could be provisioned to provide a virtual link for the packet layer. In addressing this request, the optical-layer PCE consults a TED for the optical-layer network.

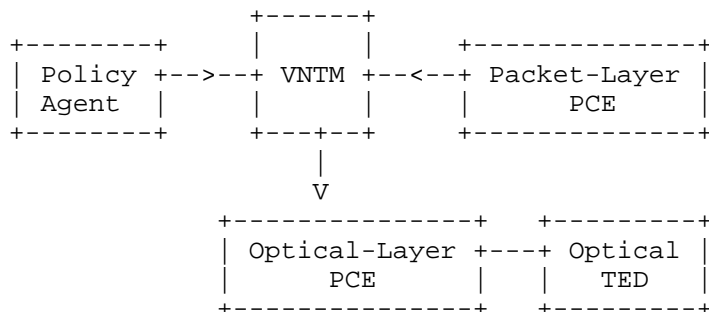


Figure 11: Invocation of VNTM and Optical-Layer Path Computation

4. Provisioning in the Optical Layer

Once a path has been found across the optical-layer network, it needs to be provisioned. The options follow those in step 3 of Section 3.1. That is, provisioning can be initiated by the optical-layer PCE or by its user, the VNTM. The command can be

sent to the head end of the optical LSP (P3) so that the control plane (for example, GMPLS RSVP-TE [RFC3473]) can be used to provision the LSP. Alternatively, the network resources can be provisioned directly, using any of the mechanisms described in Section 2.3.2.6.

5. Creation of Virtual Topology in the Packet Layer

Once the LSP has been set up in the optical layer, it can be made available in the packet layer as a virtual link. If the GMPLS signaling used the mechanisms described in [RFC6107], this process can be automated within the control plane; otherwise, it may require a specific instruction to the head-end router of the optical LSP (for example, through I2RS).

Once the virtual link is created as shown in Figure 12, it is advertised in the IGP for the packet-layer network, and the link will appear in the TED for the packet-layer network.




Figure 12: Advertisement of a New Virtual Link

6. Path Computation Completion and Provisioning in the Packet Layer

Now there are sufficient resources in the packet-layer network. The PCE for the packet layer can complete its work, and the MPLS LSP can be provisioned as described in Section 3.1.

7. Verification and Notification of Service Fulfillment

As discussed in Section 3.1, the ABNO Controller will need to verify that the end-to-end LSP has been correctly established before reporting service fulfillment to the Application Service Coordinator.

Furthermore, it is highly likely that service verification will be necessary before the optical-layer LSP can be put into service as a virtual link. Thus, the VNTM will need to coordinate with the OAM Handler to ensure that the LSP is ready for use.

3.2.1. Data Center Interconnection across Multi-Layer Networks

In order to support new and emerging cloud-based applications, such as real-time data backup, virtual machine migration, server clustering, or load reorganization, the dynamic provisioning and allocation of IT resources and the interconnection of multiple, remote Data Centers (DCs) is a growing requirement.

These operations require traffic being delivered between data centers, and, typically, the connections providing such inter-DC connectivity are provisioned using static circuits or dedicated leased lines, leading to an inefficiency in terms of resource utilization. Moreover, a basic requirement is that such a group of remote DCs can be operated logically as one.

In such environments, the data plane technology is operator and provider dependent. Their customers may rent lambda switch capable (LSC), packet switch capable (PSC), or time division multiplexing (TDM) services, and the application and usage of the ABNO architecture and Controller enable the required dynamic end-to-end network service provisioning, regardless of underlying service and transport layers.

Consequently, the interconnection of DCs may involve the operation, control, and management of heterogeneous environments: each DC site and the metro-core network segment used to interconnect them, with regard to not only the underlying data plane technology but also the control plane. For example, each DC site or domain could be controlled locally in a centralized way (e.g., via OpenFlow [ONF]), whereas the metro-core transport infrastructure is controlled by GMPLS. Although OpenFlow is specially adapted to single-domain intra-DC networks (packet-level control, lots of routing exceptions), a standardized GMPLS-based architecture would enable dynamic optical resource allocation and restoration in multi-domain (e.g., multi-vendor) core networks interconnecting distributed data centers.

The application of an ABNO architecture and related procedures would involve the following aspects:

1. Request from the Application Service Coordinator or NMS

As shown in Figure 13, the ABNO Controller receives a request from the Application Service Coordinator or from the NMS, in order to create a new end-to-end connection between two end points. The actual addressing of these end points is discussed in the next section. The ABNO Controller asks the PCE for a path between these two end points, after considering any applicable policy as defined by the Policy Agent (see Figure 1).

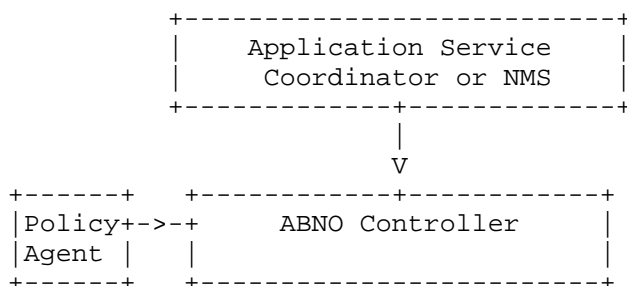


Figure 13: Application Service Coordinator Request Management

2. Address Mapping

In order to compute an end-to-end path, the PCE needs to have a unified view of the overall topology, which means that it has to consider and identify the actual end points with regard to the client network addresses. The ABNO Controller and/or the PCE may need to translate or map addresses from different address spaces. Depending on how the topology information is disseminated and gathered, there are two possible scenarios:

2a. The Application Layer Knows the Client Network Layer

Entities belonging to the application layer may have an interface with the TED or with an ALTO Server allowing those entities to map the high-level end points to network addresses. The mechanism used to enable this address correlation is out of scope for this document but relies on direct interfaces to other ABNO components in addition to the interface to the ABNO Controller.

In this scenario, the request from the NMS or Application Service Coordinator contains addresses in the client-layer network. Therefore, when the ABNO Controller requests the PCE to compute a path between two end points, the PCE is able to use the supplied addresses, compute the path, and continue the workflow in communication with the Provisioning Manager.

2b. The Application Layer Does Not Know the Client Network Layer

In this case, when the ABNO Controller receives a request from the NMS or Application Service Coordinator, the request contains only identifiers from the application-layer address space. In order for the PCE to compute an end-to-end path, these identifiers must be converted to addresses in the client-layer network. This translation can be performed by the ABNO Controller, which can access the TED and ALTO databases allowing the path computation request that it sends to the PCE to simply be contained within one network and TED. Alternatively, the computation request could use the application-layer identifiers, leaving the job of address mapping to the PCE.

Note that in order to avoid any confusion both approaches in this scenario require clear identification of the address spaces that are in use.

3. Provisioning Process

Once the path has been obtained, the Provisioning Manager receives a high-level provisioning request to provision the service. Since, in the considered use case, the network elements are not necessarily configured using the same protocol, the end-to-end path is split into segments, and the ABNO Controller coordinates or orchestrates the establishment by adapting and/or translating the abstract provisioning request to concrete segment requests by means of a VNTM or PCE that issues the corresponding commands or instructions. The provisioning may involve configuring the data plane elements directly or delegating the establishment of the underlying connection to a dedicated control plane instance responsible for that segment.

The Provisioning Manager could use a number of mechanisms to program the network elements, as shown in Figure 14. It learns which technology is used for the actual provisioning at each segment by either manual configuration or discovery.




Figure 14: Provisioning Process

4. Verification and Notification of Service Fulfillment

Once the end-to-end connectivity service has been provisioned, and after the verification of the correct operation of the service, the ABNO Controller needs to notify the Application Service Coordinator or NMS.

3.3. Make-before-Break

A number of different services depend on the establishment of a new LSP so that traffic supported by an existing LSP can be switched with little or no disruption. This section describes those use cases, presents a generic model for make-before-break within the ABNO architecture, and shows how each use case can be supported by using elements of the generic model.

3.3.1. Make-before-Break for Reoptimization

Make-before-break is a mechanism supported in RSVP-TE signaling where a new LSP is set up before the LSP it replaces is torn down [RFC3209]. This process has several benefits in situations such as reoptimization of in-service LSPs.

The process is simple, and the example shown in Figure 15 utilizes a Stateful PCE [Stateful-PCE] to monitor the network and take reoptimization actions when necessary. In this process, a service request is made to the ABNO Controller by a requester such as the OSS. The service request indicates that the LSP should be reoptimized under specific conditions according to policy. This allows the ABNO Controller to manage the sequence and prioritization of reoptimizing multiple LSPs using elements of Global Concurrent Optimization (GCO) as described in Section 3.4, and applying policies across the network so that, for instance, LSPs for delay-sensitive services are reoptimized first.

The ABNO Controller commissions the PCE to compute and set up the initial path.

Over time, the PCE monitors the changes in the network as reflected in the TED, and according to the configured policy may compute and set up a replacement path, using make-before-break within the network.

Once the new path has been set up and the network reports that it is being used correctly, the PCE tears down the old path and may report the reoptimization event to the ABNO Controller.

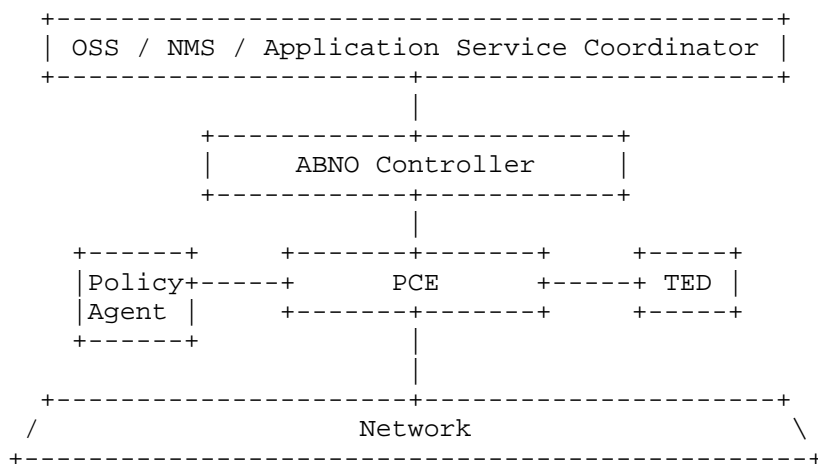


Figure 15: The Make-before-Break Process

3.3.2. Make-before-Break for Restoration

Make-before-break may also be used to repair a failed LSP where there is a desire to retain resources along some of the path, and where there is the potential for other LSPs to "steal" the resources if the

failed LSP is torn down first. Unlike the example in Section 3.3.1, this case addresses a situation where the service is interrupted, but this interruption arises from the break in service introduced by the network failure. Obviously, in the case of a point-to-multipoint LSP, the failure might only affect part of the tree and the disruption will only be to a subset of the destination leaves so that a make-before-break restoration approach will not cause disruption to the leaves that were not affected by the original failure.

Figure 16 shows the components that interact for this use case. A service request is made to the ABNO Controller by a requester such as the OSS. The service request indicates that the LSP may be restored after failure and should attempt to reuse as much of the original path as possible.

The ABNO Controller commissions the PCE to compute and set up the initial path. The ABNO Controller also requests the OAM Handler to initiate OAM on the LSP and to monitor the results.

At some point, the network reports a fault to the OAM Handler, which notifies the ABNO Controller.

The ABNO Controller commissions the PCE to compute a new path, reusing as much of the original path as possible, and the PCE sets up the new LSP.

Once the new path has been set up and the network reports that it is being used correctly, the ABNO Controller instructs the PCE to tear down the old path.

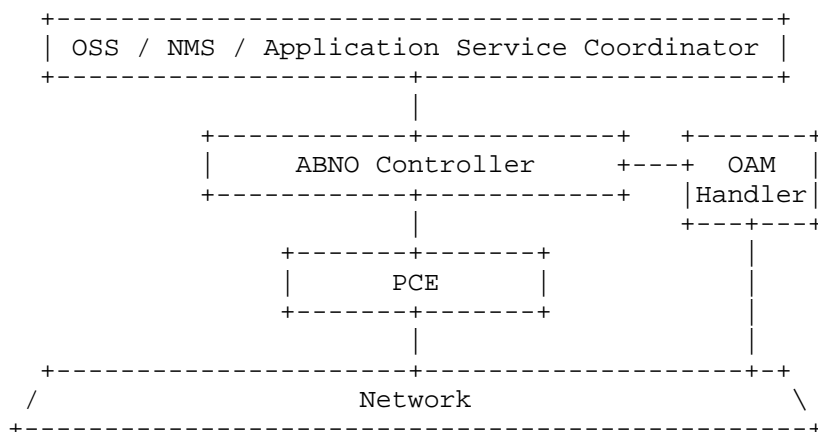


Figure 16: The Make-before-Break Restoration Process

3.3.3. Make-before-Break for Path Test and Selection

In a more complicated use case, an LSP may be monitored for a number of attributes, such as delay and jitter. When the LSP falls below a threshold, the traffic may be moved to another LSP that offers the desired (or at least a better) quality of service. To achieve this, it is necessary to establish the new LSP and test it, and because the traffic must not be interrupted, make-before-break must be used.

Moreover, it may be the case that no new LSP can provide the desired attributes and that a number of LSPs need to be tested so that the best can be selected. Furthermore, even when the original LSP is set up, it could be desirable to test a number of LSPs before deciding which should be used to carry the traffic.

Figure 17 shows the components that interact for this use case. Because multiple LSPs might exist at once, a distinct action is needed to coordinate which one carries the traffic, and this is the job of the I2RS Client acting under the control of the ABNO Controller.

The OAM Handler is responsible for initiating tests on the LSPs and for reporting the results back to the ABNO Controller. The OAM Handler can also check end-to-end connectivity test results across a multi-domain network even when each domain runs a different technology. For example, an end-to-end path might be achieved by stitching together an MPLS segment, an Ethernet/VLAN segment, another IP segment, etc.

Otherwise, the process is similar to that for reoptimization as discussed in Section 3.3.1.

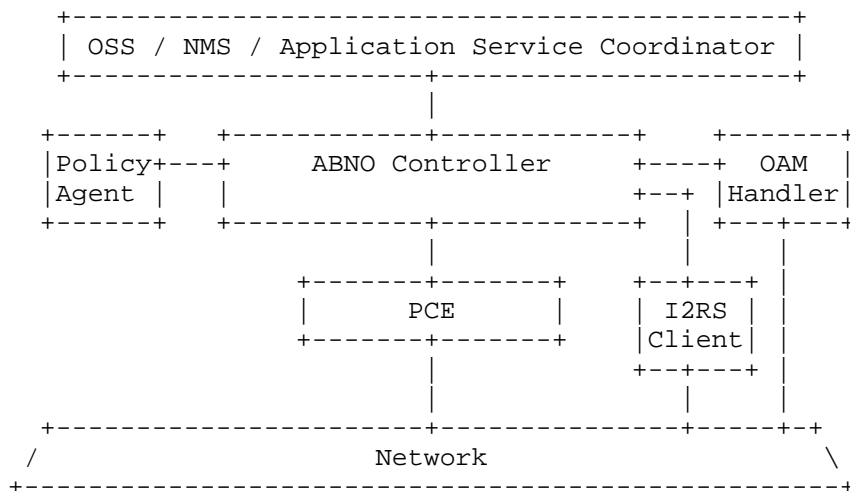


Figure 17: The Make-before-Break Path Test and Selection Process

The pseudocode that follows gives an indication of the interactions between ABNO components.

OSS requests quality-assured service

```
:Label1
```

DoWhile not enough LSPs (ABNO Controller)

Instruct PCE to compute and provision the LSP (ABNO Controller)

Create the LSP (PCE)

EndDo

```
:Label2
```

DoFor each LSP (ABNO Controller)

Test LSP (OAM Handler)

Report results to ABNO Controller (OAM Handler)

EndDo

Evaluate results of all tests (ABNO Controller)

Select preferred LSP and instruct I2RS Client (ABNO Controller)

```
Put traffic on preferred LSP (I2RS Client)
```

DoWhile too many LSPs (ABNO Controller)

Instruct PCE to tear down unwanted LSP (ABNO Controller)

Tear down unwanted LSP (PCE)

EndDo

```
DoUntil trigger (OAM Handler, ABNO Controller, Policy Agent)
  keep sending traffic (Network)
  Test LSP (OAM Handler)
EndDo

If there is already a suitable LSP (ABNO Controller)
  GoTo Label2
Else
  GoTo Label1
EndIf
```

3.4. Global Concurrent Optimization

Global Concurrent Optimization (GCO) is defined in [RFC5557] and represents a key technology for maximizing network efficiency by computing a set of traffic-engineered paths concurrently. A GCO path computation request will simultaneously consider the entire topology of the network, and the complete set of new LSPs together with their respective constraints. Similarly, GCO may be applied to recompute the paths of a set of existing LSPs.

GCO may be requested in a number of scenarios. These include:

- o Routing of new services where the PCE should consider other services or network topology.
- o A reoptimization of existing services due to fragmented network resources or suboptimized placement of sequentially computed services.
- o Recovery of connectivity for bulk services in the event of a catastrophic network failure.

A service provider may also want to compute and deploy new bulk services based on a predicted traffic matrix. The GCO functionality and capability to perform concurrent computation provide a significant network optimization advantage, thus utilizing network resources optimally and avoiding blocking.

The following use case shows how the ABNO architecture and components are used to achieve concurrent optimization across a set of services.

3.4.1. Use Case: GCO with MPLS LSPs

When considering the GCO path computation problem, we can split the GCO objective functions into three optimization categories:

- o Minimize aggregate Bandwidth Consumption (MBC).
- o Minimize the load of the Most Loaded Link (MLL).
- o Minimize Cumulative Cost of a set of paths (MCC).

This use case assumes that the GCO request will be offline and be initiated from an NMS/OSS; that is, it may take significant time to compute the service, and the paths reported in the response may want to be verified by the user before being provisioned within the network.

1. Request Management

The NMS/OSS issues a request for new service connectivity for bulk services. The ABNO Controller verifies that the NMS/OSS has sufficient rights to make the service request and apply a GCO attribute with a request to Minimize aggregate Bandwidth Consumption (MBC), as shown in Figure 18.




Figure 18: NMS Request to ABNO Controller

- 1a. Each service request has a source, destination, and bandwidth request. These service requests are sent to the ABNO Controller and categorized as GCO requests. The PCE uses the appropriate policy for each request and consults the TED for the packet layer.

2. Service Path Computation in the Packet Layer

To compute a set of services for the GCO application, PCEP supports synchronization vector (SVEC) lists for synchronized dependent path computations as defined in [RFC5440] and described in [RFC6007].

- 2a. The ABNO Controller sends the bulk service request to the GCO-capable packet-layer PCE using PCEP messaging. The PCE uses the appropriate policy for the request and consults the TED for the packet layer, as shown in Figure 19.

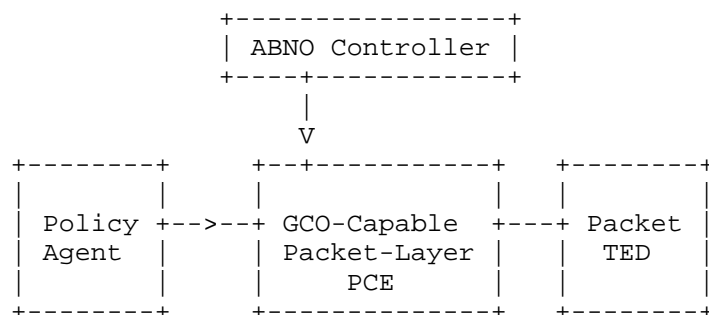


Figure 19: Path Computation Request from GCO-Capable PCE

- 2b. Upon receipt of the bulk (GCO) service requests, the PCE applies the MBC objective function and computes the services concurrently.
- 2c. Once the requested GCO service path computation completes, the PCE sends the resulting paths back to the ABNO Controller. The response includes a fully computed explicit path for each service (TE LSP).

3. The concurrently computed solution received from the PCE is sent back to the NMS/OSS by the ABNO Controller as a PCEP response, as shown in Figure 20. The NMS/OSS user can then check the candidate paths and either provision the new services or save the solution for deployment in the future.




Figure 20: ABNO Sends Solution to the NMS/OSS

3.5. Adaptive Network Management (ANM)

The ABNO architecture provides the capability for reactive network control of resources relying on classification, profiling, and prediction based on current demands and resource utilization. Server-layer transport network resources, such as Optical Transport Network (OTN) time-slicing [G.709], or the fine granularity grid of wavelengths with variable spectral bandwidth (flexi-grid) [G.694.1], can be manipulated to meet current and projected demands in a model called Elastic Optical Networks (EON) [EON].

EON provides spectrum-efficient and scalable transport by introducing flexible granular traffic grooming in the optical frequency domain. This is achieved using arbitrary contiguous concatenation of the optical spectrum that allows the creation of custom-sized bandwidth. This bandwidth is defined in slots of 12.5 GHz.

Adaptive Network Management (ANM) with EON allows appropriately sized optical bandwidth to be allocated to an end-to-end optical path. In flexi-grid, the allocation is performed according to the traffic volume, optical modulation format, and associated reach, or following user requests, and can be achieved in a highly spectrum-efficient and scalable manner. Similarly, OTN provides for flexible and granular provisioning of bandwidth on top of Wavelength Switched Optical Networks (WSONs).

To efficiently use optical resources, a system is required that can monitor network resources and decide the optimal network configuration based on the status, bandwidth availability, and user service. We call this ANM.

3.5.1. ANM Trigger

There are different reasons to trigger an adaptive network management process; these include:

- o Measurement: Traffic measurements can be used in order to cause spectrum allocations that fit the traffic needs as efficiently as possible. This function may be influenced by measuring the IP router traffic flows, by examining traffic engineering or link state databases, by usage thresholds for critical links in the network, or by requests from external entities. Nowadays, network operators have active monitoring probes in the network that store their results in the OSS. The OSS or OAM Handler components activate this measurement-based trigger, so the ABNO Controller would not be directly involved in this case.
- o Human: Operators may request ABNO to run an adaptive network planning process via an NMS.
- o Periodic: An adaptive network planning process can be run periodically to find an optimum configuration.

An ABNO Controller would receive a request from an OSS or NMS to run an adaptive network manager process.

3.5.2. Processing Request and GCO Computation

Based on the human or periodic trigger requests described in the previous section, the OSS or NMS will send a request to the ABNO Controller to perform EON-based GCO. The ABNO Controller will select a set of services to be reoptimized and choose an objective function that will deliver the best use of network resources. In making these choices, the ABNO Controller is guided by network-wide policy on the use of resources, the definition of optimization, and the level of perturbation to existing services that is tolerable.

This request for GCO is passed to the PCE, along the lines of the description in Section 3.4. The PCE can then consider the end-to-end paths and every channel's optimal spectrum assignment in order to satisfy traffic demands and optimize the optical spectrum consumption within the network.

The PCE will operate on the TED but is likely to also be stateful so that it knows which LSPs correspond to which waveband allocations on which links in the network. Once the PCE arrives at an answer, it returns a set of potential paths to the ABNO Controller, which passes them on to the NMS or OSS to supervise/select the subsequent path setup/modification process.

This exchange is shown in Figure 21. Note that the figure does not show the interactions used by the OSS/NMS for establishing or modifying LSPs in the network.

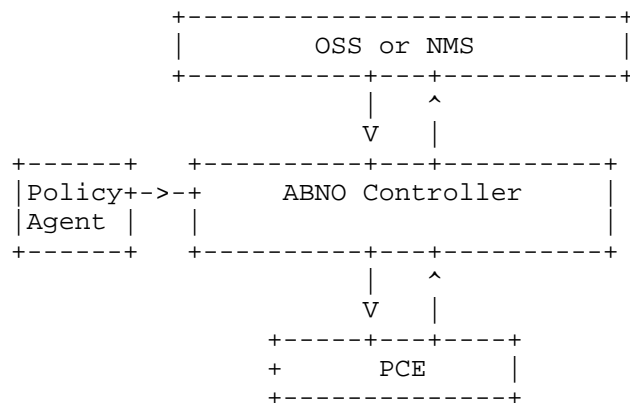


Figure 21: Adaptive Network Management with Human Intervention

3.5.3. Automated Provisioning Process

Although most network operations are supervised by the operator, there are some actions that may not require supervision, like a simple modification of a modulation format in a Bit-rate Variable Transponder (BVT) (to increase the optical spectrum efficiency or reduce energy consumption). In this process, where human intervention is not required, the PCE sends the Provisioning Manager a new configuration to configure the network elements, as shown in Figure 22.

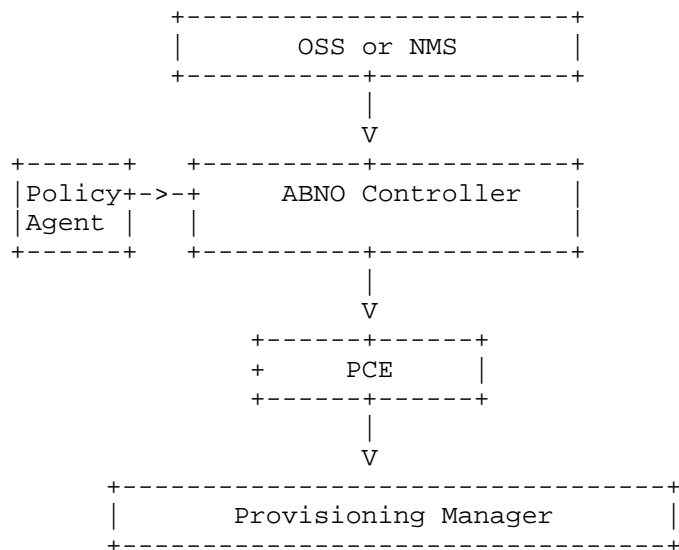


Figure 22: Adaptive Network Management without Human Intervention

3.6. Pseudowire Operations and Management

Pseudowires in an MPLS network [RFC3985] operate as a form of layered network over the connectivity provided by the MPLS network. The pseudowires are carried by LSPs operating as transport tunnels, and planning is necessary to determine how those tunnels are placed in the network and which tunnels are used by any pseudowire.

This section considers four use cases: multi-segment pseudowires, path-diverse pseudowires, path-diverse multi-segment pseudowires, and pseudowire segment protection. Section 3.6.5 describes the applicability of the ABNO architecture to these four use cases.

3.6.1. Multi-Segment Pseudowires

[RFC5254] describes the architecture for multi-segment pseudowires. An end-to-end service, as shown in Figure 23, can consist of a series of stitched segments shown in the figure as AC, PW1, PW2, PW3, and AC. Each pseudowire segment is stitched at a "stitching Provider Edge" (S-PE): for example, PW1 is stitched to PW2 at S-PE1. Each access circuit (AC) is stitched to a pseudowire segment at a "terminating PE" (T-PE): for example, PW1 is stitched to the AC at T-PE1.

Each pseudowire segment is carried across the MPLS network in an LSP operating as a transport tunnel: for example, PW1 is carried in LSP1. The LSPs between PE nodes may traverse different MPLS networks with the PEs as border nodes, or the PEs may lie within the network such that each LSP spans only part of the network.




Figure 23: Multi-Segment Pseudowire

While the topology shown in Figure 23 is easy to navigate, the reality of a deployed network can be considerably more complex. The topology in Figure 24 shows a small mesh of PEs. The links between the PEs are not physical links but represent the potential of MPLS LSPs between the PEs.

When establishing the end-to-end service between Customer Edge nodes (CEs) CE1 and CE2, some choice must be made about which PEs to use. In other words, a path computation must be made to determine the pseudowire segment "hops", and then the necessary LSP tunnels must be established to carry the pseudowire segments that will be stitched together.

Of course, each LSP may itself require a path computation decision to route it through the MPLS network between PEs.

The choice of path for the multi-segment pseudowire will depend on such issues as:

- MPLS connectivity
- MPLS bandwidth availability
- pseudowire stitching capability and capacity at PEs
- policy and confidentiality considerations for use of PEs




Figure 24: Multi-Segment Pseudowire Network Topology

3.6.2. Path-Diverse Pseudowires

The connectivity service provided by a pseudowire may need to be resilient to failure. In many cases, this function is provided by provisioning a pair of pseudowires carried by path-diverse LSPs across the network, as shown in Figure 25 (the terminology is inherited directly from [RFC3985]). Clearly, in this case, the challenge is to keep the two LSPs (LSP1 and LSP2) disjoint within the MPLS network. This problem is not different from the normal MPLS path-diversity problem.

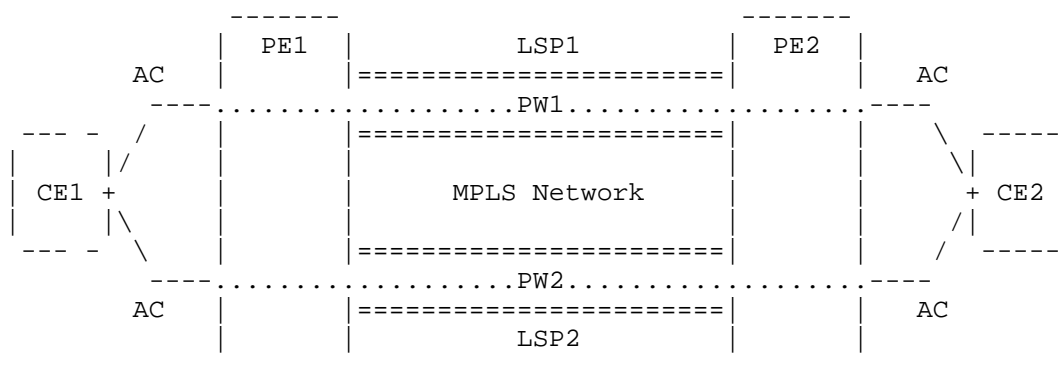


Figure 25: Path-Diverse Pseudowires

The path-diverse pseudowire is developed in Figure 26 by the "dual-homing" of each CE through more than one PE. The requirement for LSP path diversity is exactly the same, but it is complicated by the LSPs having distinct end points. In this case, the head-end router (e.g., PE1) cannot be relied upon to maintain the path diversity through the signaling protocol because it is aware of the path of only one of the LSPs. Thus, some form of coordinated path computation approach is needed.

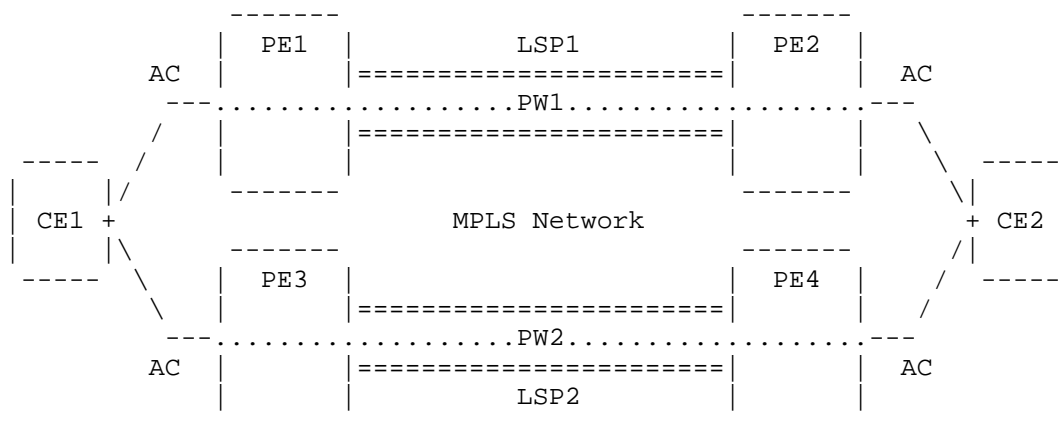


Figure 26: Path-Diverse Pseudowires with Disjoint PEs

3.6.3. Path-Diverse Multi-Segment Pseudowires

Figure 27 shows how the services in the previous two sections may be combined to offer end-to-end diverse paths in a multi-segment environment. To offer end-to-end resilience to failure, two entirely diverse, end-to-end multi-segment pseudowires may be needed.

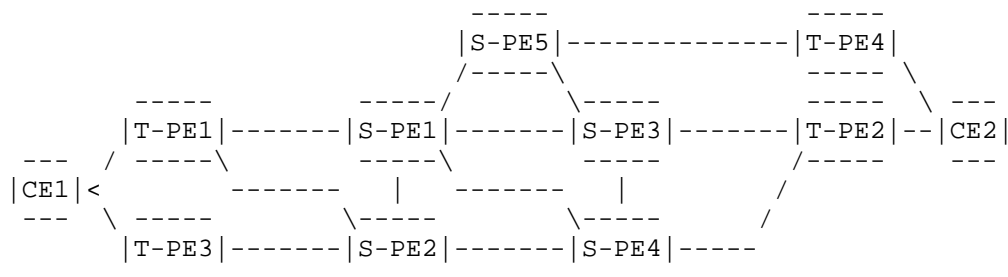


Figure 27: Path-Diverse Multi-Segment Pseudowire Network Topology

Just as in any diverse-path computation, the selection of the first path needs to be made with awareness of the fact that a second, fully diverse path is also needed. If a sequential computation was applied to the topology in Figure 27, the first path CE1,T-PE1,S-PE1,S-PE3,T-PE2,CE2 would make it impossible to find a second path that was fully diverse from the first.

But the problem is complicated by the multi-layer nature of the network. It is not enough that the PEs are chosen to be diverse because the LSP tunnels between them might share links within the MPLS network. Thus, a multi-layer planning solution is needed to achieve the desired level of service.

3.6.4. Pseudowire Segment Protection

An alternative to the end-to-end pseudowire protection service enabled by the mechanism described in Section 3.6.3 can be achieved by protecting individual pseudowire segments or PEs. For example, in Figure 27, the pseudowire between S-PE1 and S-PE5 may be protected by a pair of stitched segments running between S-PE1 and S-PE5, and between S-PE5 and S-PE3. This is shown in detail in Figure 28.




Figure 28: Fragment of a Segment-Protected Multi-Segment Pseudowire

The determination of pseudowire protection segments requires coordination and planning, and just as in Section 3.6.5, this planning must be cognizant of the paths taken by LSPs through the underlying MPLS networks.

3.6.5. Applicability of ABNO to Pseudowires

The ABNO architecture lends itself well to the planning and control of pseudowires in the use cases described above. The user or application needs a single point at which it requests services: the ABNO Controller. The ABNO Controller can ask a PCE to draw on the topology of pseudowire stitching-capable PEs as well as additional information regarding PE capabilities, such as load on PEs and administrative policies, and the PCE can use a series of TEDs or other PCEs for the underlying MPLS networks to determine the paths of the LSP tunnels. At the time of this writing, PCEP does not support

path computation requests and responses concerning pseudowires, but the concepts are very similar to existing uses and the necessary extensions would be very small.

Once the paths have been computed, a number of different provisioning systems can be used to instantiate the LSPs and provision the pseudowires under the control of the Provisioning Manager. The ABNO Controller will use the I2RS Client to instruct the network devices about what traffic should be placed on which pseudowires and, in conjunction with the OAM Handler, can ensure that failure events are handled correctly, that service quality levels are appropriate, and that service protection levels are maintained.

In many respects, the pseudowire network forms an overlay network (with its own TED and provisioning mechanisms) carried by underlying packet networks. Further client networks (the pseudowire payloads) may be carried by the pseudowire network. Thus, the problem space being addressed by ABNO in this case is a classic multi-layer network.

3.7. Cross-Stratum Optimization (CSO)

Considering the term "stratum" to broadly differentiate the layers of most concern to the application and to the network in general, the need for Cross-Stratum Optimization (CSO) arises when the application stratum and network stratum need to be coordinated to achieve operational efficiency as well as resource optimization in both application and network strata.

Data center-based applications can provide a wide variety of services such as video gaming, cloud computing, and grid applications. High-bandwidth video applications are also emerging, such as remote medical surgery, live concerts, and sporting events.

This use case for the ABNO architecture is mainly concerned with data center applications that make substantial bandwidth demands either in aggregate or individually. In addition, these applications may need specific bounds on QoS-related parameters such as latency and jitter.

3.7.1. Data Center Network Operation

Data centers come in a wide variety of sizes and configurations, but all contain compute servers, storage, and application control. Data centers offer application services to end-users, such as video gaming, cloud computing, and others. Since the data centers used to provide application services may be distributed around a network, the decisions about the control and management of application services, such as where to instantiate another service instance or to which

data center a new client is assigned, can have a significant impact on the state of the network. Conversely, the capabilities and state of the network can have a major impact on application performance.

These decisions are typically made by applications with very little or no information concerning the underlying network. Hence, such decisions may be suboptimal from the application's point of view or considering network resource utilization and quality of service.

Cross-Stratum Optimization is the process of optimizing both the application experience and the network utilization by coordinating decisions in the application stratum and the network stratum. Application resources can be roughly categorized into computing resources (i.e., servers of various types and granularities, such as Virtual Machines (VMs), memory, and storage) and content (e.g., video, audio, databases, and large data sets). By "network stratum" we mean the IP layer and below (e.g., MPLS, Synchronous Digital Hierarchy (SDH), OTN, WDM). The network stratum has resources that include routers, switches, and links. We are particularly interested in further unleashing the potential presented by MPLS and GMPLS control planes at the lower network layers in response to the high aggregate or individual demands from the application layer.

This use case demonstrates that the ABNO architecture can allow cross-stratum application/network optimization for the data center use case. Other forms of Cross-Stratum Optimization (for example, for peer-to-peer applications) are out of scope.

3.7.1.1. Virtual Machine Migration

A key enabler for data center cost savings, consolidation, flexibility, and application scalability has been the technology of compute virtualization provided through Virtual Machines (VMs). To the software application, a VM looks like a dedicated processor with dedicated memory and a dedicated operating system.

VMs not only offer a unit of compute power but also provide an "application environment" that can be replicated, backed up, and moved. Different VM configurations may be offered that are optimized for different types of processing (e.g., memory intensive, throughput intensive).

VMs may be moved between compute resources in a data center and could be moved between data centers. VM migration serves to balance load across data center resources and has several modes:

- (i) scheduled vs. dynamic;
- (ii) bulk vs. sequential;
- (iii) point-to-point vs. point-to-multipoint

While VM migration may solve problems of load or planned maintenance within a data center, it can also be effective to reduce network load around the data center. But the act of migrating VMs, especially between data centers, can impact the network and other services that are offered.

For certain applications such as disaster recovery, bulk migration is required on the fly, which may necessitate concurrent computation and path setup dynamically.

Thus, application stratum operations must also take into account the situation in the network stratum, even as the application stratum actions may be driven by the status of the network stratum.

3.7.1.2. Load Balancing

Application servers may be instantiated in many data centers located in different parts of the network. When an end-user makes an application request, a decision has to be made about which data center should host the processing and storage required to meet the request. One of the major drivers for operating multiple data centers (rather than one very large data center) is so that the application will run on a machine that is closer to the end-users and thus improve the user experience by reducing network latency. However, if the network is congested or the data center is overloaded, this strategy can backfire.

Thus, the key factors to be considered in choosing the server on which to instantiate a VM for an application include:

- The utilization of the servers in the data center
- The network load conditions within a data center
- The network load conditions between data centers
- The network conditions between the end-user and data center

Again, the choices made in the application stratum need to consider the situation in the network stratum.

3.7.2. Application of the ABNO Architecture

This section shows how the ABNO architecture is applicable to the cross-stratum data center issues described in Section 3.7.1.

Figure 29 shows a diagram of an example data center-based application. A carrier network provides access for an end-user through PE4. Three data centers (DC1, DC2, and DC3) are accessed through different parts of the network via PE1, PE2, and PE3.

The Application Service Coordinator receives information from the end-user about the desired services and converts this information to service requests that it passes to the ABNO Controller. The end-users may already know which data center they wish to use, or the Application Service Coordinator may be able to make this determination; otherwise, the task of selecting the data center must be performed by the ABNO Controller, and this may utilize a further database (see Section 2.3.1.8) to contain information about server loads and other data center parameters.

The ABNO Controller examines the network resources using information gathered from the other ABNO components and uses those components to configure the network to support the end-user's needs.

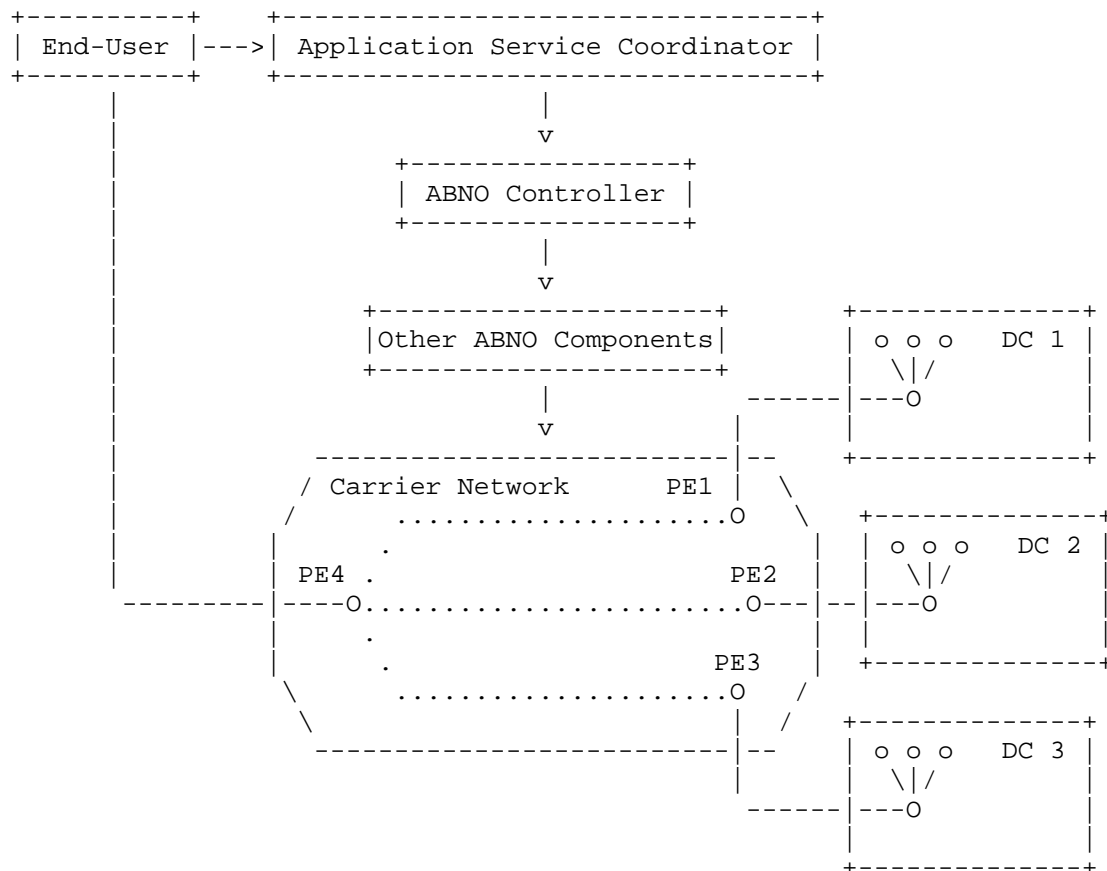


Figure 29: The ABNO Architecture in the Context of Cross-Stratum Optimization for Data Centers

3.7.2.1. Deployed Applications, Services, and Products

The ABNO Controller will need to utilize a number of components to realize the CSO functions described in Section 3.7.1.

The ALTO Server provides information about topological proximity and appropriate geographical location to servers with respect to the underlying networks. This information can be used to optimize the selection of peer location, which will help reduce the path of IP traffic or can contain it within specific service providers' networks. ALTO in conjunction with the ABNO Controller and the Application Service Coordinator can address general problems such as the selection of application servers based on resource availability and usage of the underlying networks.

The ABNO Controller can also formulate a view of current network load from the TED and from the OAM Handler (for example, by running diagnostic tools that measure latency, jitter, and packet loss). This view obviously influences not just how paths from the end-user to the data center are provisioned but can also guide the selection of which data center should provide the service and possibly even the points of attachment to be used by the end-user and to reach the chosen data center. A view of how the PCE can fit in with CSO is provided in [CSO-PCE], on which the content of Figure 29 is based.

As already discussed, the combination of the ABNO Controller and the Application Service Coordinator will need to be able to select (and possibly migrate) the location of the VM that provides the service for the end-user. Since a common technique used to direct the end-user to the correct VM/server is to employ DNS redirection, an important capability of the ABNO Controller will be the ability to program the DNS servers accordingly.

Furthermore, as already noted in other sections of this document, the ABNO Controller can coordinate the placement of traffic within the network to achieve load balancing and to provide resilience to failures. These features can be used in conjunction with the functions discussed above, to ensure that the placement of new VMs, the traffic that they generate, and the load caused by VM migration can be carried by the network and do not disrupt existing services.

3.8. ALTO Server

The ABNO architecture allows use cases with joint network and application-layer optimization. In such a use case, an application is presented with an abstract network topology containing only information relevant to the application. The application computes its application-layer routing according to its application objective. The application may interact with the ABNO Controller to set up explicit LSPs to support its application-layer routing.

The following steps are performed to illustrate such a use case.

1. Application Request of Application-Layer Topology

Consider the network shown in Figure 30. The network consists of five nodes and six links.

The application, which has end points hosted at N0, N1, and N2, requests network topology so that it can compute its application-layer routing, for example, to maximize the throughput of content replication among end points at the three sites.




Figure 30: Raw Network Topology

The request arrives at the ABNO Controller, which forwards the request to the ALTO Server component. The ALTO Server consults the Policy Agent, the TED, and the PCE to return an abstract, application-layer topology.

For example, the policy may specify that the bandwidth exposed to an application may not exceed 40 Mbps. The network has precomputed that the route from N0 to N2 should use the path N0->N3->N2, according to goals such as GCO (see Section 3.4). The ALTO Server can then produce a reduced topology for the application, such as the topology shown in Figure 31.




Figure 31: Reduced Graph for a Particular Application

The ALTO Server uses the topology and existing routing to compute an abstract network map consisting of three PIDs. The pair-wise bandwidth as well as shared bottlenecks will be computed from the internal network topology and reflected in cost maps.

2. Application Computes Application Overlay

Using the abstract topology, the application computes an application-layer routing. For concreteness, the application may compute a spanning tree to maximize the total bandwidth from N0 to N2. Figure 32 shows an example of application-layer routing, using a route of N0->N1->N2 for 35 Mbps and N0->N2 for 30 Mbps, for a total of 65 Mbps.




Figure 32: Application-Layer Spanning Tree

3. Application Path Set Up by the ABNO Controller

The application may submit its application routes to the ABNO Controller to set up explicit LSPs to support its operation. The ABNO Controller consults the ALTO maps to map the application-layer routing back to internal network topology and then instructs the Provisioning Manager to set up the paths. The ABNO Controller may re-trigger GCO to reoptimize network traffic engineering.

3.9. Other Potential Use Cases

This section serves as a placeholder for other potential use cases that might get documented in future documents.

3.9.1. Traffic Grooming and Regrooming

This use case could cover the following scenarios:

- Nested LSPs
- Packet Classification (IP flows into LSPs at edge routers)
- Bucket Stuffing
- IP Flows into ECMP Hash Bucket

3.9.2. Bandwidth Scheduling

Bandwidth scheduling consists of configuring LSPs based on a given time schedule. This can be used to support maintenance or operational schedules or to adjust network capacity based on traffic pattern detection.

The ABNO framework provides the components to enable bandwidth scheduling solutions.

4. Survivability and Redundancy within the ABNO Architecture

The ABNO architecture described in this document is presented in terms of functional units. Each unit could be implemented separately or bundled with other units into single programs or products. Furthermore, each implemented unit or bundle could be deployed on a separate device (for example, a network server) or on a separate virtual machine (for example, in a data center), or groups of programs could be deployed on the same processor. From the point of view of the architectural model, these implementation and deployment choices are entirely unimportant.

Similarly, the realization of a functional component of the ABNO architecture could be supported by more than one instance of an implementation, or by different instances of different implementations that provide the same or similar function. For example, the PCE component might have multiple instantiations for sharing the processing load of a large number of computation requests, and different instances might have different algorithmic capabilities so that one instance might serve parallel computation requests for disjoint paths, while another instance might have the capability to compute optimal point-to-multipoint paths.

This ability to have multiple instances of ABNO components also enables resiliency within the model, since in the event of the failure of one instance of one component (because of software failure, hardware failure, or connectivity problems) other instances can take over. In some circumstances, synchronization between instances of components may be needed in order to facilitate seamless resiliency.

How these features are achieved in an ABNO implementation or deployment is outside the scope of this document.

5. Security Considerations

The ABNO architecture describes a network system, and security must play an important part.

The first consideration is that the external protocols (those shown as entering or leaving the big box in Figure 1) must be appropriately secured. This security will include authentication and authorization to control access to the different functions that the ABNO system can perform, to enable different policies based on identity, and to manage the control of the network devices.

Secondly, the internal protocols that are used between ABNO components must also have appropriate security, particularly when the components are implemented on separate network nodes.

Considering that the ABNO system contains a lot of data about the network, the services carried by the network, and the services delivered to customers, access to information held in the system must be carefully managed. Since such access will be largely through the external protocols, the policy-based controls enabled by authentication will be powerful. But it should also be noted that any data sent from the databases in the ABNO system can reveal details of the network and should, therefore, be considered as a candidate for encryption. Furthermore, since ABNO components can access the information stored in the database, care is required to ensure that all such components are genuine and to consider encrypting data that flows between components when they are implemented at remote nodes.

The conclusion is that all protocols used to realize the ABNO architecture should have rich security features.

6. Manageability Considerations

The whole of the ABNO architecture is essentially about managing the network. In this respect, there is very little extra to say. ABNO provides a mechanism to gather and collate information about the network, reporting it to management applications, storing it for future inspection, and triggering actions according to configured policies.

The ABNO system will, itself, need monitoring and management. This can be seen as falling into several categories:

- Management of external protocols
- Management of internal protocols
- Management and monitoring of ABNO components
- Configuration of policy to be applied across the ABNO system

7. Informative References

- [BGP-LS] Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", Work in Progress, draft-ietf-idr-ls-distribution-10, January 2015.
- [CSO-PCE] Dhody, D., Lee, Y., Contreras, LM., Gonzalez de Dios, O., and N. Ciulli, "Cross Stratum Optimization enabled Path Computation", Work in Progress, draft-dhody-pce-cso-enabled-path-computation-07, January 2015.
- [EON] Gerstel, O., Jinno, M., Lord, A., and S.J.B. Yoo, "Elastic optical networking: a new dawn for the optical layer?", IEEE Communications Magazine, Volume 50, Issue 2, ISSN 0163-6804, February 2012.
- [Flood] Project Floodlight, "Floodlight REST API", <<http://www.projectfloodlight.org>>.
- [G.694.1] ITU-T Recommendation G.694.1, "Spectral grids for WDM applications: DWDM frequency grid", February 2012.
- [G.709] ITU-T Recommendation G.709, "Interface for the optical transport network", February 2012.
- [I2RS-Arch] Atlas, A., Halpern, J., Hares, S., Ward, D., and T. Nadeau, "An Architecture for the Interface to the Routing System", Work in Progress, draft-ietf-i2rs-architecture-09, March 2015.
- [I2RS-PS] Atlas, A., Ed., Nadeau, T., Ed., and D. Ward, "Interface to the Routing System Problem Statement", Work in Progress, draft-ietf-i2rs-problem-statement-06, January 2015.

- [ONF] Open Networking Foundation, "OpenFlow Switch Specification Version 1.4.0 (Wire Protocol 0x05)", October 2013.
- [PCE-Init-LSP] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", Work in Progress, draft-ietf-pce-pce-initiated-lsp-03, March 2015.
- [RESTCONF] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", Work in Progress, draft-ietf-netconf-restconf-04, January 2015.
- [RFC2748] Durham, D., Ed., Boyle, J., Cohen, R., Herzog, S., Rajan, R., and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000, <<http://www.rfc-editor.org/info/rfc2748>>.
- [RFC2753] Yavatkar, R., Pendarakis, D., and R. Guerin, "A Framework for Policy-based Admission Control", RFC 2753, January 2000, <<http://www.rfc-editor.org/info/rfc2753>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC3292] Doria, A., Hellstrand, F., Sundell, K., and T. Worster, "General Switch Management Protocol (GSMP) V3", RFC 3292, June 2002, <<http://www.rfc-editor.org/info/rfc3292>>.
- [RFC3412] Case, J., Harrington, D., Presuhn, R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3412, December 2002, <<http://www.rfc-editor.org/info/rfc3412>>.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003, <<http://www.rfc-editor.org/info/rfc3473>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.

- [RFC3746] Yang, L., Dantu, R., Anderson, T., and R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework", RFC 3746, April 2004, <<http://www.rfc-editor.org/info/rfc3746>>.
- [RFC3985] Bryant, S., Ed., and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005, <<http://www.rfc-editor.org/info/rfc3985>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, February 2008, <<http://www.rfc-editor.org/info/rfc5150>>.
- [RFC5212] Shiimoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, July 2008, <<http://www.rfc-editor.org/info/rfc5212>>.
- [RFC5254] Bitar, N., Ed., Bocci, M., Ed., and L. Martini, Ed., "Requirements for Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)", RFC 5254, October 2008, <<http://www.rfc-editor.org/info/rfc5254>>.
- [RFC5277] Chisholm, S. and H. Trevino, "NETCONF Event Notifications", RFC 5277, July 2008, <<http://www.rfc-editor.org/info/rfc5277>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, December 2008, <<http://www.rfc-editor.org/info/rfc5394>>.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009, <<http://www.rfc-editor.org/info/rfc5424>>.
- [RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009, <<http://www.rfc-editor.org/info/rfc5440>>.

- [RFC5520] Bradford, R., Ed., Vasseur, JP., and A. Farrel,
"Preserving Topology Confidentiality in Inter-Domain Path
Computation Using a Path-Key-Based Mechanism", RFC 5520,
April 2009, <<http://www.rfc-editor.org/info/rfc5520>>.
- [RFC5557] Lee, Y., Le Roux, JL., King, D., and E. Oki, "Path
Computation Element Communication Protocol (PCEP)
Requirements and Protocol Extensions in Support of Global
Concurrent Optimization", RFC 5557, July 2009,
<<http://www.rfc-editor.org/info/rfc5557>>.
- [RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel,
"Framework for PCE-Based Inter-Layer MPLS and GMPLS
Traffic Engineering", RFC 5623, September 2009,
<<http://www.rfc-editor.org/info/rfc5623>>.
- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic
Optimization (ALTO) Problem Statement", RFC 5693,
October 2009, <<http://www.rfc-editor.org/info/rfc5693>>.
- [RFC5810] Doria, A., Ed., Hadi Salim, J., Ed., Haas, R., Ed.,
Khosravi, H., Ed., Wang, W., Ed., Dong, L., Gopal, R., and
J. Halpern, "Forwarding and Control Element Separation
(ForCES) Protocol Specification", RFC 5810, March 2010,
<<http://www.rfc-editor.org/info/rfc5810>>.
- [RFC6007] Nishioka, I. and D. King, "Use of the Synchronization
VECTor (SVEC) List for Synchronized Dependent Path
Computations", RFC 6007, September 2010,
<<http://www.rfc-editor.org/info/rfc6007>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for
the Network Configuration Protocol (NETCONF)", RFC 6020,
October 2010, <<http://www.rfc-editor.org/info/rfc6020>>.
- [RFC6107] Shiimoto, K., Ed., and A. Farrel, Ed., "Procedures for
Dynamically Signaled Hierarchical Label Switched Paths",
RFC 6107, February 2011,
<<http://www.rfc-editor.org/info/rfc6107>>.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence
Protocol (XMPP): Core", RFC 6120, March 2011,
<<http://www.rfc-editor.org/info/rfc6120>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed.,
and A. Bierman, Ed., "Network Configuration Protocol
(NETCONF)", RFC 6241, June 2011,
<<http://www.rfc-editor.org/info/rfc6241>>.

- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", RFC 6707, September 2012, <<http://www.rfc-editor.org/info/rfc6707>>.
- [RFC6805] King, D., Ed., and A. Farrel, Ed., "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, November 2012, <<http://www.rfc-editor.org/info/rfc6805>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013, <<http://www.rfc-editor.org/info/rfc7011>>.
- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, September 2014, <<http://www.rfc-editor.org/info/rfc7285>>.
- [RFC7297] Boucadair, M., Jacquenet, C., and N. Wang, "IP Connectivity Provisioning Profile (CPP)", RFC 7297, July 2014, <<http://www.rfc-editor.org/info/rfc7297>>.
- [Stateful-PCE]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", Work in Progress, draft-ietf-pce-stateful-pce-10, October 2014.
- [TL1] Telcorida, "Operations Application Messages - Language For Operations Application Messages", GR-831, November 1996.
- [TMF-MTOSI]
TeleManagement Forum, "Multi-Technology Operations Systems Interface (MTOSI)", <<https://www.tmforum.org/MTOSI/2319/home.html>>.
- [YANG-Rtg] Lhotka, L. and A. Lindem, "A YANG Data Model for Routing Management", Work in Progress, draft-ietf-netmod-routing-cfg-17, March 2015.

Appendix A. Undefined Interfaces

This appendix provides a brief list of interfaces that are not yet defined at the time of this writing. Interfaces where there is a choice of existing protocols are not listed.

- o An interface for adding additional information to the Traffic Engineering Database is described in Section 2.3.2.3. No protocol is currently identified for this interface, but candidates include:
 - The protocol developed or adopted to satisfy the requirements of I2RS [I2RS-Arch]
 - NETCONF [RFC6241]
- o The protocol to be used by the Interface to the Routing System is described in Section 2.3.2.8. The I2RS working group has determined that this protocol will be based on a combination of NETCONF [RFC6241] and RESTCONF [RESTCONF] with further additions and modifications as deemed necessary to deliver the desired function. The details of the protocol are still to be determined.
- o As described in Section 2.3.2.10, the Virtual Network Topology Manager needs an interface that can be used by a PCE or the ABNO Controller to inform it that a client layer needs more virtual topology. It is possible that the protocol identified for use with I2RS will satisfy this requirement, or this could be achieved using extensions to the PCEP Notify message (PCNtf).
- o The north-bound interface from the ABNO Controller is used by the NMS, OSS, and Application Service Coordinator to request services in the network in support of applications as described in Section 2.3.2.11.
 - It is possible that the protocol selected or designed to satisfy I2RS will address the requirement.
 - A potential approach for this type of interface is described in [RFC7297] for a simple use case.
- o As noted in Section 2.3.2.14, there may be layer-independent data models for offering common interfaces to control, configure, and report OAM.

- o As noted in Section 3.6, the ABNO model could be applied to placing multi-segment pseudowires in a network topology made up of S-PEs and MPLS tunnels. The current definition of PCEP [RFC5440] and associated extensions that are works in progress do not include all of the details to request such paths, so some work might be necessary, although the general concepts will be easily reusable. Indeed, such work may be necessary for the wider applicability of PCEs in many networking scenarios.

Acknowledgements

Thanks for discussions and review are due to Ken Gray, Jan Medved, Nitin Bahadur, Diego Caviglia, Joel Halpern, Brian Field, Ori Gerstel, Daniele Ceccarelli, Cyril Margaria, Jonathan Hardwick, Nico Wauters, Tom Taylor, Qin Wu, and Luis Contreras. Thanks to George Swallow for suggesting the existence of the SRLG database. Tomonori Takeda and Julien Meuric provided valuable comments as part of their Routing Directorate reviews. Tina Tsou provided comments as part of her Operational Directorate review.

This work received funding from the European Union's Seventh Framework Programme for research, technological development, and demonstration, through the PACE project under grant agreement number 619712 and through the IDEALIST project under grant agreement number 317999.

Contributors

Quintin Zhao
Huawei Technologies
125 Nagog Technology Park
Acton, MA 01719
United States
EMail: qzhao@huawei.com

Victor Lopez
Telefonica I+D
EMail: vlopez@tid.es

Ramon Casellas
CTTC
EMail: ramon.casellas@cttc.es

Yuji Kamite
NTT Communications Corporation
EMail: y.kamite@ntt.com

Yosuke Tanaka
NTT Communications Corporation
EMail: yosuke.tanaka@ntt.com

Young Lee
Huawei Technologies
EMail: leeyoung@huawei.com

Y. Richard Yang
Yale University
EMail: yry@cs.yale.edu

Authors' Addresses

Daniel King
Old Dog Consulting

EMail: daniel@olddog.co.uk

Adrian Farrel
Juniper Networks

EMail: adrian@olddog.co.uk

