

A Terminology for Control Models at Optical Exchanges

Freek Dijkstra¹, Bas van Oudenaarde², Bert Andree¹, Leon Gommans¹, Paola Grosso¹, Jeroen van der Ham^{1,3}, Karst Koymans¹, Cees de Laat¹

¹ Universiteit van Amsterdam, Kruislaan 403, Amsterdam, The Netherlands
fdijkstra@science.uva.nl

² Finalist IT Group, Postbus 1354, Rotterdam, The Netherlands

³ TNO Defense, Security and Safety, The Hague, the Netherlands

Abstract. Optical or lambda exchanges have emerged to interconnect networks, providing dynamic switching capabilities on OSI layer 1 and layer 2. So far, the only inter-domain dynamics have occurred on layer 3, the IP layer. This new functionality in the data plane has consequences on the control plane. We explain this by comparing optical exchanges with current Internet exchanges.

Descriptions of optical exchanges have appeared in the literature, but discussions about these exchanges have been hampered by a lack of common terminology. This paper defines a common terminology for exchanges. Discussion in the community revealed four different meaning for the term “open exchange”. We list them in this paper.

We classify the different kind of exchanges based on the interactions between the domains at the control plane. We use these control models to distinguish between different types of interconnection points.

1 Introduction

1.1 Overview

The main function of Interconnection points, such as exchanges, is to facilitate traffic flows between the connected domains. Besides regular Internet-based exchanges, new types of exchanges are emerging. A wide variety of names has been proposed for these new exchanges, including optical exchange, transport exchange, grid exchange, GLIF open lambda exchange (GOLE), optical interconnection point and lightpath exchange.

The goal of this paper is to create a generally usable terminology for exchanges, both optical and Internet exchanges. The novelty in our work comes from the fact that we do so by looking at the control plane rather than the data plane, we identified conflicting definitions, and we are the first to compare optical and internet exchanges in detail.

Section 2 gives a classification of existing and new exchanges, and defines our terminology. Where possible, existing terminology is re-used. Other termi-

nology, in particular the term open exchange, draws upon discussions in the GLIF community [1,2]⁴.

A distinguishing factor for exchanges is the ability or inability of connected domains to influence the state of the core network. To this end, we define a total of three control models for exchanges in Sect. 3. This categorization will aid the discussion about the design of new exchanges. Section 4 maps these control models to each type of exchange.

The paper concludes with future trends and conclusions.

We refer to the extended version of this paper for a discussion about advanced network services on the data plane, like the conversion of data between different formats (interworking) and layers (elevator services), or on the control plane, like automated provisioning of network elements, policy based authorization, broker services, and index servers [3,4].

1.2 Related Work

This work builds on experience and discussions in the GLIF community, a collaboration of mostly National Research and Education Networks (NRENs). Here the need for high bandwidth circuits led to hybrid networks offering both routed and circuit switched network connections. Interconnections between NRENs are often made at new optical exchanges, like NetherLight, StarLight, ManLan, T-Lex, HK Light, UKLight and NorthernLight.

We rely as much as possible on existing terminology. In particular, the ownership terminology in Sect. 2.3 builds upon the management layers in Telecommunication Management Network (TMN) [5] and current practice in economic and legal communities [6].

This paper deals with the network node interface (NNI) of networks connected to an exchange, and is by no means the first to discuss this interface. The Optical Interworking Forum specified the network to network interface between domains (E-NNI) based on RSVP messaging [7]. Recent work comes from the L1VPN [8] workgroup in the IETF, which deals with the NNI for GMPLS [9].

The work provided in this paper is complimentary because it specifically deals with the network interface for an exchange rather than a transit network. This paper deals with a high level overview of the relation between the different actors, rather than specifying a practical signaling protocol.

2 Terminology

In this section we introduce a concise definition of terms like *domain*, *administrative control*, as well as *open* and *automated*.

⁴ The only exception is that we use the term “optical exchange”. The GLIF community currently uses the term “GOLE”, and the authors personally prefer the term “transport exchange”, but we felt that “optical exchange” was most widely recognized in all communities.

2.1 Peering

Traffic between separate networks is often exchanged at geographically clustered locations, called *interconnection points* or *peering points* [10,11]. For the regular Internet, the Internet service providers (ISPs), can interconnect using either *transit* or *peering* [12]. Peering, in most literature, is limited to providing connectivity to each others networks and to the customers of the other network, but not to other destinations. Transit on the other hand implies that traffic for any destination can be handled by the party providing the connectivity, usually for a fee.

In this paper we do not distinguish between peering and transit. In our terminology **peers** are network owners who connect to an interconnection point and **peering** is the concept of exchanging traffic between peers, regardless of the economic model.

2.2 Types of Interconnection Points

The most trivial interconnection point is a co-location that only provides rack space and power. This already gives the ability to initiate bilateral peerings between peers at the same facility. We are interested in exchanges, which are interconnection points with one or more core networks in place, dedicated to the exchange of traffic between peers.

Classification. We currently observe four types of interconnection points, based on the function, rather than the technical implementation:

- Internet exchanges
- mobile roaming exchanges
- optical exchanges
- points of presence

Internet exchanges, also known as Internet exchange points (IXP) or Network access points (NAP), serve as an interconnection points to exchange packet data between individual peers. The peers have one or a few physical connections to a central core infrastructure. The core network can be Ethernet LAN, ATM, or MPLS-based. The first variant is stateless, while the other two are stateful and require that the individual peers set up a path between them. Such a path is a channel in the physical connection.

Mobile roaming exchanges, such as GPRS roaming exchanges (GRX) [13] and UMTS exchanges, exchange packet data for respectively 2.5th and 3rd (3G) generation mobile telephony. In telecommunications, however, the term “exchange” is different from our usage and refers to a transit provider rather than an interconnection point. An exchange point between mobile roaming exchanges is technically not different from a packet-based⁵ Internet exchange.

⁵ GPRS and UMTS are packet based. The older CSD system is circuit switched.

Optical exchanges⁶, also known as lambda exchanges, grid exchange points, transport exchanges or GLIF open lambda exchanges, are interconnection points where peers exchange traffic at OSI layer 1 or layer 2 [3]. GMPLS Internet exchanges as defined by Tomic and Jukan [14] share the concept of circuit-switched interconnection points, but have not been implemented yet.

We use the term **Transport Exchange** to refer to circuit-switched exchanges, like current-day optical exchanges.

Unlike exchanges, **points of presence** (POP) are interconnection points where the peers are unequal. Access networks connect with an upstream network provider at a POP. In this case, the peers are unequal since the upstream provider accepts transit traffic from the customer, but the reverse is not true.

Internet versus Optical Exchanges. Table 1 highlights the differences between Internet exchanges and optical exchanges. Peers at an Internet exchange interconnect to exchange IP traffic. The core of an Internet exchange contains exactly one circuit per peering relation. In contrast, an optical network supports circuits between end-users, so at an optical exchange there is a circuit between peers for each end-to-end connection that goes through the exchange. The table further emphasizes that for an optical exchange these circuits can carry any layer 1 or layer 2 traffic. Differences in function and purpose lead to different choices in technology between Internet exchanges and optical exchanges. Finally, the table highlights that an optical exchange may offer more advanced services than an Internet exchange.

Table 1. Functional differences between Internet exchanges and current optical exchanges.

| | Internet Exchange | Optical Exchange |
|---------------------|---|---|
| OSI Layer | Transports traffic at layer 2, peers connect with layer 3 devices | Transports traffic at layer 1 or layer 2, peers connect at that same layer. |
| Traffic type | IP traffic only | Any packet data or any data at a specific framing or bit rate |
| End-points | Connection between two peering networks | Connections are part of a larger circuit between two end-hosts |
| Dynamics | Stateless, or state changes only when peering relations change | State changes for each data transport |
| Technology | Often packet switched, sometimes label-switched (with virtual circuits like MPLS and ATM) | Circuit or virtual-circuit switched (e.g. using SONET or VLANs) |
| Services | Only data transport | Data transport and other services, like the conversion of data between different formats and layers |

⁶ Optical does not imply that the exchange itself is purely photonic.

There is no clear boundary between the different interconnection points since each interconnection point may take multiple roles. We expect that the differences listed in Table 1 will change over time, as new technologies become available and are implemented. For example, customers at a POP may also directly peer with each other, a function typically seen at exchanges. Circuit switching is typically associated with optical exchanges, but not a technical necessity: ATM- and MPLS-based Internet exchanges are also circuit switched and it might be possible to create a non-circuit switched optical exchange using optical burst switching (OBS) [15].

2.3 Ownership

Owner, Administrator and Users. We distinguish between legal owner, economic owner, administrator and user(s) for each network element⁷. The **legal owner** of a network element is the entity that purchased the device and the **economic owner** is the entity that acquired the usage rights from the legal owner. We base these terms on current practice in economic and legal communities [6].

The **economic owner** determines its policy of the network. This entity carries the responsibility for the behavior of a device and has the final responsibility in case of hazards and abuse. In addition, each network element can also have a separate **administrator**, the person, organization, or software component that configures and administers the device on behalf of the economic owner. The economic owner determines the policy for a network element; the administrator enforces this policy. Finally, the **users** may use or invoke an element, if their request is in compliance with the active policy.

We assume that each network element has exactly one legal owner, one economic owner, and one administrator, but may have multiple users over time (though typically only one at a specific time).

Domains. We define a **domain** as a set of network elements⁸. An **administrative domain** is a set of network elements with the same administrator. An **owner domain** is a set of network elements with the same economic owner.

A **core network** is an administrative domain within an interconnection point that is able to exchange traffic between at least three peers. Core networks are of special interest throughout this paper and we use the term **core** to refer to a core network and its administrator.

Examples. Often the legal owner, economic owner, and administrator of a network element are the same entity. For example, in the Internet, a transit provider is typically owner and administrator of its network. But this is not always the case.

⁷ Network element is a generic term to include network devices, links, interfaces and hosts.

⁸ Including non-disjoint sets. Note that a domain does not necessarily have to be an AS-domain.

An organization leases a trans-oceanic fiber from a carrier for a year, the carrier is the legal owner, while the other organization is the economic owner.

If an organization outsources the maintenance of its network, the economic owner and administrator of this network are different entities.

In the next subsection we explain the concept of open control, where the exchange is both the legal owner as well as the administrator of a specific interface, while the peer is the economic owner of this interface.

2.4 Open Exchanges

We found that in the the GLIF community, the use of “open” in “open exchanges” was ambiguous. It could refer to at least four different meanings, as described below. We recommend that it is only used in the now prevalent meaning of *open control*. For other meanings, we suggest alternative wording.

Open Control Model. In a **closed** interconnection point, the economic owner domain is equal to the administrative domain: everyone both decides upon and enforces the policy of their network elements. In particular, the core ultimately decides on the policy for each interface in the core network.

In the open control model, the core of an open exchange delegates the policy decision of each external interface to the peer that connects to that interface. Therefore, peers of an open exchange have the ability to configure “their” interfaces in the core network and thus can decide who connects to their networks.

Figure 1 shows an optical exchange consisting of an optical cross connect at the core. The exchange has three peers: Anet, Bnet and Cnet. If Anet wants to connect to Cnet, it signals that request to the exchange. A closed exchange would autonomously decide to grant or deny that request, and choose interface 4 or 5. An open exchange outsources this policy decision to Cnet which has policy control over interface 4 and 5, even though this policy is enforced in the optical cross connect which is legally owned and administrated by the exchange.

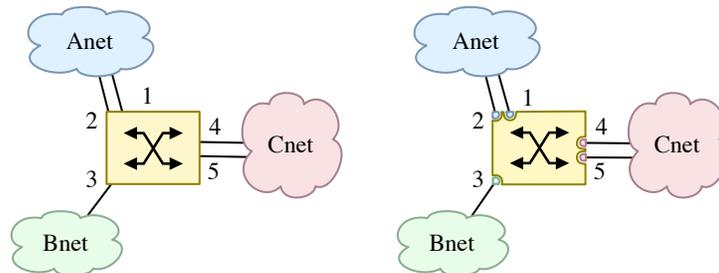


Fig. 1. Example of an optical exchange. On the left the administrative domains are shown, which are equal to the owner domains for the closed control model. On the right, the owner domains for the open control model are shown.

In the open control model, the core does not define an acceptable use policy (AUP) for its peers, and is thus *AUP free*.

Business Model. We use the word “**public**” or “**neutral**” to refer to an interconnection point with an open business model. An open business model requires that an interconnection point must have a published, clear policy for new peers to join, and has a reasonable and non-discriminatory (RAND) policy⁹ towards its peers.

A non-public interconnection point is called “**private**” or “**non-neutral**”. An open exchange can still be non-neutral. For example, an exchange economic owner may decide to only connect business partners as peers, but not others, and have the partners then decide on the policy for connections. Similarly, a neutral exchange may not be open. Hypothetically, an exchange may decide to allow every peer to connect to the core network, but grant path setup requests depending on an arbitrary decision.

Service Exposure. The term “**service exposure**” can be used to refer to the ability by peers to look in the inner workings of the exchange. The opposite of service exposure is “**service overlay**”. An exchange with a service overlay would behave like a black box. While peers can make requests to a black box, they do not know what exact devices, interfaces or other equipment are used to fulfill the request.

Automated Exchange. An exchange is called “**automated**” if peers are able to set up circuits between each other and invoke other services from the exchange without manual intervention from the economic owner of the core network.

3 Control Models

In this section, we define three different control models for interconnection points: the autonomous, federated and distributed control models. The autonomous control model is the simplest model. The federated and the distributed control model respectively extend the autonomous and the federated control models.

These models make a clear distinction between administrative control (policy enforcement) and owner control (policy decision) of the network elements. We consider a few administrative domains on the transport plane, each operated by a specific administrator. For each model, we explain how owner domains control network elements, and in particular how peers decide on the business policy for some network elements in the core network.

It is only possible to control network elements in another administrative domain if the administrators work together by sending messages to each other. It should be noted that we do not assume that these messages are automated.

⁹ This may seem to imply equal access rights to all peers. However, a distinction can be made based on the service level, as long as the service level is achievable by all peers on non-discriminatory conditions. E.g., if they pay a certain fee.

3.1 Autonomous Control Model

In the autonomous control model, there is exactly one core network, which is owned and administrated by a single entity. Peers can connect their network to the interconnection point, but there is no interaction between the peers and the core network on the control plane. Peers may interact with each other, but that is not relevant to this model.

Figure 2 shows an example of the autonomous control model. In this figure, the transport plane shows five distinct administrative domains: core, A, B, C and D, each operated by a administrator on the control plane. On the transport plane, each box represents an administrative domain, interconnected by links. On the control plane, each square represents a separate controller. There is no communication between the peers and the core on the control plane.

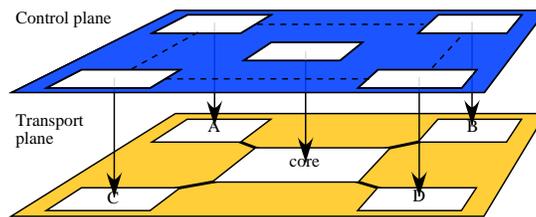


Fig. 2. Example of the autonomous control model. Squares represent administrative domains.

The economic owner of a core network determines a limited number of policies. Peers either accept the policies or take their business elsewhere.

The peers of a LAN-based Internet exchanges exchange control messages using an external routing protocol, but not with the exchange itself. So these exchanges are examples of the autonomous control model.

The autonomous control model is always closed.

3.2 Federated Control Model

In the federated control model, the interconnection point has exactly one core network. The core offers services to each peer, including the ability to interconnect with other peers.

The inner workings of the core network may be unknown to the peers (making it a black box), but peers can still check information about the state of some resources. For example, a peer can still inquire about the availability of a certain resource or get the status of a circuit it established earlier.

Figure 3 shows an example of the federated control model. The transport plane is the same as in Fig. 2, but the control plane is different: here the controller of each peer exchanges messages with the controller of the core network.

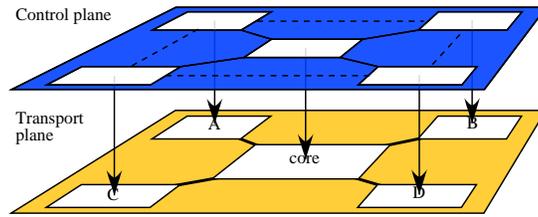


Fig. 3. Example of the federated control model.

When a peer wants to use a certain service, it invokes the administrator of the core network, which may delegate parts of the request to other parties. For example, if peer D sends a request to set up a circuit from B to D, the core economic owner checks if the requested resources in the core itself are available and contacts the economic owner(s) of the resources involved. In the case of open control, the core asks peer B if this request must be honored. If that is true, the core administrator then creates the requested circuit.

3.3 Distributed Control Model

In the distributed control model there can be multiple federations, each controlling a different core network. Every party can bring in its own equipment, e.g. fibers, and most important: its own services (and control software). Each peer exposes its own services to the rest of the community, possibly without revealing the inner details. A broker may combine multiple services and expose this combination as a single service.

The idea is that each peer still administratively controls its own network elements, but interacts with other administrators, or partially delegates its policy control, forming collaborations. Each peer can partner in multiple collaborations.

It is possible to regard one instance of the distributed control model as multiple interconnected instances of the federated control model. However, the distributed control model highlights the intelligence that is required to make all parts work together. This intelligence is not always necessary in the federated model.

Figure 4 shows an example of the distributed control model. The figure shows how peers can dedicate part of their network resources to form a dedicated core network. For example, A may expose some network elements to the other peers, which can be used by B or D to interconnect, either to A, or between each other through the core network of A. Also, C and D may decide to put some network resources in a pool, forming another, joint, core network. Typically, a core network formed by multiple peers is exposed as one single core network by a broker, which then delegates incoming requests to the individual administrators of the peers.

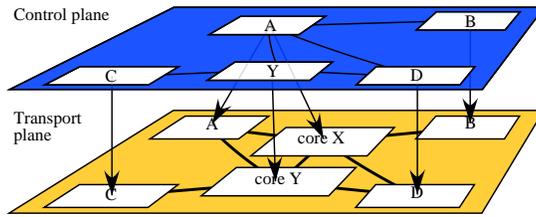


Fig. 4. *Example of the distributed control model.*

4 Model Usage

In Table 2, we give a list of viable mappings between the current interconnection points to the models described, based on our observation of current exchanges.

Table 2. Applicable models for each type of interconnection point.

| | Internet Exchange | Mobile Exchange | Optical Exchange | Point of Presence |
|---------------------------|-------------------|-----------------|------------------|-------------------|
| Autonomous control model | ✓ | ✓ | | ✓ |
| Federated control model | ✓ | ✓ | ✓ | ✓ |
| Distributed control model | | | ✓ | |

Stateless Internet and mobile exchanges use the autonomous control model, since no request needs to be sent to the core network administrator to exchange data between peers. If the Internet or mobile exchange is stateful, it can be either of these two models.

A POP typically uses the autonomous control model, because the configuration is mostly static and peers have no direct control over the inner working of the facility. However, if peers of a POP can decide on the policy, the federated control model is used.

If optical exchanges offer multiple services, standardized service discovery and service invocation are required. Both the federated and distributed control models offer this feature in a scalable way using pluggable services (a service oriented architecture). The distributed control model is more complex than the other models, and thus harder to deploy, because there is no longer a single entity that acts as a broker broker.

5 Future Trends

Large data transport on long distances is most efficient over the lowest possible layers, and peers and their users demand more flexibility to set up circuits with known quality of service (QoS) between domains. Interconnection points down in the protocol stack can offer this flexibility.

Technologies change over time, just as the requests from the users. We have reasons to believe that the current optical (transport) exchanges and Internet exchanges converge into optical exchanges that support all the required services. First there is a tendency for current optical exchanges to provide network services, and a future service might be multiparty peering like in a LAN-based Internet exchange. Secondly, Internet exchanges tend to offer more services which are now regarded as optical exchange functions, like private circuits between two peers¹⁰. Third, there is a tendency to build Internet exchanges and optical exchanges at the same locations¹¹, which indicates a possible economic advantage of combining exchanges on the same physical location.

Open control is a mind shift compared to most current exchanges. With closed control, peers sometimes have the ability to change the state of one or more network elements in a core network, but their requests are evaluated against the policy set by the exchange. With open control on the other hand, the peers decide on the policy and the exchange enforces it for them. Even if peers are in control, they do not experience it that way unless their requests are promptly answered by an automated ensemble. Thus, automation of exchanges is a necessity for this paradigm change to happen.

We also recognize a trend to let end users control the network resources as they want. For example UCLP supported by CANARIE is a control mechanism driven by users. Whether the exposition of network elements and network services will continue is yet unclear. If low layer network connections are exposed to users, authorization becomes more important to prevent abuse. Monitoring is important for peers and end-users to check if and where failures occur. This is part of our future research direction.

6 Conclusion

Formerly, discussions about optical or lambda exchanges have been hampered by a lack of common terminology. In this paper we identified ambiguous terms, in particular on “open exchanges”, and presented a consistent terminology, based on experiences in the GLIF community. We introduced multiple models for exchanges that we offer to use as reference points to the community. We did show that the terminology can be used to classify the existing exchanges according to the models that we introduced. While we are confident that the models are workable, we hope they are found as fruitful to others as they are to use in discussions on the difference between Internet exchanges and optical exchanges.

Acknowledgment

Part of this research is done under the GigaPort Next Generation project led by the Dutch National Research Network (SURFnet), and the Interactive Collabo-

¹⁰ For example, the Amsterdam Internet Exchange AMS-IX already provides private interconnects and closed user groups.[16]

¹¹ For example, Chicago, New York and Amsterdam.

rative Information Systems (ICIS) project. Both projects are supported by the Dutch Ministry of Economic Affairs, grant numbers BSIK03020 and BSIK03024.

The authors wish to thank John Vollbrecht from Internet2, Henk Steenman and Job Witteman of the AMS-IX, and members of the GLIF community for their discussions and proof-reading.

References

1. Global Lambda Integrated Facility, <http://www.glif.is/>
2. Terminology discussion in the GLIF community, Sep. 2005, <http://www.glif.is/list-archives/tech/msg00019.html>
3. F. Dijkstr, C.T.A.M. de Laat, "Optical Exchanges", GRIDNETS conference proc., Oct. 2004
4. F. Dijkstra, et al, extended version of "A Terminology for Control Models at Optical Exchanges", <http://www.science.uva.nl/~fdijkstr/>
5. "Considerations for a telecommunications management network", ITU recommendation M.3013, Feb. 2000
6. The terms "economic ownership" and "legal ownership" are common economic terms. E.g. Anne Harrison in "Definition of economic assets", Jan. 2006, <http://unstats.un.org/UNSD/nationalaccount/AEG/papers/m4EconAssets.pdf>
7. Lyndon Y. Ong (ed), et al., "Intra-Carrier E-NNI Signaling Specification", OIF specification OIF-E-NNI-Sig-01.0, Feb. 2004
8. Tomonori Takeda, et al., "Framework and Requirements for Layer 1 Virtual Private Networks", draft-ietf-l1vpn-framework, Jan. 2007, Work in Progress
9. Lou Berger, et al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, Jan. 2003
10. Bilal Chinoy and Timothy Salo, "Internet Exchanges: Policy-Driven Evolution", Harvard Workshop On Co-Ordination Of The Internet, J.F. Kennedy School Of Government, Sep. 1996
11. Geoff Huston, "Interconnection, Peering, and Settlements", Proc. of Inet'99, Jun. 1999
12. William Norton, "Internet Service Providers and Peering", Proc. of NANOG 19, May 2001
13. K.J. Blyth, A.R.J. Cook, "Designing a GPRS roaming exchange service", Second International Conference on 3G Mobile Communications Technologies, Mar. 2001
14. Slobodanka Tomic, Admela Jukan, "GMPLS-Based Exchange Points: Architecture and Functionality", Chapter 8 in "Emerging Optical Network Technologies Architectures, Protocols and Performance", Springer, ISBN 0-387-22582-X, Oct. 2004
15. Chunming Qiao, Myungsik Yoo, "Optical Burst Switching (OBS) - A New Paradigm for an Optical Internet", J. of High-Speed networks, pp. 69-84, 1999
16. Amsterdam Internet Exchange, "Services provided by the AMS-IX", <http://www.ams-ix.net/services/>