Benjamin P. Niven-Jenkins

Network Convergence onto an All-Optical Backbone

This article gives a high-level overview of how multiple legacy and future networks could be converged onto a single shared all-optical backbone network. This convergence could lead to savings in capital and operating expenditure for the service provider (or providers) that own and operate the shared all-optical backbone and its client networks. Several benefits that result from several client networks converging on a shared all-optical backbone are outlined.

Introduction

Large service providers often operate and manage several different networks, each using their own dedicated backbone infrastructure. It is not unusual for each network to be used to provide only a limited set of services; for example, a service provider may provide Internet access over one network and provide IP virtual private network (VPN) services over a different distinct network. There are several reasons why service providers choose to operate multiple distinct networks. For example, the service provider may wish to maintain distinct networks to avoid connecting what it considers a private network (for example, an IP VPN service) to a public network (for example, the Internet). In addition, some of the networks are often used to provide legacy services, for example frame relay and asynchronous transfer mode (ATM), that are still used by some of the service provider's customers but cannot be provided natively over more recent networking technologies.

In order to reduce costs, service providers are increasingly looking to decrease the number of dedicated backbone networks that they have to maintain and manage by converging some or all of their existing networks onto a single shared backbone network. In order to achieve this they require a solution that supports the graceful migration of their existing infrastructure and services onto a shared backbone network while still providing the stability, quality of service and reliability expected by, or guaranteed to, the service provider's customers.

All-Optical Networks

In a conventional optical network, data is transmitted over optical fibres that interconnect network nodes (commonly called *routers* or *switches*). As data traverses the network it will traverse one or more network nodes. At each network node, data is received on an input port, converted from an optical signal to an electrical signal, processed in the electrical domain, converted from an electrical to an optical signal and transmitted on the relevant output port. This conversion from an optical signal to an electrical signal and back again is called *optical-electrical-optical* (OEO) conversion. This OEO conversion is expensive to implement and acts as a bottleneck in the network because the maximum speed at which data can be switched from an input port to an output port depends on how fast the underlying electronics in each network node can operate.

In an all-optical network, data is also transmitted over optical fibres that interconnect network nodes. However, at each network node, data is received on an input port, switched in the optical domain and transmitted on the relevant output port without any OEO conversion taking place. The lack of OEO conversion in all-optical networking nodes (more commonly called *optical cross-connects* (OXCs)) results in the cost of networking nodes being reduced as the OEO conversion used in conventional optical nodes is relatively expensive. A side effect of eliminating the OEO conversion within a networking node is that the node is not able to examine the data being transmitted because all-optical nodes switch entire wavelengths as opposed to individual packets (or time-slots) as in a conventional optical network. In an all-optical network, all forwarding decisions are therefore made based on the physical input port and the wavelength the optical signal was received on. It is usual in an all-optical network for several different wavelengths to be used simultaneously on each fibre by utilising wavelength division multiplexing (WDM).

It is also worth mentioning that although the majority of OEO conversion is removed in an all-optical network, some OEO is still required, usually in two main places within the all-optical network nodes. Firstly, OEO is required in any nodes on the edge of the all-

---

The Author: Benjamin P. Niven-Jenkins is with BT Exact

optical network to terminate and frame client connections so they are in an acceptable format to be transported across the all-optical network. Secondly, a small amount of OEO conversion is required on every all-optical node so that control and management information can be properly decoded and processed by the all-optical node.

Converging onto an All-Optical Backbone

In the architecture described below a single all-optical backbone network (the server network) is shared between several client networks. Each client network no longer uses its own dedicated backbone infrastructure. Instead all the client networks use the server network as a single shared backbone. Requests are made from client networks (or manually by an operator) to the server network for end-to-end optical connections to be provisioned across the server network according to the client network's topology and traffic demands. These requests are processed and provisioned by the server network and are then treated as point-to-point links by the client network that requested them.

Figure 1 shows an example topology where a single all-optical backbone network is shared between two client networks. The nodes labelled PE A, PE B and PE C (and their attached CE nodes) together form a client network; for example, an IP VPN. The nodes labelled PE D and PE E (and their attached CE nodes) together form another client network, for example a frame relay network. Only two client networks are used to keep the diagram as simple as possible. The all-optical network in the centre acts as a server network to, and the backbone network of, the various client networks.

End-to-end all-optical connections are requested by a client network (or manually by an operator) and provisioned across the all-optical network, usually between two provider edge (PE) devices in the client network; for example between the nodes labelled PE A and PE C. The physical details of this all-optical connection (and the all-optical network it traverses) are transparent to both PE A and PE C, each PE believing that it is connected directly to the other PE via a dedicated point-to-point link. The optical connections themselves are not advertised into the client network but they provide the underlying infrastructure to create point-to-point links that connect the two PEs together. These point-to-point links are advertised into the appropriate client network's routing protocol.

Each point-to-point link between two PEs in a client network maps directly onto a single end-to-end all-optical connection in the server network. This means that each client link using an underlying all-optical connection has its own dedicated bandwidth in the same way that it would have if the client network was provisioned using its own dedicated backbone infrastructure.

Each client network is logically distinct from the others; in other words any links advertised into one client network will not be visible in any of the other client networks.

Figure 1 **Example topology where two client networks are converged onto a single all-optical backbone**

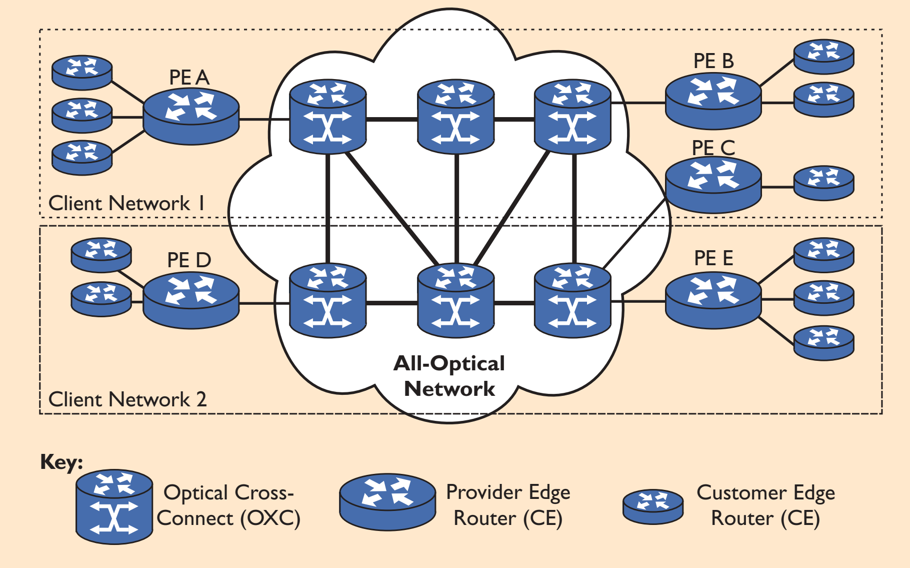

Transitioning Client Networks onto the All-Optical Backbone

Client networks can be transitioned from their existing core infrastructure to the shared all-optical infrastructure in incremental stages to minimise or avoid any disruption to customers using that particular client network. Individual links in the client network can be transitioned over to using all-optical connections one at a time with no adverse impact on the other links in the client network or the other client networks sharing the all-optical backbone.

Before any transitioning can take place it is first necessary to identify which of the links currently being used by the client network should be replaced by all-optical connections. In most, if not all, cases the high-capacity trunk links that interconnect the provider edge devices in the client network are the most suitable to be replaced with all-optical connections. In this way the client network's backbone infrastructure can be replaced with a shared all-optical backbone.

Once the links to be replaced have been identified, a new link would have to be installed connecting each networking device that has links to be replaced in the client network to at least one OXC in the all-optical server network. Then, one by one, all of the client links identified earlier would be replaced with all-optical connections. The process described below of how an existing client link could be replaced with an all-optical connection is shown in Figure 2.

Before the existing link between two networking devices (labelled PE A and PE B in Figure 2) can be replaced with an all-optical connection, the networking devices must first be connected to the all-optical network. An appropriate all-optical connection must then be requested and provisioned across the all-optical network. The details of how optical connections are requested and provisioned is outside of the scope of this article. At this point there are two parallel links between PE A and PE B, one across the original client backbone infrastructure and one across the all-optical network. Once the new all-optical connection is up and working and PE A and PE B are directly connected via a new point-to-point link across the all-optical network, the new link between PE A and PE B is advertised into the client network and it can start to carry data traffic. The data being carried on the link being replaced can be transitioned onto the new optical connection and the link being replaced can be decommissioned. This process is repeated

Figure 2 How a client network is transitioned onto the all-optical backbone

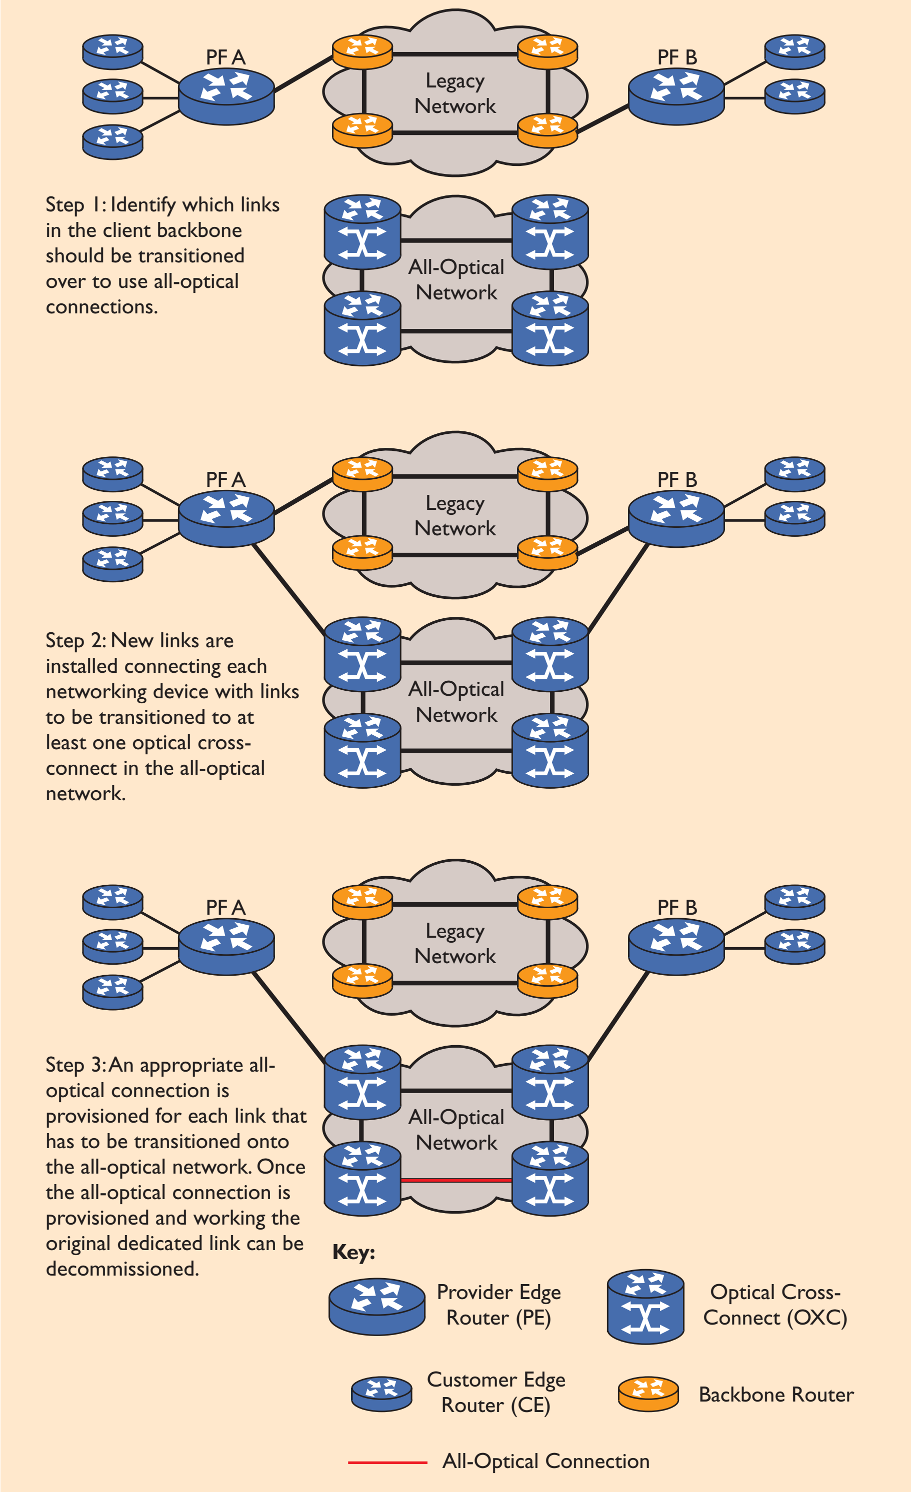

Step 1: Identify which links in the client backbone should be transitioned over to use all-optical connections.

Step 2: New links are installed connecting each networking device with links to be transitioned to at least one optical cross-connect in the all-optical network.

Step 3: An appropriate all-optical connection is provisioned for each link that has to be transitioned onto the all-optical network. Once the all-optical connection is provisioned and working the original dedicated link can be decommissioned.

for every link identified for replacement at the start of the transitioning process.

Network Convergence is More than Just the Backbone

Achieving network convergence involves more than just having a converged backbone network. There is good reason to believe that a service provider can make larger savings by converging other parts of its network; for example by using multi-service network devices that support both emerging and legacy services. Howeverm convergence on a single shared all-optical backbone can be used as a stepping stone towards full network convergence, with the all-optical backbone being used to provide a simple and flexible mechanism to provision bandwidth as point-to-point links across the all-optical backbone. This allows the service provider to deploy or converge their client networks as they see fit with the knowledge that any spare backbone network capacity can be utilised by any of the backbone's client networks depending on their individual traffic demands. As several client networks are converged with each other into a single client network utilising multi-service network devices to provide multiple services over a single client network, the service provider can simply provision additional bandwidth from the shared all-optical server backbone for the converged client network as its traffic demands change. Any bandwidth that was used across the all-optical backbone by the individual client networks before they were converged is released and can be reallocated according to the traffic demands of the remaining client networks.

It is not unusual for a service provider to overlay several client networks onto a single physical fibre infrastructure; for example, a service provider may use the same long haul wavelength division multiplex (WDM) infrastructure to carry traffic from several different client networks. In this situation there are still benefits to converging on an all optical backbone. Typically conventional optical networks offer only a limited set of static protection options (1 + 1 or 1:N protection), whereas an all-optical network gives the service provider the ability to protect its all-optical connections using static as well as dynamic protection options such as mesh restoration. An all-optical backbone also gives the service provider the flexibility to choose which protection type is most suitable on a per-connection basis so that network capacity is only used for protection as and when required.

Benefits to Using a Shared All-Optical Backbone

In the architecture outlined above no client network has any visibility of the other client networks and no two client networks share the same optical connections across the all-optical server network. This separation between client networks means that any data transmitted within one client network will always remain within that client network. There is no way for data from one client network to leak into a different client network while the data is traversing the shared all-optical backbone. This leads to several benefits, some examples of which are outlined below.

If one client network is heavily congested, any adverse affects caused by that congestion are not seen by (and do not affect) the other client networks sharing the all-optical backbone. The number of optical connections (and therefore the bandwidth) provisioned and advertised into a client network can easily be increased or decreased to reflect any medium-to-long-term growth or reduction in traffic demands. Once a client network no longer requires an optical connection, the resources used by that connection can be reused by one of the other client networks if required. In this way the all-optical infrastructure can be shared among the client networks according to their actual traffic demands and needs, rather than being constrained to any predictions made when the all-optical network was first deployed. Therefore the allocation of available resources within the all-optical server network can change according to the changing demands of its client networks. For example, the amount of IP VPN traffic being carried by a service provider may be constantly growing, while in contrast the amount of frame relay traffic being carried is static or shrinking. If a service provider's frame relay and IP VPN networks were converged on an all-optical backbone then over time any optical connections no longer required by the frame relay client network would be made available to any other client networks and could be reallocated to the IP VPN client network.

Each client network can protect its optical connections across the all-optical backbone against failures in a range of ways depending on the requirements of the particular client network being provisioned. For example, a service provider may decide that its IP VPN service requires each optical connection to have a dedicated backup connection associated with it (1 + 1 protection) because the service provider's IP VPN customers are given service level agreements that penalise the service provider when the client's service is degraded or interrupted. In contrast the same service provider may decide that its Internet service is only provided on a best-effort basis and therefore mesh, or no, restoration is acceptable. Each client network is also not restricted to using a single type of protection for all of its optical connections. It is possible for some optical connections used by a client network to have dedicated backup connections and for others to rely on shared mesh, or no, restoration.

Just as any congestion in one client network can not affect the other client networks, any malicious attack on one client network is confined to that client network alone and will not affect any of the other client networks. This means that if a network device in one client network is compromised by a malicious user, it cannot be used as a stepping stone to launch malicious attacks on any of the other client networks. Therefore a service provider could choose to use the same shared all-optical backbone for all its private and public networks knowing that it would not be possible for its private networks to be attacked via its public networks and vice versa.

Provided that client networks are not directly connected together, then each client network can use any addressing scheme that the service provider chooses. Each client network could use the same addressing scheme and the same overlapping addresses. Alternatively each client network could use a unique addressing scheme and a set of unique addresses. If two or more client networks are connected together, then those client networks that are connected together must not use addresses that overlap with each other.

It is not unusual for a service provider to have a separate (and possibly completely different) operational support system for each network backbone that they operate and manage. This is obviously costly and inefficient as many tasks and skills are duplicated as a result of each network backbone being managed independently. Converging several client networks onto an all-optical backbone enables a service provider to reduce the amount of inefficiency in their operational support systems by converging any duplicated functionality onto a single operational support system to manage the shared all-optical backbone. The cost of running a given service is largely dependent on the cost of managing and running the network over which that service is implemented. Therefore any reduction in the cost of operating and managing a network has a direct impact on the profitability of the services that are implemented using that network.

Summary

This article has presented a high-level overview of an architecture that could be used to transparently carry multiple legacy and future networks over a single shared all-optical backbone network. This architecture gives service providers the potential to reduce the number of backbone networks that they are required to maintain and operate. Consequently, a service provider can also reduce the number of network management systems that they are required to maintain and operate as well as being able to remove the duplication of skills and functionality that are often the result of a service provider having several separate network management systems. When these reductions are combined with the reduced cost of all-optical switches compared to traditional optical switches and any reductions in duplication and complexity that the service provider may achieve, reductions in capital and operating expenditure could be obtained over the medium to long term.

The architecture allows for client networks to be easily transitioned from their existing dedicated backbone network to the shared all-optical network in an incremental manner. This allows the service provider to

Glossary

Customer Edge (CE) router. A router that is part of a customer's network and that is connected and interfaces with a provider edge router in the service provider's network.

Optical Cross-Connect (OXC). An optical switch that does not contain optical-electrical-optical conversion. Data is switched (or cross-connected) optically with a single wavelength being the finest level of switching granularity available.

Optical-Electrical-Optical (OEO) conversion. The process of converting from an optical signal to an electrical signal and back again. Optical-electrical-optical conversion is not used in an optical cross-connect but it is used in conventional switches and routers to allow the switch or router to process data electronically.

Provider Edge (PE) router. A router that is part of a service provider's network and that is connected and interfaces with a customer edge router in the customer's network.

Virtual Private Network (VPN). A networking solution that enables several geographically dispersed customer sites to be interconnected via a network backbone that is shared with the VPN's other customers. Traffic separation between customers is maintained so that each VPN customer is only able to access its own data.

Wavelength Division Multiplexing (DWDM). Technology which allows several different wavelengths (frequencies) of light to be carried on a single optical fibre simultaneously.

minimise any consequences and disruption to their existing customers and services as a result of the transitioning process.

Although the various client networks all share a single all-optical backbone network, they are kept logically distinct from one another by the all-optical backbone network. This means that no client network has any visibility of the other client networks and no two client networks share the same optical connections across the all-optical backbone network. Data from one client network can not therefore leak into a different client network while the data is traversing the shared all-optical backbone.

The majority of the pieces required to build the architecture outlined in this article either exist today or are the subject of ongoing research within the industry. The architecture outlined will not be suitable for all service providers, as its suitability depends on a number of factors too complex to detail in this article. However, the architecture presented is one possible way to utilise an all-optical network to converge several client networks while maintaining the separation necessary to enable service providers to produce and offer fair, meaningful and secure services to their customers just as if each client network had its own dedicated backbone network.

Biography

Benjamin Niven-Jenkins
BT Exact

Benjamin Niven-Jenkins is currently working for BT Exact as an IP networking specialist. He graduated with a B.Sc. (Hons.) in Computer Science and Software Engineering from the University of Birmingham, England. His current research interests are routing algorithms and control planes for all-optical networks. His other interests include multi-protocol label switching (MPLS) and virtual private networks (VPNs).

benjamin.niven-jenkins@bt.com