Optical Layer Signaling: How Much Is Really Needed?

Ori Gerstel, Nortel Networks

ABSTRACT

This article challenges the emerging industry trend of adopting Internet-style distributed network-control with its full complexity for the optical transport network. Instead, we argue that an extensive telecom-style network management interface augmented with a minimal control plane and a service layer interface between management systems is more appropriate for the real needs of the optical layer. This approach will allow more flexibility in extending the interoperability between vendors and carriers as our understanding of these networks grows, increase the reliability of the network, and be a better fit for the telecom service provider. On the other hand, the simplicity of use and automation the Internet control plane promises can just as easily be achieved with our proposal.

INTRODUCTION

Optical networking technology has become one of the main promises for core transport network architectures. As the technology matures, from the simple point-to-point line systems of five years ago to full-function optical cross-connects a year from now, the optical network elements (ONEs) support more and more functionality and the need for sophisticated network management and control (NM&C) grows.

Two main trends have evolved over the years for management and control of large networks. These approaches, which we call telecom-style and Internet-style NM&C, were suited to the needs of the technologies and customers for whom they were meant: traditional telecom transport services on one hand and the Internet on the other. Both the telecom and Internet communities have tried to extend their NM&C paradigms to the optical layer. This is a very reasonable approach since these paradigms were refined over years of experience, and as a result are much more robust than some totally new approach. Even more important, there are huge costs associated with changing the mode of operation, training the craft personnel, the operational software, and so on. In this spirit the American National Standards Institute (ANSI) and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) have drafted many proposals on functional models for ONEs and how they should be managed, based on the synchronous optical network/synchronous digital hierarchy (SONET/SDH) approach; see [1] for a survey on the topic. At the same time, the Internet Engineering Task Force (IETF) has proposed how to extend multiprotocol label switching (MPLS) and related protocols to encompass the optical layer as another type of MPLS switch [2]. Variants of this approach have been adopted by the different optical layer standardization bodies (including OIF and ODSI) and are discussed in other articles in this issue.

Definitions used in this article

Carrier	The company operating the network (sometimes called service provider or network operator). This could be either a telco or an ISP
Control	The real time logic that is responsible for the operation of the NE, e.g., handling the signaling protocols
Control plane	The combined control operation across different NEs
Network design	The task of determining the nodes, the configuration per node, and the routes of lightpaths in the network. The latter part is sometimes referred to as "traffic engineering."
Network management	The slower transactions between the NEs and the management systems (as opposed to the fast control). The term "management" will be used as shorthand
Operator	The person in charge of the operation of the network through EMS and NMS GUIs. Typically stationed at a NOC (Network Operations Center).
Signaling	A fast distributed protocol between network elements. Part of the control plane
Telecom	We restrict the discussion to high-speed telecom transport equipment.
Vendor	The maker of the equipment

This article was written before the author had joined Xros and does not necessarily represent the opinion of his current or previous employers.

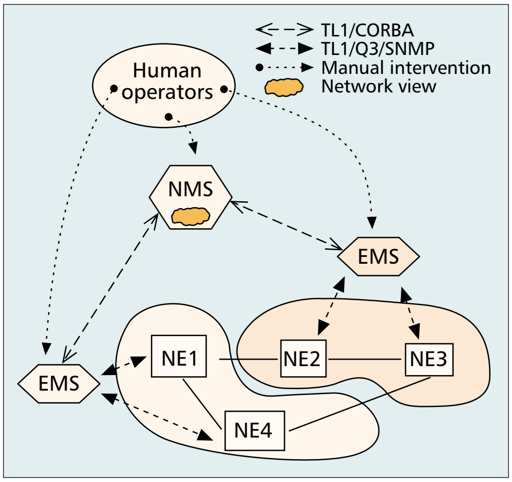

Figure 1. *Telecom-style network management*

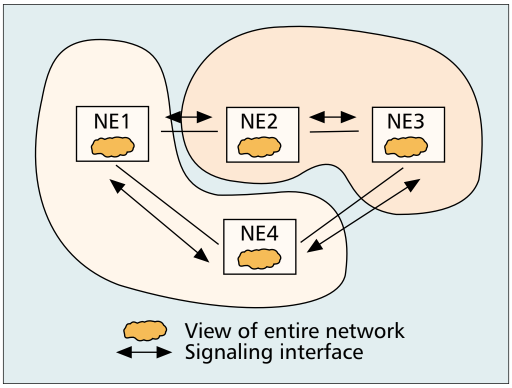

Figure 2. *Internet-style network management and control.*

The dominance of the Internet and its associated protocols as a client layer on top of the optical layer on one hand, and the failure of telecom equipment to interoperate across vendor boundaries with the same ease of Internet gear, have caused a new trend as well. The "bell heads" have started to abandon the "traditional" telecom way of managing networks through a hierarchy of management systems. Instead, they seem inclined toward a more automated Internet-style approach where the relatively dumb control plane is replaced by much smarter distributed control [3]. This trend, and the affection of the marketing and financial communities for everything related to the Internet, have pushed equipment vendors to follow suit.

This article takes a different view. After explaining the difference between these two basic approaches to NM&C, we claim that this new technology and evolving customer needs deem both approaches inadequate for the future optical network. We then propose how to augment telecom-style NM&C with some of the Internet-style features it lacks to achieve a combined approach that provides the best fit for this layer. The approach we advocate herein is not new, but it seems to be somehow missing from most discussions on optical layer interoperability standardization, which is the main motivation of this article.

Two Existing Approaches to Managing a Network

Traditionally, the telecom and Internet communities had very different views on how a network needs to operate. These differences stem from the main role of the network in both cases. Telecom networks evolved to be a symbol of reliability, allowing phone calls to go through even in the face of power outages and natural disasters; predictable quality, supporting the same voice quality regardless of network load; and optimized performance, allowing efficient use of network resources. The drawback of these networks was their inflexibility: they really only supported voice calls, and new applications were very hard to support. The Internet had a very different goal: supporting an open environment with infinite flexibility to foster new applications while keeping the network layer as simple and automatic as possible. While succeeding phenomenally on this front, the Internet neglected the other values, which where central to telecom networks.

These differences imply very different emphases on the role of NM&C subsystems. Telecom NM&C is typically much more centralized in nature than in IP networks to allow very tight human control of resources and good troubleshooting tools, as one would expect in support of the reliability and optimality requirements above. To this end they require a large number of operators at different levels.

Internet nodes emphasize distributed control, and are skinnier on their management support. They are also more automated and do not rely on human assistance for the most part, at the expense of resource overprovisioning.

In the next subsections we shall explore these issues in more detail.

Telecom-Style Management and Control

In typical telecom networks, the network elements (NEs) are not very aware of the network topology as a whole and are more focused on their own operation. Well-defined network management interfaces connect the NEs to a whole hierarchy of management systems, as defined by numerous ITU-T and Bellcore (now Telcordia) documents. These interfaces are based on the legacy TL-1 text language, more modern Common Management Information Protocol (CMIP) or Simple Network Management Protocol (SNMP), and more recently, Common Object Request Broker Architecture (CORBA). The lowest level of such managers is concerned with individual NEs; these are the element management systems (EMSs). These systems still do not have a view of the entire network and are connected to managers that manage subnetworks of elements of a similar type (subnetwork controllers, SNCs), which in turn are connected to systems that manage an entire diverse network as a whole (network management systems, NMSs). On top of these are systems that manage the various services the network provides. This hierarchy is called the telecommunications management network (TMN), and is depicted in Fig. 1. Human operators manage these systems according to the dictates of planing organizations.

Insofar as the control plane is concerned, con-

The Internet had a very different goal: supporting an open environment with infinite flexibility to foster new applications while keeping the network layer as simple and automatic as possible. While succeeding phenomenally on this front, the Internet neglected the other values, which where central to telecom networks.

Two of the main functions of the Internet control plane are not necessary for optical networks at this point in time: topology awareness at every NE and extended signaling for automated distributed setup of connections.

trary to several recent arguments on the topic (e.g., [1]), telecom networks do have a nontrivial control plane. In fact some of them, such as SONET bidirectional line switched rings (BLSRs), even have global understanding of their network topology (arguably, this represents too much of a control plane and is the reason SONET could never achieve full interoperability). It is true that telecom equipment standards try to minimize the extent of the control plane as much as possible and place as much functionality as possible in the management realm. The advantages of this approach will be articulated as part of the enhanced proposal later.

INTERNET-STYLE MANAGEMENT AND CONTROL

The Internet paradigm places heavier weight on the control than the management of the network. It is based on automatic operation of the network, with much less human intervention. When a system is put into service it typically starts its operation with comparatively little manual configuration commands (plug and play), discovers its neighbors and the rest of its domain, and is ready to ship packets in the general direction of their destination. This functionality is achieved through a set of topology discovery and routing protocols, such as Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and IS-IS, which run at each node and create an awareness of the entire network topology at each node (Fig. 2). The actual handling of packets is mainly based on connectionless IP forwarding to take the packet and figure out how it reaches the next hop en route to its destination. More recently, connection-oriented support has been added in the form of MPLS, but this has not yet changed the fundamentals of how such systems are managed.

Network management is still needed for such networks to allow operators to provision basic parameters into the systems, intervene in case of a failure, and analyze performance bottlenecks and design for the future, but these interfaces are not as comprehensive and well defined as their telecom-style counterparts. The main protocols used for those are Cisco's command language interface (CLI), SNMP, and recently Web-based HTML/Java management interfaces.

WHAT FUNCTIONALITY DOES THE OPTICAL LAYER REQUIRE?

Before delving into a discussion on how each of these styles fits the optical network, it is beneficial to examine the real needs of this layer in terms of control and management. In particular, we shall try to examine if either of the two most demanding signaling mechanisms is really needed: support for topology awareness at the NEs and signaling for automatic setup of connections:

- **Automatic discovery of the network topology:** It is important to reflect to the operator what the network looks like. Note that this can be done via the control plane, using a distributed protocol such as OSPF. On the other hand, the topology could also be auto-discovered in a central location (the NMS) as long as each node knows who its neighbors are (through exchange of hello messages) and relays this data to the NMS. The latter approach does not require the NEs to have such topology awareness.
- **Fault propagation:** The quick discovery of a failure and its dissemination to the pertinent nodes in the network must rely on distributed signaling, for protection purposes as well as fault isolation and correlation.
- **Automatic protection switching:** Clearly, the signaling to coordinate the protection switch must be fast and not rely on an MS. The setup of the protection routes can be preprovisioned into appropriate tables in the NEs at the time of connection setup, or computed on the fly when a failure occurs. The former approach does not necessitate topology awareness or fast connection setup since, upon a failure, each node reconfigures itself based on its predetermined tables. The latter approach does require such signaling; however, we claim that it is not suitable for the optical layer due to the lack of control on how the scarce protection bandwidth is utilized, and other reasons beyond the scope of this article.
- **Traffic engineering:** The need to control the routes of connections arose in MPLS networks [2] to in support for different qualities of service. This need is even stronger in the optical domain, since bandwidth resources come in large quanta (i.e., wavelengths) and the number of these wavelengths is fairly low per fiber: 100 wavelengths in some commercial systems, unlike the thousands of label-switched path (LSP) connections in MPLS. The inability of fully optical systems to change the wavelength of a connection further compounds the picture. As a result, there is a need to carefully optimize the network through sophisticated network design tools that require manual intervention.
- **Automatic setup of connections:** The ability to set a connection up in a relatively short timeframe is becoming more and more important. The bureaucratic and labor-intensive process that causes weeks of delay in current networks is clearly becoming unacceptable. At the same time, going to the other extreme, of requiring connection setup in a matter of seconds is hard to justify either. Relaxing this setup time requirement allows this automation to be just as well supported through a management system (MS). The key to shortening the delay in this case is not the use of the control plane but automation of the process. This point is further expanded below.
- **Measuring the quality of connections:** This function typically requires an in-band overhead and thus is part of the control plane.
- **Good fault isolation tools:** Due to the high bandwidth carried over the network, such tools will be needed even more for fast detection of misrouted connections and fault localization. This requires detailed management support to allow full visibility into the NE.
- **Coarse-grained interoperability:** It is theoretically conceivable to build a network where each node is of a different make, and they all interoperate to achieve optimal network functionality. We refer to this as *fine-grained interoperability*. Real-life networks, however, are

typically made up of more homogeneous domains of NEs from one vendor, with tighter coordination inside each domain and sparser interfaces between these domains.[1] This allows exploiting the proprietary innovations that distinguish one vendor from another.

One of the main perceived advantages of IP-style NM&C for the optical layer is its automatic connection setup support. We note that this function brings with it many unresolved and complex issues:

- The need for manually controlled traffic engineering, as described above, contradicts fully automated routing of connections and therefore their automatic setup.
- The need to plan for known (or estimated) future demands further complicates the design and makes automation harder.
- One can argue that the need to carefully optimize the network and the human intervention it may require can be eliminated by overprovisioning the network. However, the cost associated with each channel (including the wavelength-division multiplexing, WDM, gear and crossconnect ports) may cause this cost to be prohibitive in the short term.
- Supporting protection at the optical layer further complicates routing considerations. Even the simple case of finding diverse routes of working/protect connections is a complex algorithmic problem, and support for shared mesh protection poses additional significant constraints.
- This issue becomes even more complex if protection is supported at the client (e.g., IP) layer instead of relying on the optical layer. To this end, the optical layer will have to ascertain that the client layer links supported by it (which are mapped to optical layer connections) are sufficiently disjoint and do not all fail together. To some extent, this need has been addressed in [4] via the concept of an MPLS shared risk link group. Unfortunately, this concept only supports simple scenarios in which different links are pairwise disjoint. A realistic case that cannot be supported by this concept is presented in [5].
- Fast automatic setup becomes important if the current mode for leased lines, connections that are always up, is replaced by part-time connections set up on a need basis. However, to exploit this property for increased revenues, carriers must be able to time share the bandwidth among different customers. While this certainly works for low-speed telephone connections, it is questionable that it works for high-speed lightpaths.
- Finally, many other connection characteristics, crucial to understand before standardizing an interface, have not been considered and understood so far. An example of such parameters can be found in a proposal for an optical layer differentiated service [6].

Conclusions

From the above discussion we draw the following conclusions:

- Extensive management interfaces are needed to allow human intervention for various purposes. This is the case for both styles.
- At the same time, some level of distributed control is needed, but we claim it is very limited in scope and far more limited than what Internet-style control supports. A provisioning interface between an MS and the NEs will take care of all the rest while keeping the interface straightforward, since additional intelligence remains at the MS level and does not affect the NEs and the interface to/between them.
- Fully automated distributed setup of lightpaths is premature. Beyond the aforementioned technical issues, there are indications that customer need for such services may not exist in the near term since automated connection setup is not even offered today for much lower-speed connections (STS-1 and even T1).
- In summary, two of the main functions of the control plane are *not* necessary for optical networks at this point in time:
  –Topology awareness at every NE (via OSPF, BGP, etc.)
  –Extended signaling for automatic distributed setup of connections (i.e., using RSVP or CR-LDP to their full extent).

Are Any of The Styles a Perfect Match for the Optical Layer?

In this section we investigate whether any of the NM&C styles provides a good match for the optical layer, and come up with less than satisfactory results. This provides a motivation for the section on how to merge these NM&C styles into an appropriate NM&C layer.

What Is Broken in the Telecom Style?

The TMN hierarchy is ideal in theory for managing large and diverse networks. In practice, however, many critical components are missing or flawed:

- Network-level managers are not as ubiquitous as one would expect. Telcordia provides a substantial set of such managers, mainly for RBOCs. Most other carriers have automated their systems via homegrown tools. These systems include a monumental amount of software, and are old and cumbersome to use and very costly to upgrade in support of new applications.
- Interoperability between managers (the so-called X-interfaces in TMN) has not proven itself yet. While the ITU has drafted several Recommendations in this direction (M.3208.1 through 4), and CORBA can provide an excellent platform for such interoperability at the network or service layer, this did not yet happen in practice. As a result, setting up a connection that involves more than one carrier requires separate manual intervention for every carrier.
- For many of the new carriers, managers above the EMS level are often missing, requiring operators to repeat operations such as provisioning the same circuit through multiple EMSes. This is both error-prone and labor-intensive.
- Older systems do not even have a modern point-and-click EMS. Instead they are manually provisioned using the cryptic TL-1 text

The appropriate solution for the optical layer lies in a careful mix of the telecom and Internet management and control styles, augmented with service-level interoperability through CORBA.

[1] *Note that we refer to peer NEs at a given layer of the network. These NEs may well be connected through NEs at a lower layer, but this connectivity is transparent to them. For example, one may have a domain of same-vendor OXCs that interoperate through proprietary means, carried over a wavelength-division multiplexer layer of a different make. This is still considered coarse-grained interoperability because the interfaces between the different layers are minimal.*

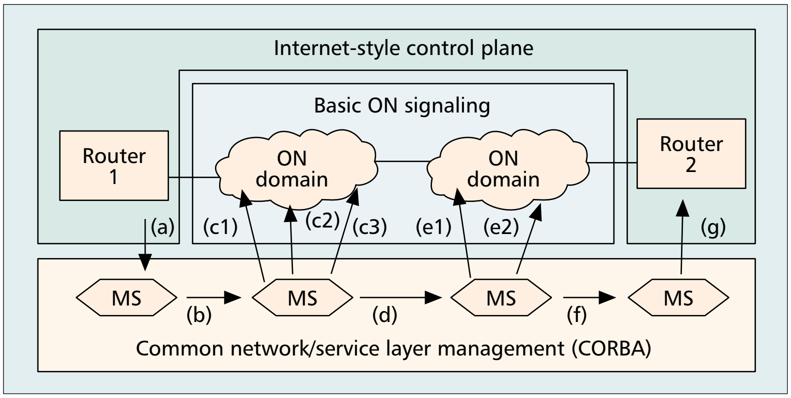

Figure 3. *A combination of telecom- and Internet-style NM&C.*

interface (similar to a UNIX/DOS command line).

Is the Internet Style Appropriate for the Optical Layer?

The Internet-style NM&C has evolved from a connectionless best-effort traffic paradigm, and in this sense is much farther apart from optical networks than the telecom-style NM&C. However, the main virtue of this style is that its developer community and their standards body (IETF) are moving extremely fast. Thus, they are expected to adapt better to the changing requirements of a new technology. Another advantage of this community is that their track record of interoperability between complex systems is much better. The Internet-style NM&C has some technical merits as well:

- The seamless integration of optical layer lightpaths into MPLS allows for a uniform control paradigm to span these different technologies and create a single consolidated network with lower operational expenses.
- The paradigm allows for fast setup of connections (within seconds), which may be needed sometime in the future.
- Interoperability may be based on existing standards and avoid "reinventing the wheel" (although the amount of proposed change in support of the optical layer may deem this point irrelevant).
- A finer granularity for interoperability is possible: each ONE could potentially be of a different make and still interoperate. In the telecom-style NM&C islands of uniform ONEs are needed to avoid having a separate MS per NE.

The challenges, however, are daunting as well and lie mainly in the reliability of the Internet control plane. Assuming that for such high-bandwidth connections the reliability requirements will not go down, it is not clear how this style will meet the 99.999 percent availability required of the transport network. The gap may currently be several orders of magnitude, as we shall see below. A few indications of the problem are:

- Routing protocols are hard to design [7] and even harder to fully test, given their complex behavior. It is known, for instance, that only a handful of implementations of dynamic routing are robust and scalable. It is also unclear how to prove their reliability, similar to what carriers require from their vendors as part of acceptance of the equipment into the network.
- Even worse, it is known from measurements on the Internet that the performance of its routing algorithms is far from perfect [8]. For example, the chances of a packet encountering a severe routing pathology (e.g., a long-lasting forwarding loop) were 3.3 percent at the end of 1995. This behavior is less of an issue in IP networks due to TCP-level retransmits, which compensate for such packet losses. A routing failure will be far more important for lightpaths since it will cause permanent creation of the wrong connection, even creation of a loop.
- While such problems can be dealt with when turning up service (via testing of the connection prior to running live traffic on it), they represent a crucial problem for protection signaling, since the protection lightpath has to work instantly.
- While the consensus seems to be that MPLS will be adapted for optical networks (see [2] and other articles in this issue), field experience with MPLS is not widespread. After all, this is a very new paradigm for the Internet community and is only deployed in a small number of networks. This may be an indication that it is too early to adopt it and expect to reap the benefits of its maturity (e.g., stability of the specification).

Our Proposal

Since none of the existing approaches provides a good fit for the optical transport layer, it is natural to try and mix them to enjoy the best of both worlds. This mix is made up of three parts, and is graphically depicted in Fig. 3 as three large rectangular shapes:

- An extensive network management interface, to allow maximum visibility into and control of the network, including:
  –Automatic discovery of node configuration and network topology at the MS level
  –Alarm reporting and fault isolation
  –External control of the routing of connections (traffic engineering), either manually or through a design tool
  –Determination of protection policies and routes
  –All other management requirements with which telecom NEs comply in the different management areas (faults, configuration, accounting, performance, and security, or FCAPS)
- A flexible service-level interoperability interface between MSs, based on a CORBA platform to allow:
  –Automatic setup of connections, possibly with sophisticated attributes for quality of service reliability, pricing, and so on
  –Alarm propagation and trouble ticketing between domains
  –Billing support
  –Service level agreement management
  –Compatibility between different versions of the interface to allow gradual upgrades of one domain at a time

- A fast, distributed control plane with minimal functionality:
  –Connection ID support for misconnection detection
  –Support for error detection and recovery: bit error rate (BER), forward error correction (FEC), and so on
  –Support for fast dissemination of failure information (which resource has failed/recovered)
  –A very rudimentary connection setup protocol for automatic protection purposes. This protocol should only specify the route of the connection, without additional advanced attributes, since the logic to determine the path is beyond the scope of the ONEs themselves and supported by the service-layer interface.

This approach carries with it a few substantial advantages:

- The tight human control this approach allows for facilitates more optimized planning of NE configurations and traffic routes. We argue that Internet networks are moving closer to this approach as well through support for traffic engineering (e.g., CR-LDP).
- It also facilitates better and quicker fault localization and troubleshooting, which becoms increasingly important even for Internet traffic with more mission-critical traffic flowing through it.
- This hierarchy allows very diverse technologies to coexist in the same network and be managed together. While each EMS can provide tailored support for a specific technology, an MS higher in the hierarchy can manage multiple technologies with an appropriate level of abstraction.
- It is easier to develop sophisticated software for management systems than for NEs. The reason is that more advanced, non-real-time operating systems are much more prevalent for the MS, better middleware (e.g., CORBA or Java) exists, and there is more support for good debugging and maintenance tools. Although the software technology for embedded systems has improved considerably over the last few years, many sophisticated tools require large amounts of memory and processing resources that many embedded systems do not have.
- It is easier to change existing MS software, in part due to the better tools, which are also more flexible for future enhancements. For example, the CORBA architecture facilitates very flexible interfaces between subsystems, which allow replacing of a subsystem on the fly without taking the rest of the software down.
- The other reason for the ease of upgrade is simply the fact that the number of management systems to upgrade is much lower than the number of NEs that must be upgraded. Not only does the sheer number make a difference; also, if the upgrade implies changes in the interface between systems (e.g., in the signaling format), it may be required to simultaneously upgrade all the NEs, or else NEs that still run the old version might not know how to handle messages in a new format.
- Another related issue is inherent flexibility of NE-MS and MS-MS interfaces vis-à-vis the more rigid NE-NE signaling interfaces, typically simpler due to the real-time requirements on them. For example, many MS interfaces use presentation-layer interfaces such as Abstract Syntax Notation 1 (ASN-1) to abstract the details of data attribute formats between systems. They may also allow new attributes to be added to one system without having to immediately upgrade the other end. This is not the case with most signaling protocols.
- CORBA is gaining momentum in many industries (http://www.corba.org), and its openness will allow reusing some of these existing tools for the telecom industry as well. For example, security tools developed for online banking, such as certificate authorities and key management, can be reused to secure automatic connection contracts [9].
- Much progress has been made defining standards for interoperability between MSs. Most notably, the TeleManagement Forum (a global consortium of over 200 groups: http://www.tmforum.org) is developing such standards as part of its SMART TMN program.

It is easier to change existing MS software, in part due to the better tools, which are also more flexible for future enhancements. For example, the CORBA architecture facilitates very flexible interfaces between subsystems.

EXAMPLE: SETUP OF A LIGHTPATH, INITIATED BY ROUTERS

In order to understand how this proposal differs from the full-scale Internet-style NM&C being proposed by IETF and other standardization bodies [2, 4], it may be helpful to consider how both approaches support a typical futuristic scenario of automatically setting up a connection. We also consider how these scenarios change if the setup is triggered manually.

Consider an optical network that provides connections to IP routers, in which a pair of routers realizes that the amount of traffic between them warrants a new lightpath to connect them directly. In this case router 1 sets up the new lightpath to router 2 through two separate domains of optical networks (which represent either different vendors or different carriers).

Our Combined Approach — Refer again to Fig. 3 for this case. The vertical arrows in the figure represent messages in support of the scenario, and the letters represent the following steps:

(a) Router 1 sends a notification to its MS requesting for the new lightpath to router 2 (without having to specify the route that the lightpath will take).
(b) The MS finds out, through the CORBA service plane, which ON domains to use, based on minimal cost, quality of service, or survivability considerations.
(c) The MS of the first domain sends messages (c1, c2, and c3) to three OXCs in it that have to be configured to support the new lightpath. The OXCs themselves need not know they are being considered to be part of the route.
(d) This MS then sends a message to the MS of the second optical domain.
(e) This MS sets up the route in its domain via management messages e1 and e2.

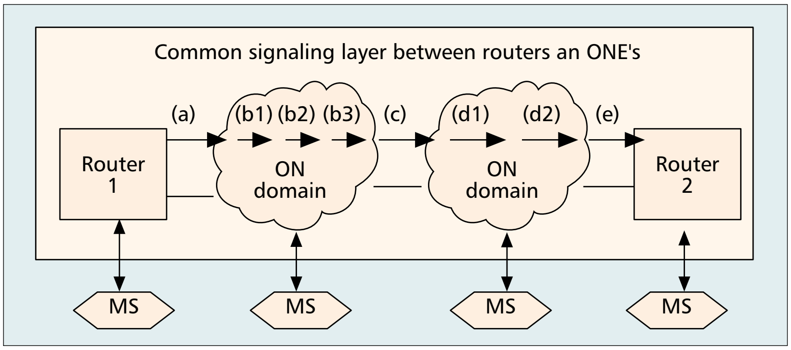

Figure 4. *Internet-style NM&C between optical networks and routers.*

(f) Now, the MS of router 2 is informed that the lightpath is ready.
(g) Finally, this MS configures router 2 to start using the new lightpath.

Internet-Style Automatic Setup — We assume that both the routers and ONEs are part of one large control plane.[2] The chain of events is depicted in Fig. 4, as follows:
(a) Router 1 sends a lightpath setup request to the adjacent ONE. Since the router knows the entire topology of the network, it can figure out independently how to route the lightpath; this is conveyed as part of the setup message (assuming, say, CR-LDP for the protocol).
(b) This message propagates from node to node in the first ON domain (messages b1–b3).
(c) The message then propagates to the next ON domain.
(d) Inside the second domain it sets up the pertinent route using messages d1 and d2.
(e) Finally, router 2 receives the setup message and the lightpath starts to be used. The information on it is disseminated to all the other nodes (say, using OSPF) so that other routers are aware of which resources are no longer at their disposal.

Manual Trigger for Setup — If, instead of an automatic event, the setup event is triggered by a human operator, both scenarios are very similar to the above, except that the initial setup request comes from an MS instead of router 1. In the telecom style this translates to a different message (a) being sent in the inverse direction, from the MS of router 1 to the router itself to notify the router of the new lightpath. The rest of the messages are the same as in Fig. 3. In the Internet style, the lightpath setup message originates at the MS of router 1 instead of the router itself. This message triggers the router to send message (a), and the rest of the scenario follows.

[2] This is commonly referred to as the "peer model". Although a different "client model" exists as well, we ignore it here to simplify the presentation. These differences are secondary with respect to the main topic of this article.

SUMMARY

In this article we have surveyed the two main trends that have evolved over the years for network management and control. These approaches, which we call telecom-style and Internet-style NM&C, were suited to the needs of the technologies and customers for which they were meant. However, new technologies and evolving customer needs may make both approaches insufficient in the future.

Indeed, in the above discussion we have tried to show that none of the traditional NM&C styles is a perfect fit for a reliable, flexible, efficient, and automated optical transport network of the future. We therefore believe that the appropriate solution lies in a careful mix of the two styles, augmented with service-level interoperability through CORBA. This mix allows keeping the NE-level software simple, and builds on the inherent flexibility and extensive feature set CORBA supports. While this approach is by no means new, it did not find its way into industry and standardization discussions for optical networks, which is the main motivation for this article.

A final disclaimer: one should not interpret the above discussion as a prediction that the industry is moving in the direction we have outlined. In fact, the opposite is more likely than not. While we have reasoned that this is not the appropriate step to take at this point in time for technical reasons mainly, other considerations such as marketing positioning are more likely to drive many vendors and customers alike in the opposite direction. A quick look at most of the other articles on this topic will prove this point.

REFERENCES

[1] A. McGuire and P. Bonenfant, "Standards: The Blueprints for Optical Networking," *IEEE Commun. Mag.*, Feb. 1998.
[2] D. Awduche *et al.*, "Multi-protocol Lambda Switching: Combined MPLS Traffic Engineering Control with Optical Crossconnects," IETF doc. draft-awduche-mpls-te-optical-00.txt, Nov. 1999.
[3] A. Greenberg, G. Hjalmtysson, and J. Yates, "Smart Routers - Simple Optics: A Network Architecture for IP over WDM," *Proc. OFC 2000*, Mar. 2000.
[4] K. Kompella *et al.*, "Extensions to IS-IS/OSPF and RSVP in Support of MPL(ambda)S," IETF doc. draft-kompella-mpls-optical-00.txt
[5] O. Crochat, J.-Y. Le Boudec, and O. Gerstel, "Protection Interoperability for WDM Optical Networks," *IEEE/ACM Trans. Net.*, June 2000.
[6] V. Paxson, "End-to-End Routing Behavior in the Internet," *IEEE/ACM Trans. Net.*, vol. 5, no. 5, Oct. 1997, pp. 601–15.
[7] Perlman and Varghese, "Pitfalls in the Design of Distributed Routing Algorithms," *SIGCOMM '88*, Aug. 1988, pp. 43–54.
[8] O. Gerstel and R. Ramaswami, "Optical Layer Survivability — An Implementation Perspective," *JSAC*, Sept., 2000.
[9] R. R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal Capacity Placement for Path Restoration in STM and ATM Mesh-Survivable Networks," *IEEE/ACM Trans. Net.*, vol. 6, no. 3, June 1998.
[10] S. Spencer and S. Woster, "Computer Industry Middleware Offers A Management Solution Without Requiring A Special Case for Telcos," *Telephony*, Aug. 1998.
[11] N. Golmie, T. D. Ndousse, and D. H. Su, "A Differentiated Optical Service Model for WDM Networks," *IEEE Commun. Mag.*, Feb. 2000, pp. 68–73.

BIOGRAPHY

ORNAN (ORI) GERSTEL (ori@ieee.org) received his B.A., M.Sc., and D.Sc. degrees from the Technion, Israel. After finishing his D.Sc. he joined the Optical Network Systems Group at IBM T. J. Watson Research Center and moved with the group to develop optical networking products with Tellabs Operations. There he served as the system and software architect for the Tellabs Optical Networking Group, building the TITAN 6100 metro DWDM product line. Recently he left Tellabs to join Xros (now Nortel), building all-optical cross-connects, where he is a senior systems architect. His research interests include network architecture, fault tolerance, and network design problems in optical networks.