NETWORK MIGRATION: EVOLUTION FROM RING TO MESH

Christopher Olszewski, John Labourdette, Frank Huang,
Daryl Chaires, Dimitrios Pendarakis
Tellium, Inc., 2 Crescent Place, Oceanport NJ 07757
colszewski@tellium.com (corresponding author)

1. Introduction

The evolution of carrier networks will likely proceed by integrating the functions of multiple traditional network elements (NEs) into single edge devices and incorporating mesh-like characteristics into the core. NE consolidation offers carriers the benefits of the latest technology in decreased operating expenses (fewer and smaller NEs taking up less space and power with reduced complexity and maintenance costs). The target architecture proposed in this paper provides both grooming at the edge and mesh networking in the core network, as well as a network management approach that makes use of automatic neighbor discovery and fast provisioning. Evolving current ring-based networks towards this next generation network (NGN) is clearly a challenge and is the subject of this paper.

Driven by increasing competition, but decreasing budgets and workforce, carriers are being pushed towards the creation of the NGN and ensuring its interoperability with their existing networks. The motivating factors of this shift are:

1. Reducing the capital expense of the bandwidth growth using fast shared-mesh restoration in the core NGN
2. Reducing operating expenses (for floor space, power, and maintenance), using fewer numbers and types of NEs in the NGN
3. Reducing configuration and provisioning errors (and the resources needed to correct them) using self-discovering NEs in the NGN
4. Accelerating revenue from new customers with shortened provisioning time using fewer hand offs in the NGN NEs and OSs
5. Maximizing the return on the current network by continuing to use embedded legacy NEs, which are expected to remain operational for several years
6. Offering new services efficiently to increasing numbers of customers, by supporting new services in a maintainable fashion in the NGN.

Existing ring-based transport networks have a variety of DCSs of different granularity (1/0, 3/1, 3/3) that groom traffic sent between offices. ADMs of different rates and types provide inter-office transport and ring protection (UPSR and/or BLSR), and also connect to large end-customer locations. An example of this network is shown in Figure 1. In the target NGN network, two NEs replace the multiple DCSs and ADMs: a multi-service transport switch (MTX) and an optical cross-connect (OXC). The MTX combines the functions of multiple ADMs for ring-type transport together with a next-generation DCS for DS3/STS1 grooming. The MTX also incorporates multi-service provisioning platform (MSPP) functionality for new services, predominantly data services like ethernet. An OXC in large offices will enact fast shared mesh protection between other OXCs, and support inter-MTX connectivity between these offices. WDM technology (coarse or dense) may also be used, particularly on highly-trafficked routes between larger offices. An example is shown in Figure 2.

Network management will also evolve in step with the deployment of the NGN. Management of the current network is carried out by several nearly independent OSs, each addressing a specific function (e.g., provisioning, fault management, etc.). The OSs are tied together (sometimes loosely and incompletely) through databases and related programs that were developed when NEs had little intelligence, and consequently had to be directed in great detail by their OSs. In the NGN, network management will be based on a layered NMS/EMS set of systems consistent with existing and evolving standards. The EMS layer allows NEs from different vendors to function in an integrated network management structure. Relying on the capability of NEs to act autonomously with network intelligence permits some NM functions to move into the NE control plane.

To efficiently use the existing infrastructure, both the current legacy network and the NGN must be connected together at a set of interfaces that support cross-network services. Provisioning service requests across both networks will be achieved by a service manager that partitions the requests appropriately between the current network and the NGN. This partitioning will use rules based on the endpoints of the service, the type of service, and the extent of the deployment of the NGN.

This paper discusses network evolution for a current-generation metro network, but the discussion may be applied to other networks as well. A real network transition plan must be prepared with an understanding of the network's architecture, intended use, and its existing management processes.

This paper is organized as follows: Section 2 discusses the current legacy network in terms of physical network infrastructure (i.e., the NEs and the network architecture), network management structure, and supported services. Section 3 presents a parallel discussion of the NGN. Section 4 discusses interworking between the current legacy network and the NGN, in terms of topology, management, and services. Section 5 offers a brief summary.

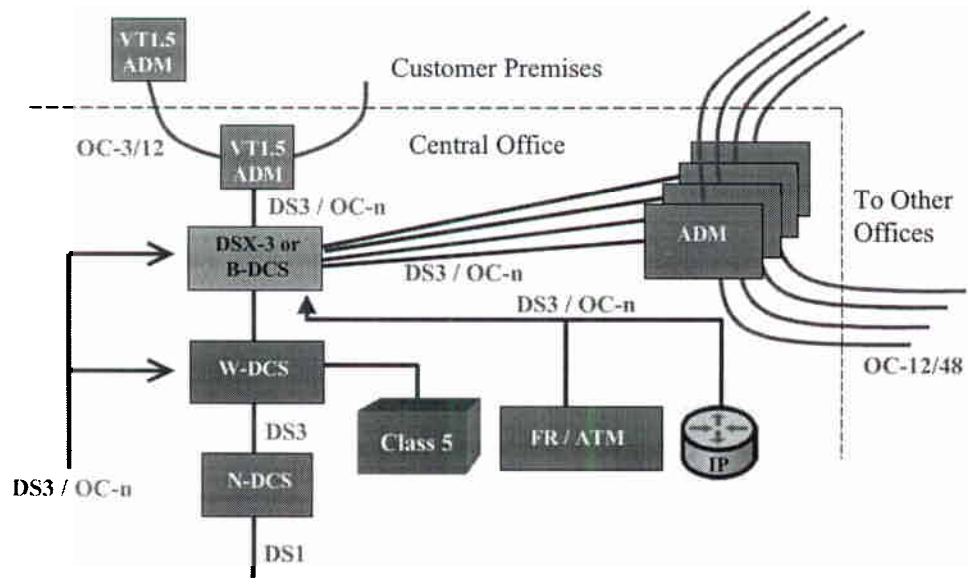

Figure 1 Current office architecture

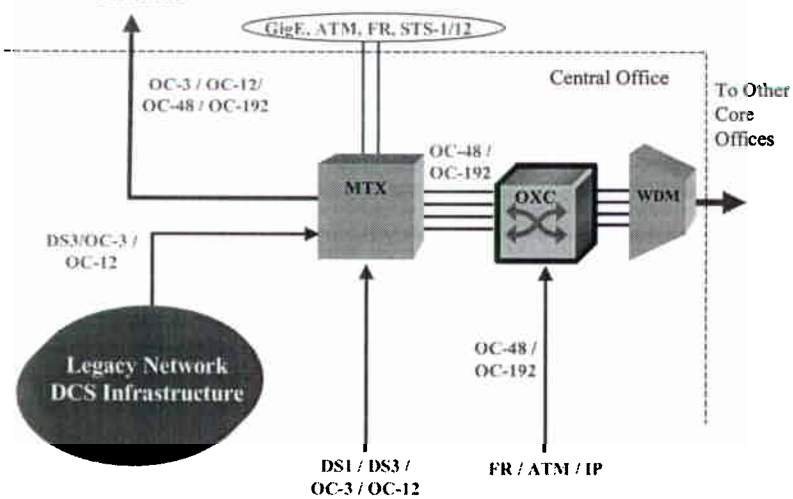

Figure 2 Next generation office architecture

2. Current Legacy Network

Physical Network Infrastructure

A schematic picture of the current transport network is given in Figure 1. Voice and private line services enter the transport architecture through the narrowband DCS (N-DCS), wideband DCS (W-DCS), and the broadband DCS (B-DCS). Patch panel frames of the appropriate signal rate provide connectivity between these DCSs.

Connections to customer locations can be through multiple means: for lower rate SONET services (OC-3, OC-12), ADMs may be used; for DS3 interfaces and below, TDM muliplexers can be used; for very high-rate services, multiplexers and dark fiber can be used. Data switches (IP routers, ATM) support customer data service interfaces directly.

The B-DCS serves as the main gateway for inter-office connectivity. Connections are sent between offices by ADMs, using BLSR or UPSR at the OC-48 line rate or above. Figure 3 shows a typical topology for the network. This figure illustrates the local office / hub office / core office roles in this network. Smaller local offices close to one another are connected together by small SONET rings. Larger offices on these rings are designated as hubs, and serve to interconnect traffic coming from the smaller rings. Hub offices with the largest amounts of traffic, aggregating traffic going to other regions of the network, are called core offices. Traffic from one local office to another may go through hub and/or core offices.

Core offices, hub offices, and local offices each have the same types of NEs in them. However, local offices may not have a B-DCS, and will likely have less equipment. Because of their small size and lesser volumes of traffic, local offices may not be upgraded as quickly as core offices or hub offices to the NGN.

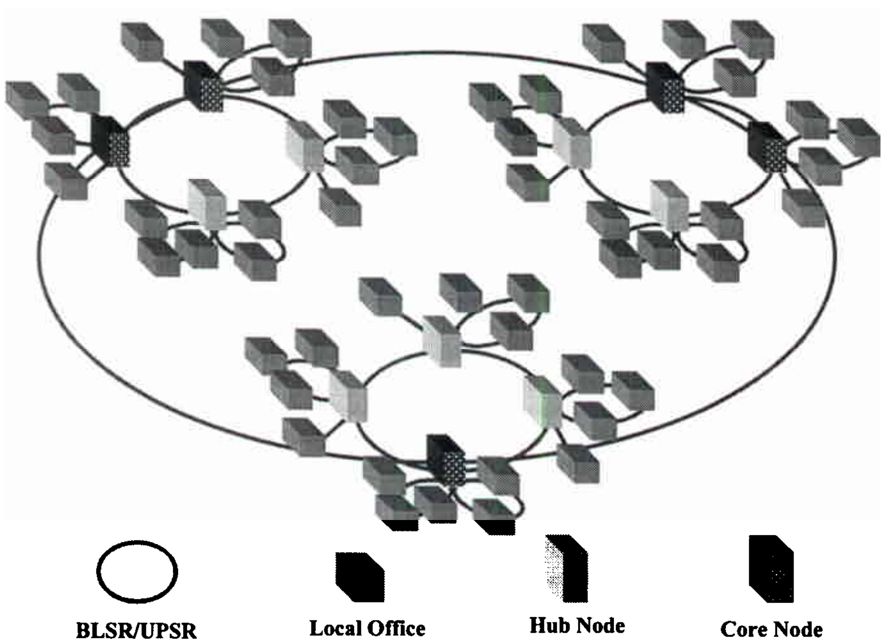

Figure 3 Interoffice connections and edge/hub/core concept of current network

Network Management Structure

The network management structure of the current network is generally complex and rigid. Because the NEs themselves have very little network awareness, the intelligence directing them comes from the management systems. Several separate systems carry out different aspects of network management (configuration management, fault management, etc.). These disparate systems work from a common view of the network, which for many carriers is based on long-enduring database programs for keeping track of trunks and circuits with their mapping and assignments. Communication from these systems to NEs can be direct (e.g., through TL1 messages), or through a layer of EMSs tailored specifically for each type of NE.

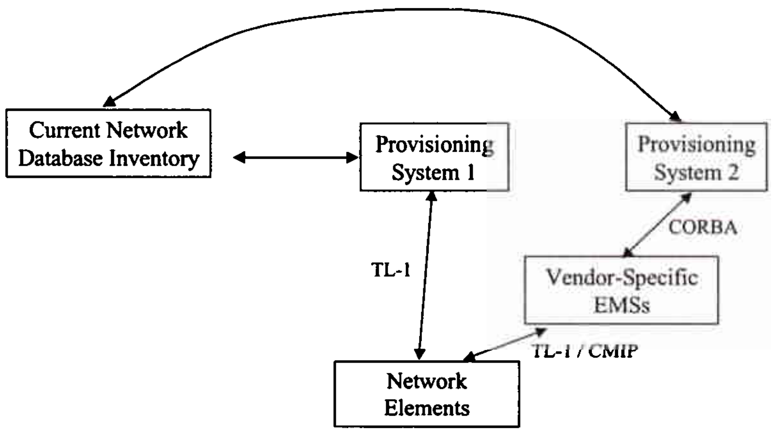

Figure 4 Current OSs for provisioning

Figure 4 shows a generic example of a current-generation OS structure for provisioning, which acts on specific work orders for new service. Provisioning system 1 interfaces directly with the NEs, while provisioning system 2 communicates with a different set of NEs through an intermediate EMS.* Multiple provisioning systems are not unusual in a company, since different types of services may require different processes and NEs, and also since telecommunications firms with different approaches acquire or merge with each other. Messages to the NEs configure the equipment for new service connections, and responses back indicate the success of these commands.

Other network management functions (fault management, capacity planning, etc.) use systems that essentially replicate the architecture in Figure 4. The network database keeps track of the configuration of the network. The database system for the current network keeps track of connections between managed transport elements, and is integrally important to many other functions. Due to their longevity and entrenchment, these legacy systems are difficult to change, are not designed to support the capabilities of the NGN equipment, and often exhibit performance limitations.

Current Network Services

The existing network supports traditional services using the array of various DCSs, together with the SONET ADMs. These traditional services include

- voice (switched) services
- private line services (DS1, DS3, and OC-N)
- data services (FR, ATM).

WDM equipment supports some wavelength services (OC-48) directly.

Service requests are made by customers, and are then followed by a complex and time-consuming process. Service requests are first verified and then passed on to the service-provisioning group. At this point, circuits are laid out in an automated but manually assisted fashion on the existing set of transport NEs, and the routes are planned out. The routes are analyzed to determine if extra equipment (e.g., port cards) is needed and in which offices. If required, orders for such equipment are placed. When all the equipment has been installed, tested, and made ready for service, the circuit provisioning orders are given to the existing OSs, and are executed. Any errors in the ordering, equipment, or database will become apparent, and will cause delays in the process. When noticed, these errors will have to be examined and fixed. Finally, when the provisioning request is successfully carried out, the customer will receive service and be billed.

Many of the separate steps in this process require manual hand-offs between separate organizational groups. Thus, the opportunity for typical errors in human-to-human communication is introduced which can produce errors in the

* One widely deployed database system for provisioning is TIRKS, developed in the late 1970s; it uses a system called TEMS to communicate with the NEs. These systems are traditionally deployed in RBOC networks.

request for service. The planning process requires human involvement in an off-line process, which can in turn slow the process down. Typical cycle times for service can range up to four months.[1]

3. Next Generation Network

Physical Network Infrastructure

A diagram of NGN NEs and their connectivity is shown in Figure 2. This diagram represents a core office as shown in Figure 3. The NGN NEs include a core optical cross connect switch (OXC), a multi-service transport switch (MTX), and wavelength-division multiplexers (WDMs). The OXCs provide switched connectivity between core offices. The MTX combines the functions of multiple layers of DCSs, multiple ADMs, and an MSPP. The WDM provides connectivity on routes between offices carrying a large amount of traffic.

Between OXCs in core offices, the NGN supports shared mesh topology with fast restoration.[2,3] Connections from the MTX, as well as wavelength services directly from customer locations, go through the OXC and thus to other core offices and their destinations. The OXCs provide a variety of protection classes based on optical mesh technology. These protection classes include diverse dedicated protection, shared protection, and preemptible protection classes. Direct high-bandwidth trunking connections between packet switches for ATM or IP services are also supported by the OXC.[4,5]

Connections to an office without an OXC will be through the MTX. The MTX may also support mesh, but will likely use only ring protection. The overall network architecture (core mesh, with peripheral rings) was found to be very efficient in terms of capital expense in previous studies, and shows similar efficiencies to two-tier networks.[6,7,8] As the network evolves, larger peripheral offices using only rings will gradually be brought into the core mesh network. The pace of this evolution will be governed by traffic growth and carrier economics. Customer locations will also connect directly to the MTX.

The MTX serves to groom incoming streams up to the line rate to be efficiently transported across the network. Having interfaces at the DS3 through OC-192 rates, the MTX will have the capability to groom traffic at the STS-1 or possibly even lower rate. The MTX thus acts as a replacement for the many DCSs in the current network. Additionally, the MTX supports connections from new and existing data services:

- Gigabit Ethernet (Gig-E)
- Frame relay (FR)
- ATM
- Etc.

Offices having insufficient available fibers between them, or offices with large geographical distances between them, are places where WDM technology will be deployed. Some services could be carried directly on the WDM systems.

The main difference between hub and core offices is the absence of the OXC in the hub offices. As mentioned above, hub offices will likely be connected to the rest of the network by rings. Rapidly growing hub offices may become core offices. Local offices may remain mostly unchanged by the evolution in the larger offices of the network. With small amounts of traffic, upgrading local nodes to the NGN may not be cost effective.

Network Management Structure

One natural and increasingly applied approach to manage the NGN is to deploy separate element management systems (EMSs) for each new type of NEs, and manage them together under a Network Management System (NMS). In this approach, a standards-based interface such as TMF 513 / TMF 814 can be used between the EMSs and the NMS to manage the entire network. [9,10] Multiple vendors of different types of NEs can then be managed in this way from a single platform, allowing end-to-end provisioning and network visibility. The challenge in this architecture is the detailed development of the NMS. Figure 5 illustrates this approach. The problem may be simplified if multiple types of NEs could be managed with one EMS, but this approach is common only within a single vendor's suite of equipment.

Control plane (NE-to-NE) interactions could help in managing a network of disparate equipment for establishing connections. NGN NEs will have the capability of carrying out and coordinating actions that in the current network are carried out purely through the management plane.[4] One instance of this coordinated activity is restoration of services using shared mesh between OXCs. In the NGN, this restoration would take place entirely through the control plane: however, NEs involved in these actions would then inform the NMS (through the management-plane) that failure and restoration events have occurred.

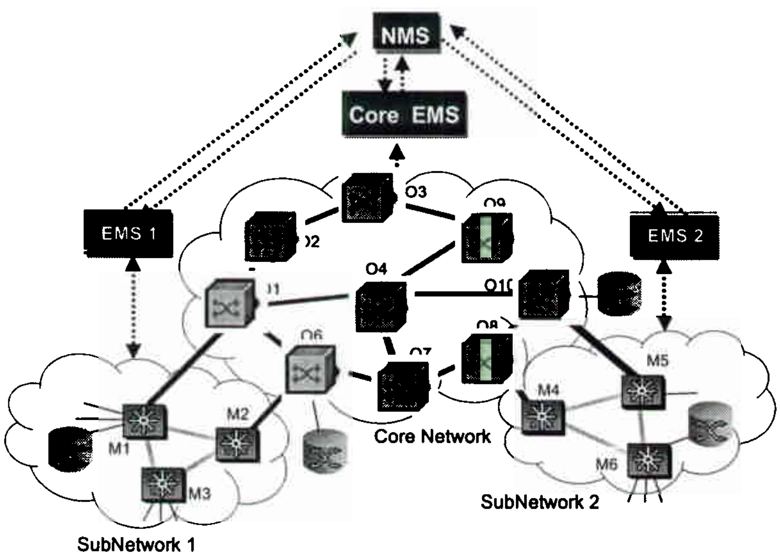

Figure 5 Next generation network management structure with NMS/EMS integration

Self-discovery of the configuration and topology of the NGN NEs maintains the accuracy and availability of the NGN's network data as the NGN grows. Network usage tracking and capacity planning for the NGN will thus be a more straightforward task than in the current generation network. Equipment and service provisioning will be faster and more accurate.

NGN Transport Services

In addition to the set of services supported on the current network, the NGN transport network will support new and existing services using the MTX and the OXC:

- Gigabit Ethernet (GigE)
- ESCON, FICON, etc.
- OVPN[11]
- Wavelength services
- ATM / FR

The exact mix of services offered will depend on the customer demand at a given location and the deployed NEs. In addition to the end-customer services, the NGN will also offer transport services (connectivity) between current generation NEs.

Orders must still be validated in the NGN transport network. However, the time required for service planning and provisioning can be decreased and manual sources of errors removed.[1] The network itself becomes the database of record, and the network (or the NMS) will be able to quickly determine the route of a service and then implement it automatically (or, under human supervision if so desired). Services could be requested directly by the user through UNI signaling with proper authorization.

As the NGN nears the limits of its installed capacity, additional equipment will need to be deployed. With accurate tracking of used and available capacity, the carrier will be made aware if the amount of available capacity falls below a pre-specified threshold.

4. Legacy Network / NGN Interworking

Physical Network Connectivity

As show in Figure 2, one of the "customers" of the NGN is the current generation network. This connectivity allows a service originating in the existing part of the network to be transported over the NGN to its destination. An equivalent representation is that the current legacy network provides the NGN connectivity to customers served by smaller offices without the necessity of replacing functional NEs in those offices. As the network grows, NGN NEs will likely be deployed to serve new customers as current NEs reach their capacity. Thus, the NGN network is a hybrid not only in the ring-mesh sense, but also in the current generation / next generation sense.

Network Management Interoperation

The current legacy network and the NGN will coexist over an extended period of time, since current operational NEs will not be replaced in bulk. Both the current and the NGN NEs will simultaneously be active in the same network, and so they will need to be managed together. The aim in managing this hybrid network is to incorporate the new capabilities of the NGN NEs, while continuing to use the current NEs where it makes sense.

In managing this heterogeneous network, various goals of this network migration are to simultaneously:

- Minimize changes to the current network management structure
- Allow use of the capabilities of the NGN NEs
- Provide a common way to manage the configuration, provisioning, and alarm information of both the current and next generation networks.

A simple approach to manage these networks together is at the service level. Below the service level the networks remain essentially separate. Each network (the current and the NGN) is responsible for managing itself. Physical connections between the current and NGN carry services going between them. Each part of the network thus sets up and monitors services within itself: the current network with its current processes, and the NGN with its more streamlined process. To allocate service requests between the two networks a *service manager* is introduced. The service manager determines, using simple rules, whether a service request can be met entirely within one network, or must enter both networks. The service manager generates separate requests on each network for each connection.

Figure 6 illustrates this concept. The functions of the service manager could be incorporated into the NMS of the NGN, although these functions would not be needed at all if the current network did not exist. Implementation of very similar concepts has already begun in standards processes and commercially.[12,13]

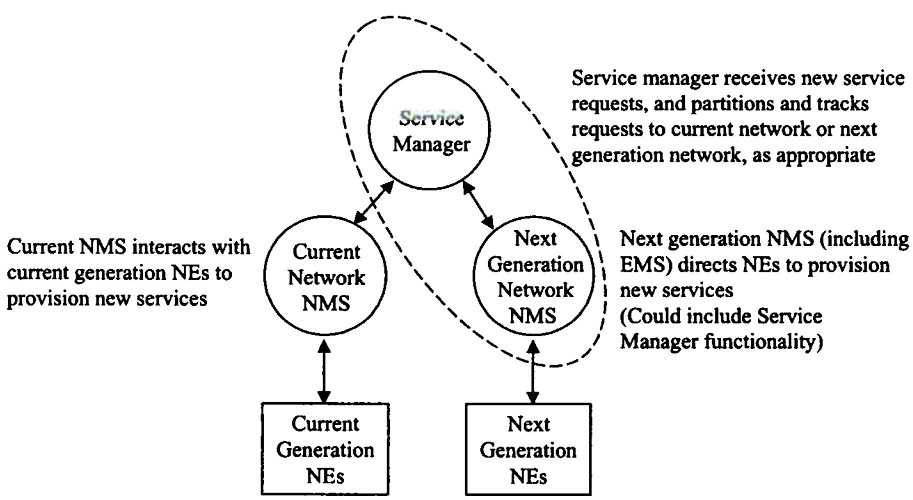

Figure 6 Service manager partitions services to current and next generation networks.

This approach minimizes changes that need to be made to current network management systems. Requests made on the current network will be carried out with the same accuracy and speed as they are now. Requests made solely on the NGN, by contrast, should be met faster and more accurately. As more and more new services are provisioned on the NGN, service delivery time and performance will improve. Large scale and presumably costly changes to the current OSs to support the NGN NEs are not needed.

Some other obvious approaches have operational complexities, or are equivalent to the proposed solution:

- The current network set of OSs could be enhanced to manage the NGN NEs. This approach ignores the increasingly intelligent behavior of the NGN NEs, and forces them to be managed as their less-capable predecessors. Neither service delivery time nor accuracy will be improved by this approach.
- A control-plane approach could be implemented to convey service requests between the current network and the NGN. Because the current equipment does not support control-plane signaling, a "proxy agent" for the current network would need to be developed to handle this communication.[1] However, the functions that would be invoked in this way are very similar to those invoked with the service manager concept. Indeed, the proxy agent would have to communicate with the current OS structure in much the same way.

The proposed approach also has the effect of quarantining the current network management from the NGN, while it allows the current network to continue to satisfy new service requests where deploying NGN NEs is not economic.

For the purpose of order tracking, trouble shooting, and billing, the service manager must be able to associate a single service request with service orders on both the current and NGN networks. The service manager will generate a unique identifier that is passed as part of the service requests in the current and NGN networks. This unique identifier can then be used to track orders in both networks. During failures, these identifiers would also be used to track problems with a circuit between networks.

Current Network/Next Generation Network Service Interoperation

The NGN's deployment will be an overlay network in the sense that both current and NGN networks will exist in the same locations. Since the NGN will be more efficient, that part of the network should grow preferentially. Some services will be available on both the current network and the NGN while other new services may only be available on the NGN, and so must wait for the deployment of NGN NEs in a nearby office.

During and after the deployment of the NGN transport network, current generation network equipment will still be operating and generating revenue. The current generation of NEs must also be available to meet new services in places where the NGN NEs are not deployed. These service requests will initiate activities in both the current and NGN to satisfy these demands. The NGN will increasingly supply the transport connectivity even between current generation NEs. In this sense the NGN will be supporting existing services in multiple ways.

Validated orders for new services will be given to the service manager. Using simple rules, this service manager will parse the order request into pieces that can be satisfied in the current network (if necessary), and pieces that can be satisfied in the NGN network.

To parse a service order and send it to the specific network to be filled, the service manager will need to know:

- The rules for dividing a complete service request into pieces
- The interface points on both the current and the NGN that are connected together and in service.

The rules for dividing up a service are based on:

- The type of service requested
- The endpoints of the service
- The availability of NGN NEs at the endpoints to meet the service needs
- The connectivity of the offices serving the endpoints to the rest of the network.

Policies governing the allocation of services to the current or NGN will be embedded in the rules themselves. For example, in certain locations where a service connection could be made through either the current legacy network or the NGN, a policy may prefer one network over the other, or distribute the load in some proportion. Note that this policy governing the assignment of services between networks is in addition to policies within the current or NGN networks.

The service manager must also have the complete set of the interface points between the current network and the NGN. These interface points form essential parts of the service orders that the service manager generates to the service ordering procedure for both the current network and NGN. Besides physical interfaces, these points may also be associated with specific types of services being offered in one network or the other (e.g., an interface to a frame relay switch in the existing network that has capacity to serve more service requests).

The part of the request sent to the NGN will be handled within that piece. The NMS of the NGN will have a complete topological view of the NGN, as well as the knowledge of which NEs can satisfy requests for a particular service. If there are sufficient resources already allocated in the NGN, the provisioning requests will be executed immediately, or at the time specified on the order. Once the NMS receives the request, the provisioning process should be carried out very quickly (minutes).

The part of the request sent to the current network will be handled through its usual existing procedures.

When both parts (current and NGN) of a multi-part request are satisfied, the service manager will initiate the procedure of notifying the customer and beginning the billing cycle.

Figure 7 gives an example of the process of service allocation: a DS3 private line service is ordered (offered on both the current and NGN networks) between an office not in the NGN (A), and an office in the NGN (C).

We assume that the service manager has the following information or rules:

1. Office A does not have NGN NEs
2. Office C does have NGN equipment able to accept DS3 service requests
3. Between office A and office C for DS3 service, the NGN network should be used as much as possible
4. Office B serves as the gateway to the NGN for office A.
5. A list of interface points between the current network and the NGN in both offices B and C (points (d) and (e) in Figure 7).

The service manager then computes least-cost routes through the hybrid network. Some routes could be entirely in the legacy network; other routes could be within the legacy network to the interface points ((d) and (e) in Figure 7), and then within the NGN between these interface points. Based on the policies above, the service manager will route

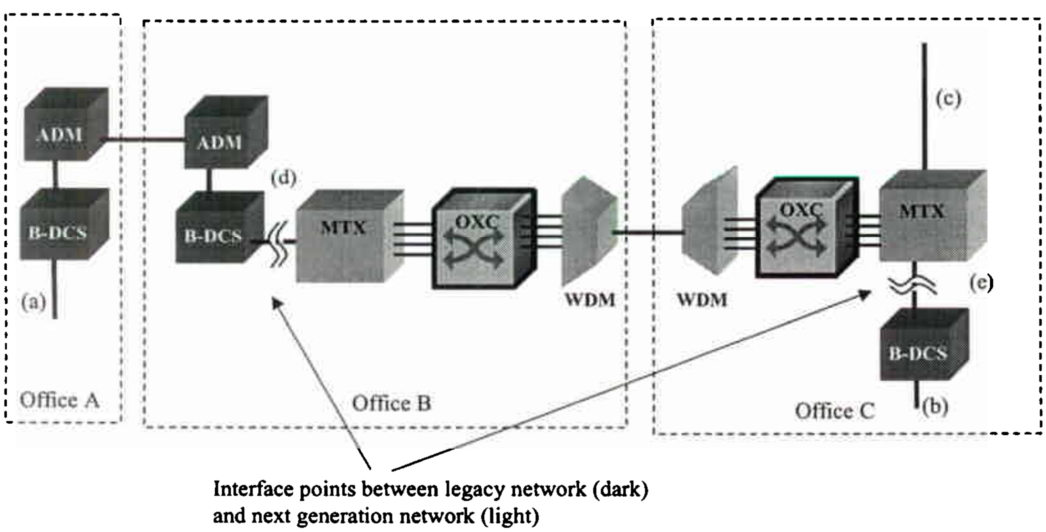

Figure 7 DS3 service crossing between current and NGN.
Points (a) and (b) represent the termination point on the current network.
Point (c) represents an alternate termination point on the NGN

the service from (a) to (b), going through the interface points (d) and (e) between the current legacy network and the NGN.

The service manager will then generate two requests on the current network:

1. A DS3 connection between (a) and (d)
2. A DS3 connection between (e) and (b).

In addition, the service manager will generate one request on the NGN:

1. A DS3 connection between (d) and (e).

If the endpoint in office C were not specified to be at point (b) on the current B-DCS, but instead were at point (c) on the MTX, the second service request on the current network would not be needed, and the request for service on the NGN would be from point (d) to point (c).

5. Summary

Incorporating NEs with new capabilities in an existing network will become increasingly essential. Driven by economic factors, carriers will need to take advantage of NGN NEs, while continuing to use the current legacy network. Not only will the growth of the physical network be challenging, so also will the management of this hybrid network. The approach proposed here is essentially to manage the current network and the NGN separately, and to use a service manager to unify service requests between them. In this way, operations within the old network do not need to be changed significantly (at significant cost), and the NGN network can operate without the constraints imposed by current network management techniques. This division of responsibilities allows the core of the network to migrate from a ring architecture to mesh.

6. References

1. "Network Element 'Secret Agents': Starting at the Bottom", Yankee Group, January 2002.
2. "Routing and Restoration Architectures in Mesh Optical Networks", G. Ellinas et al., Optical Networks Magazine, Issue 3:5, September/October 2002.
3. "Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks", J. Labourdette et al., Photonic Network Communications, vol. 4., no. 3, 2002, pp. 219-235.
4. "IP over Optical Networks: Architectural Aspects", B. Rajagopalan, et al, IEEE Communications magazine, September 2000.
5. "Role of Optical Network and Spare Router Strategy in Resilient IP Backbone Architecture", J. Labourdette, E. Bouillet, and S. Chaudhuri. DRCN 2003.
6. "Carrier Network Architectures", D. Johnson and N. Hayman. BT white paper. Available at: http://www.tellium.com/documents/BTexact_Network_Study_Whitepaper_FINAL.pdf.
7. "Capacity Requirements of Ring, 1+1, and Mesh Configurations", S. Samieian, et al., Lightwave, August 2001.
8. "Two-Tier Network Economics", C. Olszewski et al., NFOEC 2002.
9. TeleMangement Forum, TMF 513, "Multi-Technology Network Management Business Agreement: NML-EML Interface v2.0", October 2001. Available at http://www.tmfcentral.com/sdata/documents/TMFC663%20TMFC338%20TMF513v2[1].0.pdf.
10. TeleMangement Forum, TMF 814, "Multi-Technology Network Management Solution Set: NML-EML Interface v2.0", October 2001. Available at http://www.tmfcentral.com/sdata/documents/TMFC967%20TMF814v2.0.pdf.
11. Optical Network Magazine, special issue on oVPN to appear. Information at http://optical-networks.com/cfp_jukan.html.
12. TeleManagement Forum, TMF 508, "Connection and Service Management Business Agreement, v3.0", April 2001.
13. "Supporting Next Generation Networks without Replacing the Legacy OSS", D. Curtis. Wisor white paper. Available at: http://www.wisor.com/solutions/whitepapers/Next%20Generation%20Solution%20in%20a%20Legacy%20World.pdf.