Intercarrier Bandwidth Exchange: An Engineering Framework

Subir Biswas and Debanjan Saha, Tellium Inc.
Nitish Mandal, Queens University

ABSTRACT

A mismatch between demand and supply for bandwidth is common in transport carrier networks. This mismatch is generally a result of the disparity between a carrier's capacity buildout and its anticipated customer demand. A carrier with temporary bandwidth deficit or lack of presence in a geographical region and a carrier with surplus capacity in the right locations can be brought together by the emerging bandwidth exchange technology. Bandwidth exchange offers a win-win solution, in which the carrier with deficit avoids losing revenue by buying capacity from the carrier with surplus, and the latter makes additional revenue by retail sale of its excess capacity. While the concept of real-time purchase and exchange of bandwidth has attracted a lot of attention, many technical challenges stand in the way of making it a reality. The purpose of this article is to provide an engineering framework for enabling real-time bandwidth exchange with committed quality of service and service level agreement among transport carriers. Special emphasis is given to identifying the architectural requirements and the enabling infrastructure necessary for building a viable bandwidth exchage that can be used for creating revenue out of surplus stranded capacity. In-depth analysis of cross-carrier SLA specification, capacity publication, route design, and service provisioning are also provided in the article.

INTRODUCTION

Bandwidth exchange and trading as a tool for investment risk management is lately gathering a great deal of attention in the telecom industry [1, 2]. In today's competitive marketplace where the price of bandwidth is falling and attaining profitability is proving to be harder and harder, carriers are unwilling to see even an ounce of unsold capacity in their network. They have come to a realization that in addition to the traditional wholesale approach to bandwidth, its commoditization and retail sale across competing carriers can save lost revenue by complementing each other's surplus and deficit.

Capacity surplus and deficit are generally a result of the mismatch between a carrier's capacity buildout and its anticipated customer demand. A carrier with temporary bandwidth deficit or lack of presence in a customer's site, and a carrier with surplus capacity in the right places can be brought together by the emerging bandwidth exchange technology, which has the potential to offer a win-win solution to both carriers. The carrier with deficit avoids losing revenue by buying capacity from the carrier with surplus, and the latter makes additional revenue by retail sale of its excess capacity.

Figure 1 illustrates today's capacity exchange and service activation process. The slow process of manual contracting and exchanging is aggravated by the lack of any standardization in capacity availability, dynamic pricing, service level agreement (SLA), and quality of service (QoS) publication formats among carriers. The service activation cycle is also lengthened due to the lack of any just-in-time provisioning infrastructure. The result is a six-to-nine-month turnaround time for supporting services involving multiple carriers.

Excess capacity, tied up with this long selling cycle, loses its liquidity, and that translates into lost revenue opportunity for a carrier. Also, the capacity that is surplus only for a duration shorter than the exchange-activation cycle cannot be sold. In addition, the sales, general and administration (SG&A) cost [3] involved in business and legal contracting adds to the reduction in earnings.

Therefore, the clear business mandate for the industry is to shrink both the capacity purchase and service activation cycles; for that, automation is deemed to be the solution. In this article we describe an engineering framework and its necessary components for enabling automated bandwidth exchange across multiple carriers.

Automatic bandwidth exchange across two or multiple carriers would require the following:

- A standard framework for SLA specification and monitoring
- Physical infrastructure for bandwidth exchange through pooling points across multiple carriers' networks
- Availability publication and resource sharing across carriers at the operations support system (OSS) and network management layers
- Intercarrier service provisioning

0163-6804/03/$17.00 © 2003 IEEE

The purpose of this article is to address these issues and to define an engineering framework that can be used for integrating various architectural subsystems into a real-time bandwidth exchange infrastructure with committed SLA.

The rest of this article is organized as follows. We define the scope of bandwidth exchange and, using a real-life example, explain how it is manually done in today's transport networks. An exhaustive definition and characterization of the SLA is provided. The issues regarding intercarrier connectivity and in-carrier capacity surplus are explained. We present the proposed automatic exchange architecture and its components, including an intelligent exchange layer. Cross-carrier service management, including capacity advertisement, routing, and provisioning, is also presented.

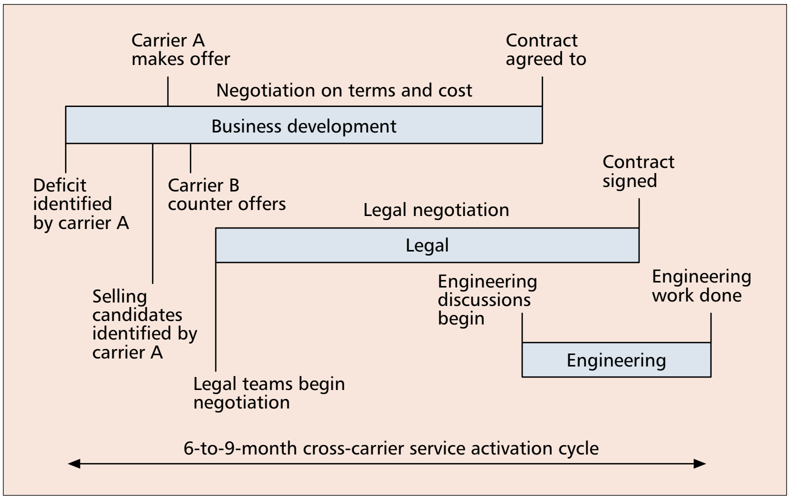

Figure 1. *Today's manual inter-carrier capacity exchange cycle.*

AUTOMATIC BANDWIDTH EXCHANGE

The process of automatic bandwidth exchange is explained using the example scenario shown in Fig. 2. In this example, carrier 1 and carrier 2 both have network presence in North America and Europe. While carrier 1 has denser presence in Europe, carrier 2 provides better geographical coverage in North America. Consider the situation in which carrier 2 receives a 2.5 Gb managed lambda service request from city A in North America to city B in Europe. Since carrier 2 does not have a point of presence (POP) in that European city, the only way it can earn revenue from this request is by spot-purchasing capacity from carrier 1 in real time. Note that the carriers share joint presence in a number of cities and POPs. These shared POPs are referred to as *collocation points* and are typically used as junction points for capacity exchange across the carriers.

Another situation in which real-time purchase will be useful is for a service request P-to-Q from a customer of carrier 1. Although carrier 1 has presence in both end cities, it does not have enough internal capacity to route the service. Therefore, a solution for carrier 1 would be to spot-purchase retail capacity from carrier 2.

For reasons outlined earlier, this spot-purchase/exchange process should be automated and the amount of manual intervention should be minimized. Our proposed automation cycle is shown in the flowchart in Fig. 3.

After a customer enters a service request with the desired protection and SLA specifications, its servicing transport service provider (TSP) checks if it has POPs at both ends of the requested service. If the servicing TSP has POPs at both ends and has enough internal capacity to satisfy the desired SLA, the service is approved and provisioned using the TSP's own resources. This is the usual flow of events most of the time. However, when the servicing TSP realizes that either it does not have POPs at one or both end cities, or the internal resources are not enough for supporting the specified SLAs, it invokes what is referred to as a *capacity exchange cycle*.

After the initiation of an exchange cycle, the servicing TSP checks an external resource database that contains capacity availability information published by other TSPs. A TSP publishes availability information about its surplus capacity with the following attributes:

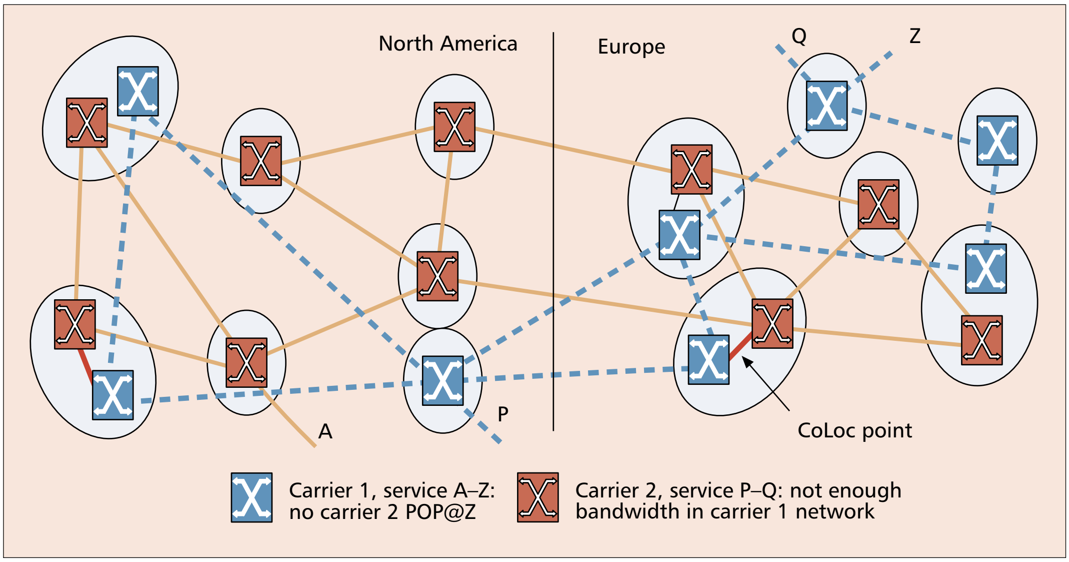

Figure 2. *An example of capacity exchange between two transatlantic carriers.*

Note that after service feasibility is confirmed, it may or may not be possible to invoke cross-carrier service activation immediately because, although the seller carrier has surplus capacity, there may not be sufficient inter-carrier physical connectivity to sell that bandwidth.

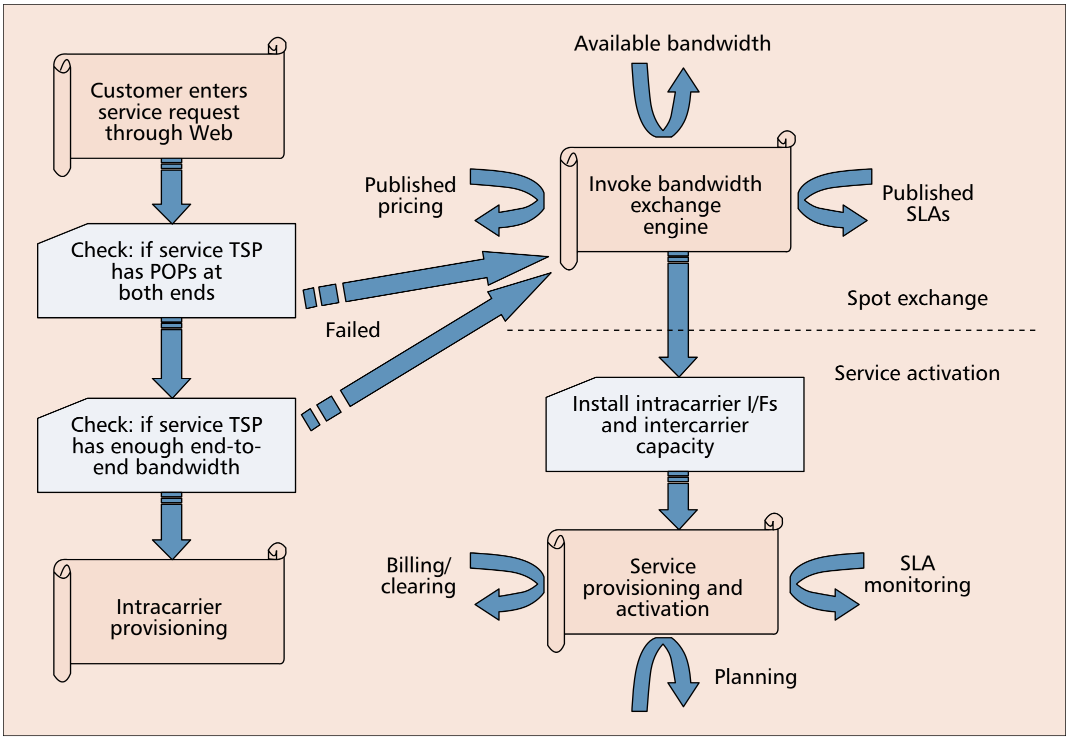

Figure 3. *Automated exchange and activation cycle for cross-carrier service provisioning.*

- Amount and topology of the available capacity that is ready to be sold
- SLA that can be supported on this capacity
- Asking price for real-time spot sell

Based on this external capacity availability information, the servicing TSP determines if it can support the requested service with specified grades of SLA. While making this decision, the servicing TSP considers its own internal resources as well as resources available for spot purchase from other participating TSPs.

If a service route is feasible, the servicing TSP initiates an intercarrier service provisioning process. More about provisioning using OSS and network management system (NMS) infrastructure will be explained later. The carriers also execute billing and clearing procedures across each other for cross-accounting purposes. Finally, after a cross-carrier service is established, the servicing TSP continues to monitor SLA and end-to-end performance on the service. This is done in order to validate the pre-exchange SLA contracts from the advertising TSPs.

Note that after service feasibility is confirmed, it may or may not be possible to invoke cross-carrier service activation immediately because, although the seller carrier has surplus capacity, there may not be sufficient intercarrier physical connectivity to sell that bandwidth. In this situation, the right physical connectivity has to be manually established through the collocation points (Fig. 2); only then can the automated provisioning start.

SERVICE LEVEL AGREEMENT

SLAs [4, 5] specify the service quality commitments of a service provider to its customers. SLAs first came into prominence in the late 1990s, when the Frame Relay Forum adopted its "Service Level Definitions Implementation Agreement" in document FRF13 [6]. Guidelines in this implementation agreement have defined acceptable parameters for several key characteristics of frame relay services, such as frame transfer delay, frame delivery ratio, data delivery ratio, service availability, and uptime. Besides service quality, SLAs also cover a variety of details, including corrective actions, penalties and incentives, dispute-resolution procedures, nonconformance, acceptable service violations, reporting policies, and rules for terminating a contract.

There are different types of SLAs covering different areas of end-to-end services. For example, SLAs can be designed for network services, hosting services, or application services. For the purpose of this article, we focus on network SLAs. A network SLA covers the characteristics of the network itself, such as network availability, throughput, data loss rates, latency, and sometimes security. In the following we outline each of these items.

Network Availability: Network availability is defined as the percentage of time the network is up and operational. While 100 percent uptime may be everyone's goal, 99.5–99.9 percent are more realistic averages. Some SLAs also list stricter standards for network availability including mean time between failures (MTBF), service restoration latency, and so on. One key element of a network SLA is specifying penalties for downtime. Penalties may vary depending on when outages occur. For instance, downtime at 2 a.m. may not disrupt a typical enterprise's business, and hence may carry a lesser penalty than an outage at 2 p.m., which may be catastrophic.

Network Throughput: Throughput is defined as sustained data rate measured over a period of time. For private line services this runs from 56 kb/s (a dialup connection) up to 10 Gb/s (OC-192/STM1-6 service). For data services throughput may vary over time and is more complicated to specify and measure. For example, for frame relay it is often specified as a long-term average rate and a short-term burst rate. For IP services, service providers often use 95th percentile throughput measured over time windows of 5 min as the service throughput.

Loss Rates: Another key part of a network SLA is data loss rate. Low data loss rate is very important for real-time applications such as voice over IP (VoIP) or interactive media that cannot operate effectively at a high loss rate; packet loss in the 5 percent range is acceptable for non-streaming applications such as Web browsing.

Data Latency: Data latency, like data loss, is critical in VoIP and multimedia applications where delays must not impact end-user performance. Real-time interactive applications require response times of 100 ms or less. Web browsing, on the other hand, remains viable at 250 ms. In practice, network providers often guarantee a round-trip delay of 80–100 ms between the routers in their core networks.

Security: Security SLAs are becoming increasingly important. A security SLA defines how data is protected while in transit over the service provider's network. Unlike the hard and fast metrics of the network SLA parameters, a security SLA is determined by customer requirements and is more subjective. Issues specific to this SLA include use of public or private encryption keys and the level of encryption such as DES or Triple-DES [7].

Some of the manual and most time-consuming parts of today's bandwidth exchange are SLA publication, negotiation, monitoring, reporting, and management. One of the prerequisites of automated bandwidth exchange is automating SLA management, and the first step of that is SLA standardization. Unfortunately, network SLA standardization development is still in its infancy. The Policy working groups in the Distributed Management Task Force (DMTF) and Internet Engineering Task Force (IETF) are working on the first steps of SLA standardization. However, any real deployment based on these standards is still a long time away. Meanwhile, private interoperability between the service providers is the only path available for moving forward.

INTERCARRIER CONNECTIVITY

Excess capacity inventory and intercarrier physical connectivity are two essential components for real-time bandwidth exchange. To understand the form of capacity surplus, one needs to look at a typical capacity deployment scenario. New capacity installation between two central offices involves:

- Fiber laying through dug-up conduits between the offices
- Connecting dense wavelength-division multiplexing (DWDM) systems at the fiber ends
- Connecting DWDM transponders to the switch interface cards of the central office switches

Once a conduit is dug, since the incremental cost of installing more fibers through it is relatively small, the carriers prefer to lay a large number of fibers even when not all the capacity is pre-sold before installation. Installation of DWDM and the switch interface cards, on the other hand, is relatively less involved and are typically done on demand.

From an exchange perspective, the unused bandwidth for which the fiber infrastructure is already installed (also known as *dark fiber*) is considered surplus. An unsold wavelength from city A to city B through an installed fiber system is considered excess even if it is not connected to the DWDM and the switch.

Exchanging surplus capacity requires intercarrier physical connectivity, which is achieved through shared central offices referred to as colocation (CoLoc) facilities [8]. Participating carriers place their switches in a CoLoc office, and the desired connectivity across carriers is achieved by connecting the switches. As shown in Fig. 4, connectivity across these switches can be established either directly or through carrier-neutral switches referred to as *pooling points*. In direct connectivity, fibers are run directly across different carriers' switches. The role of the CoLoc facility in this case is just to provide floor space, power, and other resources for hosting switches from multiple carriers.

Connections through pooling points offer more flexible intercarrier connectivity and therefore allow quick reconfiguration with new capacity exchange requirements. The flip side is the involvement of a third party pooling point provider who needs to play an intermediary role in the exchange and service provisioning process. Typically, a CoLoc facility provider also offers the pooling point infrastructure for its shared office [8].

Note that an intercarrier service always needs to be provisioned through one or multiple CoLoc points, which makes it particularly vulnerable to a failure or disaster involving the CoLoc facilities. Therefore, if an intercarrier service requires service protection, its backup service routing should always be CoLoc-divergent from the primary service route.

In a metro environment, it is also possible for carriers to be directly connected through fibers that are installed between the carriers' own central offices. This method of connectivity has the advantage of being completely independent of the CoLoc facility providers.

BANDWIDTH EXCHANGE ARCHITECTURE

We propose a capacity exchange framework that leverages the transport carriers' existing OSS and NMS infrastructure for automating cross-carrier availability publication, resource sharing, service provisioning and performance management. As illustrated in Fig. 5, a capacity exchange protocol layer is proposed on top of the existing OSS/NMS layer. For a carrier that participates in spot bandwidth exchange, this new protocol layer consists of an exchange engine, which interacts with its peer exchange engines from other carriers. An exchange engine also communicates with its local OSS/NMS/EMS layer software in order to perform capacity availability and cost publication, cross-carrier service

Unfortunately, network SLA standardization development is still in its infancy. The Policy working groups in the DMTF and the IETF are working on the first steps of SLA standardization. However, any real deployment based on these standards is still a long time away.

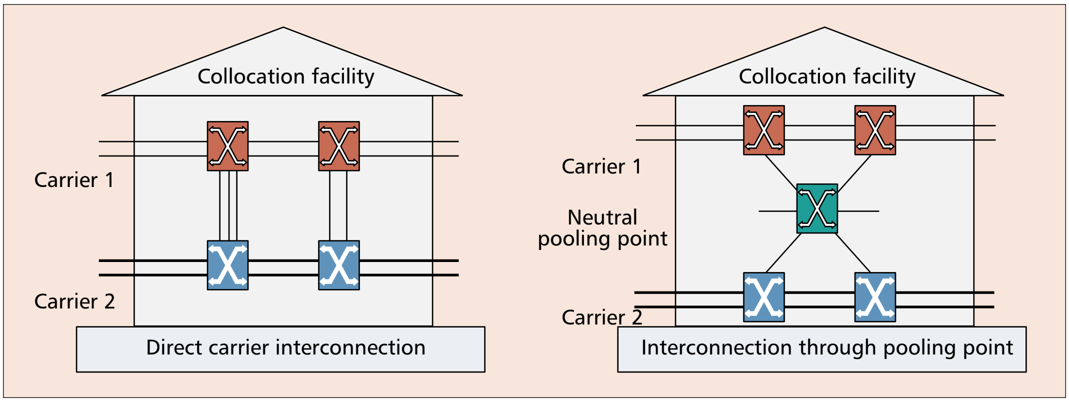

Figure 4. *Intercarrier connectivity scenarios at CoLoc facilities.*

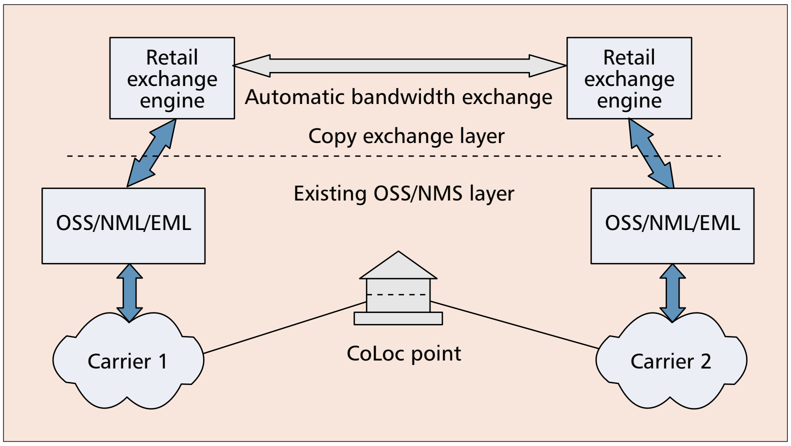

Figure 5. *Automatic capacity exchange using OSS/NMS infrastructure.*

route design, provisioning, as well as other standard fault, configuration, accounting, performance, and security (FCAPS) management operations [9, 10]. In addition, the exchange engine is also responsible for coordinating cross-carrier billing/clearing and QoS/SLA monitoring.

Details of the exchange engine, OSS/NMS layer, and their role in the exchange process are illustrated in Fig. 6. For simplicity, only the carrier 1 infrastructure is shown in this figure. Other carriers that participate in automated capacity exchange should maintain similar exchange infrastructure.

Based on geographical and administration boundaries, a carrier's network is partitioned into multiple subnetworks. In the element management layer (EML), for each subnetwork there is one element nanagement system (EMS), which is responsible for configuration, provisioning, fault, and performance management within its managed subnetwork. The network management layer (NML) is responsible for similar functions but across multiple EMSs' subnetworks within the same carrier. For example, the NML layer will handle transatlantic path provisioning within carrier 1's network in Fig. 2. In order to handle the technology and protocol heterogeneity of equipment from multiple vendors, it is desirable to develop a standardized interface between the EML and NML. Such an interface is currently being developed by the TeleManagement Forum (TMF). TMF's MTNM 513 [11] is a Common Object Request Broker Architecture (CORBA)-based interface that is widely accepted by transport carriers as well as equipment vendors, and has already been implemented within a wide range of transmission networks that comply with the TMN [12] framework.

While NML and EML perform provisioning, fault, and performance management, the OSS layer is responsible for managing the business aspects of a service. Typical OSS layer functions [12] include service fulfillment (ordering and provisioning), service assurance (QoS, SLA management), and customer billing. TMF is developing a standard interface, Connection and Service Management Information Model (CaSMIM), [13] that can be used for interaction between the OSS and NML layers. CaSMIM provides a means of translating the service definitions made in service management systems into technology- and network-specific implementations.

Components of the bandwidth exchange layer are shown within the shaded area in Fig. 6. The resource availability engine taps the inventory and configuration databases at multiple layers and maintains a database of excess inventory that can be potentially sold to other carriers. Available capacity is qualified with attributes such as pricing, QoS, and feasible SLA specifications. The availability engine also includes a policy module that enforces any resource publication policy based on intercarrier business agreements. Examples of such policies include variable intercarrier pricing, capacity visibility to selected partner carriers, and planned oversubscription to surplus capacity.

It is the exchange agent that, with appropriate security and authentication semantics, makes these published quantities available to its peer agents within participating carriers. Through this automated and mutual inventory publication process, a carrier knows how much of other carriers' bandwidth is available for real-time purchase during service creation for its own customers. Purchasing and service creation take place based on these published quantities. Note that the exchange agent of a carrier represents the only point of contact with other carriers. Cross-carrier service provisioning and management is performed through these peering agents.

The cross-carrier billing engine is responsible for billing the spot sells among the carriers. Peri-

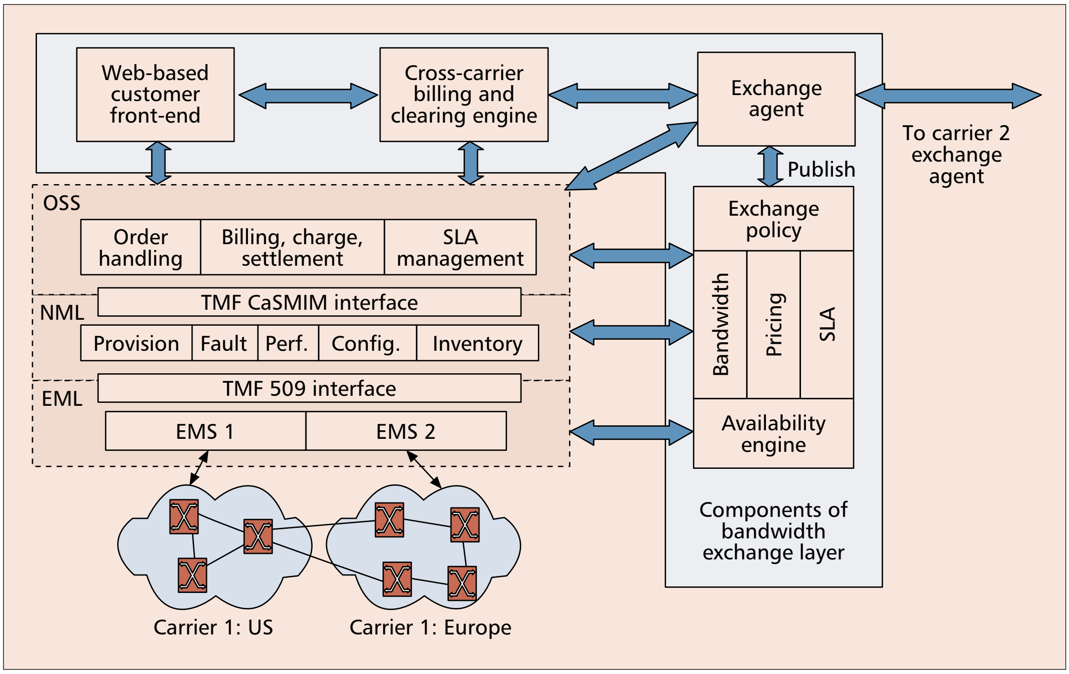

Figure 6. *Architectural components of the capacity exchange layer.*

The cross-carrier billing engine is responsible for billing the spot sells among the carriers. Periodic clearing, instead of transaction for every sell, keeps the cross-carrier accounting overhead under control.

odic clearing, instead of a transaction for every sale, keeps the cross-carrier accounting overhead under control. Note that this billing engine is different from the usual billing entity built into the OSS layer, primarily for customer billing. This cross-carrier billing engine works in synchronization with the OSS layer billing for handling the sold cross-carrier services.

CROSS-CARRIER SERVICE MANAGEMENT

AVAILABILITY PUBLICATION

A participating carrier publishes its surplus capacity through its availability engine and the exchange agent. An example of availability publication is illustrated in the example network in Fig. 7. Consider two carriers' networks interconnected through CoLoc facilities at cities A, B, and C. Between two CoLoc points i and j, the surplus capacity intended for real-time sale is designated by an availability vector $\mathbf{V}_{i,j}$. A carrier publishes its surplus capacity to another carrier in the form of a table of availability vectors as shown in Fig. 7. Carrier 1, for instance, in Fig. 7 publishes an availability table with entries for A–C, B–C, C–A, and C–B. This indicates that carrier 1 intends to sell its excess capacity between cities A and C, and B and C, but not between A and B.

In addition to the sellable bandwidth information, an availability vector contains pricing, provisioning granularity, duration of availability, and offered protection and SLA guarantees. Depending on the specific cross-carrier relationship, a carrier may publish different availability vectors for the same sellable bandwidth to different carriers. Based on the external capacity availability information in its availability tables, a servicing TSP determines if it can support a cross-carrier service with the required grades of SLA specification.

ROUTING

Cross-carrier service provisioning involves route design based on local resources as well as remote resources that appear in a servicing carrier's availability tables. We will explain the route design process using the example shown in Fig. 8. Consider a scenario in which carrier 1 intends to buy retail capacity from carrier 2 in order to support service A–Z that it otherwise cannot accommodate with its internal available capacity. To compute a cross-carrier route, carrier 1 first creates its own network graph, which is shown as the blue dashed graph. It then overlays additional links based on the availability information published by carrier 2. These overlay links and their costs are assigned based on the information published by seller carrier 2. The overlay links are depicted as solid orange lines in the graph in Fig. 8. Path computation for service A–Z is performed based on this overlaid graph consisting of the solid as well as dashed links.

After it computes an end-to-end route for the service, carrier 1 identifies its route segments. In this example, the end-to-end service A–Z has three segments among which segments A–B and C–Z are to be routed using arrier 1's capacity and segment B–C is to be routed using capacity that needs to be bought from carrier 2. After the service route is designed, carrier 1 provisions segments A–B and C–Z, and informs carrier 2 about segment B–C. Upon receiving this information, carrier 2 provisions segment B–C with committed SLA, QoS, and protection specification.

Protection offered to individual segments is typically local to the carrier. In this example, while carrier 1 provides and executes protection/restoration protocols for segments A–B and C–Z, carrier 2 protects segment B–C. Note that the protection commitment for B–C is part of the availability information published by carrier 2, and the protection type is considered while computing the end-to-end path.

In some situations, it may make perfect sense for a servicing carrier to route a service entirely through purchased bandwidth from carriers that offer low-price deal to sell off their surplus capacity. This will help the servicing carrier in better servicing its internal capacity only to the local customers.

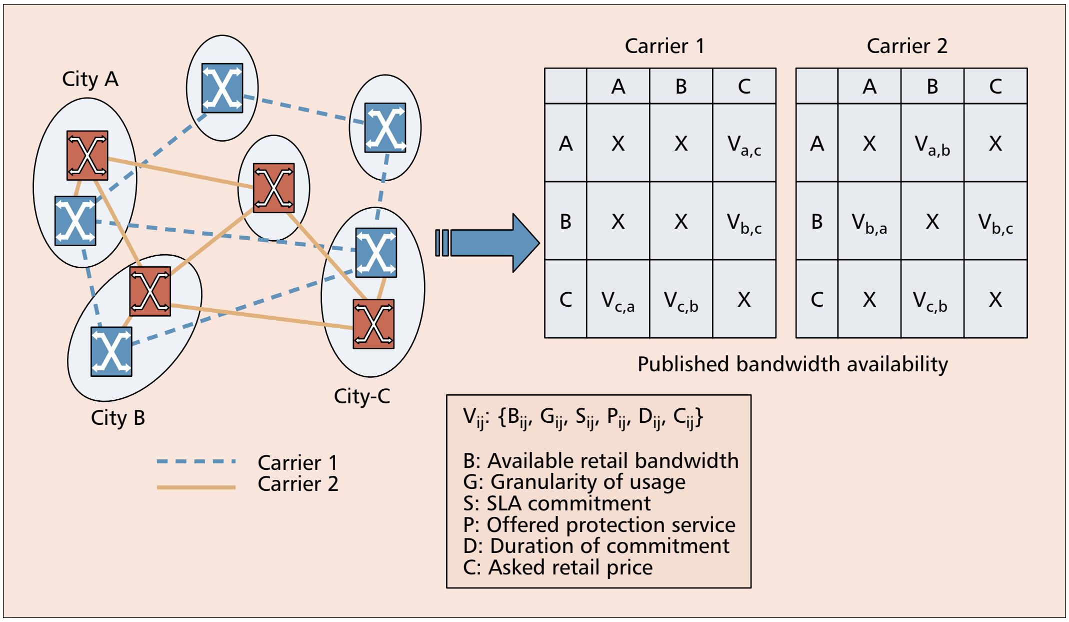

Carrier 1

	A	B	C
A	X	X	$V_{a,c}$
B	X	X	$V_{b,c}$
C	$V_{c,a}$	$V_{c,b}$	X

Carrier 2

	A	B	C
A	X	$V_{a,b}$	X
B	$V_{b,a}$	X	$V_{b,c}$
C	X	$V_{c,b}$	X

V_{ij}: $\{B_{ij}, G_{ij}, S_{ij}, P_{ij}, D_{ij}, C_{ij}\}$

B: Available retail bandwidth
G: Granularity of usage
S: SLA commitment
P: Offered protection service
D: Duration of commitment
C: Asked retail price

Figure 7. *Capacity availability publication vector.*

PROVISIONING

The exchange, provisioning, and billing cycle for a cross-carrier service is explained through the timing diagram shown in Fig. 9. After a carrier's OSS/NML/EML software decides it cannot support a service with its own available capacity, it decides to buy retail bandwidth from one or multiple participating partner carriers. Prospective seller selection is performed based on the availability databases available in the buyer's availability engine. After a seller is chosen, the buyer's exchange agent queries the seller's exchange agent about the proposed spot purchase and potential cross-carrier service provisioning. The agents exchange a number of messages for mediating any unsynchronized availability information before a transaction is agreed upon.

At this point, the exchange engines on both sides check if the necessary interconnectivity exists between the carriers in question. If not present, work orders are generated to connect them through appropriate CoLoc POPs, through either direct connectivity or third party pooling switches. The seller's exchange engine also checks if the published excess capacity is *connected* or *unconnected*. If it is unconnected, the seller's exchange agent generates work orders for installing appropriate DWDM and switch interface cards for converting the involved capacity status from unconnected to connected. It should be noted that this part of the exchange cycle is manual and, if needed, can slow the provisioning process in certain situations.

Once capacity installation is performed, the buyer carrier's NML and EML layers perform route computation based on the internally available and purchased external capacity. Different segments of the multicarrier service route are provisioned as described in the previous section. After provisioning is completed, the exchange agents of both carriers initiate billing and clearing processes. Also, SLA and QoS objects are created with which the buyer carrier starts monitoring the health of the provisioned service. Monitored parameters are relayed back to the billing server so that the billing and clearing can be performed based on the committed SLA and QoS compliance from the seller's side.

In this example, cross-carrier service provisioning was illustrated with only two participating carriers. The same principles can be extended for scenarios where a buyer carrier spot purchased bandwidth from multiple carriers and an end-to-end service is routed through all the involving carriers' networks. The exact business model and pricing for the spot exchange will decide the nature of the cross-carrier service route chosen by a servicing carrier. In some situations, it may make perfect sense for a servicing carrier to route a service entirely through purchased bandwidth from carriers that offer low-price deals to sell off their surplus capacity. This will help the servicing carrier to better service its internal capacity only to local customers.

SUMMARY AND CONCLUSIONS

This article provides an engineering framework for enabling automatic cross-carrier bandwidth exchange with committed QoS and SLAs. It surveys today's manual exchange technology and demonstrates how its lengthy business, legal, and engineering contracting processes are responsible for limited asset liquidity and subsequent revenue loss in transport carriers' networks. It is shown that automated bandwidth exchange, coupled with near-real-time service activation, has the ability to translate surplus capacity into short-term revenue.

We have proposed an architecture in which an intelligent capacity exchange layer is added on top of carriers' existing OSS/NMS/EMS layers. This new exchange layer offers a number of exchange-related functions, including cross-carrier capacity availability publication, billing, exchange peering, and cross-carrier provisioning. The role of OSS/NMS/EMS and their integration with this new layer was addressed.

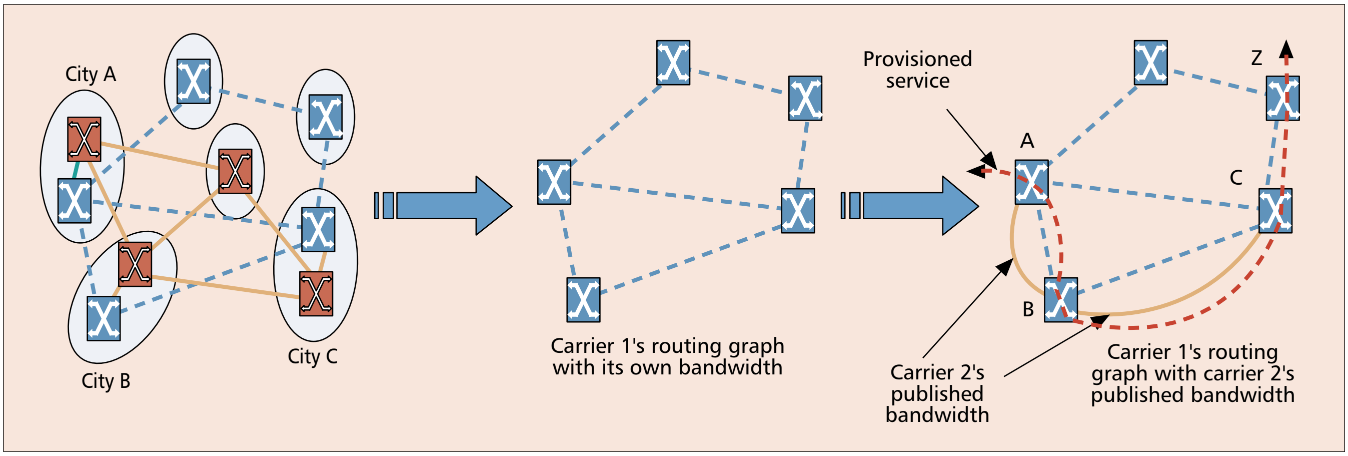

Figure 8. *Route computation for cross-carrier services.*

The article identifies a set of standardization requirements on which the telecommunication industry and carriers have to agree in order to support the proposed exchange framework. Some of these standards have already started and are in a reasonably mature state of development. The role of these standards in the exchange process and their current developmental state are also addressed in the article.

REFERENCES

[1] "Bandwidth Trading: The Long and the Short of It," *Yankee Group Rep., Energy Commun.*, vol. 3, no. 7, Nov. 2000.
[2] P. V. Kohut, "Understanding the Need for Bandwidth Trading," *Capacity 2001*.
[3] J. Altmann, B. Rupp, and P. Varaiya, "Internet Demand under Different Pricing Schemes," *ACM Conf. Elect.Commerce 1999*.
[4] D. Verma, *Supporting Service Level Agreements on IP Networks*, 1st ed., Macmillan., 1999.
[5] Ri. Sturm, W. Morris, and M. Jander, *Foundations of Service Level Management*, 1st ed., 2000.
[6] FRF.13, "Frame Relay Forum Service Level Definitions (SLD) Implementation Agreement," 1998.
[7] R. Atkinson, "Security Architecture for the Internet Protocol," IETF RFC 1825, 1995.
[8] "Pooling Points: Building the Ubiquitous Network, One Port at a Time," *Yankee Group Rep., Energy Commun.*, vol. 3, no. 15, Dec. 2000.
[9] R. D. Doverspike, S. J. Phillips, and J. R. Westbrook, "Transport Network Architectures in an IP World," *Proc. INFOCOM 2000*, vol. 1, pp. 305–14.
[10] D. Medhi *et al.*, "A Network Management Framework for Multi-Layered Network Survivability: An Overview," *IEEE/IFIP Conf. Int. Net. Mgmt.*, Seattle, WA, May 2001, pp. 293–96.

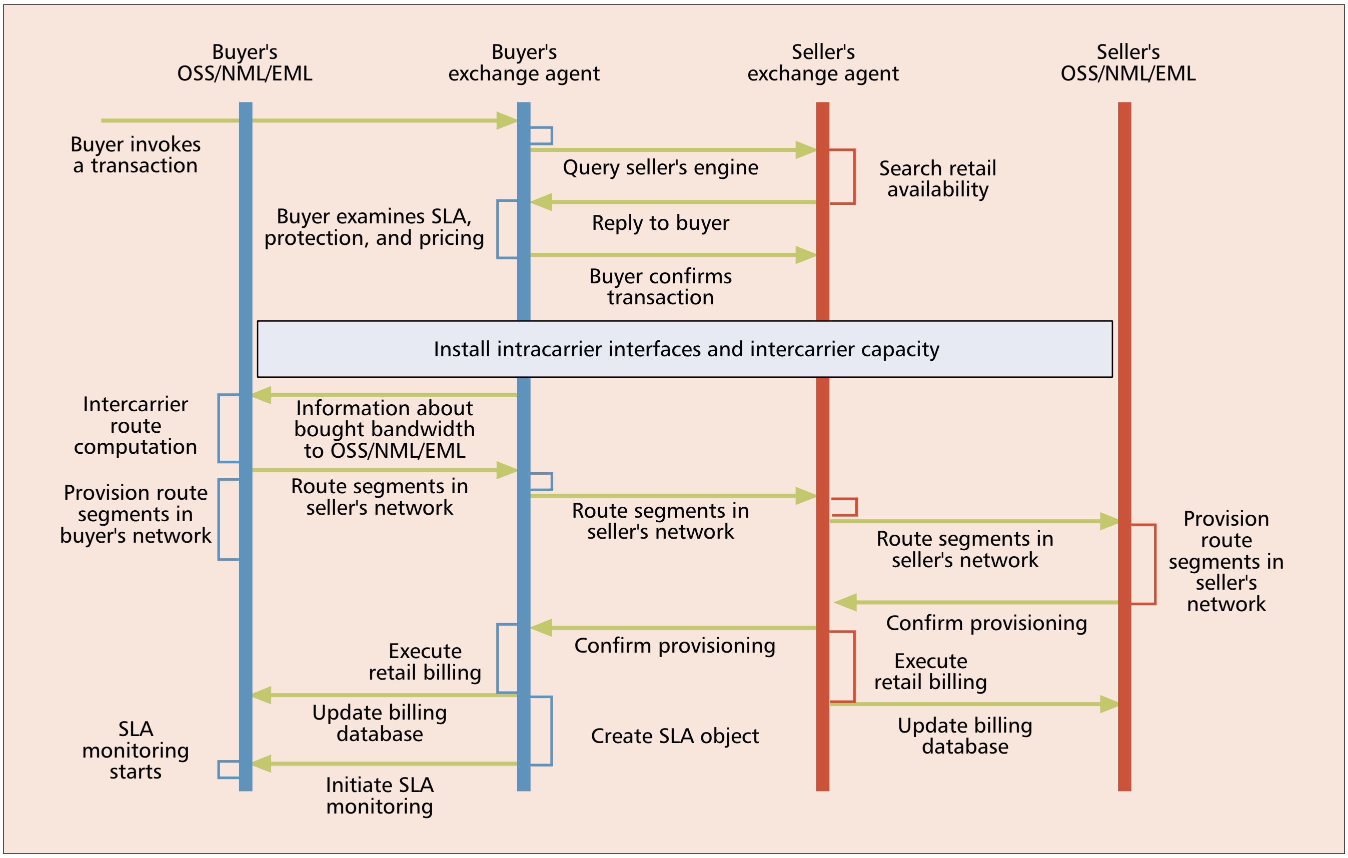

Figure 9. *A timing diagram for a cross-carrier capacity exchange and provisioning cycle.*

Some of the standards identified in this article have already started and are in a reasonably mature state of development.

[11] "Multi-Technology Network Management Business Agreement, NML-EML Interface Version 2.0," TMF 513, Aug. 2001.
[12] "Telecommunications Management Network (TMN)," white paper by Vertel; http://www.iec.org/tutorials/tmn/index.html
[13] "Connection and Service Management Information Model (CaSMIM) Business Agreement," TMF508 v3.0, Sept. 2001.

BIOGRAPHIES

SUBIR BISWAS, Ph.D. (subir_biswas@ieee.org), is a principal architect at Tellium. In this role, he is the lead architect for optical network management and advanced application development using Tellium's Aurora™ family of optical switches and StarNet Software Suite™. Prior to joining Tellium, he worked for NEC's C&C Research Laboratories, Princeton, New Jersey, where he worked on wireless ATM, IP multicast, and IP traffic engineering problems. He holds a Ph.D. from the University of Cambridge, United Kingdom. He has more than 10 years of experience in data and transmission networking industry and has published around 40 journal/conference articles and book chapters and registered several patents in these areas. His areas of interest include next-generation optical and wireless network protocols, Internet security, 3G wireless services, and mobile commerce.

DEBANJAN SAHA, Ph.D. (debanjan@acm.org), heads the Business Applications group at Tellium. In this role he consults with Tellium customers, some of the largest service providers in the world, on IP and transport network planning, design, and deployment. Previous to his tenure at Tellium, he spent several years at IBM Research and Lucent Bell Laboratories where he designed and developed protocols for IP routers, multiservice switches, and Internet servers. He is well known in the industry for expertise in the area or IP and optical networks. He is actively involved with various standards bodies, most notably IETF and OIF. He also serves as editor of international journals and magazines, and technical committee member of workshops and conferences. He is a notable author of numerous technical articles on various topics of networking and is a frequent speaker at academic and industrial events.

NITISH MANDAL (nmandal@mbast.queensu.ca) is currently pursuing an M.B.A. at Queens University, Canada. He recieved his M.A.Sc. in electrical engineering from Simon Fraser University. He has over nine years of experience in the telecom industry, and has held various positions in product development, systems engineering, and marketing of core networking products. His area of interest is in management of technology and operations.