Ericsson Review
284 23-3174 | Uen

Transport networks in the cloud age

Architecture simplification in the move to cloud service provider.

ELISA BELLAGAMBA, PAUL LÜSSE, DAVID ALLAN AND RICCARDO MARTINOTTI

Network convergence, the rise of content distribution networks and the consolidation of service hosting enabled by the cloud model are significantly altering traffic patterns being observed by service providers. These developments, along with changes in the business environment, are driving the choices for technology.

The cloud phenomenon

Several factors have come together to form a virtuous circle that is changing the way people use the internet and fixed and mobile-broadband services. Ubiquitous broadband access, the reduced cost and increased number of mobile devices, the spread of content-delivery networks, and the virtualization of web services are all causing a massive shift in user behavior – so that traffic patterns bear almost no relation to those seen just a few years ago.

Broadband and mobility have driven up demand for bandwidth and network-hosted applications. Cloud computing and the massive virtualization it enables are influencing the way this demand will be met and what network architecture will be needed to support applications that are bandwidth-hungry, storage-demanding and computing-intensive.

Cloud computing is a framework for hierarchical virtualization of resources, and leads to extreme consolidation of computing capacity and hosting facilities. Wholesaling interfaces are exposed at every level of the hierarchy – ranging from infrastructure as a service (IaaS) to software as a service (SaaS) – with every layer being continually expanded to the point where the term currently used is anything as a service (XaaS).

Cloud computing offers increased efficiency and resource elasticity, which results in significant savings in network hosting costs – primarily through data-center consolidation, massive virtualization, deployment of a uniform resource base (identical servers in large quantities, for example) and, ultimately, significant economies of scale.

Impact on the traffic matrix

The cloud supports an extensive range of applications that exhibit a variety of characteristics and place different demands on the network. As a result, several hosting scenarios are possible, each one having a different impact on the network.

Storage-demanding applications are extremely sensitive to latency and packet loss, and have a high ratio of back-end bandwidth to bandwidth consumed by users. Studies[1] indicate that the bandwidth consumed within the data center (DC) exceeds user-to-DC bandwidth by about a factor of 4, and when inter-data-center traffic is factored in, the ratio of DC-to-user traffic is about five to one. The ratio of user-to-DC traffic is expected to remain constant for the next few years due to better storage networking, and improved workload management by large-scale applications through the use of reduction techniques. One such technique – MapReduce – splits user requests into smaller, mutually exclusive, problems and then distributes the smaller pieces across several servers

BOX A Terms and abbreviations

AES	Advanced Encryption Standard
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
CLE	customer-located equipment
CL-PS	connectionless packet switching
CO-CS	connection-oriented circuit-switched
CO-PS	connection-oriented packet-switched
DC	data center
DoS	denial-of-service
EDGE	Enhanced Data rates for GSM Evolution
E-LAN	Ethernet LAN service
E-Line	Ethernet line service
FC	Fiber Channel
FR	frame relay
GFP	generic framing procedure
IaaS	infrastructure as a service
inter-DC	inter-data-center
IP	Internet Protocol
IPsec	IP Security
IP-VPN	Internet Protocol Virtual Private Network
ISP	internet service provider
LAN	local area network
MPLS	multi-protocol label switching
OAM	operations, administration and maintenance
OTN	optical transport network
OTT	over-the-top
POTP	packet optical transport platform
PSTN	public switched telephone network
QoS	quality of service
SaaS	software as a service
SAN	storage area network
SDH	synchronous digital hierarchy
SDN	Software-Defined Networking
SLA	Service Level Agreement
SME	small and medium-sized enterprise
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/ Internet Protocol
TDM	time division multiplex
TLS	Transport Layer Security
VLAN	virtual local area network
VM	virtual machine
VPN	virtual private network
WAN	wide area network
WDM	wavelength division multiplexing
XaaS	anything as a service

working in parallel. Employing such techniques can cut both internal latency and the bandwidth-distance product required to support a given application, which in turn lowers transport costs and maximizes usage of cloud-computing resources.

The effect on the network is clear. Current storage-demanding applications in combination with consolidation of hosting services are shaping the traffic matrix into a hub-and-spoke pattern – where the hubs are large data centers utilizing cloud technologies. Just a few years ago, carriers were concerned about the rise of peer-to-peer networking, which was causing unpredictable shifts in the network-traffic matrix. This is in stark contrast to the current situation, where a significant amount of user-generated traffic is routed to comparatively few destinations rather than fanning out to a large number of end points. As more and more private applications switch to public clouds, the hub-and-spoke trend is likely to continue.

As traffic patterns become more predictable, overall network operations become less complicated as the dilation factor – the additional bandwidth needed to add burst tolerance to an unpredictable any-to-any traffic matrix – applied to network design can be reduced. This in turn enables improved utilization of network assets, while cloud technologies create efficiencies in computer processing.

The growth in traffic to the cloud is expected to continue in parallel with the process of consolidation, as cloud-based applications proliferate and the consumption rates of bandwidth-hungry subscribers continue to rise.

FIGURE 1 Cloud services for enterprises – structural relations

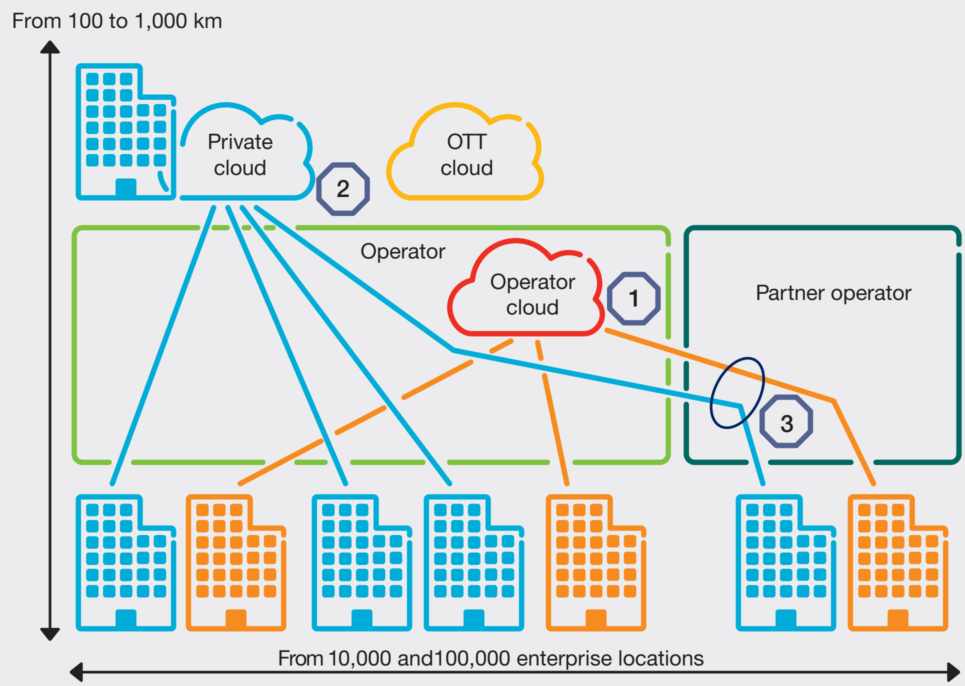

Requirements

The basic system requirements and many of the fundamental technologies for data centers remain the same, despite the transition to the cloud. Scale, virtualization and the associated security considerations, however, have all changed.

Inter-DC connectivity places a number of new requirements on the network in terms of latency, bandwidth and packet loss. The need for connectivity between data centers is primarily driven by geo-redundancy – which enhances service availability – and thereby the need to synchronize multiple storage locations in real time.

Replicating data between data centers in a synchronized manner is usually not possible due to the extensive physical distance between data centers. Concurrent replication over long distances is not feasible in a reasonable time frame, without creating processing delays or race conditions among applications. As a result, replication usually occurs in an asynchronous way, which brings latency into play.

Consolidation of application hosting into larger data centers using cloud techniques leads to a proportional growth in the bandwidth of connection pipes. The need for lossless SAN interconnect typically imposes requirements for non-blocking WAN connectivity. The trend toward large cloud-based data centers generates a similar need, but on a much larger scale, and creates the possible necessity for dedicated inter-DC connectivity.

In the immediate future, owing to the variety of cloud and virtualization technologies available, the network will have to support the common denominator technology: Ethernet VLAN bridging. This technology enables Layer 2 (L2) mobility among computing resources – virtual machines (VMs) can be migrated to different sites in the cloud – which requires some form of L2 overlay in the LAN and WAN to facilitate scaling. Layer 3 (L3) overlays are also under discussion, but are in their infancy in terms of being a practical replacement for bridging because they require other techniques, such as SDN, to be standardized.

The transition over time of computing and storage resources to cloud-enabled data centers will give rise to a number of interconnect arrangements. These arrangements will impose additional security requirements when it comes to accessing applications and where public and private clouds are interconnected – in what are known as hybrid cloud scenarios.

Hybrid cloud scenarios are not the only development necessitating new security requirements. Wholesale models and hosting consolidation into ever-larger data centers places strong requirements on tenant isolation in

FIGURE 2 **Cloud-service connectivity – the transport operator's point of view**

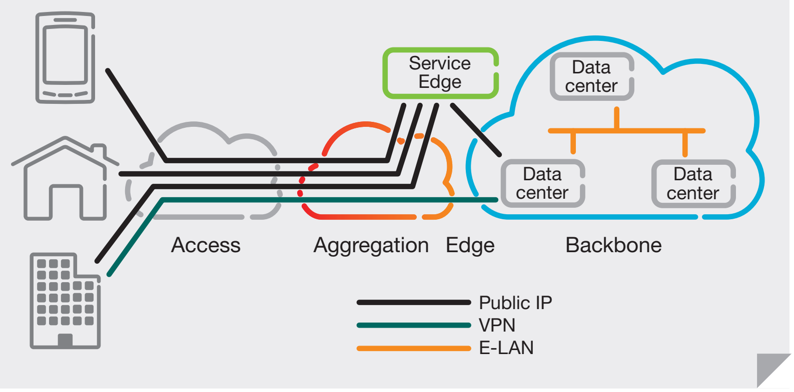

the LAN, SAN and WAN. The current best-practice method of meeting these requirements is VLAN L2 isolation – a technique that will probably continue to be used not only as a security mechanism, but also as the most practical form of resalable network virtualization. By bridging VLANs with an overlay that eliminates geographic restrictions, network connectivity can be provided to usefully sized communities or closed user groups without the complication of further defining virtualized networking components (switching and routing, for example). Here, the absence of a virtualized control plane makes the architecture simpler.

Business environment

The expectations of private subscribers and enterprise users for seamless network access, regardless of the connection method, are rising constantly. The need to meet these expectations is driving the convergence of fixed and mobile networks – minimizing transport infrastructure costs by maximizing resource utilization. The enabler of such convergence is considered to be IP, as it provides the connectivity logically required to support all the services of the Networked Society.

To enable the cloud age, the implementation of the simpler hub-and-spoke networks that will ultimately replace distributed PSTNs will not only take time, but will also present a set of challenges relating to the ability to support and migrate legacy services during the transition period.

Operators are facing topology changes in their networks and must constantly adapt to a rapidly evolving business environment. One of the most obvious current trends is the effort being made by some OTT players to integrate their business models, which comprise social networks, communication services, mobile devices and content – all dependent on huge data centers. To be successful, the services provided by OTT models need to be able to reach the user, and thus the ability to integrate these models depends entirely on the connectivity provided by network operators. If OTT players and traditional service providers are to be able to cooperate or compete, several business models will need to be established, especially as large multinational service providers begin to launch cloud-based services in direct competition with OTT players. In another model, tier-2 and tier-3 operators might partner with OTT players. The trend of consolidation and partnership is both broadening the competitive landscape and leading to greater cost pressure.

To serve residential users, operators are establishing a service layer that makes use of cloud-computing principles to increase flexibility and reduce time-to-market for their offerings. To effectively serve enterprises, stringent SLA requirements that address concerns such as security are needed. Services provided by cloud operators can be attractive to small and medium-sized enterprises (SMEs) because they enable these businesses to buy multiple services, including connectivity, from a single supplier and avoid the role of system integrator (point 1 in **Figure 1**). Large multinational enterprises and government organizations are tending toward consolidation of their services in private clouds (point 2 in Figure 1). In such cases, operators have to provide connectivity between locations, often using only L1 or L2. The wholesale model – where operators resell connectivity provided by a wholesale operator to enterprises – adds yet another dimension (point 3 in Figure 1).

Network requirements in the cloud era are changing fundamentally – tens of millions of residential subscribers and tens of thousands of enterprise locations require wide-area connectivity, supported by a small number of data centers spanning geographical areas easily exceeding 100,000sq km.

A balanced approach to WAN infrastructure is needed: MPLS will satisfy flexibility requirements; digital transport technologies such as optical transport network (OTN) and data-plane-driven OAM tools will effectively monitor and troubleshoot large-scale networks for cloud applications.

Convergent transport

The service-experience expectations of today's users are becoming more difficult to meet. Residential and business users want ubiquitous and seamless access to both fixed and mobile networks. Network operators are trying to converge their transport infrastructure to optimize service provisioning, network resources and total cost of ownership. Cloud service providers can take advantage of this convergence, promising simpler connectivity.

Service requirements

Cloud-based services place a number of requirements on network operators in terms of user-to-data-center connection and data-center interconnection. The cloud-service connectivity needed from the point of view of the transport operator is shown in **Figure 2**.

Note that traffic in a single data center is not included in the scope of this section, as it is not relevant to WANs. The challenge is to extend the data-center operational environment into the inter-data-center arena in a seamless fashion.

User-to-Data-Center connection

How users access DCs and cloud services is an important consideration when it comes to maintaining user-perceived performance. Users tend to access cloud services through the public internet – either through a fixed or mobile broadband connection. Such access requires IP connectivity from the user to their ISP, and from the ISP to the DC. It is the responsibility of the cloud service provider to ensure that access-related security requirements are met. For example, some providers use Transport Layer Security (TLS) or Advanced Encryption Standard (AES) over TCP/IP sessions.

Some SMEs require traffic separation and more stringent security measures. To provide the most attractive QoS characteristics in such cases, a variety of virtual private interconnection services are available. Choices include today's commonly used Ethernet E-Line and Ethernet LAN services (E-LAN) and IP VPNs such as L3 VPN based on Border Gateway Protocol (BGP) – which usually runs over MPLS.

Security can be enhanced by using IPsec, which encrypts upper-layer communication. For VPN access, the cloud service provider needs to set up a dedicated virtual connection from the user to the nearest DC (or data centers in case of redundant access). The distance covered by such a connection can vary greatly and can include several international links, and both of these factors affect the cost of providing this service.

DC-interconnection requirements

Interconnecting several data centers to enable the exchange of both IP traffic and storage traffic can be achieved in different ways; E-LAN is the most common service used for this purpose today. Each data-center-interconnection may carry all the traffic related to a specific cloud service provider.

Application traffic is usually based on IP, while SAN traffic is typically Fiber Channel (FC) and can be encapsulated over IP, Ethernet or the physical layer directly. Direct encapsulation in the physical layer can be implemented when large amounts of data are replicated over short distances (up to about 10km). In this case, storage protocols can be transported via multiple parallel wavelength division multiplexing (WDM) channels, with OTN framing to increase link reliability.

FIGURE 3 Migration of services and simplification of network layers

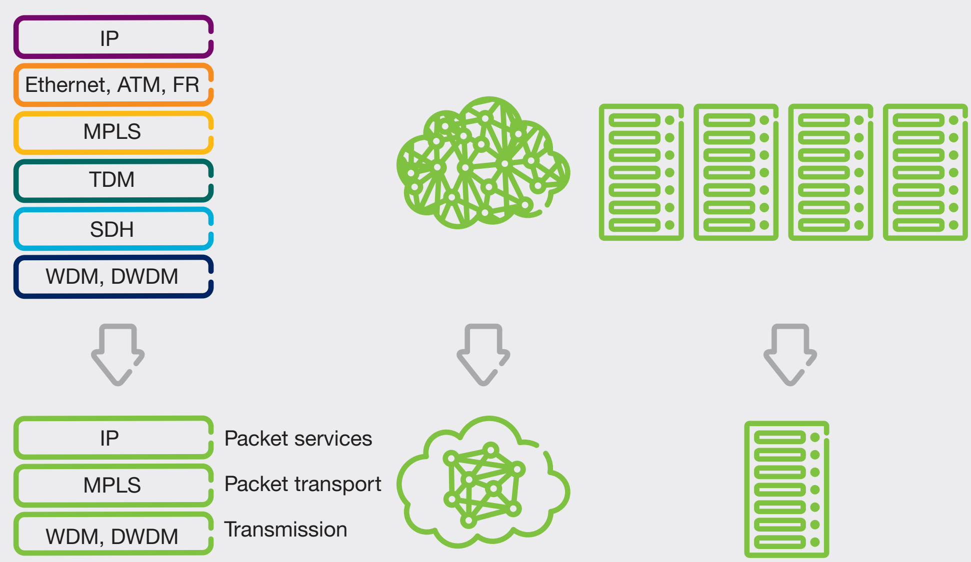

Both the amount of traffic that can enter and exit the VLAN on each interface and the performance targets in terms of loss, delay, and delay variation are specified by E-LAN services. These services are agnostic to the underlying transport layer and as such can be provided in different ways including packet-based, MPLS, OTN or TDM-based.

Different service levels

Cloud service providers may want to provide a variety of service levels to meet the varying needs of a broad subscriber base and sell premium services at the right price.

A transport network should be able to support different service levels for connectivity and availability. In this context, connectivity is defined in terms of bandwidth and QoS – with related targets for loss, delay, and delay variation – and availability is defined in terms of reduced congestion and latency.

Operational requirements

The continual ad-hoc addition of new services to networks over the years has resulted in today's typical networks being owned by an established operator and consisting of multiple, separate networks of equipment – each providing a small number of services (and in some cases, only one). Each of these service-specific networks was surely built on sound economic and business decisions at the time of deployment, and during its lifetime it very probably made a valuable contribution to the operator's revenue. Often, however, such systems remain in place long after becoming redundant – well beyond the point at which their services should have been withdrawn and the network shut down.

The dominant information-payload carried by networks today is IP, which is in turn the carrier for a multitude of inner service payloads and as such represents a global, scalable service convergence layer. The growth in IP traffic, outpacing legacy network traffic, has given network operators a financial incentive to optimize their networks to accommodate IP – the dominant traffic form – and remove the other service-specific network technologies.

Optimizing a network for simplicity, bandwidth and service availability usually requires three layers:

- a connectionless packet switching (CL-PS) service that can be provided by IP or Ethernet;

FIGURE 4 Client-server relations for the type of public IP connection shown in Figure 2

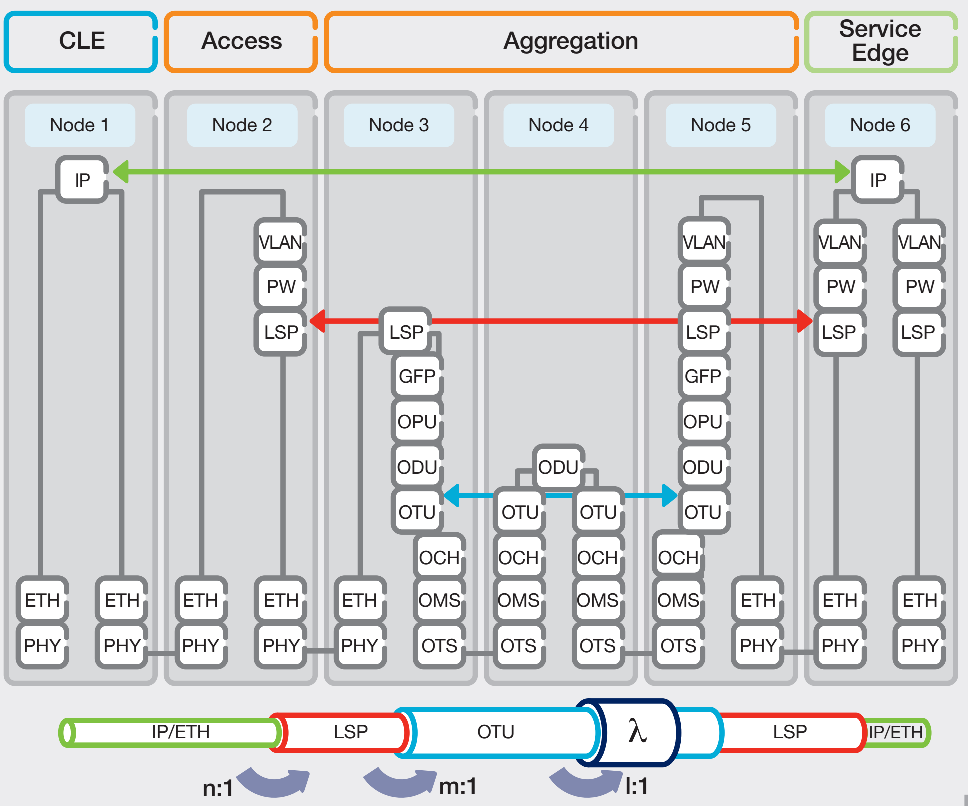

- a connection-oriented packet-switched (CO-PS) service that can be provided by MPLS; and;
- a connection-oriented circuit-switched (CO-CS) service that can be provided by OTN/WDM.

The establishment of a client-server relationship between these layers is shown in **Figure 3** and **Figure 4**.

The IP/Ethernet packet-services layer could be a public packet service (the internet) based on multipoint (connectionless) IP or multiple closed (private) packet services – either multipoint (connectionless) using IP or Ethernet, or point-to-point (connection-oriented) Ethernet. These services are provided by routing and switching service nodes, interconnected by point-to-point transport links. The packet transport layer provides point-to-point connections with a range of throughputs from a few Mbps to tens of Gbps. Service traffic passes into and out of the transport layer over Ethernet interfaces.

In general, transport links are delivered through a hierarchy of point-to-point connections provisioned to satisfy long-term bandwidth estimates of demand between network nodes. The transport layer of a large carrier network should maximize the utilization of the network's resources over several years to support many (client) services while minimizing the total cost of transport. Hence, the attributes of carrier-class transport can be summarized as:

- connections provisioned to satisfy long-term anticipated demand;
- scalable overall network connectivity – deploying a connection hierarchy effectively limits the number of individual connections controlled by each network node;
- transparent and non-intrusive delivery of client information;
- traffic-content agnostic;
- deterministic throughput of client signals between points on the server layer – where performance is defined in terms of bandwidth, latency (average and variation) and the probability of information loss;
- automatic protection of some or all client connections against faults in the server layer (the long-standing gold standard is <50ms); and
- a set of operational tools that minimize operational-staff costs by automating recurring activities.

The optimal internal forwarding plane technology of the transport layer is usually considered to be MPLS, especially given its transport profile – which is operationally very similar to traditional SDH, and hence minimizes the associated opex.

The characteristics of carrier-class transport networks are equally valid when it comes to cloud networking. The inter-data-center traffic of certain applications needs to be promptly redirected to ensure synchronization so as not to disrupt the application in use. Consequently, of the characteristics listed above, the most important ones when it comes to cloud networking are sub-50ms traffic restoration and performance monitoring. The OAM can issue SLA reports directly, and the transport profile of MPLS and OTN can provide performance indicators.

At the transmission level, the transport layer sends information over an optical medium. The developments taking place in transmission technology have led to increases in speed and range, while media costs are dropping. Forwarding individual MPLS packets across a switching node is expensive compared to forwarding the traffic as a continuous bit stream over an optical transmission medium.

In this context, OTN switching – even if considered at a lower hierarchical layer than MPLS – is a circuit-switched (CO-CS) approach that offers similar characteristics to that of – MPLS in terms of OAM capabilities.

Lower layers of the transport hierarchy, closer to the edge of the network, prepare packet traffic for bulk OTN transport through aggregation and grooming so the packets can be forwarded in bulk to a single downstream node. Specifically:

- aggregation ensures that the bit streams passing through the high-speed nodes are reasonably full of packets; and

grooming ensures that all packets are routed to the same destination.

Service impact

Simplifying the network layer structure does not necessarily imply that existing services have to be withdrawn – especially given that many of them will continue to generate significant revenue. Such services can be carried (emulated) over the packet transport network, but operational cost pressures will drive service portfolio rationalization.

Best technology choices

The general and scalability requirements posed by inter-data-center connectivity, transport and aggregation of a large number of end points can be met by two transport technologies: MPLS and OTN.

The characteristics of carrier-class transport apply to both MPLS and OTN, but each of these technologies also has its own distinct characteristics. Neither solution is the best – a variety of factors should be considered before choosing the most appropriate technology in each case.

Through statistical multiplexing and differentiated QoS handling, MPLS supports efficient use of network resources. Because the bandwidth for any given service can span from just a few Mbps to several Gbps, MPLS provides very fine granularity.

Full transparency and higher traffic isolation for client traffic is supported by OTN. Given the hierarchical nature of this technology, it is inherently robust to denial-of-service (DoS) and man-in-the-middle attacks. The granularity of the traffic is defined via the ODU structure, which is either a multiple of 1Gbps (ODU0) or 2.5Gbps (ODU1), thereby covering 10Gbps (ODU2), 40Gbps (ODU3) and 100Gbps (ODU4).

An example of an enterprise location transported to the operator's service edge is shown in Figure 4. In the example, a bandwidth of 100Mbps is required, and the IP traffic is taken over by the service provider with CLE (1). In the access layer, the traffic of n subscribers is aggregated via an MPLS-based node (2) connected via a 10-gigabit Ethernet to a node (3). POTP nodes are deployed in the aggregation part of the network (Nodes 3, 4 and 5), further aggregating the traffic up to l x ODU4 links (each 100Gbps). A 96-channel WDM system is therefore able to transmit nearly 10Tbps via a single fiber.

A smaller number of uncontended connections at the high end of the leased-line market start from 1Gbps capacity with high security requirements. These connections will be provided directly via an Ethernet (or IP)-over-OTN connection. For such services, requirements for uncontended speed and security indicate that OTN is preferable to MPLS.

Conclusion

Seamless access to broadband and applications is raising the stakes for both private and commercial users, regardless of the type of connection method used. Access is driving the convergence of fixed and mobile networks to minimize transport infrastructure costs and maximize resource utilization.

At the same time, the cloud age, fuelled by massive growth in traffic and the consolidation of data centers, is changing the shape of the network-traffic matrix, making it less complex than before and making it easier for operators to predict traffic patterns.

The challenge for the operator is to ensure that investments made to simplify transport architecture are protected in the move to cloud service provider.

As operators make the transition to cloud service provider, the process of network simplification can be successfully carried out. The characteristics of carrier-class transport networks are equally valid when it comes to cloud networking, and apply to both MPLS and OTN. Each of these technologies has their own characteristics, which ultimately determines the best solution in each case.

References

1. Cisco, White paper, 2011, Cisco Global Cloud Index: Forecast and Methodology, 2010-2015, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns1175/Cloud_Index_White_Paper.html

Elisa Bellagamba

works in the Product Area IP & Broadband, as portfolio strategy manager. She is currently involved in the innovation of the IP and Broadband portfolio including strategies for new products, partnerships, standardization and solutions. She holds an M.Sc. in computer science engineering from University of Pisa, Italy.

Paul Lüsse

is currently responsible for business development of converged networks. As part of Ericsson's business intelligence community he has a strong focus on operator business models and strategies. He holds a Dr.Ing. (Ph.D.) in electrical engineering from Technische Universität Braunschweig, Germany.

David Allan

is a distinguished engineer at DUIB Systems and Technology, Business Unit Networks at Ericsson Silicon Valley in San Jose, California. He joined Ericsson in 2009 and his current role focuses on carrier and cloud infrastructure based on MPLS and Ethernet. He co-chairs the End-to-End Architecture committee at the Broadband Forum and was recently honored as a Distinguished Fellow. He holds a B.Eng. from Carleton University, Ottawa, Canada.

Riccardo Martinotti

joined Ericsson in 2006 and is a senior specialist at DUIB PDU Optical and Metro System and Technology, Business Unit Networks. His published work includes patents and papers on packet transport networks. He holds an M.Sc. in electronic engineering from University of Genoa, Italy.