PennWell

LIGHTWAVE®

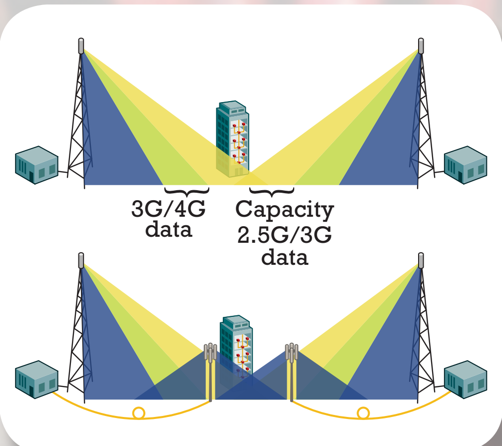

EDITORIAL GUIDE

Fiber support of small cells

As mobile network operators move to small cell architectures, the demands for both traditional mobile backhaul and emerging mobile fronthaul become more complex. Operators are likely to increasingly choose fiber for these applications, as these articles describe.

SPONSORED BY:

JDSU

Now you can monitor service quality to the edge and see the network the way your customers experience it.

JDSU

Only with PacketPortal™ from JDSU.

The result? New and innovative revenue opportunities, faster time-to-market for new services, and mean-time-to-repair measured in minutes—not to mention a dramatically improved customer experience.

PacketPortal: see the network the way your customers experience it.

Watch the video to find out more.

NORTH AMERICA 1 855 275-5378

LATIN AMERICA +1 954 688 5660

ASIA PACIFIC +852 2892 0990

EMEA +49 7121 86 2222

PacketPortal™

Originally published May 1, 2012

Better backhaul with MPLS to the cell site

by **NIR HALACHMI**

LAST YEAR LEFT North American mobile operators breathless. A wide adoption of smartphones, an exponential increase in smart devices (phones, tablets, game consoles, etc.), and rapid growth in smartphone applications compounded to make the mobile economy very dynamic. This mobile data usage and LTE deployment trials drove operators to address bandwidth shortages for mobile backhaul. More important, carriers had to face the realization that there is a true need for network change to support the troika of the new mobile economy tsunami: the convergence of mobile broadband, smart-device growth, and applications adoption.

The first quarter of 2012 was no different, with exciting applications like Instagram added to Android (later acquired by Facebook) or Apple announcing iPad3 with LTE connectivity. The rapid rise of smart devices and applications acceptance has been staggering – and it's a global phenomenon (see Figure 1).

This rising data use has forced many mobile operators to massively invest in the network infrastructure to remain competitive and minimize churn, despite the fact that they can't link this capital investment to increased revenues. Since data traffic is taking the higher share of the operator networks, there is a need to migrate mobile backhaul links to technologies that are more efficient in delivering these services as well as supporting the exponential growth in demand.

Change is in the air

Change is expected in almost all aspects of the network (perhaps beside the OSS/BSS) as a result of mobile services trends. To address the mobile broadband arena, carriers are adapting their networks through a number of mechanisms.

:: RAN technology is moving forward with HSPA+, progressing to LTE and later LTE Advanced to accommodate up to 1 Gbps of downlink bandwidth. It is

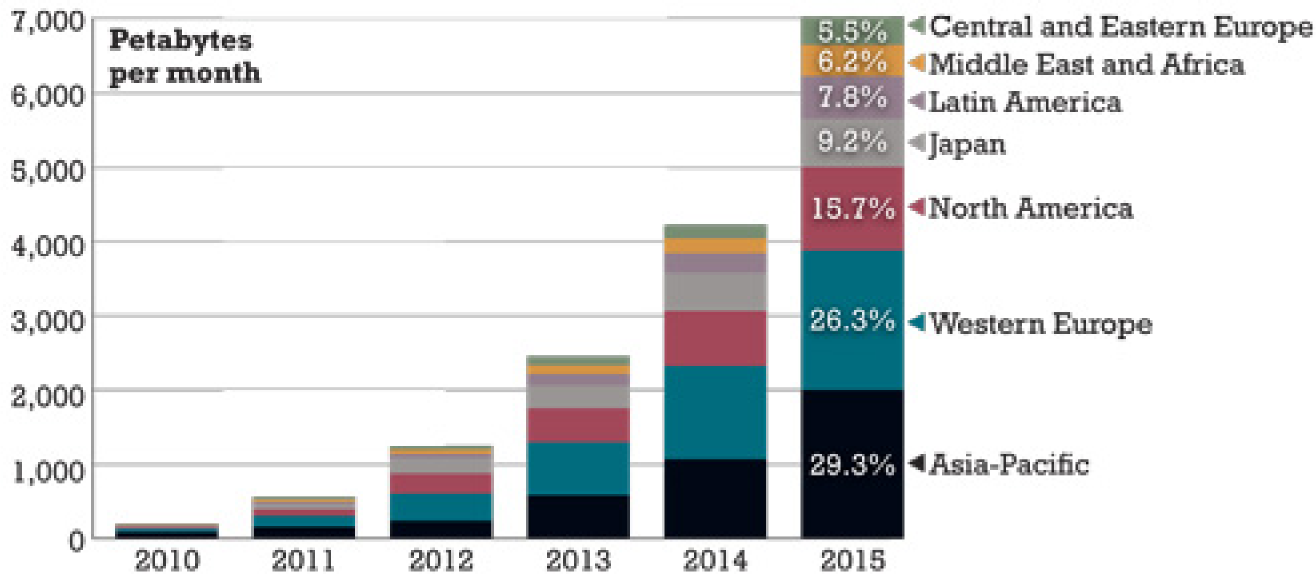

FIGURE 1. *The explosion in mobile bandwidth demand is a worldwide phenomenon, with a projected 92% CAGR from 2010 through 2015. Source: Cisco VNI Mobile, 2011*

ironic that needs are changing so fast that while LTE is hardly commercialized, the next "advanced" generation is already being introduced. The fast pace of changing technologies may cause some operators to skip some technological generations while others will have an even bigger mix of technologies in their network.

- Mobile architecture is changing, with new concepts entering the market, such as "small cell" and "Cloud RAN." The Evolved Packet Core (EPC) concept of flat, all-IP-based network also has caught on as LTE offerings mandate an end-to-end IP service. Such architecture will enable easier introduction and creation of services to support new business models, partnerships, and deployment options.

As a result of these changes, the mobile backhaul space is evolving as well. But unlike the RAN and packet core, which have been well defined by the 3GPP standards body, mobile backhaul traditionally hasn't been defined at all, leaving operators with multiple technology options that offered different values and disadvantages.

Fortunately, two standard bodies have noticed this problem and taken action. The Metro Ethernet Forum (MEF) now offers MEF22.1 and the new CE2.0 initiative, which promotes assured services; operations, administration, and maintenance;

and network-to-network interconnection for Carrier Ethernet in the role of transport technology. Meanwhile, the Broadband Forum has created the TR-221 specifications for MPLS use in mobile backhaul networks.

Mobile backhaul requirements and options

Regardless of who provides mobile backhaul – the mobile operator itself, carrier's carriers, utilities, bandwidth wholesalers, etc. – the requirements remain complex. For example, the typical mobile network combines multiple mobile technology generations like 2G, 3G, and 4G – all of which may coexist in the same cell or in different cells. Therefore, any backhaul network's technology must offer a seamless migration path from TDM- to packet-based transport. Since each mobile macro cell will serve a large number of customers and may serve multiple base stations, operators must implement a transport protocol(s) that can provide high resiliency with sub-50-msec recovery time.

While there are few technological options for mobile backhaul, there is one common denominator: The most viable options are packet-based. Two options stand out: IP/MPLS and Layer 2 Carrier Ethernet. But as the network grows with more cell sites (both large and small), scalability can become a limiting factor in the use of Layer 2 Ethernet. Therefore, mobile backhaul providers should consider the benefits of IP/MPLS to the cell site for mobile backhaul.

Why extend MPLS to the cell site?

MPLS was created to combine the best of two worlds: ATM switching and IP routing. MPLS decouples the data plane from the control plane; it is a connection-oriented technology, so the connection has to be established prior to the data's delivery. The MPLS control plane establishes the connection by signaling through each hop along the path. MPLS has significant traffic engineering capabilities that can be used to provide end-to-end service-level-agreement assurance. The MPLS data plane switches the packets based on MPLS labels that are carried inside a 32-bit MPLS header.

IP/MPLS is the de facto standard in the core today. While most edge and access networks are Layer 2, rapid changes due to the dynamic nature of mobile connectivity have forced operators to consider extending MPLS to the access and aggregation layers for easier control, resiliency, redundancy, and scalability.

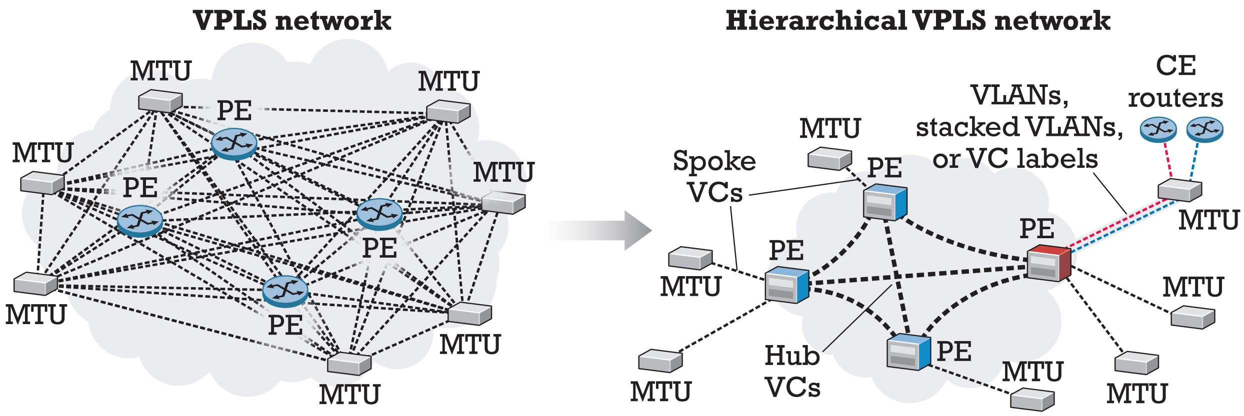

FIGURE 2. *Virtual private line services (VPLS) network vs. hierarchical VPLS network.*

MPLS at the edge of the network for mobile backhaul provides multiple advantages.

Maximizing scalability. MPLS is highly scalable. The 20-bit label enables more than one million label-switched paths (LSPs) per node. With each node changing the label and reusing labels, practically infinite LSPs can be supported. By using virtual private wire/line services (VPWS/VPLS), such a network can support thousands of customers and each customer can have a different logical topology. Hierarchical VPLS (H-VPLS) technology further increases scalability by segmenting the network into several fully meshed partitions, each concentrating into the VPLS hub (see Figure 2).

In contrast, Ethernet's 12-bit VLAN tags support only 4,000 VLANs per switch. VLAN stacking (Q-in-Q) enables 4,000 customer VLANs to be carried in 4,000 provider VLANs. Since each customer is likely to use multiple VLAN IDs, the number of customers that can be supported is quite limited.

Dynamic path creation. As mentioned, MPLS is a connection-oriented technology where control plane protocols (namely LDP and RSVP variants) handle path creation, starting from the source label edge router (LER), traversing the label switch routers all the way through the destination LER. These protocols base their path creation on the dynamic routing information exchanged between peers. The dynamic nature of MPLS minimizes service creation time while increasing network scalability since most of the work is done by dynamic protocols.

When path creation can be accomplished by configuring only the end devices, manageability of the network becomes even easier.

Traffic engineering capabilities. MPLS provides strong traffic engineering capabilities embedded in the MPLS control protocols. The control plane can check and reserve bandwidth when establishing a path only after assuring the required committed information rate is available throughout the proposed path. The control plane also can mandate that traffic pass through specific nodes using strict rules or provide other protocols full dynamic control to ease the operational management of the network as changes occur.

Support of TDM and other legacy services. Because it is essentially a tunneling protocol, MPLS supports the transport of any service available today – TDM, Ethernet, Frame Relay, ATM, IP, etc. These services and protocols are encapsulated with MPLS labels, then switched to the destination, which may be another customer site or a service within the provider network. Taking into account the very strict timing requirements of TDM-based mobile technologies, traffic engineering can be used to assure the proper delivery of these services concurrently, combined with other less sensitive data services. MPLS with traffic engineering thus can guarantee dedicated bandwidth for the TDM-based mobile elements still in service to minimize the dreaded "iPhone dropped call" problems when bandwidth-hungry applications usurp the link capacity.

Designing the network for resiliency. As each cell site supports a large amount of end users, downtime translates into large revenue loss. Redundancy therefore must be part of the network design from the cell site, across the access and aggregation networks, and through the core. Primary and backup paths using VPLS technologies enable two levels of protection – at the tunnel level and service level – while restricting the use of backup bandwidth for failure scenarios only. If a failure occurs, MPLS Fast ReRoute and Bidirectional Forwarding Detection provide sub-50-msec switchover, using local repair techniques and signaling across the services to identify and initiate rerouting.

Seamless network improves service creation time

With the increasing deployment of small-cell technologies, the number of cells will grow exponentially. This trend has caused scalability issues and service

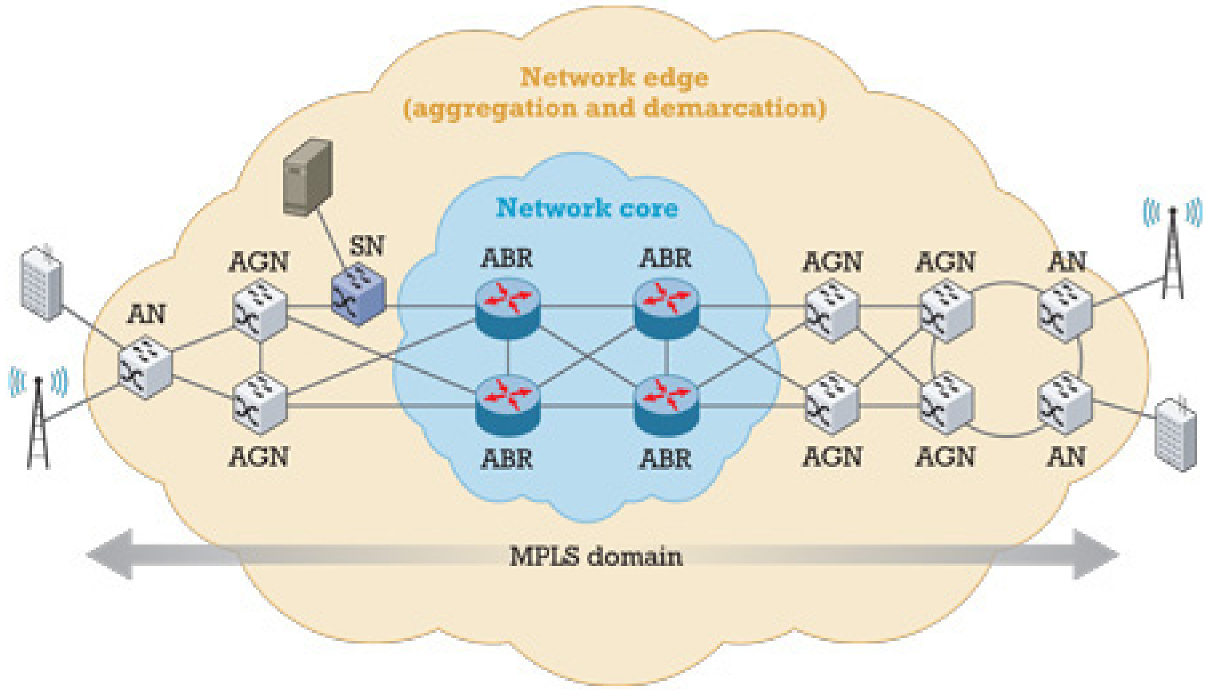

FIGURE 3. *A "seamless" MPLS architecture decouples the service layer from the transport layer. This decoupling increases the flexibility to define and introduce new services by enabling service nodes to be placed at optimal locations in the network.*

creation challenges. The best way to improve service creation/delivery time is to have a network that operates independent of the services yet can support any service-deployment scenario. The network should not have transport boundaries that limit access to services.

A "seamless" MPLS architecture inherently has no boundaries and hence decouples the service layer from the transport layer. This decoupling increases the flexibility to define and introduce new services by enabling service nodes to be placed at optimal locations in the network rather than at the "boundary nodes."

Although both the service and transport layers use the same MPLS packet formats, the difference is in the use of the MPLS control plane. Using the MPLS control plane end-to-end enables a management system to select the endpoints of the service then trigger signaling to set up the services across the network between the endpoints (see Figure 3).

Winning move

Moving MPLS to the cell site or aggregation point integrates access and aggregation networks with the core onto a single MPLS-managed network to create significant operational advantages. This network architecture is decoupled from the service architecture and incorporates intelligent switching closer to the cell site to optimize network resources and improve the network's overall performance in an increasingly dynamic mobile-focused world. It also enables true customizable services, because quality of service parameters can be incorporated end-to-end. MPLS is manageable and scalable and can support any legacy services required to enable smooth migration to a pure IP network.

Factoring in the cost savings for bandwidth efficiency and network resiliency in an increasingly dynamic, bandwidth-hungry environment, MPLS enables additional revenues from customized services and cost savings through improved service creation. Combined, all these factors make a strong business case for driving MPLS to the cell site.

NIR HALACHMI is product manager, mobile backhaul products, at Telco Systems. He is responsible for the design and development of Telco Systems's mobile backhaul offerings focusing on both cellular and wireless technology as well as QoS, data security, and communications.

Originally published July 2, 2012

Optical transport for cell backhaul: SONET, Carrier Ethernet, and cell-site fiber challenges

by **JON ANDERSON**

WITH THE ADVENT of LTE/4G technology, our communication industry frequently hears and sees that copper T1 service to cell tower sites is quickly becoming inadequate. Wireless carriers continue to increase the number of new cell sites and to upgrade existing 3G sites to 4G/LTE. These carriers are more frequently requesting a minimum of 50-Mbps Ethernet initial service to these new tower sites and Ethernet upgrades to existing T1 services — often with follow-on commitments to add 150-Mbps Ethernet service with just a three- or four-month notice.

This surge in cell backhaul bandwidth demand threatens to overwhelm facilities sooner rather than later. For example, historically the typical cell site might have been served with eight T1s. More recently, SONET ring networks were extended to include cell sites to facilitate dropping off a mixture of T1, DS3, and Ethernet type circuits. Now, not only is the number of new cell tower sites growing, but these sites are being designed to accommodate as many as six wireless carriers each. The result of this explosion in bandwidth is the need for fiber-fed connectivity based on Carrier Ethernet.

The emergence of Carrier Ethernet transport

The SONET unidirectional path-switched ring (UPSR) architecture continues to be very popular for mobile backhaul given its huge embedded base, variety of interfaces, and scalable bandwidths across OC-3/12/48/192 backbone rates. In addition, SONET offers proven reliability with less than 50-ms ring switching time.

Yet there is a new technology trend in network architectures to support cell-site tower locations – Carrier Ethernet transport. Typically these active platforms support Gigabit Ethernet to 10-Gigabit Ethernet backbone optics and are very scalable.

The drivers behind this new service provider model are quite logical: the ubiquity of the Ethernet interface (whether copper RJ-45 or optical 10/100/1000 Mbps), the advancement of ITU-T G.8031/2 standards for ring protection switching (also sub-50 ms), and five-9s of reliability. These technological advancements will enable Carrier Ethernet to become the predominant technology for serving the ever-growing demand for cell backhaul. In addition, the sheer volume of Ethernet chip sets across the application landscape has facilitated lower silicon component costs, greater availability, and reliability improvements.

At the heart of the acceptance of this network topology is Ethernet Protection Ring Switching (EPRS). EPRS was defined by the ITU and Metropolitan Ethernet Forum (MEF), is widely accepted, and continues to evolve into more complex network architectures with the recent announcement of Carrier Ethernet 2.0 by the MEF. ERPS began at ITU-T as part of the G.8032 Recommendation to provide sub-50-ms protection and recovery switching for Ethernet traffic in a ring topology while ensuring there are no loops formed at the Ethernet layer. G.8032v1 supported a single-ring topology and G.8032v2 supports multiple ring/ladder topologies.

Additional Carrier Ethernet service definitions are expected as new standards-based features are created, implemented in silicon, and deployed in active systems.

All of this bodes well for the continued use of Carrier Ethernet to meet 4G/LTE requirements for bandwidth increases and to expedite the push for fiber ring deployments to cell sites. In some sense cell backhaul is fast becoming the FTTx of the "Mobile Device Generation."

Don't neglect the Physical Layer

Yet with all of the technological advances on the active platform side of equation, service providers often neglect the challenges related to the Physical Layer until the installation and service due dates are almost upon them. With up to six cell carriers per new cell tower site and thousands of new tower sites popping up across the nation, service providers are understandably seeking deployment improvements, efficiencies, and best methods for delivering fiber handoffs to multiple carriers.

The cell-site location often represents a harsh environment in which to land these small- to medium-count fiber cables and involves special requirements for separate (non-shared) fiber facilities, separate demarcation points, lockable access fiber cabinets, outdoor NEMA 4 rated fiber cabinets, as well as hut-based fiber cross-connect panel/frame equipment. Existing fiber cable routes near cell sites may be small count, thereby limiting bandwidth without expensive fiber cable overbuilds. Wireless carriers may require multimode fiber cross-connect panel fiber connector/terminations to accommodate the lower-cost Gigabit Ethernet SFPs in their active platforms.

The fiber distribution portfolio used in these environments must address the basic challenges inherent in the central office/hub/outside plant/cell-site environments to gracefully improve fiber deliverability and fiber management for cell backhaul. The point of every component within a fiber management system — from the cladding on the fiber, cable jacketing, optical component packaging, and the route paths within them — is to protect and reduce the risk of fiber damage. Period.

Fiber distribution and management equipment that does not accomplish this in an easy and intuitive way is over-thought and costs you money. Fiber management should be approached with three simple goals in mind:

1. The first and most important objective is to minimize your fiber risk in the cable plant.
2. The second goal is to attempt to eliminate deployment and maintenance headaches.
3. Third is to reduce the cost of broadband deployment via careful attention to not only lower capital equipment prices, but also lower operational costs.

Perhaps the biggest key to achieving these three objectives is to reduce risk by eliminating as much interaction with fiber jumpers and the fiber tail as possible. Fiber management equipment that integrates fiber distribution and slack storage within a small footprint enables the service provider to quickly and conveniently deploy the fiber as well as access it at a later time if necessary.

In addition, the costs of delivering the fiber to the site should not be overlooked. A simple in-ground drop cable would be convenient, but often is not possible. New developments in ruggedized microduct that enable service providers to push the fiber through existing conduit — even environments previously considered exhausted — are being brought to market and should be investigated.

Bringing together the whole package

While consumer demands for bandwidth will drive the use of Carrier Ethernet, careful system engineering and plant design will enable the service provider to create a backhaul network that economically meets the needs of the wireless carrier. Careful consideration of not only the active electronics but the physical layer as well as will ensure Carrier Ethernet is an economical business driver for the entire industry.

JON ANDERSON is an applications engineer at Clearfield Inc. He joined Clearfield's Market Segment Application Engineering team with nearly 40 years of experience in the telecommunications and data networking industry with companies such as Fujitsu, Teltrend, Lynch Communications, and Alcatel. He has direct engineering and application experience with a variety of optical platforms and technologies such as SONET (TDM), DWDM, Ethernet, and FTTx. Jon studied at Seattle Pacific University and Clover Park Technical College where he earned his degree in telecommunications.

Originally published January 1, 2013

Demarcation of Carrier Ethernet 2.0 services

Carrier Ethernet 2.0 promises to make Ethernet-service deployment easier than ever. Proper demarcation is an important element in fulfilling this promise.

by **TY ESTES**

THE METRO ETHERNET Forum (MEF) announced Carrier Ethernet 2.0 (CE 2.0) as the next generation in the evolution of Ethernet services. This next generation is defined as "networks and services that enable multiple Classes of Service (multi-CoS) and manageability over interconnected provider networks." These three service attributes – multi-CoS, manageability, and interconnect – reflect major new capabilities in E-Line and E-LAN services as well as the addition of new E-Tree and E-Access services.

The first generation of Carrier Ethernet provided the standardization that enabled the growth of Ethernet to a ubiquitous service available in most metropolitan markets worldwide (see Table). CE 2.0 addresses the issues that have delayed global deployments of Ethernet services across multiple networks with efficient use of network resources and cohesive management platforms. MEF CE 2.0 standards and certified-compliant equipment are the next step in the evolution of Carrier Ethernet and will open new opportunities in cloud services, mobile backhaul, and multipoint services.

The MEF recently announced the first equipment vendors to achieve CE 2.0-certified compliance. CE 2.0-certified compliant network interface devices (NIDs) are a critical element in the demarcation and delivery of these next generation services as they traverse from the enterprise network across one or more service-provider networks on a regional, national, or global scale.

TABLE: Carrier Ethernet 1.0 and 2.0 standards comparison from the Metro Ethernet Forum

Carrier Ethernet generation	CE 1.0	CE 2.0
Characteristics	**Standardized**	**Multi-CoS, managed, interconnected**
Services	E-line, E-LAN	E-Line, E-LAN, E-Tree, E-Access
Services Specs and IAs	MEF 6	MEF 6.1 and 6.1.1, 33 MEF 22.1
SUPPORTING WORK		
Attributes IAs	MEF 10	MEF 10.2 and 10.2.1, 26.1 MEF 13, 20, 23.1
Management	MEF 7, 15	MEF 7.1, 16, 17, 30, 31
ENABLED APPLICATIONS		
MBH	2G/3G migration	4G MBH migration, MBH optimization
Business services	Metro, regional	Local, regional, national, global Application- and distance-oriented, private cloud
Wholesale	–	Buy/sell access services

IAs = Implementation agreements

Multi-CoS

Ethernet services can deliver several data flows, including voice, video, and business-critical data. Multi-CoS enables services to be differentiated, prioritized, and assigned unique bandwidth profiles for application delivery in cloud services, mobile backhaul, and business services.

For the service provider, multi-CoS optimizes bandwidth use and improves the quality of service (QoS) by segregating voice, video, and data flows, especially for services like mobile backhaul with high-priority and bursty traffic. Service providers can avoid over-subscription of network resources (over-building networks with full-throughput ports) and deploy services to align with bandwidth needs. Multi-CoS can also conserve provider virtual LANs (VLANs) by classifying traffic within Ethernet virtual circuits (EVCs) instead of using multiple EVCs to deliver applications with service multiplexing.

For the enterprise, multi-CoS enables lower cost per assured bit with granular bandwidth profiling and service-level-agreement (SLA) metrics to verify guaranteed performance. Enterprises can achieve predictable application delivery with multi-CoS standards that account for application types and distances.

Carrier Ethernet has many types of data flows:

:: **Per user-to-network interface (UNI) or port.** All data received by the UNI from the subscriber network is accepted, bundled, and delivered as a single service. The data is not segregated, and it shares the same bandwidth profile and performance parameters for SLA assurance.

:: **Per EVC.** An EVC is a logical connection between at least two UNIs. Data received by the UNI from the subscriber is mapped to the EVC based on their CE-VLAN ID. A UNI can also support multiple EVCs with service multiplexing, and multiple EVCs are segregated by unique service-provider VLAN IDs. Any data received by the UNI that's not mapped to EVCs is filtered by the UNI. Each EVC of a UNI can have unique bandwidth profile and performance parameters.

:: **Per CoS.** Each EVC contains multiple CoS flows for different data types or applications, and each CoS for an EVC can have unique bandwidth profile and performance parameters.

CoS can be identified by priority code point (Layer 2 VLAN priority), by L2CP (Layer 2 Control Protocol identification based on destination MAC address and Ethernet protocol), or by differentiated services code point (Layer 3 priority). In Figure 1, each CoS flow is differentiated with a priority and bandwidth profile for a private cloud E-Line service: CoS 6 is for synchronization data, which has a lower data rate, or committed information rate (CIR), but the highest CoS priority because it requires the lowest possible latency; CoS 5 is for voice data, which has a lower data rate, or CIR, but a higher CoS priority because it requires low latency for voice quality; CoS 3 is for video data, which requires a higher CIR for streaming video and a higher CoS priority because latency and dropped frames can affect video quality; and CoS 1 is for web and background email data and has a high CIR and low priority.

MEF 23.1 introduces the industry's first standardized multi-CoS performance objectives (MPOs) with new metrics for specific applications, including mobile backhaul, VoIP, videoconferencing, and financial trading. In addition to the ITU Y.1731 performance metrics of frame delay (latency), inter-frame delay variation (jitter), and frame loss ratio, MEF 23.1 adds mean frame delay and frame loss range. These MPOs are defined to enable precise SLA metrics for application-specific delivery.

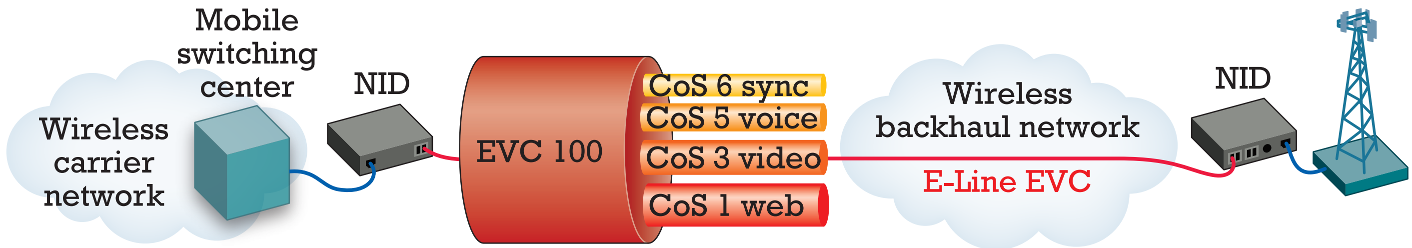

FIGURE 1. *Multi-CoS in a mobile backhaul service.*

The MEF further defines MPOs with performance tiers that allow service providers and enterprises to predict SLA performance metrics based on network distances. These four distance-related performance tiers adjust MEF 23.1 metrics for metro (205-km), regional (1,200-km), continental (7,000-km), and global (27,500-km) point-to-point services. MEF CE 2.0-certified compliant NIDs enable multi-CoS at the UNI and the external network-to-network interface (ENNI) using traffic shaping and policing with bandwidth profiles and deliver application-specific MEF 23.1 MPOs.

NIDs can also provide ITU-T Y.1564 service testing and Y.1731 performance monitoring for each CoS, MPO, or distance-related performance tier. Y.1564 service testing ensures each CoS meets SLA requirements and can run multiple CoS test flows simultaneously. For example, when turning up multiple services with multiple flows, Y.1564 can ensure that data traffic does not affect voice traffic. Once the service is activated, NIDs provide Y.1731 real time performance monitoring for SLA assurance of each CoS.

Manageability

The CE 2.0 manageability attribute refers to service OAM management functions throughout the lifecycle of an Ethernet service across one or more networks. The lifecycle includes service provisioning and verification, performance monitoring for SLA assurance, and fault monitoring and isolation. These capabilities are achieved via ITU-T Y.1564 service activation testing, ITU-T Y.1731 performance monitoring, and IEEE 802.1ag connectivity fault management. CE 2.0 extends traffic management to include both ingress and egress granular bandwidth profiles that can be applied per UNI, per EVC, and per CoS.

Management of global CE 2.0 services requires a standardized management interface for management systems – including element management systems (EMS), network management systems (NMS), and operational support systems –

to easily interoperate with a variety of different network devices. To address this requirement, the MEF has developed the Service OAM Management Information Bases (SOAM MIBs). A MIB is a network communication protocol used for SNMP management of the entities in a network. The MEF has ratified specifications and implementation agreements for the management of IEEE 802.1ag functionality with MEF 30 and 31 Service OAM Fault Management MIBs (SOAM FM MIBs).

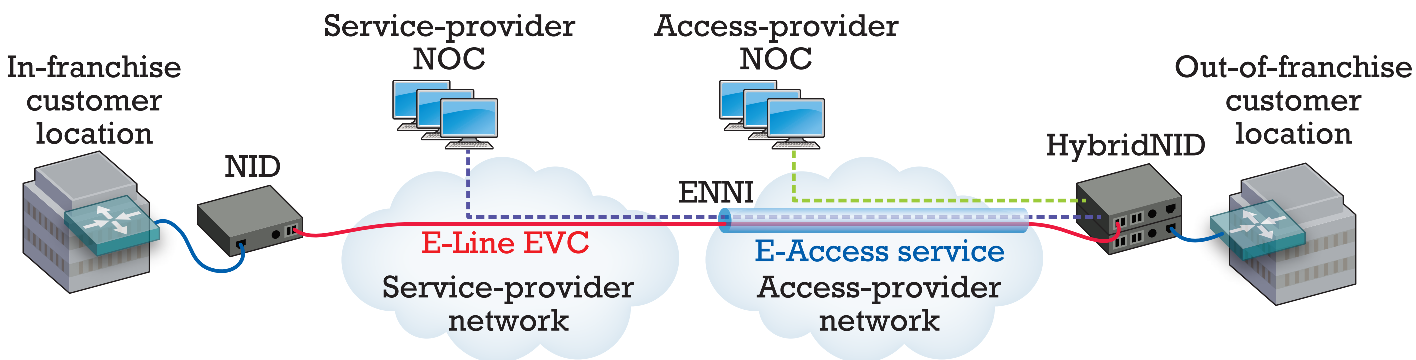

FIGURE 2. *E-Access service delivery with a "HybridNID."*

The MEF is also publishing service OAM standards for performance monitoring (SOAM PM MIB) and service management (service MIB project). It's important to note that the MEF clearly defines the attributes of CE 2.0 within the context of a linear timeline, similar to a technology release, which means that CE 2.0 attributes were frozen at the time of release in early 2012. Although there are other developments in MEF management standards directly related to CE 2.0, they're not officially included in the definition. Thinking of CE 2.0 as an incremental process enables the MEF and other organizations to develop standards that cover SOAM PM MIBs, service management MIBs, and protocols like YANG and NETCONF.

There's also a joint effort among the MEF, the Broadband Forum, and the TeleManagement Forum on Ethernet management to align terminology and requirements. The goal is to produce a plan to identify standards gaps and determine resources to complete a unified Carrier Ethernet management standard.

Standardized management interfaces and protocols are essential to industry interoperability. Service providers spend millions of dollars integrating equipment into EMS/NMS software. Industry support for management standards will reduce the costs associated with deploying service portals and EMS/NMS and simplify management and SLA assurance of multipoint E-LAN and E-Tree services.

CE 2.0-compliant NIDs with comprehensive support of the Y.1564, Y.1731, and 802.1ag SOAM standards streamline service activation, enable SLA reporting, and provide the tools for rapid fault isolation. These capabilities improve customer satisfaction throughout the entire lifecycle of CE 2.0 services. NIDs that support the MEF SOAM MIB specifications and implementation agreements reduce the costs of EMS integration and the time-consuming complexities of multivendor deployments.

Interconnect

Carrier Ethernet has truly become a global service, and delivering services across multiple regions and operators requires wholesale arrangements between service providers. Interconnecting to out-of-franchise enterprise customer locations and cell towers is a business necessity. The MEF standardized the service-provider interconnect with the MEF 26 ENNI specification. The MEF further defined the interconnect with MEF 26.1 to include specifications for multipoint services across multiple networks.

The challenge is that the wide variety of Carrier Ethernet services makes interconnecting these services a complex and time-consuming process. Aligning the different SLAs and CoS parameters, along with QoS and bandwidth between service providers at the ENNI, can take weeks of negotiation to create a single end-to-end service.

The MEF has addressed this challenge within CE 2.0 via two new standards: MEF 23.1 CoS includes guidance for aligning CoS at the ENNI and MEF 33 E-Access simplifies the service interconnection. MEF 33 E-Access streamlines the deployment of wholesale services with a single CoS for the ENNI-to-ENNI or ENNI-to-UNI connection. The end-to-end EVC can provide multi-CoS, so the bandwidth profile and SLA metrics of the E-Access CoS will meet or exceed the highest CoS of the EVC.

E-Access is a service composed of a data tunnel that enables transport of EVCs. Most NIDs are capable of delivering E-Access services, but multiple NIDs currently are deployed at the out-of-franchise location for demarcation of the E-Line and

E-Access services. A “HybridNID” enables wholesale E-Access services with virtual NID functionality and allows the service provider and access provider to manage and provision the service to one demarcation device (see Figure 2). That enables the service provider to “trust but verify” the E-Access SLA without having to deploy another NID at the out-of-franchise customer location.

An MEF survey shows that 93% of service providers will buy more wholesale services and 90% will sell more wholesale services. The wholesale Ethernet services market is booming, and MEF 33 E-Access will further accelerate this growth. E-Access will open new opportunities with mobile backhaul services to cell-tower sites, distributed cloud services, and global enterprise services.

Although not officially part of MEF CE 2.0, there’s a logical correlation between interconnect and equipment interoperability. Interoperability events such as those held by EANTC and CableLabs provide a venue for vendors to interconnect equipment in a multi-operator network environment. These interoperability events enable dozens of vendors to collaborate in a complex network environment and provide tested and verified interoperability.

CE 2.0 is more than the latest technical specifications; it brings standardized functionality that reduces operational costs and simplifies service deployments. Multi-CoS optimizes bandwidth and enables services with predictable QoS, standardized management simplifies EMS integration for global networks, and E-Access streamlines wholesale service deployments. CE 2.0 opens new business opportunities that will enable the projected 100% growth of global Ethernet services by 2016.

TY ESTES is marketing communications director at Omnitron Systems.

Company Description:

JDSU serves markets that are at the heart of what's important, making virtually every network in the world faster and more reliable. JDSU leads the fastest-growing segments of the optical-networking market: tunable XFPs, transport blades, and reconfigurable optical add/drop multiplexers (ROADMs). From development and planning, through deployment and turn-up testing, to optimization and assurance, JDSU test instruments, software, and expertise ensure that all networks—xDSL, fiber, cable, and wireless—are always working at their best.

JDSU also applies its sophisticated optical technologies to a range of markets that use light to achieve a unique purpose, from fighting counterfeiting and protecting consumer brands to enabling satellites.

JDSU diverse markets include:

Broadband, Mobile, and Enterprise Networks — JDSU components and test solutions are part of every major network in the world

- Components, Modules, and Systems
- Network and Service Enablement
- Network Intelligence and Analytics

Anti-Counterfeiting — protecting over 100 of the world's top currencies

- Security Pigments and Banknote Threads

Aerospace and Defense — part of every manned NASA spacecraft vehicle since 1961 and ensuring the security and reliability of mission-critical communications

- Custom Optics
- Network and Service Enablement

Gesture Recognition — driving new heights of gaming and natural interfaces

- Optical Thin-Film Coatings and Components
- Diode Lasers

Lasers — changing the face of precision manufacturing

- Solid-State, Diode, Gas and Fiber Lasers

LINKS:

- Learn more about how JDSU enables small-cell deployments
- Accelerate backhaul deployments
- Small cell assurance solution
- Watch video – small cell assurance solution
- Small Cells and the Evolution of Backhaul Assurance