JDSU

White Paper

Next-Generation Cellular Backhaul

Today's most critical real-world issues and the methodologies used to combat them

Introduction—Wireless Backhaul the New Critical Service

Wireless standards and spectrum auctions enable incredible bandwidth availability. The advances in modulation schemes (multiple input/multiple output [MIMO], orthogonal frequency division multiplexing [OFDM]) and system spectral efficiencies combined with the significant amount of spectrum now dedicated to the 3.5G/4G services leads to an aggregate radio bandwidth available to/from cell sites between 50 and 300 Mbps. Delivering this sort of bandwidth with physical T1/E1s is not only impractical, the combined bandwidth will likely exceed the theoretical bandwidth available on all the copper pairs delivered to the tower.

The evolution of wireless backhaul is very similar to what happened in the mid-to-late 90s with wavelength division multiplexing (WDM) and Erbium doped fiber amplifiers (EDFAs); the invention of which together caused exponential improvements in long-distance optical transport. Providers were able to carry several terabits on a fiber and could go thousands of kilometers between expensive electrical regeneration points. While this increase in capacity required tremendous investment in the core, it eventually focused the attention and the money toward investments in the other more limiting elements in the value chain: switching, metro, and local loop technologies. Similarly for wireless carriers, the focus is shifting toward backhaul rather than radio access bandwidth, with backhaul becoming the most crucial resource that wireless operators must manage.

Today most wireless backhaul traverses copper-fed T1s/E1s at a cost of approximately $150-$400 per month. Wireless operators currently must provide primarily voice services with defined revenues requiring small amounts of bandwidth—$0.10/minute using approximately 13 kbps of bandwidth per call—revenue for a wireless operator tracks directly to bandwidth usage and ultimately backhaul charges.

Service	Cost to Customer	Rate	Monthly Usage	Monthly Revue	Revenue Tracks Usage
Voice	$0.10/min	13 kbps	~50 Mbyte (500 min)	$50	Yes
Data	$50/month unlimited	100k+	~2 Gbytes (40x voice)	$50	No

Table 1. Comparison of Voice and Data Usage and Revenue

As data services become more widely adopted, utilization of bandwidth for these services will increase while revenue will not. This decoupling of backhaul bandwidth usage and customer revenue yields a fundamental business problem for operators—reducing the cost of backhaul, which is already their single largest operational cost. Furthermore, adoption of data and video streaming services will likely drive significant increases in bandwidth needs while disproportionally contributing to revenue. With 3G service usage already significant, and 4G services on the horizon, wireless carriers have no choice but to seek more cost-effective backhaul technologies that also provide flexibility in reaching 100s of Mbps thus delivering on the promise of 4G while avoiding bankruptcy.

This paper assesses technology options available for wireless backhaul and examines the pros and cons of the various backhaul architectures, the Ethernet services options, and the existing operation, administration, and management (OAM) strategies and standards. In conclusion, the paper offers a practical strategy for implementing a wireless backhaul service assurance strategy based on the assessment provided.

Technology Overview

Cell Site Backhaul Network Service

Backhaul Network Service[1] (BNS) provides the high-speed point-to-point connectivity between the Base Station Controller (BSC) and an aggregation point or hub operated by the wireless provider. The BNS (Ethernet, Time Division Multiplexing [TDM], Internet Protocol [IP], and asynchronous transfer mode [ATM]) can be provided over a multitude of technologies deployed over microwave, fiber, or copper.

All backhaul networks include a mobile switching center (MSC)/Data Center Hub that peers with the core network or the Public Switched Telephone Network (PSTN). Some operators also deploy a second layer of microwave hubs that combines multiple tower uplink connections into a single link before handing off to a fiber- or copper-based backhaul link. Figure 1 shows a typical backhaul scenario. Note that a wireless operator could have multiple third-party backhaul providers or own the backhaul network outright.

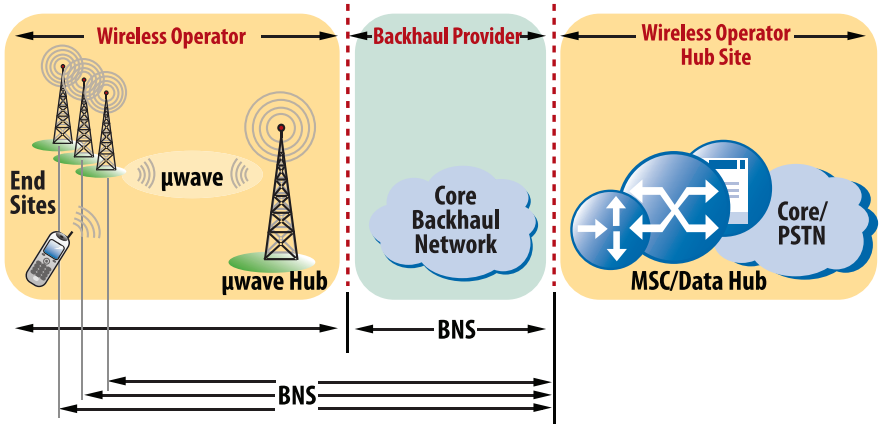

Figure 1 Typical BNS

Wireless operators are typically very proficient at managing tower services and other RF-type facilities. They manage, or directly oversee the operation of, the microwave portion of any backhaul service. Unless the wireless operator is part of an integrated landline carrier, they will typically go to a third-party backhaul provider for the connections between the microwave hubs or single-tier end sites and the MSC/Data Hub.

[1] Unmanaged services such as Wi-Fi hot-spot backhaul, or best-effort mesh networking technologies, are technically backhaul but not in the scope of this analysis. The different economic paradigm of unmanaged services drive decisions for architecture and service assurance based on a consumer model of network deployment rather than a more traditional business services model. Therefore this analysis considers only managed BNS used to deploy premium mobile voice, data, and/or video services.

Fiber to the Tower—The New Monopoly

Regardless of the service carried over the physical medium, wireless operators must make fundamental choices regarding the backhaul architecture. First, it is important to note that only microwave and fiber have a practical path to attain hundreds of Mbps. While many initial deployments may have tens of megabits aggregate bandwidth; they are just within the reach of copper-based services. Anything deployed over copper is a temporary solution viable for only a few years at sites with spectrums in the tens of megahertz. Copper-based offerings maintain solid footing for providing business services that allow telecom providers to offer affordable Ethernet services to the 85 percent of office buildings not yet served by fiber. However, copper-based services do not carry the same flexibility, quality of service (QoS), or capacity as fiber. Hence, its utilization should only be in areas of low projected usage or during periods of initial ramp-up, until fiber- and microwave-based solutions become available.

In North America, where incumbent local exchange carrier (ILEC) plesiochronous digital hierarchy (PDH) backhaul charges are more affordable (now as low as $200 for 1.5 Mbps/month), the division between fiber/copper/microwave backhaul tends to favor a star copper fed T1 architecture. In EMEA, where backhaul charges are considerably more expensive (~$800 for 2 Mbps), more of the backhaul is aggregated with microwave first and then brought over a backhaul core. Figure 2 shows a breakdown of backhaul by transport technology for Europe, Middle East, and Africa (EMEA) and North America.

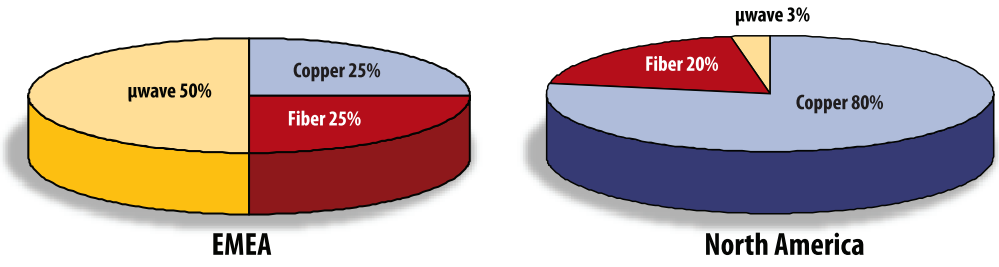

Figure 2 Backhaul Transport Technology Breakdown

As with communication services to office buildings, the first operator that delivers fiber to a tower will likely own the advanced services to that tower. Service pricing and flexibility that fiber operators obtain providing Ethernet services along with the ability to deliver traditional TDM services will likely result in a monopoly. That is, it will not make economic sense for other operators to invest in delivering fiber or microwave services to that tower thus becoming a second network operator. Rather, the fixed line operator with the closest fiber runs and the lowest engineering and digging expense will deliver the fiber and own the tower.

Backhaul Transport Architectures Analysis

Alternative backhaul providers must consider implementing a metro network architecture that allows for raising their return on investment (ROI) by delivering other business services as well as wireless backhaul over the same infrastructure. It is likely that the majority of the services new providers will sell to business and wireless operators will be Layer 1 circuit emulation (T1s), Layer 2 (Ethernet) or Layer 3 (IP virtual private network [VPN]). However, these services can be carried over a multitude of architectures that each offer advantages for providers. The choices for metro network architectures fall into these categories:

– Layer 1: Synchronous digital hierarchy/synchronous optical network (SDH/SONET)[2] Transport—G.707

– Layer 2: Ethernet with stacked virtual local area networks (VLANs) (Q-in-Q)—802.1ah

– Layer 2: Provider Backbone Bridges (PBB or 'MAC in MAC') Ethernet—802.1ad

– Layer 3: IP/Multiprotocol Label Switching (MPLS) over Ethernet using MPLS Pseudowires and Transport MPLS (T-MPLS)

– Layer 2: Provider Backbone Bridges with Traffic Engineering (PBB-TE) Ethernet—802.1Qay

– Layer 1/Layer 2 (L1/L2) Hybrid: Combinations of SDH/SONET Transport with a Layer 2 Core Capability

While it may seem incongruous to compare transport technologies from various layers, the underlying protection and multiplexing protocols used by each of the categories fundamentally affects solution costs and efficiency in carrying assorted traffic types. This paper also addresses certain hybrid architectures that attempt to mix some of the Layer 2 technologies with those in Layer 1 from SONET to garner the benefits of both forms of transport.

Understanding the differing capabilities of the above-mentioned architectures requires familiarity with connections as well as circuit and packet switching. The different network categories above provide connectionless and connection-oriented networks that operate with circuit or packet switched core elements. Connectionless networks do not operate end-to-end with engineered paths, rather one device forwards traffic to another without direct control or knowledge of whether the services are delivered to the final destination. These networks are simple to set up and operate because they are typically auto-configured and auto-recovered. Whereas, connection-oriented networks operate end-to-end with engineered links and can generally guarantee delivery as well as implement planned recovery paths, which requires additional complexity in the operations side to set up each connection. Circuit switching allows for the set up of paths beforehand and reduces the overhead costs associated with small amounts of traffic. Circuit switching typically takes longer in establishing a session but results in lower overhead during the session. On the other hand, packet switching allows more dynamic forwarding on a packet-by-packet basis.

Layer 1—SDH/SONET Transport—Legacy Transport

Using SDH/SONET (ITU G.707) requires an adaptation and encapsulation of the Ethernet traffic. On top of SONET path mapping, Virtual Concatenation (VC) allows SONET to bond multiple Virtual Tributary (VT) (~1.7 Mbps) and synchronous transport signal (STS) (~52 Mbps) groups assigned to carry a particular Ethernet payload and hence provide for a more efficient and flexible bandwidth transport. Link Capacity Adjustment Scheme (LCAS, ITU G.7042) is an enhancement to VC that allows changing bandwidth—dynamically allocating capacity. Adjustments to bandwidth can occur in increments of a VT up to an STS level and then in increments of STSs above that. Finally, Generic Framing Protocol (GFP, ITU G.7041) or Point-to-Point Protocol (PPP) is used to encapsulate the Ethernet in the virtual container.

[2] SDH and SONET according to ITU G.707/Y.1322 are almost completely compatible. SONET includes a lowest native mapping of a single STS-1 at 51.84 Mbps, while the lowest native container for SDH is an STM-1 (STS-3) at 155.52 Mbps.

SONET/SDH transport allows operators to natively carry PDH services (T1/E1), which preserves strict stratum clock transport allowing the easy delivery of customer and cell site timing. Once Ethernet services are mapped into the transport, additional network hops do not add appreciable jitter or latency other than distance. SONET/SDH provides protection below 50 ms and 100 percent non-blocking availability for all services. The solutions are available today, scalable up to arbitrarily large networks, and can operate up to OC-768 (40 G) rates.

However, SONET/SDH transport is more expensive per transported bit than Ethernet or IP/MPLS equipment, as it does not allow for statistical multiplexing of traffic in the core. This effectively raises the required capacity per bit delivered. SONET/SDH is inherently a point-to-point service and does not natively support local area network (LAN) services. Additionally, SONET does not allow for different classes of service. All services are protected the same way and at the same expense, also raising the effective cost per bit served.

Layer 2—Ethernet Transport—802.1ad Provider Bridging or 'Q-in-Q'

As a transport infrastructure, Ethernet has had a varied history. Unlike SONET which was designed from the ground up as a reliable transport technology, Ethernet evolved from the corporate LAN environment. Initially, it was unsuitable for carrier transport. The first adaptation of Ethernet for transport operators was Q-in-Q or stacked VLANs (IEEE 802.1ad) where the customers' VLANs were stacked upon a provider-supplied VLAN (or P-VLAN), which allowed the provider to enable LAN-type connectivity to up to 4,094 customers while preserving their already applied VLAN tags. Protection and reconfiguration was supplied by the Spanning Tree Protocol (STP), which allows networks to entirely reconfigure themselves in event of a failure without a separate management application directing the effort. STP could provide restorations times as low as 150 ms; however, the upper bound is unspecified and only limited by network configuration—hardly the QoS guarantees expected in many wireless backhaul configurations.

This network implementation did not see widespread carrier deployment due to the limitation on the number of customers supported, limitations in offering high-QoS services and other scalability concerns. Instead, most operators' first wide scale packet-based transport infrastructure was based on Layer 3 IP/MPLS routers.

Layer 2—Ethernet Transport—802.1ah, MAC-in-MAC or PBB

To address the scalability limitation of Q-in-Q, MAC-in-MAC Ethernet (802.1ah) was introduced as a connectionless network technology that uses the MAC address from the provider to encapsulate the MAC address of the users while allowing the network to accomplish forwarding and addressing for a large set of customers.

While MAC-in-MAC allowed the operator to scale beyond 4,094 supported customers to nearly a limitless number (~1 M), MAC-in-MAC still uses the STP to accomplish network reconfiguration. Similar to Q-in-Q, MAC-in-MAC suffers from the same nondeterministic network reconfiguration, limiting its application in high QoS environments. For this reason, it may be deployed to deliver services that do not require strict restoration requirements, however, it is not an acceptable technology for the deployment of wireless backhaul requiring <50 ms restoration time.

Layer 3—IP/MPLS Transport including T-MPLS

IP/MPLS networking is the most common packet-based transport used for high QoS Ethernet services today. IP/MPLS allows operators to engineer links with full redundancy end-to-end while allowing the remainder of the traffic to operate in a connectionless mode. IP/MPLS networks can implement Virtual Private Wire Service (VPWS), LAN, and VPLS using point-to-point MPLS tunnels over the IP network. While the original justification of MPLS labels and wire speed forwarding have long past as a need, they now allow core networks to control complexity and cost when providing virtual network services.

With Fast Reroute (FRR), VPLS services can have lower 50 ms protection. Additionally, stacked MPLS tags allow a hierarchical operation (H-VPLS) for a nearly unlimited number of user connections across a common shared network.

While widely deployed and proven in the core packet networks of most global carriers, IP/MPLS is largely considered overkill for wireless backhaul. It adds additional cost and layers of complexity when deploying simpler Layer 2 services such as wireless backhaul point-to-point circuits and other metro Ethernet services. OAM for these networks alone operates at Layer1 (SONET or Ethernet), Layer 2 (Ethernet typically), Layer 2.5 (MPLS), Layer 3 (IP), and finally the Ethernet service. All these elements increase network complexity and ultimately add implementation expense.

To address these concerns, IP/MPLS has been simplified as T-MPLS for certain Layer 2 applications. T-MPLS discards the connectionless operation of the IP network and introduces a single end-to-end OAM along with protection mechanisms to simplify operation and lower costs. T-MPLS benefits from large support coming from the vendor base that is already supplying IP/MPLS equipment as well as from the strength of having come out of the already carrier class core.

Layer 2—Ethernet Transport—802.1Qay, PBB-TE

The final base transport technology and the latest to arrive in the networking world is PBB-TE, which keeps the same frame structure as PBB, or MAC-in-MAC, but adds a connection orientation to enable high levels of QoS. Users can select PBB-TE engineering on a PBB network that enables select switched services with fully redundant connections, guaranteed QoS, and fast restoration. STP continues to provide protection of the less critical services while the high QoS paths can be enabled with SONET-like fully redundant paths end to end. Unlike IP/MPLS, PBB-TE builds protection and scalability into the Layer 2 transport making the solution more economical to implement and simpler to administer.

PBB-TE provides less than 50 ms protection on engineered links and natively supports both point-to-point and LAN services. PBB-TE can scale to millions of supported customers on large networks and achieves the lowest cost per bit transported through Ethernet transport and statistical multiplexing. Inherently, it is less expensive and easier to administer than IP/MPLS Layer 3 networks.

However, PBB-TE is still not widely deployed in existing customer networks, still has interoperability and equipment availability issues, and needs careful timing distribution for some wireless backhaul applications (GSM).

Layer 1/Layer 2 Hybrid Networks—SONET/SDH Plus Ethernet in One Platform

Some equipment vendors have released solutions that attempt to combine the benefits of native Layer 1 SONET/SDH transport for T1/E1 services while adding Ethernet Layer 2 transport in a separate virtual container for native packet switched services. These solutions will typically rely on the more expensive SONET/SDH transport interfaces similar to what ATM networks have done, but they ultimately provide the operator with the accompanying timing and protection mechanisms.

Summary of Backhaul Transport Architectures

In summary, SONET/SDH, MPLS, T-MPLS, PBB-TE, and the L1/L2 hybrid are all viable network architectures for providers seeking to offer wireless backhaul services along side Ethernet and Layer 3 VPN services over the same network. Q-in-Q and MAC-in-MAC architectures fail to provide either the scaling or the QoS guarantees to make them suitable choices. Providers must then select from the remaining choices based on other requirements including: equipment availability, timing of the rollout, network management expertise, current and future traffic mix, required service level agreements (SLAs), QoS, and finally cost. Table 2 summarizes these considerations.

Ethernet Services for BNS

Regardless of the core network infrastructure chosen, the next generation of services being deployed for wireless backhaul and business services is Ethernet. Unfortunately, the core PDH or MPLS OAM that provides sectionalization and performance monitoring of the core does not map to the services carried over it. Hence, PDH or MPLS OAM does not provide the precise performance measurements of the end-to-end service being delivered. The following section reviews the elements of Ethernet services and the Ethernet fault and performance management specifications that are critical to delivering wireless backhaul or high SLA Ethernet business services.

	SONET/SDH Transport	Ethernet Q in Q	Ethernet MAC in MAC	IP/MPLS over Ethernet	T-MPLS over Ethernet	PBB-TE	L1/L2 Hybrid
Availability	✔	✔	✔	✔	Few Vendors	Few Vendors	Few Vendors
<50 ms restoration	✔	✗	✗	✔	✔	✔	✔
Stratum Timing	✔	← Evolving Standards Support →					✔
Stat Mux in Core	✗	✔	✔	✔	✔	✔	On L2 only
PDH Transport	✔	← Circuit Emulation Services (CES) →					✔
Multipoint LAN Serv	✗	✔	✔	✔	✔	✗ (very limited)	✔
Differentiated COS	✗	✔	✔	✔	✔	✔	✔
Scalability	✔	✗	✔	✔	✔	✔	✔
Cost Effectiveness	$$$$$	$	$	$$$$	$$$	$$	$$$$$
Suitable for Backhaul Transport?	✔	✗	✗	✔	✔	✔	✔
Chosen for what reason?	• High mix of T1 transport • Easy to operate • Available now • Timing transport			• Shared with core infrastructure • Availability	• Availability • End-to-End QoS	• Cost Effectiveness	• High mix of T1/E1 transport • Flexability
Relevant Standards	• G.707	•IEEE 802.1q • IEEE 802.1ad	• 802.1ah	• RFC 3031	• G.8110.1 • G.842/14 • G.8121/31/32	• 802.1 Qay	• N/A

Table 2 Comparison of Transport Architectures

User to Network Interfaces, Ethernet Virtual Connections, and Network to Network Interfaces

The Metro Ethernet Forum[3](MEF) has provided much of the structure for making Ethernet carrier class. The name "Metro Ethernet" belies the true mission of the MEF—making Ethernet the global transport and service technology of choice. The MEF member companies work together to establish guidelines, pre-standardization efforts, and implementation agreements that can be used to progress toward this goal. Seventeen technical specifications (TS) and implementation agreements (IA), MEF 1 through MEF 17, have been released covering aspects of Ethernet services, architecture, management and test and measurement.[4]

Fundamentally, the MEF has an overarching mandate for defining Ethernet and the services carried over it to help ensure flexibility and compatibility among implementations. As such, the MEF has defined Ethernet services in terms of the following fundamental parameters:

– User-to-Network Interface (UNI)—Physical interface between the subscriber and the carrier, for example 10/100 Gigabit Ethernet (GigE) and 10 GigE.

– Ethernet Virtual Connection (EVC)—Logical description of the service as defined by the association between two or more UNIs, for example 20 Mbps Committed Information Rate (CIR) or 2 ms latency. Termed virtual because many EVCs may exist between a single pair of UNIs to carry traffic of different priority.

– Network-to-Network Interface (NNI)—Interface that defines the inter-carrier services when multiple carriers are used to deliver that service.

NNIs are considered a crucial part of making Ethernet a widespread national and global business service by allowing Ethernet providers in the core to connect to wholesale local Ethernet providers. Areas currently under study include common service definitions, protection requirements, OAM mechanisms and SLA management. NNIs are not generally involved in BNS, because the networks exist on a metro-only level and do not involve a second Ethernet carrier.

EVC Service Types—Ethernet Virtual Private Line, Private Lines, and E-Line

In the MEF 6 Technical Specification, the MEF defines the first fundamental property of an EVC as the Ethernet Service Type. It can be:

– E-Line—Point-to-point EVC also known as E-pipe where the network operator delivers bits and does not dynamically switch content.

– E-LAN—Multipoint-to-multipoint EVC where the network operator provides a LAN or switching function.

The E-Line services type forces the operator to deliver a pipe, which is the typical transport option for carriers of Layer 1 T1/E1 services. E-LAN services have connectivity to a cloud with accessibility to all locations. Wireless backhaul relies almost exclusively on E-Line point-to-point services between BSCs and MSC/Data Hubs. However, many of the other business services customers the backhaul operator will also serve with the same infrastructure will demand E-LAN type services to connect multiple locations.

[3] http://metroethernetforum.org/page_loader.php?p_id=29 MEF 6 Full specification

[4] The Metro Ethernet Forum allows access to the published specifications. Access to the draft work is limited to members.

On top of the E-Line service type, the MEF also defines two relevant Ethernet services:

– Ethernet Private Line (EPL)
  - Only one EVC per UNI and guaranteed transparency
  - Does not allow Excess Information Rate (EIR) services
  - Used when the goal is to provide a simple pipe
– Ethernet Virtual Private Line (EVPL)
  - Multiple EVCs per UNI with near transparency
  - Allows EIR services
  - Used to provide flexibility for managing different Ethernet Class of Service (CoS), and multiple services over a single UNI.

The two Ethernet services types above (EPL and EVPL) are similar except the EVPL service can carry multiple EVCs and, therefore, is limited to providing transparency across the EVCs. Both the provider and customer must agree at the UNI on how the EVCs are mapped and carried. The EPL service is completely transparent and has no restrictions as to what is carried, so the provider and customer do not have to coordinate for the EPL EVC to provide 100-percent transparency.

Because towers may initially be provisioned with a single EVC provided to the wireless carrier, it is possible for the service to be an EPL. However, as the diversity of the tower grows, separate EVCs may be needed at a single tower to provide:

– Low latency CIR traffic for voice codec voice over IP (VoIP)
– Inexpensive high bandwidth EIR services for data applications
– Low latency, low frame loss CIR traffic for PDH emulation services

With the eventual need for multiple EVCs and to avoid service downtime and expense during upgrades, the BSN should be designed and deployed to support EVPL service. Providing EPL services in the short term will necessitate a disruption in service when upgrades are required to multiple EVCs over the single UNI.

[5] For a full discovery of progress on E-NNI, see the member area of MEF.

[6] EIR services are throughput services that are offered conditionally based on existing network capacity. For example, a link might have an EVC defined with 10 Mbps of Committed Information Rate (CIR) and 10 Mbps of EIR traffic. Capacity allowing, the customer could carry 20 Mbps over the link but may be limited at any time to its CIR of 10 Mbps.

Key UNI and EVC Attributes

For each service type, the MEF EPL and EVPL service definitions set many of the UNI and EVC Service Attributes.[7] Some highlights of the key UNI parameters include:

Per UNI Parameter	Ethernet Private Line—MEF 6 Section 6.1	Ethernet Virtual Private Line—MEF 6 Section 6.2
Service Multiplexing	No, only one EVC per UNI	Supported – multiple EVCs supported per UNI
Customer Edge VLAN (CE-VLAN) preservation	Yes, Customer's VLANs tags are guaranteed to be carried through	Optional transparency – typically preserved.
Number of EVCs per UNI	1, Only one service per UNI interface	1 or more services allowed over the same UNI interface
CIR and Committed Burst Size (CBS)	Supported per UNI, per EVC or per CoS	Supported per UNI, per EVC or per CoS
EIR and Excess Burst Size (EBS)	Not supported—No ability to specify the maximum capacity of the network to carry traffic that the network may optionally transport	Optional. Must also specify EBS. Service performance parameters do not apply to EIR traffic.[8]

Key EVC parameters include:

Per EVC Parameter	Ethernet Private Line—MEF 6 Section 6.1	Ethernet Virtual Private Line—MEF 6 Section 6.2
EVC Type	Point to Point	Point to Point
Service Performance Specification (SPS)	Only one allowed – specified for the EVC and includes: Frame Delay (FD) Frame Delay Variation (FDV) Frame Loss Ratio (FLR)	Optionally per CoS ID includes: FD FDV FLR

[7] See Sections 6.1 and 6.1 in MEF 6 for complete EVC and UNI parameter tables defined by the services and service types. They are not all included in this white paper because they do not have strategic impact on other aspects of cell backhaul.

[8] Determined by how the UNI bandwidth profile per Ingress UNI is specified. For traffic that is not considered in compliance, the Service Level Specification (SLS) performance parameters of frame loss, delay, and frame delay variation do not apply.

Case for Ethernet for Wireless Backhaul

Now that Ethernet and the relevant protocols and specifications have been explained, the pros and cons can be analyzed. First, these reasons explain why Ethernet should NOT be used:

– Cost-effective PDH services—Current backhaul providers may respond to the competitive threat posed by alternative backhaul providers and significantly lower the rates on existing PDH services. An attractive option as it will hold-off and reduce the incentive to switch, saving money now and not incurring extra engineering or capital expense.

– Difficulty conveying stratum timing—While Network Time Protocol (NTP), Precision Time Protocol (PTP), and IEEE 1588 v2 all make advances conveying timing over packet networks, Global System for Mobile Communications (GSM) wireless carriers may be unable to avoid obtaining timing from a PDH link.

– Long-term communications contracts—Do not switch service when in the second year of a 5-year, non-cancellable contract.

However, from the material in the prior sections, the advantages of migrating to Ethernet are also clear:

– Lowest cost per bit for a switched service—Through a combination of statistical multiplexing, economies of scale from the enterprise, and low-cost transceivers and switching elements, Ethernet packet services results in the price leader for switched services.

– Flexibility of interface—Ethernet leads the industry in flexibility with its ability to support multiple services and classes of service over a single UNI, along with the ability to support CIR and EIR.

– Transport over any medium—Ethernet can be carried over any medium including fiber, copper, and microwave. Additionally, Ethernet is carried on all major transport network architectures including SONET/SDH, IP/MPLS and PBB-TE networks.

– Scalability—From 1 to 1000 Mbps over one interface. While sector speeds are around 10 to 30 Mbps for 3G, 4G will enable 100s of megabits. Requires capability to deliver those speeds over future interfaces.

Ethernet Fault and Performance Management Challenges

Role of OAM

OAM allows operators to remotely manage networks. While enterprises can generally have a physical presence to test services with portables, carriers do not have that luxury and require tools that are built into the protocols and network elements to support remote management. The OAM protocols are well established in PDH, Frame Relay, and ATM communication services and generally provide the ability to:

– Proactively monitor the network and services—Includes performance reporting or fault detection on key network elements, links, service, and devices as well as the end-to-end service fault and performance.

– Promptly isolate faults in the network and service—The first bullet addresses detection, but the OAM protocols must also provide the ability to isolate further, so that the operator can issue the proper provisioning command or send a technician to the proper location equipped to solve the problem.

– Measure key SLA parameters—The OAM protocols must include the ability to proactively and reactively measure such key parameters as throughput, delay, jitter, permitted burst size, and availability.

– Measure usage for billing or capacity management—With packet-based services, usage varies. With properly designed billing, operators can maximize revenue; and with optimal capacity management, they can lower the capital expenditure needed to provide service.

Ethernet OAM Standards

Since the 1980s Ethernet has lacked OAM because of its enterprise roots impeding it from becoming a high QoS carrier service.

To enable Ethernet to achieve high QoS, several OAM standards have been released including:

– IEEE 802.3ah (2005)—Ethernet in the First Mile—This OAM specification, found in section 5/57 of the IEEE 802.3 2005 standard, covers link level OAM used to help sectionalize the most troublesome link in the network—between the customer and the provider. It operates between the Provider Edge (PE) switch and the Customer Edge (CE) switch. The main feature of the specification is providing the dying gasp messages between the CE and PE that inform the operator why the CE device may no longer be reachable.

– RFC 2544 (1999)—Benchmarking Methodology for Network Interconnect Devices. Originally intended as a specification on load testing network elements, this standard is now used to benchmark packet-based services—Ethernet, IP, Fiber Distributed Data Interface (FDDI), and others. It is a collection of tests that verify throughput, latency, frame loss rate, as well as ability to tolerate back-to-back frames and frames of different sizes. It does this by sending a varying load into the device or network under test and measuring the characteristics of the passed traffic. Practical adaptations also typically include enhancements for jitter or Fault Detection Verification (FDV) testing as well. RFC 2544 testing is intrusive and hence does not help operate the network. Instead, it is used to verify the settings for traffic shaping and policing at service turn-up. It may also be used after a repair event to verify proper operation.

– IEEE 802.1ag (2007)—End-to-End Service Connectivity Fault Management

– ITU Y.1731 (2006)—OAM Functions and Mechanisms for Ethernet-Based Networks

Y.1731 and 802.1ag cover mechanisms to perform quick fault detection, protection switching, discovery and fault isolation of Ethernet networks. Y.1731 covers additional performance protocols useful for measuring delay, delay variation, and loss rate.

The OAM detailed in both specifications can operate at up to 10 different simultaneous levels but nominally include the following:

– Customer OAM—Operates from CE switch to the opposite CE switch. Covers the entire service from the customer's point of view.

– Provider OAM—Operates from PE switch to the opposite PE switch. Covers the entire service from the retail operator's point of view.

– Operator OAM—Operates over local, metro, or core segments between UNI and NNI devices. It is used to help sectionalize service issues among different network segments.

There are several key protocol elements present in both Y.1731 and 802.1ag including:

– Continuity Check Messages (CCMs)—Periodic frames sent from a particular source MAC to a destination MAC used for protection switching and performance monitoring. The two MAC devices are called Managed End Points (MEPs) and they exchange CCMs in both directions at a rate between 300 frames per second and 1 frame every 6 hours. The absence of the frames, typically three or more missing frames, indicates the presence of a fault and notification and protection switching events would occur.

– Link Trace Message (LTM) and Link Trace Reply (LTR)—On-demand protocols sourced from one node and sent to other nodes on the EVC for purposes of discovery. One node sends out an LTM and all other nodes on the EVC reply with an LTR associated with their MAC. Initially, this protocol can be used to discover which devices are present on an EVC; and, if that information is benchmarked, it can also be used at a later time to perform fault isolation. The protocol will discover all nodes on the EVC including the MEPs and the other intermediate 802.1ag/Y.1731 aware nodes called Managed Interior Points (MIPs).

– Loop Back Message (LBM)/MAC Ping and Loop Back Reply (LBR)—On-demand protocols used to identify fault location. The originating managed end point (MEP)/managed interior point (MIP) sends the LBM to another MEP/MIP discovered in the LTM process. The destination MEP/MIP will respond with a LBR message back to the source if connectivity is present.

Many other protocols in 802.1ag and Y.1731 not mentioned in this paper augment the aforementioned protocols to provide a rich fault detection and isolation toolset for the Ethernet operator.

– ITU Y.1731 OAM—As mentioned earlier, Y.1731 alone contains certain protocols that assist in conducting Ethernet service performance monitoring including:

  - One-Way Delay Measurement (1DM)/Frame Delay (FD)—Command may be executed periodically or on demand and can be used with end-point synchronized clocks to measure one-way delay in the network. Synchronizing clocks at different points in the network can be achieved with global positioning system (GPS) timing, NTP, and PTP-IEEE 1588 v2.

  - One end sends the 1DM message with a precise time stamp attached, and the receiving end subtracts that measurement from its own internal clock to calculate the delay.

  - DMM/DMR Message, Round-Trip Delay, One-Way Delay, and FD Measurement—The originating MEP/MIP sends a DMM message and the destination MEP/MIP responds with a DMR that contains the time measurement that it took for the MEP/MIP to respond. The originating MEP then takes the total time from the DMM transmit to DMR receipt and subtracts the hold time for a true measurement of the round-trip delay, which is then divided in half to give an estimated one–way delay, or FD.

- Frame Delay Variation (FDV)—Measures the variation on FD calculated by either of the above methods. No separate protocol measures FDV. This measurement is always available on a one-way basis as only relative time measurements and not absolute synchronizations are needed to calculate the variations in FD.

  The industry has two distinct interpretations of FDV. The first is put forth by the Internet Engineering Task Force (IETF) and the MEF and is basically the same as RFC 3393—Inter Packet Delay Variation (IPDV). It measures the variation in delay between two adjacent packets, commonly mislabeled as jitter.

  The second methodology for calculating FDV is put forth by the ITU. This measurement is the variation between any two packets during a specified time interval such as a PM collection interval. While PDH purists would not call this jitter, it has the same basic interpretation and is typically referred to as jitter.

  The ITU method calculates the maximum difference in FD over the time period and results in a significantly different number from the MEF/IEFT metric.

  The reason for the discrepancy can be traced to the motivation for RFC 3393—IPDV. It was specified to give an easy to calculate metric that did not require any filtering or stable clocks to make a statistical measurement on the jitter present on multimedia streams. While the number does not directly represent jitter, it does give an easy to calculate metric that does relate to jitter.

  However, when PDH circuit emulation requirements are placed on an EVC, the hard requirements relate to jitter so that the de-jitter buffer of the PDH circuit emulation service (CES) device will not drop frames. Because it is the CES application, and not VoIP, that is driving the jitter requirement, JDSU believes the ITU version will ultimately be adopted for wireless backhaul performance metrics.

- Frame Loss Rate (FLR)—FLR measurement also does not have a specific protocol associated with it. FLR between two MEP/MIPS may be calculated continuously or on demand. If calculated continuously, the MEPs can either:

  - Make the measurement synthetically on CCMs only—The percentage of missing CCMs is extrapolated and reported as the estimated, synthetic, FLR. This methodology works for E-LAN or E-Line service types and is generally compatible with how packet performance metrics, such as IP SLA, were calculated in the past.
  - Use CCMs to measure actual customer frame loss—The already active CCM sessions contain measurements of transmitted and received packets that are sent to a paired MEP. These measurements are then compared with the received counts from and an absolute count of frame loss and the frame loss rate can be calculated.
  - Ethernet Loss Measurement (ETH-LM)—On-demand protocol than can be executed to make single or multiple FLR measurements and is similar to the actual customer FLR measurement above.

Managing Service Level Agreements

You should now have an understanding of the definition of an EVC and the different options of measuring the key performance parameters for Ethernet either at turn-up or during service operation. It is these definitions that are key to monitoring your conformance to SLAs.

It is important to note that Tier 1 commercial services or wholesale services operators tend to get paid from SLA conformance. SLA Contracts are signed that specify values for performance parameters such as:

– Delay, Jitter, and Frame Loss Rate—Parameters directly measured from the Y.1731 performance parameters. Care must be taken to ensure the FLR measurement and the one-way or two-way delay measurements truly fulfill the customer requirements.

– CIR, CBS, EIR, EBS—Parameters generally only verifiable at turn-up or at acceptance after a repair using RFC 2544 and similar procedures. Proprietary tools are available within the industry to measure the CIR parameter in operation.[9]

– Protection/Restoration Time—Verified by inspecting the transport architecture and/or observation of actual protection switching events.

– Mean Time to Repair (MTTR)—Typically measured from the ticketing system and is the time between a ticket opening and closing.

– Availability—Percentage of time the circuit was available for service while meeting or exceeding the SLA metrics. This metric has been well specified in the past with PDH services through standards such as M.2100, T1.231 and G.826. With the flexibility offered by Ethernet, it is natural that the definition of available time be dynamic and settable by the customer. For example, a 0.05 percent FLR may represent near perfect service for a business Ethernet customer while it would be intolerable to a wireless backhaul provider trying to do circuit emulation.

  For this reason, it is expected that customers should be able to specify a FLR and an associated time period (such as 5.0 x 10-7 over 1 hour) to define what they consider available time. ITU Y.1731 does not specify the rate or the time interval for availability.

  To remain congruent with the legacy PM standards above, it is desirable to have availability measurements at 1 second intervals. This allows a more accurate representation of available time but does force the hardware to collect and measure available time.

[9] Refer to Accedian EtherNID line of Ethernet Demarc Devices. They contain an 'RFC 2544 like' routine that fills the EVC with traffic up to the CIR while not impacting customer traffic.

Ethernet Service Assurance Challenges

The above-mentioned strategies for assuring that Ethernet meets the stringent requirements for Wireless backhaul are not without their challenges, including:

– Lack of Support for Y.1731 PM in Network Elements—The standards were only recently approved and full implementation is not simple. MEP/MIP and 802.1ag capabilities may be available soon, but bulk availability of Y.1731 PM statistics on core Ethernet switches will likely take upgrades all the way down to the chip sets. This will lengthen the amount of time it will take operators to release this capability and may require hardware upgrades to actually implement.

  However, add-on devices such as UNI NIDs and NNI NIDs are available in the market to book-end the EVCs and enable Y.1731 PM without upgrading the infrastructure.

– Seamless transition from a book-end approach to Network Element (NE) Supported Model—While this does not present an immediate challenge, ensuring the stability of operations, consistent SLA performance monitoring and continuity of tool investment as the operator migrates to the final NE model is key.

– Wireless carriers will need to manage a diversity of backhaul operators—The Wireless carriers will need standard PM and SLA reporting as well as flow through trouble resolution to ensure they do not get operationally mired in dealing with 30 different backhaul providers. Key business processes include consistent reporting, known Y.1731 standardized capabilities in any backhaul provided NID, and automated flow through capabilities.

Practical implementation of Ethernet Backhaul Assurance

Once the OAM measurement and loopback capabilities mentioned earlier become available, operators must follow a basic process to ensure that service is successful.

– Turn-up the service right the first time—Typically, a technician performs an RFC 2544 test in the field using a test probe in the core to verify the EVC parameters. Steps for service turn-up include:
  - Perform quick provisioning check with technician's portable and the remote system
  - Release of the technician
  - Activate the central probe system to conduct a long-term loopback test to verify high QoS.

Service can now be properly tested at the traffic shaper ingress, burned in over a long period of time, and the amount of technician time reduced, because the long-term test is now automated.

– Continuously monitor the network—Operators monitor key statistics for long-term trending and performance violations.
  - Network link statistics—Runts, code violations
  - Element CPUs and buffer stats
  - Ring and/or path performance

Engineers use summarized reports and Month-to-Date, Top-N-offender-type reports to proactively address issues and raise the level of QoS of all services on the network.

– Continuously Monitor the Service—Operators monitor the fulfillment of SLAs to discover issues before the SLA is violated. This depends heavily on the Y.1731 PM capabilities in the circuit. Monitoring the service and providing that data to the SLA management system is ultimately how providers get paid!

– Enable Quick Fault Isolation—Have the same turn-up tools available on demand or, better yet, on a flow-through basis, to confirm troubles and properly dispatch technicians.

Conclusions

The question remains, however, it's not a matter of "If" but rather "When" to perform Ethernet wireless backhaul. The promise of 4G makes Ethernet, microwave, and fiber backhaul a necessity in providing affordable, flexible high bit rate BNS. Legacy performance management models have proven that the MEF for Ethernet service PM proposal will ultimately be adopted; it is just a matter of time. Wireless backhaul services remain the critical Ethernet service pushing the industry and forcing the adoption and progression of these standards.

Acronyms

1DM	One-Way Delay Measurement
3G	Third Generation
4G	Fourth Generation
ATM	Asynchronous Transfer Mode
BNS	Backhaul Network Service
BSC	Bay Station Controller
CBS	Committed Burst Size
CE	Customer Edge
CE-VLAN	Customer Edge Virtual Local Area Network
CES	Circuit Emulation Service
CIR	Committed Information Rate
CoS	Class of Service
CCM	Continuity Check Message
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
EIR	Excess Information Rate
E-LAN	Multipoint-to-Multipoint Emulated Local Area Network
E-LINE	Point-to Point Ethernet Virtual Circuit (pipe or line)
EDFA	Erbium Doped Fiber Amplifier
EMEA	Europe, Middle East, and Africa
EPL	Ethernet Private Line
ETH-LM	Ethernet Loss Measurement
EVC	Ethernet Virtual Circuit
EVPL	Ethernet Virtual Private Line
FD	Frame Delay
FDDI	Fiber Distributed Data Interface
FDV	Frame Delay Variation
FLR	Frame Loss Ratio
FRR	Fast Re-Route
GFP	Generic Framing Protocol
GPS	Global Positioning System
GSM	Global System for Mobile Communications
H-VPLS	Hierarchical Virtual Private Local Area Network Service
IA	Implementation Agreements
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ILEC	Incumbent Local Exchange Carrier
IP	Internet Protocol
IPDV	Inter Packet Delay Variation
L1-L5	Layers 1 through 5
LAN	Local Area Network
LBM	Loop Back Message
LBR	Loop Back Reply
LCAS	Link Capacity Adjustment Scheme
LTM	Link Trace Message
LTR	Link Trace Reply
MAC	Media Access Control
Mbps	Megabits per Second
MEF	Metro Ethernet Forum
MEP	Managed End Point
MIMO	Multiple Input Multiple Output
MIP	Managed Interior Point
MPLS	Multi Protocol Label Switching
MSC	Mobile Switching Center
MTTR	Mean Time to Repair
NE	Network Element
NID	Network Interface Device
NNI	Network-Network Interface
NTP	Network Time Protocol
OAM	Operations, Administration, and Maintenance
OFDM	Orthogonal Frequency-Division Multiplexing
P-VLAN	Provider-supplied Virtual Local Area Network
PBB	Provider Backbone Bridges
PBB-TE	Provider Backbone Bridges
PDH	Plesiochronous Digital Hierarchy
PE	Provider Edge
PPP	Point-to-Point Protocol
PSTN	Public Switched Telephone Network
PTP	Precision Time Protocol
PWE	Pseudowire
Q-in-Q	Layer 2—Ethernet Transport—802.1ad Provider Bridging
QoS	Quality of Service
ROI	Return On Investment
SDH	Synchronous Digital Hierarchy
SLS	Service Level Specification
SONET	Synchronous Optical Networking
STP	Spanning Tree Protocol
STS	Synchronous Transport Signal
T-MPLS	Transport Multi Protocol Label Switching
TDM	Time Division Multiplex
TS	Technical Specifications
UNI	User-Network Interface
VC	Virtual Concatenation
VLAN	Virtual Local Area Network
VoIP	Voice over Internet Protocol
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VPWS	Virtual Private Wire Service
VT	Virtual Tributary
WDM	Wavelength Division Multiplexing
Wimax	Worldwide Interoperability for Microwave Access

Test & Measurement Regional Sales

NORTH AMERICA
TEL: 1 866 228 3762
FAX: +1 301 353 9216

LATIN AMERICA
TEL: +55 11 5503 3800
FAX: +55 11 5505 1598

ASIA PACIFIC
TEL: +852 2892 0990
FAX: +852 2892 0770

EMEA
TEL: +49 7121 86 2222
FAX: +49 7121 86 1222

WEBSITE: www.jdsu.com

 30149499 000 0808 NXTGNCELLULAR.WP.SAS.TM.AE August 2008