Low-Cost Next Generation Metro Architectures for Carrier Ethernet Services

by Mark Showalter, Director of Service Provider Marketing

5/6/2009

Carriers around the world are transitioning their legacy circuit-based network architectures to next-generation packet architectures to broaden their service mix and expand their revenue streams. Legacy architectures rely on a separate network for each service � time-division multiplexing for voice, frame relay for private data, hybrid fiber-coax for video, CDMA for mobile, and IP/MPLS for high-speed Internet access. The broader trend of convergence in telecommunications means carriers need to offer an array of services to remain competitive. Operating and maintaining disparate networks is expensive and limits profitability for carriers with legacy architectures.

The challenge facing metro network planners is to find the most cost-effective way to deliver profitable services in the near-term while deploying a network flexible enough to support an array of new services in the future. Next generation network architectures first are built to profitably meet the needs of an initial market such as business Ethernet services. Once this infrastructure is in place, the goal is to leverage the investment to deliver additional services including Ethernet mobile backhaul and residential triple-play aggregation in order to maximize the return on their investment.

Infonetics Research forecasts carrier revenues for Ethernet services will approach $30B annually by 2013. Much of this growth will be at the expense of legacy services like T1, frame relay and ATM. Carrier Ethernet services defined by the Metro Ethernet Forum (MEF) have gained wide acceptance by service providers and equipment vendors around the world, providing a framework for customers to make educated comparisons between various offerings.

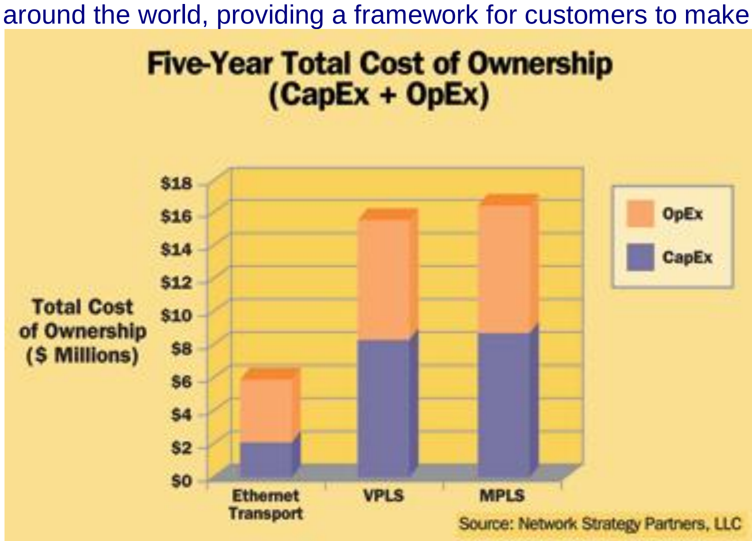

Some of these MEF service definitions are modeled after the popular frame relay service definitions, although at much higher data rates. As with frame relay, the term User Network Interface, or UNI, is used to define the demarcation point between the carrier and the subscriber. Frame relay UNIs are typically available with 56 Kbps to 45 Mbps data rates, whereas Ethernet UNIs may be provisioned with data rates from 10Mbps to 10 Gigabits per second. Also as with frame relay, Committed Information Rate (CIR) and Excess Information Rate (EIR) settings may be used to provision fractional rate services, such as fractional 10 Gigabit service available in 1 Gbps increments. This article compares next-generation metro architectures that enable MEF E-LINE and ELAN services at the UNI, while providing a solid foundation for additional services in the future.

E-LINE services provide a point-to-point connection between two Ethernet UNI interfaces. These services are alternatives to T1/E1, T3/E3 and SONET/SDH leased line services. ELAN services provide multi-point service that connects three or more locations and serves as an alternative to packet-switched services such as frame relay.

There are several alternative metro network architectures for delivering this array of services. The most common approaches may be generally categorized as metro Virtual Private LAN Service (VPLS), metro Multi-Protocol Label Switching (MPLS), and metro Ethernet Transport. All of these approaches enable service providers to combine an array of services onto a common metro network to reduce operational complexity and costs relative to legacy architectures.

MPLS has been successfully deployed to address complex routing issues at the core of IP networks. The metro MPLS approach extends a rich set of IP routing features across the metro, where Ethernet switches with 1 GE optical uplinks aggregate traffic onto routers with MPLS line cards at the serving central offices and at the metro core.

VPLS was introduced to enable a robust set of Ethernet private LAN services above and beyond what MPLS was designed to support, while taking advantage of MPLS core network services. The metro VPLS approach extends these services to the edge by deploying routers with VPLS optimized silicon at the serving central offices and at the metro core.

In contrast, the **metro Ethernet Transport** approach relies primarily on Ethernet technologies across the metro, including VLAN translation, IEEE 802.1ad Provider Bridging (PB), IEEE 802.1ah Provider Backbone Bridges (PBB), and IEEE 802.1Qay Provider Backbone Bridges - Traffic Engineering (PBB-TE).

With each of these approaches, the metro network connects to an MPLS carrier core network which extends Ethernet services beyond the metro for those businesses whose offices are spread around the world. The Ethernet Network-to-Network Interface (ENNI) serves as the demarcation point between the different carriers enabling global services to be designed across an archipelago of metro networks.

Although these metro differing approaches all provide the tools to deliver E-LINE, E-LAN and other services, each approach involves trade-offs that impact the total cost of ownership (TCO) over the life of a metro network. Network Strategy Partners (NSP), management consultants to the telecommunications industry, released a white paper detailing their study comparing the TCO of metro Ethernet, metro VPLS and metro MPLS approaches for residential triple-play, business Ethernet, and wireless backhaul services. Figure 1 shows the relative TCO of the three approaches over five years as reported by NSP. The study shows the TCO for Ethernet Transport is 61% and 63% lower than for the metro VPLS and metro MPLS alternatives, respectively. These savings result from both lower capital expenditures (CapEx) and operational expenses (OpEx).

According to the study, "The VPLS and MPLS alternatives have much higher capital expense than the Ethernet Transport alternative. This cost difference is due to the VPLS and MPLS alternatives use of IP/MPLS embedded features in their line-cards as compared to the widely deployed Ethernet L2/L3 features of the Ethernet Transport line-cards. The volume of production of Ethernet line-cards is two orders of magnitude greater than that of the VPLS and MPLS line-cards. Consequently Ethernet Transport vendors can profitably support price levels that are less than half of those of the highly specialized (and relatively low volume) VPLS and MPLS line-cards." Additionally, NSP reports that VPLS and MPLS switches have lower 10GE port densities, thus requiring more switches to satisfy the same amount of metro traffic.

On the OpEx side of the equation, the study shows the top five contributors to OpEx include service contract expenses, engineering/facilities/installation expenses, network care expenses (which include the cost of provisioning, monitoring, and maintenance of the network), training expenses, and testing/certification operations expenses. Because many of these contributing factors are a fixed percentage of CapEx (which is higher for both VPLS and MPLS), Ethernet Transport enables the lowest relative OpEx. Service contacts are the biggest contributor to OpEx, typically costing 10-15% of the CapEx. Other OPEX savings accrue from lower power consumption and subsequent reduction in heating, ventilation, and air-conditioning (HVAC) costs.

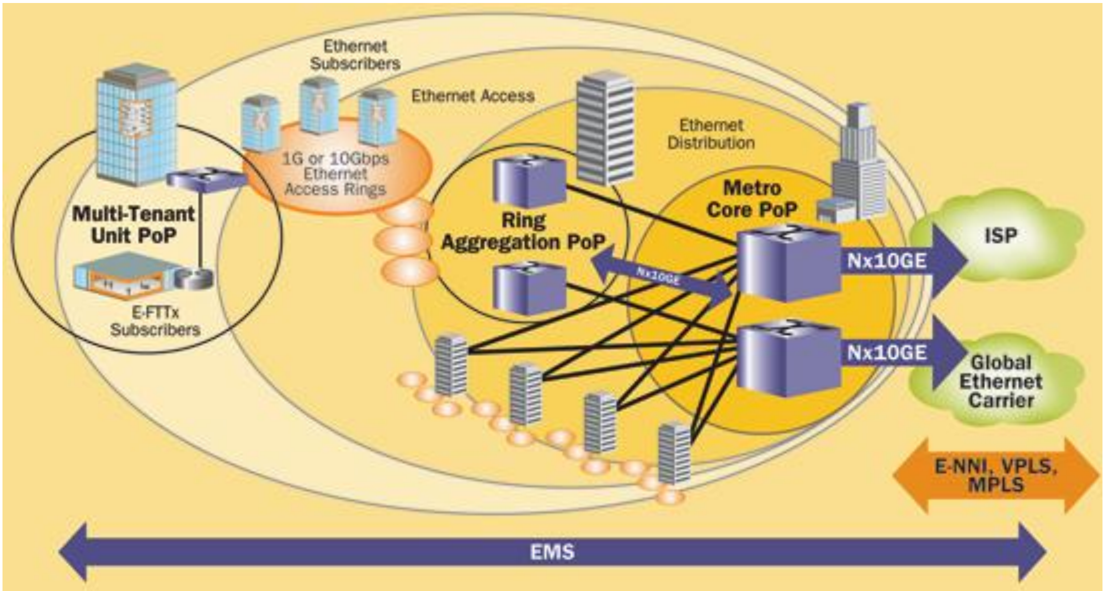

Figure 2

Figure 2 is an Ethernet FTTX (E-FTTX) architecture representing the Metro Ethernet Transport solution that meets these requirements. It is composed of multiple tiers including the metro core point-of-presence (POP), the provider ring aggregation POP and the multi-tenant unit (MTU) POP and can gracefully scale from hundreds to hundreds of thousands of UNIs per metro core POP. Access switches are placed at the MTU-POP and provide the first level of aggregation in the metro network. Aggregation switches provide the next level of aggregation and are placed at the ring aggregation POP, collecting traffic from multiple access switches. The metro core provides high-density 10 GigE connectivity between the ring aggregation sites.

To deliver services, the 802.1Q (VLAN Tagging) and 802.1ad (Provider Bridging) protocols are utilized to support highly scalable point-to-point (P2P) and multi-point to multi-point (MP2MP) connections.

A ring topology is the most redundant type of network topology but depending on geographical characteristics, existing cabling, and redundancy requirements, star or meshed configurations can also be supported. On the subscriber side, a customer premise equipment (CPE) is provided and managed by the provider, or the subscriber may choose to procure their own.

Granular quality of service (QoS) control and high levels of bandwidth capacity provide the capability to support a wide range of services over the same network, each with its own performance requirements. Network resiliency is provided using sub-50ms recovery from failures on metro fiber rings.