Connection-Oriented Ethernet – Completing the Ethernet Revolution

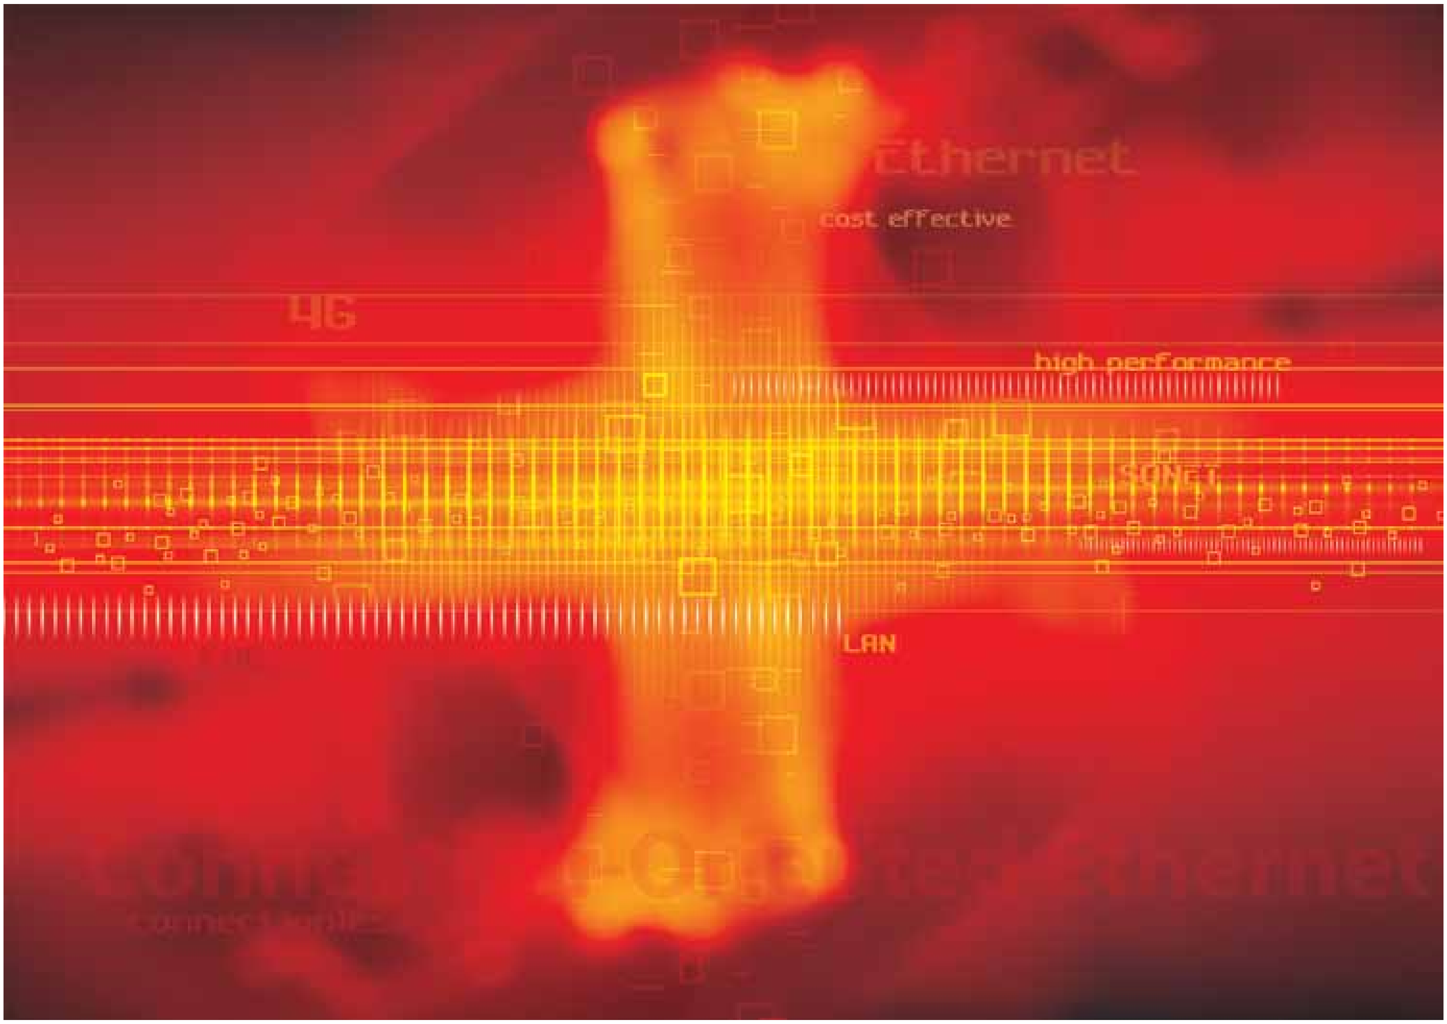

Ethernet is the hottest thing going in metro networking. Ethernet enterprise services are forecasted to grow at a compound annual rate of over 25% per year and first-generation triple play networks are all based on Ethernet.

To satisfy these demands, service providers have deployed a wide variety of Ethernet technologies, depending upon the specific application. Broadly speaking, these have been either connectionless native Ethernet technologies or Ethernet encapsulated into connection-oriented networking technologies like SONET or OTN. The connectionless networking approaches have been vital for multipoint-to-multipoint E-LAN enterprise services applications, but have also been used for aggregation of Ethernet traffic although the service quality has been lower. The connection-oriented SONET and OTN encapsulations have been vital for providing high-quality Ethernet point-to-point services and infrastructure transport, but they cannot provide any kind of statistical multiplexing or the port aggregation necessary to scale the network economically.

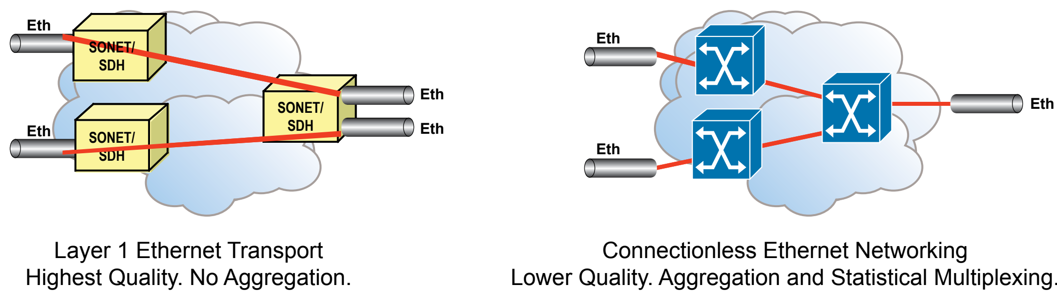

Figure 1: Layer 1 Ethernet Transport and Connectionless Ethernet Characteristics

As Ethernet plays an increasingly pervasive role in metro networking for 4G mobile backhaul, enterprise services, continued growth in residential broadband, and as an access technology for IP services, network providers must be able to deploy a single, general-purpose Ethernet aggregation and transport environment for all these applications. This new infrastructure must deliver the robustness and performance guarantees of Layer 1 SONET/SDH/OTN encapsulation, along with the aggregation and statistical multiplexing economies of native Ethernet. To create a general-purpose infrastructure, Ethernet must evolve beyond simple connectionless networking. Connection-oriented Ethernet is the only technology that provides both efficient aggregation and performance guarantees, thereby completing the Ethernet revolution by accomplishing for Ethernet what SONET/SDH accomplished for TDM.

Connection-oriented Ethernet is the only technology that provides both efficient aggregation and performance guarantees, thereby completing the Ethernet revolution by accomplishing for Ethernet what SONET/SDH accomplished for TDM.

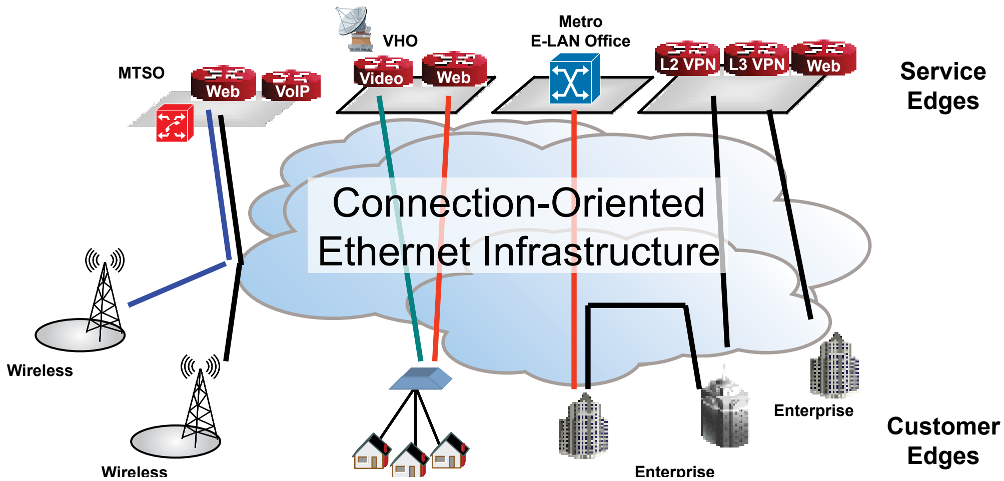

Figure 2: General-Purpose Connection-Oriented Ethernet Aggregation Infrastructure

Making Ethernet Connection-Oriented

Historically, Ethernet has been a connectionless technology by design. In classic LAN environments, the connectionless capabilities of Ethernet MAC bridging and CSMA/CD provided considerable flexibility, simplicity, and economy in networking latency-insensitive traffic within a single, well-bounded administrative domain.

To bring the cost and flexibility benefits of Ethernet into the public network, the industry has made enormous modifications, enhancements, and extensions to classic Ethernet protocols. The result is "Carrier Ethernet." Most of these enhancements have focused on extending the connectionless classic Ethernet protocol to a service provider environment. This has brought about new capabilities in:

- OAM (IEEE 802.1ag, IEEE 803.2ah, ITU-T Y.1731)
- Survivability (ITU-T G.8031, IEEE 802.3ad, IEEE 802.1s)
- Scalability (IEEE 802.1ad, IEEE 802.1ah)
- Speed and distance (e.g., IEEE 802.3ae)

In addition to the above standards, equipment vendors "beefed up" enterprise-class Ethernet hardware and software platforms to provide a number of carrier-class features.

The next step forward is to make Ethernet connection-oriented. Connection-oriented Ethernet is a high-performance implementation of Carrier Ethernet that is optimized for aggregation infrastructure and EVPL, EPL, and E-Tree services.

Connection-oriented Ethernet is a high-performance implementation of Carrier Ethernet that is optimized for aggregation infrastructure and EVPL, EPL, and E-Tree services.

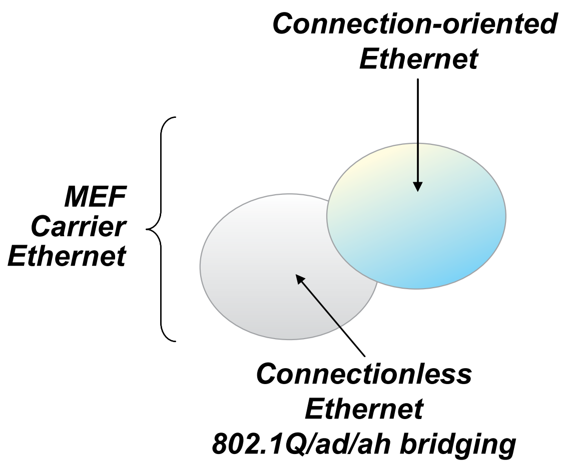

Figure 3: Connection-Oriented Ethernet: A High-Performance Implementation of Carrier Ethernet

The Three Essential Functions of Connection-Oriented Ethernet

The three essential functions that make Ethernet connection-oriented are:

- Predetermined EVC paths
- Resource reservation and admission control
- Per-connection traffic engineering and traffic management

The ability to predetermine the EVC path through the Ethernet network is fundamental to making Ethernet connection-oriented. In classic connectionless Ethernet bridging, Ethernet frames are forwarded in the network according to the MAC bridging tables in the learning bridge. If a destination MAC address is unknown, the bridge floods the frame to all ports in the broadcast domain. Spanning tree protocols like IEEE 802.1s are run to ensure there are no loops in the topology and to provide network restoration in the event of failure. Depending upon the location and sequence of network failures, the path EVCs take through the network may be difficult to predetermine.

Predetermining the EVC path—either through a management plane application or via an embedded control plane—ensures that all frames in the EVC pass over the same sets of nodes. Therefore, intelligence regarding the connection as a whole can now be imparted to all nodes along the path.

Resource reservation and CAC is the next critical function. Now that the EVC path through the network has been explicitly identified, the actual bandwidth and queuing resources required for each EVC are reserved in all nodes along the path. This is vital to ensure the highest possible levels of performance in terms of packet loss, latency, and jitter. CAC ensures that the requested resource is actually available in each node along the path prior to establishing the EVC.

Once the path has been determined and the resources allocated, the traffic engineering and traffic management functions ensure that the requested connection performance is actually delivered. After packets have been classified on network ingress, a variety of traffic management functions must be provided in any packet-based network. These include:

- Policing
- Shaping
- Queuing
- Scheduling

These traffic management functions are briefly outlined below.

Packet classification is the processes of identifying to which EVCs the incoming frames belong. The ingress equipment can examine a variety of Ethernet and IP layer information to make this decision. Once the incoming frame is classified, policing is then applied to ensure that all frames coming into the network conform to the traffic contract, known as the bandwidth profile, agreed to upon connection setup. Two-level, three-color marking allows incoming frames that conform to the CIR to be admitted to the network. Frames that exceed even the EIR are discarded immediately, and frames that exceed the CIR, but not the EIR, are marked for possible discard later, should the network become congested. An EVC can be subject to a single such policer if the bandwidth profile is applied to the entire EVC. EVCs can also include bandwidth profiles for each of many CoS classes within the EVC. In this case, a single EVC can be subject to multiple policers.

The Heart of Connection-Oriented Ethernet

The shaping, queuing, and scheduling granularity determines whether each individual EVC enjoys the significant performance benefits of connection-oriented Ethernet, or is essentially connectionless from a performance point of view.

Every time an Ethernet frame in any Ethernet packet network waits for a transmission opportunity on an egress port, it is queued along with other connections that are also bound for the same egress port and a scheduler determines which frame goes next and possibly there is a shaping function. The critical question is how these frames are queued, scheduled and shaped and at what level of granularity. Several possibilities exist within a connection:

- Per card
- Per port
- Per EVC connection
- Per CoS

On many Carrier Ethernet traffic management implementations, Ethernet frames from many EVCs are placed into a single set of egress queues and therefore visibility into individual EVCs becomes lost. When this happens, the network begins to play "priority roulette" where, essentially, all EVCs in a single priority class get random access to transmission opportunities; the service quality of individual connections consequently goes down.

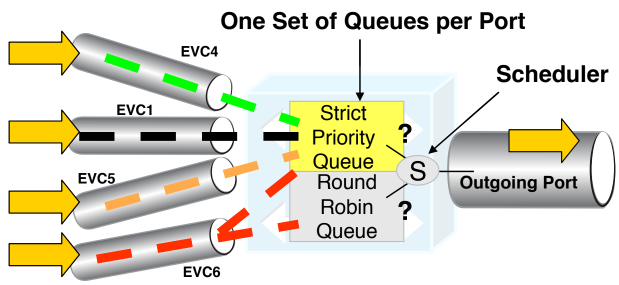

Figure 4: "Priority Roulette": Individual EVC Visibility Lost to Scheduler in Per-Port Queuing

Connection performance can, therefore, only be guaranteed by providing policing, scheduling, and hierarchical egress shaping functions down to the granularity of an individual CoS class within an EVC.

Connection performance can, therefore, only be guaranteed by providing policing, scheduling, and hierarchical egress shaping functions down to the granularity of an individual CoS class within an EVC.

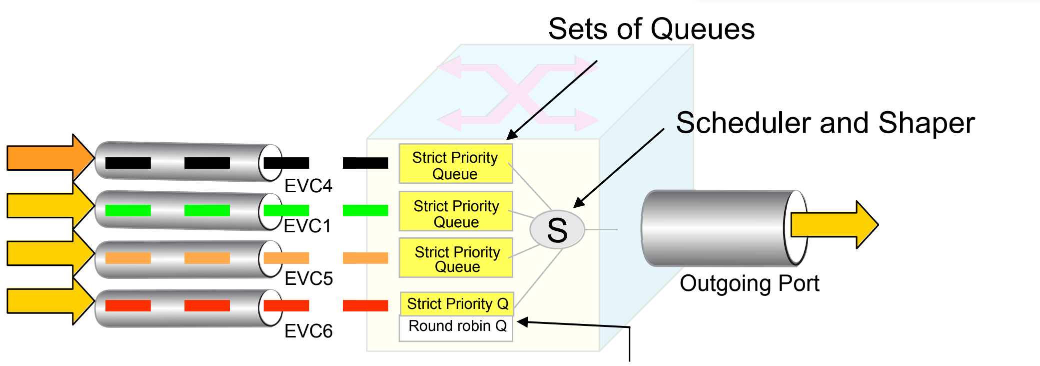

Figure 5: Per-EVC Queuing, Scheduling and Shaping

By providing the previously discussed three essential functions of connection-oriented Ethernet, Ethernet connections can now enjoy levels of service quality on a par with Ethernet over SONET/SDH—along with the aggregation flexibility of native Ethernet.

A Word About Protection

Once Ethernet frames flow in a connection-oriented manner across the network, it becomes possible to provide dedicated, deterministic, automatic protection switching functionality equivalent to that provided by SONET/SDH—with 10 ms failure detection and 50 ms protection switching speed.

ITU-T G.8031 and IEEE 802.1Qay standards both offer methods to create dedicated protection resources with the same deterministic characteristics as the working path resources (bandwidth profile and loss, latency and jitter performance). Upon Ethernet layer or lower-layer failure or degradation, protection switching logic (located in the originating and terminating nodes of a protection domain) performs rapid 10 ms failure detection and switches to the protection resource in 50 ms. This allows service providers to provide availability guarantees similar to TDM, over arbitrary Ethernet network topologies and large distances.

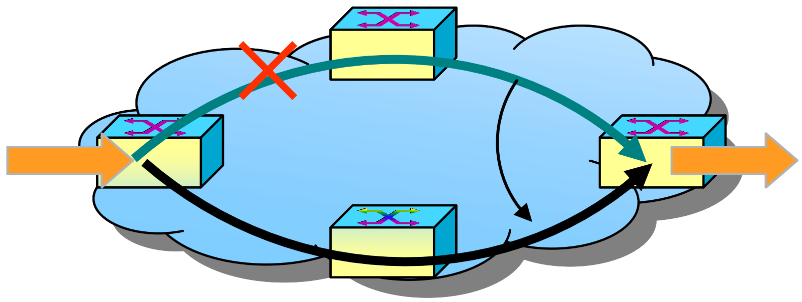

Figure 6: Automatic Protection Switching for Connection-Oriented Ethernet

The Technology Choices

Several existing and emerging technologies can realize connection-oriented Ethernet. These technologies include:

- VLAN switching
- 802.1Qay (PBB-TE)
- MPLS-TP
- T-MPLS
- PW/MPLS

There are two important considerations when choosing a connection-oriented Ethernet approach for metro aggregation and transport infrastructure:

- How many layers are required in the network?
- Is a routed IP data plane and associated complexity required?

Ethernet-only approaches, such as VLAN switching and PBB-TE, do not require additional pseudowire or MPLS layers. They are therefore simpler to integrate into OSS systems and require fewer layers of OAM for fault and performance management.

Furthermore, Ethernet-only approaches do not require deployment of a software intensive IP-routed data plane deep into the metro aggregation network. This is vital, because there are often thousands of remotely deployed elements in this network; the introduction of an IP-routed control plane would greatly increase the complexity of software maintenance, troubleshooting and provisioning, and therefore raise the operational cost of the network.

For a detailed discussion of the benefits of Ethernet-only approaches to connection-oriented Ethernet, see the white paper "Connection-Oriented Ethernet: Deployment and Operational Considerations."

Important Applications for Connection-Oriented Ethernet

There are several important applications that will drive deployment of connection-oriented Ethernet. These include:

- Wireline backhaul for 4G mobile wireless networks
- Infrastructure scaling for enterprise Ethernet services networks
- Private-line-equivalent quality EPL services
- Triple play network optimization
- General Ethernet infrastructure

Some of these are discussed in the sections that follow.

Backhaul for 4G Mobile Wireless Networks

WiMAX and LTE mobile wireless networks will roll out in 2009 and 2010 respectively. Both technologies enable significant advances in mobile multimedia applications and, as a consequence, drive significantly more bandwidth than previous mobile technologies. 4G mobile technologies specify native Ethernet interfaces from the base station to the wireline network, in contrast to 3G technologies that specified T1 base station interfaces.

In an anticipated 4G mobile backhaul network, there may be 50 to 100 cell towers, each bringing back hundreds of Mbps of bandwidth to a set of mobile switching offices. This traffic needs to be funneled down to a handful of Ethernet ports at the mobile switching office. This requires a large amount of efficient Ethernet aggregation that does not compromise the latency and jitter of certain voice and signaling applications. This combination of large-scale Ethernet aggregation and tight latency and jitter performance requirements, make connection-oriented Ethernet a must for these new networks. Connection-oriented Ethernet deployed in an optical backhaul network enables the necessary aggregation, provides network protection, and allocates resources for all the Ethernet connections, while allowing oversubscription for interactive and background applications.

This combination of large-scale Ethernet aggregation and tight latency and jitter performance requirements, make connection-oriented Ethernet a must for these new networks.

Scaling of Enterprise Ethernet Services Networks

Many enterprise Ethernet services deployments have employed small form-factor platforms designed to provide Ethernet-based access to Internet or other IP services, or to provide point-to-point or multipoint connectivity services. As these networks grow, it becomes necessary to improve the efficiency with which these small form-factor platforms are interconnected to the IP service edge, or to provide connectivity among embedded Ethernet platforms.

Ethernet cross-connect or "DCS" machines with connection-oriented Ethernet provide a solid, scalable foundation for this interconnection function. Ethernet DCS with a large-scale grooming fabric consolidates EVCs or VLANs from multiple Fast Ethernet and Gigabit Ethernet ports to highly filled 10GE pipes and interconnect traffic efficiently among Ethernet access networks. Being able to manipulate, translate, and push/pop tags flexibly in this Ethernet DCS provides an incredible amount of networking and interconnection power. These platforms offer economic payback at fairly low levels of traffic and provide considerable management flexibility.

Because the Ethernet DCS utilizes connection-oriented Ethernet, the SLA requirements for any type of service can be met or surpassed through the device. Even with simple CIR provisioning, there is considerable economic and management benefit. Adding statistical multiplexing capability provides additional efficiencies.

If connection-oriented Ethernet switching is integrated into the electronic switching fabric of a Packet ONP that may already exist in the office for SONET grooming, then the Ethernet DCS common-equipment cost is effectively eliminated, and the service provider can derive an immediate economic benefit by deploying Ethernet grooming.

If connection-oriented Ethernet switching is integrated into the electronic switching fabric of a Packet ONP that may already exist in the office for SONET grooming...service provider can derive an immediate economic benefit by deploying Ethernet grooming.

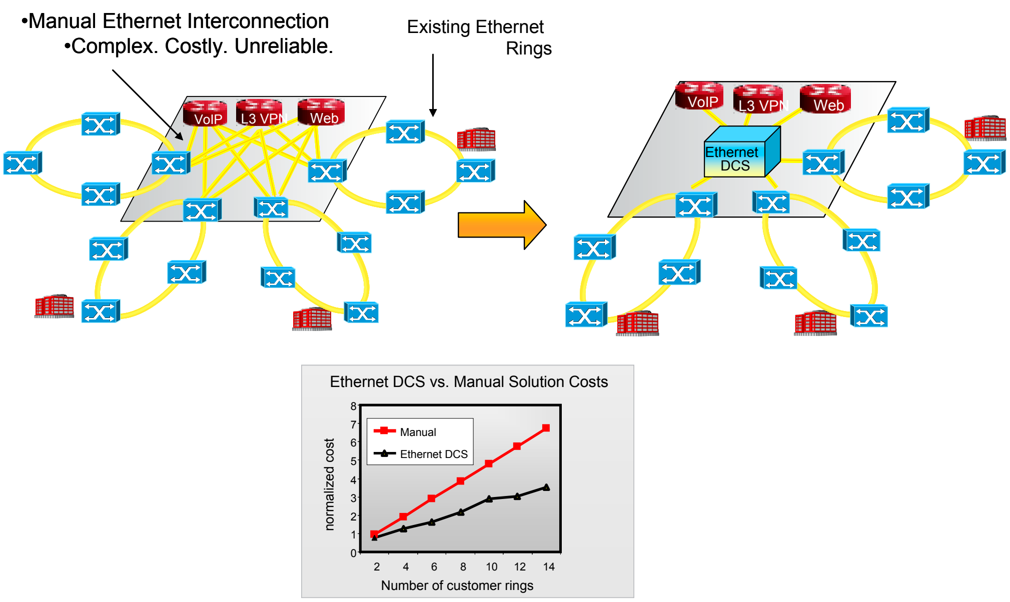

Figure 7: Cost and Managability Improvements of Connection-Oriented Ethernet DCS

Point–to–Point Services

The MEF has defined two general types of point-to-point Ethernet services—EPL and EVPL. EPL services have a single EVC connecting one ingress port to one egress port across the network. This type of service is delivered over both connectionless native Ethernet networks as well as connection-oriented SONET/SDH networks.

EVPL services are more flexible in that they have more than one EVC per port, so a single Ethernet port at a hub location can have multiple EVCs that connect to different remote locations.

When EVPL services are provided over connectionless Ethernet networks, the service quality is often not as good as with traditional TDM private-line circuits. This is because of the hop-by-hop nature of the traffic management, the need to account for unexpected Ethernet flooding, and reliance upon spanning tree protocols for restoration. In addition, connectionless Ethernet networks often place significant requirements and restrictions upon CPE. These restrictions include such items as requiring the CPE device to shape traffic, requiring the CPE device to be a router, and limiting the amount of latency-sensitive traffic that can be introduced onto the network over a customer UNI.

With connection-oriented Ethernet technology, EVPL services with EPL quality can be delivered along with 50 ms protection switching and all significant restrictions can be removed from the CPE.

When connection-oriented Ethernet is integrated directly into the optical access environment in a Packet ONP, there are additional opportunities for cost-saving by offloading EVPL services from the connectionless E-LAN services platform.

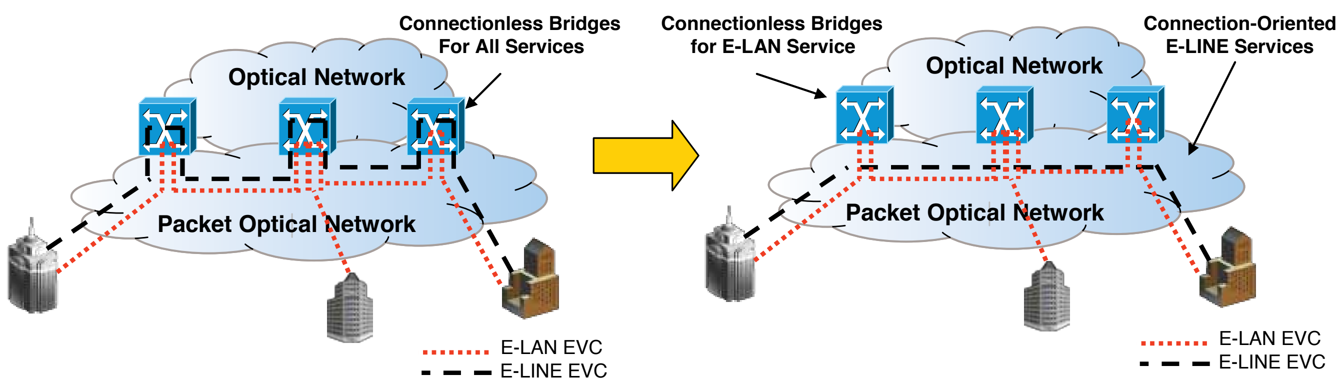

Figure 8: Offloading EVPL Services onto the Connection-Oriented Packet Optical Network

Conclusion

Service providers delivered enormously profitable TDM services because they had deployed a single common aggregation and transport infrastructure based on connection-oriented SONET/SDH technology. As demand for Ethernet continues to grow, and as service providers are faced with providing higher quality Ethernet networking for all applications, there is again a need for a universal aggregation and transport infrastructure that connects all types of end users to each other and to every type of higher layer service.

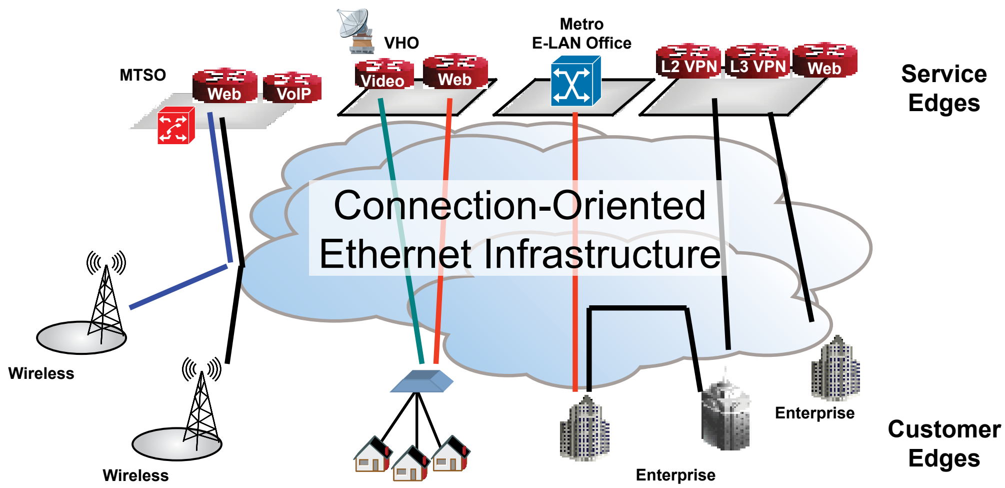

Figure 9: General-Purpose Connection-Oriented Ethernet Aggregation Infrastructure

This aggregation infrastructure, comprising Ethernet-only connection-oriented Ethernet, delivers the cost-effectiveness of statistical multiplexing and large-scale Ethernet aggregation together with the performance quality and manageability of SONET. It allows service providers to deliver the full range of network applications most economically for mobile, enterprise, and residential users.

Acronym	Definition
CAC	Connection Admission Control
CD	Collision-Detecting
CIR	Committed Information Rate
CoS	Class of Service
CPE	Customer-Premises Equipment
CSMA	Carrier Sense Multi-Access
DCS	Digital Cross-Connect
EIR	Excess Information Rate
E-LAN	Ethernet LAN
EPL	Ethernet Private Line
EVC	Ethernet Virtual Circuit
EVPL	Ethernet Virtual Private Line
IP	Internet Protocol
LTE	Long Term Evolution
MAC	Media Access Control
MEF	Metro Ethernet Forum
MPLS	Multiprotocol Label Switching
MTSO	Mobile Technology Switching Office
MTU	Maximum Transfer Unit
OAM	Operations, Administration and Maintenance
OTN	Optical Transport Network
Packet ONP	Packet Optical Networking Platform
PBB-TE	Provider Backbone Bridging with Traffic Engineering
PW-MPLS	Pseudowire MPLS
SDH	Synchronous Digital Hierarchy
SLA	Service-Level Agreement
SONET	Synchronous Optical Networking
TDM	Time-Division Multiplexing
UNI	User-Network Interface
WiMAX	Worldwide Interoperability for Microwave Access

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom