

Metro Ethernet Networks - A Technical Overview

Summary

The target audience for this comprehensive white paper is broad, including analysts and technologists unfamiliar with Ethernet deployment in the metro network and the work of the Metro Ethernet Forum.

The paper presents an overview of key issues pertaining to deployment of Ethernet in the metro arena, and addresses the following:

What is a Metro Ethernet Network (MEN)?

Why use Ethernet in the Metro?

What are the current limitations and possible solutions for using Ethernet in the metro?

The role of the MEF in addressing these limitations

What is a Metro Ethernet Network (MEN)?

A Metro Ethernet Network is generally defined as the network that bridges or connects geographically separated enterprise LANs while also connecting across the WAN or backbone networks that are generally owned by service providers. The Metro Ethernet Networks provide connectivity services across Metro geography utilizing Ethernet as the core protocol and enabling broadband applications.

Ethernet is a widely deployed cost-effective and well-known technology, and Ethernet interfaces are available on a plethora of data communication/telecommunication devices. Standards-compliant interfaces are available for 10/100/1000 Mbps and the standard for 10 Gbps Ethernet was ratified in the IEEE in 2002.

In Metropolitan Area Networks (MANs), Ethernet has the potential to cost-effectively increase network capacity and offer a wide range of service offerings in a scalable, simple, and flexible manner. An Ethernet based MAN is generally termed a Metro Ethernet Network (MEN). Some service providers have extended the MEN-like technology for the Wide Area Network (WAN) as well.

In enterprise networks, Ethernet currently has two key service applications that are garnering lots of attention and growth: Connectivity to the public Internet and connectivity between geographically separate corporate sites- LAN extension. The latter service extends the functionality and reachability of corporate networks, as shown in the following diagram.

The diagram and associated call-out text highlights the key areas of interest for a MEN, and sets the scene for the rest of this white-paper.




Figure 1

Links are primarily point-to-point and can be any speed of Ethernet.

Nodes can be either switches or routers, depending on:

- Their location in the MEN
- The nature of the services being provisioned
- The level of service resilience (network protection)

Nodes are meshed to whatever degree necessary to provide the desired connectivity, services and protection

WAN links connect MEN's together across large distances.

Ethernet Services can be topologically classified into either E-Line (point-to-point as shown), or E-LAN (multipoint-to-multipoint). Services can be further classified according to the bandwidth provisioned and which can be exclusive or shared across multiple users. Bandwidth can be provisioned on demand e.g. from 1Mbps to 1Gbps.

Varying degrees of service resilience is obtained by implementing a combination of network protection techniques which can be end-to-end (as shown) or nodal

Quality of Service (QoS) is obtained by using a combination of techniques, providing both 'hard' and 'soft' bandwidth and packet-loss guarantees. QoS can be end-to-end (as shown) nodal. QoS is visible from an enterprise end-customer perspective as a technical/operational Service Level Specification (SLS), which is underwritten by a commercial Service Level Agreement (SLA)

Why use Ethernet in the Metro anyway?

Today, 98 percent of all data traffic in all enterprise LANs start and end on an Ethernet port. It is the most dominant standard protocol in the networking industry and has 30 years of great history in the Enterprise LANs. As enterprises are looking for network connectivity beyond the walls of their LANs, Metro Ethernet becomes an obvious choice both a technical and cost perspective.

When corporate networks are connected to, or interconnected within a MAN, a bottleneck can be created. Indeed, the so-called 'bandwidth glut' of which analysts often speak is difficult to find in the Metro – as here, connections are dominated by often traditional T1/E1, T3/E3 or ATM links. Over the last decade bandwidth has increased over 300 fold in the backbone and 100 fold in the access, but only a meager 16 fold in the metro producing a significant Metro bottleneck. Enterprise customers are pushing service providers to connect their sites via metro networks, yet T1 lines often cannot provide the flexible metro bandwidth capacity that enterprises know Optical Ethernet can provide.

Deploying Gigabit Ethernet based platforms in the metropolitan areas is a compelling and commercially proven approach to break the metro bandwidth bottleneck for the following reasons:

Cost Effectiveness

Infrastructure equipment costs for Ethernet are significantly less than frame relay (FR) or ATM costs. This is due to:

- (a) The economies-of-scale arising from the existing installed base of Ethernet that ensures lower material and developments costs
- (b) The relative technical simplicity of Ethernet.

Further, the economics on Ethernet have held true over the 30 year lifetime of Ethernet – where the 'rule of thumb' has been: *"Roughly ten times the performance of the preceding generation, at about three times the cost."*

Provisioning costs (mainly operational and planning related) are also significantly less than for TDM (T1/E1, T3/E3, SONET / SDH) with comparatively effortless adoption and higher available data rates as an added bonus.

Rapid Provisioning on Demand

From a service provider perspective, service velocity is a key competitive differentiator. The present lack of customer-centric flexibility as well as the coarseness of bandwidth granularity of TDM and ATM legacy systems are seen to be major impediments to providing promising revenue generating services such as described in section 6. Today, the prevalent connection choices are DS1, DS3, OC3, etc - and the flexibility and fine granularity required by many IT managers running corporate networks, is often simply not available.

On the other hand, Ethernet access services offer a wide range of speeds, from 1Mbps to 1Gbps, in increments of

1Mbps or less, which can be provided on-demand and quickly – possibly even by the customer through a web-based tool.

Packet-based

Ethernet is an asynchronous frame-based technology that has particular flexibility advantages over its more rigid cell-based or synchronous competitors. With suitable rate-limiting functions to manage the available resources and with sufficiently large trunk capacity, Ethernet can provide rapid bandwidth-on-demand.

Ease of Interworking

Removing the inter-working or protocol translation issues between different platforms and environments makes service provisioning and activation simpler. In other words, a layer of complexity (e.g., ATM and SONET) is removed from the WAN access. This ‘plug and play’ feature of Ethernet that reduces configuration requirements also enables a simple migration path from low to high speeds. Consequently, it is relatively simple to integrate and interface end-customer IT systems to Ethernet metro services.

Ubiquitous Adoption

Ethernet has been the dominant technology in enterprise and campus LANs for many years. Standard interfaces are readily available for 10/100/1000/10000Mbps. The implications extend beyond economy-of-scale and include less-obvious benefits such as being easy-to-learn compared to the complexities of SONET and ATM, which impacts employee hiring capability and availability.

Many of these advantages are the result of the inherent simplicity of Ethernet. However, as is often the case, fundamental laws of conservation apply here too, and these benefits are not without their price. In the next section, the shortcomings of deploying Ethernet in the metro are defined.

What are the current limitations of using Ethernet in the MEN?

When speaking about Ethernet in the Metro, there is a justifiable tendency to draw comparisons with familiar ATM and Frame Relay services.

However, in spite of the following limitations, it is important to note that many incumbent and new-entrant service providers are presently providing Ethernet-based services on a profitable basis to a wide range of satisfied customers.

Using ATM and Frame Relay as a frame of reference then, one can identify the following limitations of ‘pure’ or ‘native’ Ethernet as a Metro transport medium. (Possible solutions to these limitations are discussed in section 5)

1) End-to-end QoS Guarantees

The key words here are ‘end-to-end’, because any intermediate element affects the customer perception of service at the end-point. Ethernet needs mechanisms for:

- *Connection admission for new service requests.*
i.e. At short timescales, can a demand request with a specific QoS requirement be admitted, without compromising the service performance of already accepted demands?
- *Scheduling and policing to maintain fair access*
i.e. At the packet level, can fair access to the available bandwidth be ensured in the face of competing demands, and can it be guaranteed during times of congestion?
- *Ensuring optimal path establishment through the network*
i.e. Path establishment via the spanning tree algorithm means non-optimal paths exist, which could introduce loss, jitter and delay.
- *‘Packet coloring’*
i.e. To be able to mark the packets for the purpose of prioritization, scheduling and policing for example, like DiffServ. (Although IEEE 802.1Q does make provision for three user priority bits)

2) Protection Mechanisms

Given its origins in the LAN space, Ethernet has traditionally had limited protection mechanisms in the event of network disruptions:

Slow failure recovery

Link breaks in an Ethernet environment are handled via the spanning tree algorithm, which can take up to tens of seconds to converge, depending on the size of the network. Compared to the 50msec automatic protection switching (APS) available in a SONET network, this is orders of magnitude too long for voice/video/mission critical applications.

Lack of fault isolation capability

Similarly, Ethernet has no 'built-in' alarms such as LOS, RDI etc found in SONET, which allow fault isolation to the malfunctioning section/line/path

3) In-service OAM

Ethernet has no overhead capability to monitor in-service bit-error rate (BER) such as the BIP-8 overhead bytes found in SONET. Such capability is particularly useful at service demarcation points, where it is used for loop-back and reachability testing.

4) Scalability and Network Resource Utilization

One of the advantages of Ethernet in an enterprise domain is the ability to logically partition distinct user groups over the same physical network, using the virtual LAN (VLAN) concept. Scaling VLANs to the Metro, where such user groups would be different companies poses challenges:

Limited VLAN tag space

The IEEE 801.Q standard defines an address space of only 4096 available tags. This is insufficient for a large service provider.

Spanning tree issues

A single spanning tree allows only one loop-free path. This implies uneven load distribution and potential bottlenecks. Multiple spanning trees are under consideration by the IEEE 802.1s working group in an attempt to address this shortcoming.

With such limitations, the reader will by now be wondering how Ethernet service is currently being successfully offered, given the preceding limitations. This is defined in the next section.

What are possible solutions to these limitations?

It is important to reiterate that both network equipment manufacturers and service providers have already found different and innovative ways around some of the preceding limitations. Each of these has specific implementations but broadly speaking, the following possible solutions exist.

1) End-to-End QoS Performance

One way to provide 'reasonable' service performance is via over-provisioning a best-effort service. Given the economics of Ethernet in infrastructure compared to say, legacy SONET, this may be attractive. However over-provisioning is not an economically viable long-term solution as the customer base expands. Multi Protocol Label Switching (MPLS) has the potential to be a long-term solution towards addressing some QoS limitations. This is because it:

- Provides Traffic Engineering capability (via the 'explicit route' object)
- Provides guaranteed-bandwidth 'pipes'
- Has an inherent packet coloring capability (via the EXP bits).
- Seamlessly runs on Ethernet (or any) transport (via insertion of the MPLS 'shim')

2) Protection Mechanisms

Slow failure recovery.

It is important to appreciate that different degrees of protection are warranted depending on the nature of the customer service and on the SLA terms. Several (possibly co-existent) technology options exist here and are being implemented in the industry and/or being proposed in the standards bodies

- *IEEE 802.1s (Multiple Spanning Trees)* is a proposal to allow more than one loop-free path in a VLAN environment. The benefit from a protection viewpoint is a redundant path, at the expense of management complexity.
- *IEEE 802.1w (Rapid Reconfiguration Spanning Tree)* implements a faster-convergence algorithm, taking about one second to converge (which may indeed be suitable for a particular class of service). However, this still doesn't approach the SONET-like restoration times of 50 msec which may be required for other classes of service.

- *IEEE 802.3ad (Link Aggregation)* provides sub-second fail-over on trunk groups. This approach has merit in certain applications, say, where geographic path diversity is needed. It also has the added benefit of allowing load-balancing techniques to be used.
- *The IEEE 802.17 (Resilient Packet Ring)* working group is advocating a ring protection protocol which gives sub-50msec failover.
- *MPLS* offers flexible and scalable solutions via capabilities such as backup LSP's, fast-reroute and LSP pre-emption. This flexibility paves the way to allowing Service Providers to sell different levels of protection at different price-points.

Lack of fault isolation capability

Currently, this aspect is being addressed as part of Ethernet OAM discussions in different standards bodies with recent discussions in ITU-T Q3/13.

3) In-service OAM

Currently, this aspect is being addressed in different standards bodies including IEEE, MEF and ITU-T Q3/13. IEEE 802.3ah Ethernet in the First Mile (EFM) working group is specifying Ethernet link OAM using a frame-based approach. MEF and ITU-T are addressing end-to-end in-service OAM for Ethernet. Some of the OAM functions being addressed include performance monitoring and fault management.

4) Scalability and Network Resource Utilization

Certain network equipment manufacturers currently supply equipment that increases the available VLAN tag space through various proprietary 'tag-stacking' schemes. This approach works to a degree, but does introduce network management complexity as well as interoperability considerations. For example, what does a non-tag-stacking switch do, when it receives stacked frames?

Given the limitations of spanning tree implementations, various MPLS techniques currently in draft within the IETF hold promise as long-term scalable solutions that maximize network resource utilization.

By now, the observant reader will have appreciated the need for an industry body such as the MEF to provide some alignment among the proponents of the various solutions. The role and current efforts of the MEF is discussed in the following section.


The role of the MEF in addressing these limitations

Compelling economic and operational reasons exist to deploy Ethernet in the metro and several possible solutions exist to address some of the identified technical limitations.

It should be evident that imperatives exist to:

- Ensure alignment and technical interoperability among the various implementations currently being used or considered
- Investigate ways to address those limitations receiving minimal or no attention

To facilitate these imperatives, the current MEF technical committee is currently focusing their work on five key areas of technical development work. These areas are: Management, Architecture, Protocols/Transport, Services and Testing. Each of these areas have sub-teams focused on delivering documents, frameworks, agreements and interoperability content to support the advanced adoption of Metro Ethernet as the network of choice on a global basis.



Architecture Area: Mandated to develop an architectural reference model and set of common linguistic tools for the technical teams

- **Services Area:** Mandated to define Ethernet services models, definitions, traffic management, service parameters and attributes
- **Protocols and Transport Area:** Mandated to define methods, procedures and protocols to support the defined Metro Ethernet services
- **Management Area:** Mandated to develop Metro Ethernet operations, administration and maintenance (OAM) requirements, models and definitions
- **Testing Methods:** Newly developed area within MEF to address the market demands for standardized interoperability and compliant solutions.

Note that the MEF is not a standards body, but an industry forum. It will draw on work developed in other standards organizations as far as possible, defining its own standards only as a last resort.

MEF technical work areas as at August 2003

The technical committee has organized its work under four key broad areas of technical development work. Considerable progress has been made in reaching consensus, as indicated through many approved projects moving to MEF working draft status. This section summarizes the primary approved projects:

1) Architecture Area Projects

To develop an architectural reference model and set of common linguistic tools for the technical teams, architecture area has made significant progress. The approved projects that have progressed to MEF working draft status include:

- MEF Architecture Framework: Part 1 Generic Framework
- MEF Architecture Framework: Part 2 Ethernet Services Layer
- UNI Requirements & Framework
- UNI Type 1 Implementation Agreement
- E-NNI Technical Specification

MEF Architecture Framework: Part 1 Generic Framework

The MEF Generic Architecture Framework describes the layer network decomposition model for a Metro Ethernet Network (MEN). The model reuses the native Ethernet frame structure and the architectural constructs created to describe connection and connection-less oriented transport networks in ITU-T Recommendations G.805 and G.809 respectively.

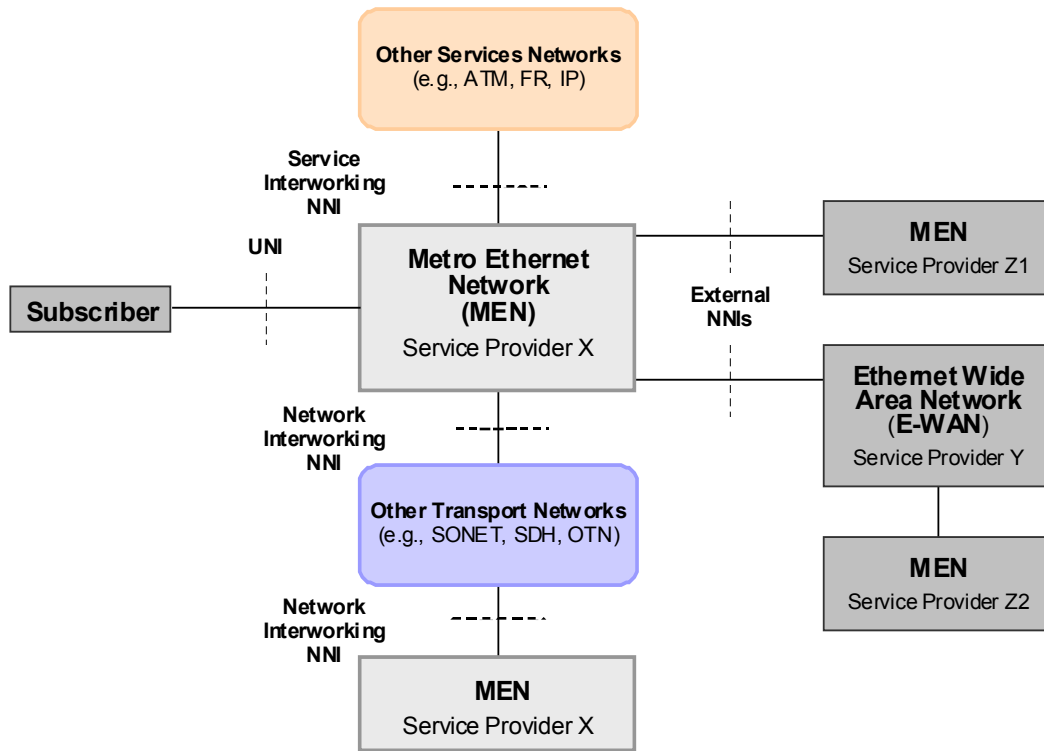


Figure 2

Figure 2 shows external architectural components to a MEN, their associated interfaces and reference points. External MEN elements include a) the subscribers to the MEN service, b) other MEN networks, and c) other (non-Ethernet) Transport and Service Networks. Subscribers connect to a MEN at a User-Network Interface reference point. Typically, internal network elements (NE) are interconnected via Internal Network-to-Network interfaces. Two autonomous MEN may interconnect at an external NNI (E-NNI) reference point. A MEN may interconnect with other Transport and Service networks at a Network Interworking NNI (NI-NNI) or a Service Interworking NNI (SI-NNI) reference point.

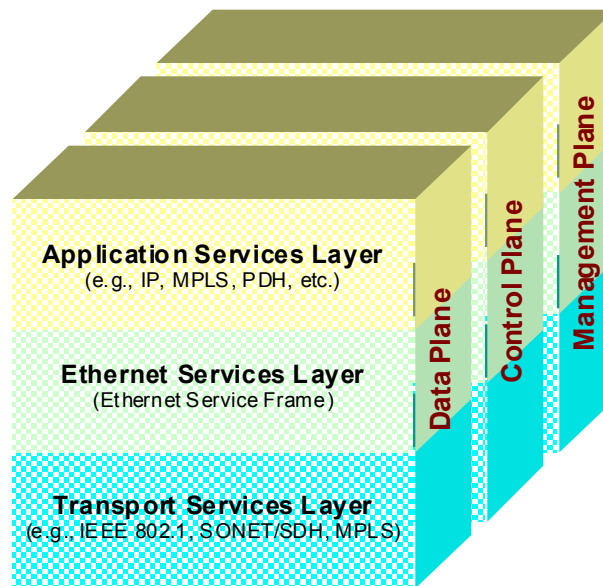


Figure 3

The MEN layer network model in Figure 3 defines the MEN in terms of three layered network components: the Ethernet Services (ETH) Layer supporting basic layer 2 Ethernet data communication services; a set of one or more supporting Transport Service (TRANS) Layer(s) and, optional Application Services (APP) Layer supporting applications carried on the basic L2 Ethernet services. In addition, each of these layer networks may be further decomposed into their data, control and management plane components.

MEF Architecture Framework: Part 2 Ethernet Services Layer

The Ethernet Services (ETH) Layer is a specific layer network, within a MEN responsible for the instantiation of Ethernet-centric connectivity services and delivery of Ethernet PDUs presented across well-defined internal and external interfaces. The ETH layer shall be responsible for all service aware aspects associated with interconnect, operations, administration and management capabilities required to request and support basic layer 2 Ethernet connectivity services.

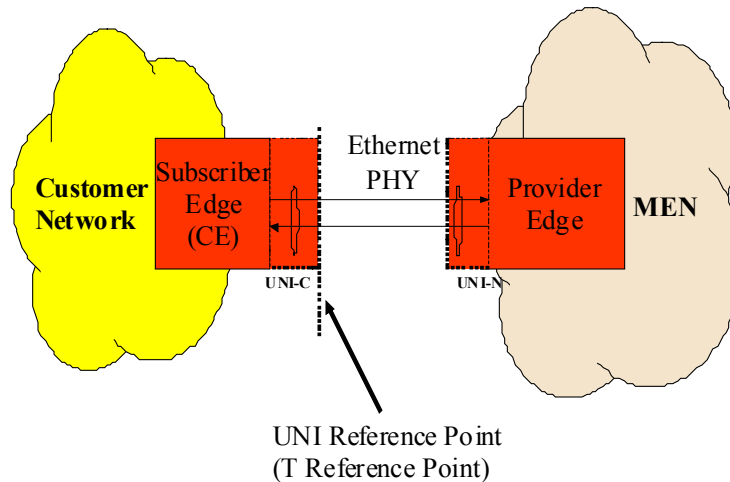
The Ethernet Virtual Connection (EVC) is an important logical construct in the ETH layer used to enable end-to-end subscriber services instances across one or more MEN service providers. An EVC represents a single instance of an ETH layer service. Subscriber flows may be mapped into one or more EVCs though subscriber flows corresponding to a single service instance get mapped into a single EVC.

UNI Requirements & Framework

UNI requirements are specified for three phases of UNI deployment. The first phase of the UNI, called, UNI Type 1, focuses on existing Ethernet deployment with UNI-Cs that require no changes and use existing IEEE Ethernet PHY and MAC functionality. UNI Type 2 requires static service discovery functionality with auto-discovery and OAM capabilities. UNI Type 3 requires dynamic connection setup such that EVCs can be setup and/or modified from a UNI-C.

UNI defines a reference point for all interactions between subscribers of MEF defined services and the MEN service provider. The UNI reference point (also called as the T reference point) provides the demarcation between a subscriber and the MEN boundaries.

UNI Reference Point



The UNI functionality is split between the subscriber’s equipment (UNI-C) and the MEN equipment (UNI-N). The UNI-N may be separated from the T reference point by a transport access network other than a directly attached native IEEE 802.3 PHY. The UNI Functional Model consists of three planes:

Data plane: Defines the means of transporting information across the UNI reference point.

Control plane: Defines the means for the subscriber and the MEN provider to make use of the data plane.

Management plane: Controls the operation of the UNI data and control planes

UNI Type 1 Implementation Agreement

UNI Type 1 IA specifies how MEF UNI Type 1 will operate in manual configuration operations. This allows existing Ethernet devices (switch, router, workstations, etc) acting as CEs to be instantly compliant to this IA with no additional

software or hardware upgrades. This IA provides bare minimum connectivity services with no control plane capabilities and a few data and management plane functions.

E-NNI Technical Specification

The purpose of E-NNI specification is to develop the requirements for the E-NNI, an interface between two MENs where each MEN is under the control of a distinct administrative authority. It is also intended to specify sufficient details of the protocol operating at the E-NNI to ensure interoperability.

A component Ethernet Virtual Connection (CEVC) is a building block for multi-MEN EVCs. While EVC is defined as an association of two or more UNIs, a CEVC is defined as an association of at least one UNI and at least one E-NNI, or an association of two or more E-NNIs attached to a single MEN.

2) Services Area Projects

One of the primary priorities of the MEF is to define services for Metro Ethernet Networks. In line with this priority, the Services area has also made significant progress with the following projects moving to MEF working draft status:

- Ethernet Services Model - Phase 1
- Ethernet Services Definitions – Phase I
- Traffic Management Specification - Phase I
- CES Requirements
- PDH Implementation Agreement

Ethernet Service Model - Phase 1

Ethernet Service Model specifies a model for specifying Ethernet services, which are modeled from the point of view of the subscriber’s equipment (CE) that is used to access the service. Besides the basic elements of Ethernet services, a number of service attributes are defined that may be offered as part of an Ethernet Service including the definition of Service Level Specification (SLS). It also specifies an Ethernet Service framework that provides the definition and relationship between attributes and their associated parameters used to create an Ethernet Service.




Figure 4

The scope of Phase I is limited to those services that are based on Ethernet, and where from the subscriber equipment point of view, protocol operating at the UNI between CE and MEN is a standard Ethernet protocol (PHY and MAC layers). The services are limited to those between two or more UNIs.

Ethernet Services Definitions

Ethernet Services Definitions use service attributes and parameters defined in Ethernet Services Model – Phase 1 and Traffic Management Specification – Phase 1, to define two generic services Ethernet Service types. These service types along with their associated service attributes and parameters can be used to create point-to-point and multipoint-to-multipoint services. The services are instantiated at an Ethernet UNI, are agnostic of the underlying network infrastructure and may be supported over different transport technologies in the service provider’s network such as SONET, RPR, WDM, MPLS, etc.

The two Ethernet service types, namely, Ethernet Line (E-Line) and Ethernet LAN (E-LAN), are shown in Figure 5 and 6.




Figure 5




Figure 6

E-Line provides a point-to-point EVC between two UNIs (subscribers). E-LAN provides a multipoint-to-multipoint EVC between two or more UNIs. Subscriber data sent from one UNI can be received at one or more of the other UNIs.


Traffic Management Specification - Phase 1

Traffic Management Specification defines the traffic and performance parameters that may be specified as part of an Ethernet SLS. These parameters are defined to enable the subscriber to quantitatively and qualitatively measure the service being provided. Consistent use of these parameters in the SLS will allow subscribers to compare various offerings. Additionally, service providers may use these parameters as guidelines for specifying CoS-based Ethernet services.

The scope of Phase I is to provide sufficient technical details for network equipment vendors and service providers to implement the minimum set of traffic management capabilities to support CoS-based SLSs for Ethernet Services.

CES Requirements

Circuit Emulation Services (CES) is the emulation of TDM circuits over a MEN. CES is a useful technique to allow MEN service providers to offer services to customers who are outside the reach of the MEN itself or who already have TDM-based equipment. As a result, MEN service providers can extend their reach and addressable customer base. As an example, the use of CES enables metro Ethernet transport networks to connect to PBXs on customer premises and deliver TDM voice traffic along side of data traffic on Ethernet.


Figure 7

CES requirements cover PDH services (e.g. Nx64 kbps, T1, E1, T3, E3) and SONET/SDH services (STS1, STS3, STS3c, STS12, STS12c and European equivalents). References are made to specifications produced by other standards organizations (notably the ITU, ANSI, IETF PWE3 and ATM Forum), which are adapted to address the specific needs of MEN.

PDH Implementation Agreement

The scope of PDH Implementation Agreement is limited to a) the essential agreements needed to create inter-operable equipment to reliably transport PDH circuits across MEN, and b) the required performance of the underlying MEN to enable the provisioning of circuit emulated PDH services that meet the existing TDM standards, as defined by ITU-T and ANSI.

3) Protocols & Transport Area Projects

The projects under the Protocol and Transport area that have progressed to MEF working draft status include:

- Protection Requirements
- Framework for Metro Ethernet Protection
- MPLS Protection Implementation Agreement
- TMF Technical Specification
- Ethernet Interworking Function (E-IWF)
- EoS Implementation Agreement

Protection Requirements

The Protection Requirements provide requirements to be satisfied by the protection and restoration mechanisms in Ethernet services-enabled MENs. The requirements have been based on the interpretation of SLS for Ethernet services (e.g. availability, mean time to restore, mean time between failures, etc.) to reflect into network protection requirements (e.g. connectivity restoration time, protection resource allocation etc.) This mandates that the protection offered by the network is directly related to the services supplied to the user and the requirements are derived from the need to protect the services provided to the user.

The scope of these requirements includes protection control and signaling, QoS during protection, interaction of such mechanisms with transport and application specific protection mechanisms etc.

Protection Framework

To deliver protection to Ethernet services implemented over MEN, a reference model has been created to allow description of a unified protection structure. The Protection Reference Model (PRM) allows a consistent description of protection capabilities applied on these services across various transport technologies, network topologies and policies, thereby enabling the description of protection of ETH layer services.

The elements are described in a “bottom-up” approach. Each layer may contain protection capabilities, and may run above a lower layer that might contain protection capabilities as well. The entire protection scheme is controlled according to the application protection constraint policy.

The following styles of network protection mechanisms are currently under consideration:

- Aggregated Line and Node Protection (ALNP) service

- End-to-End Path Protection (EAPP) service
- MP2MP protection service
- Link Protection based on Link Aggregation

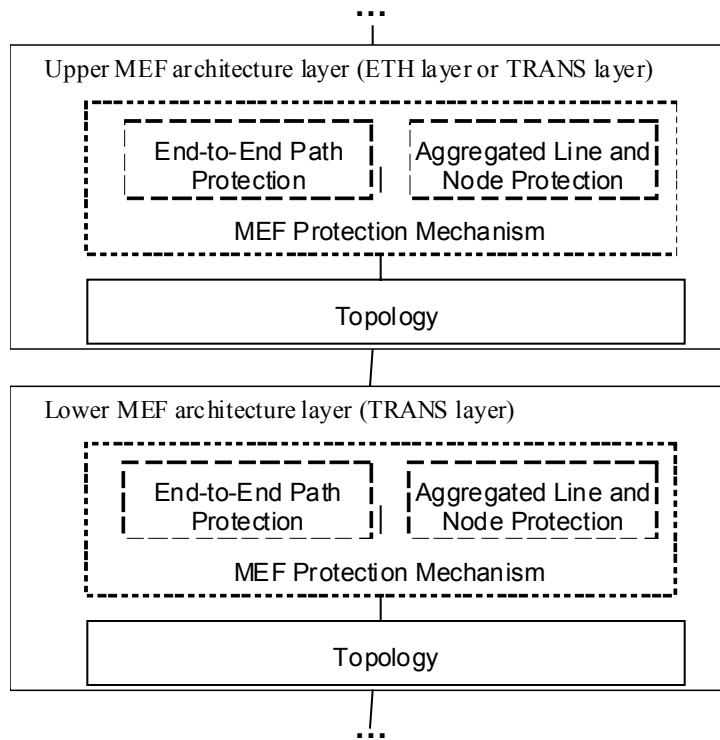


Figure 8

Transport Multiplexing Function (TMF) Technical Specification

TMF is a functional element within the TRANS layer of the MEF architecture. The purpose of the TMF is to provide a mapping of multiple individual UNI links to one multiplexed TRANS Link as shown in Figure 9. The individual component links comprise a stream of Ethernet Frames received from (and transmitted to) a subscriber/customer endpoint termination (CE).

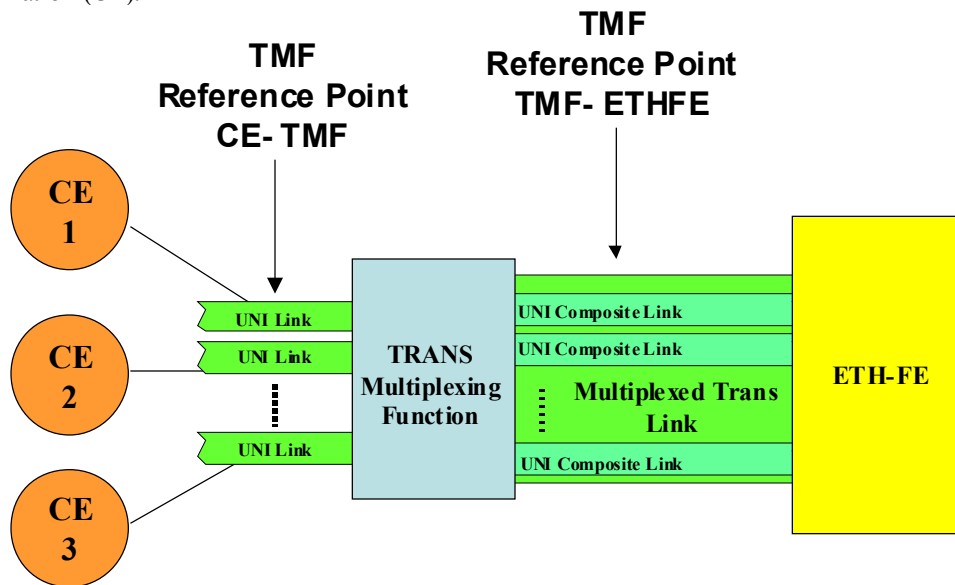


Figure 9

TMF provides a generic access link aggregation function. As such the TMF may be used with any type of transport facilities present in the MEN.

Ethernet Interworking Function (E-IWF)

Ethernet Interworking Function is defined as inter-working between Ethernet and other types of layer 2 circuits such as ATM and FR. This document addresses how new Ethernet services can be augmented to existing topologies built with ATM and FR technologies for Layer 2 VPN services. E-IWF’s scope is limited to point-to-point EVCs that are cross connected to virtual circuits of type ATM or FR. The multipoint EVCs are considered for interworking functions only when the payload is in the bridged Ethernet format over ATM or FR circuits.

4) Management Area Projects

This area is addressing specifically identified OAM&P (Operations, Administration, Maintenance and Provisioning) issues in MENs. Projects that have moved to draft status include:

- EMS Requirements
- EMS-NMS Information Model
- Ethernet Services OAM
- Performance Monitoring

EMS Requirements

EMS Requirements identifies the requirements for EMS(s) deployed by MEN service providers. An EMS is a system that can be used to manage one or more network elements (NEs).

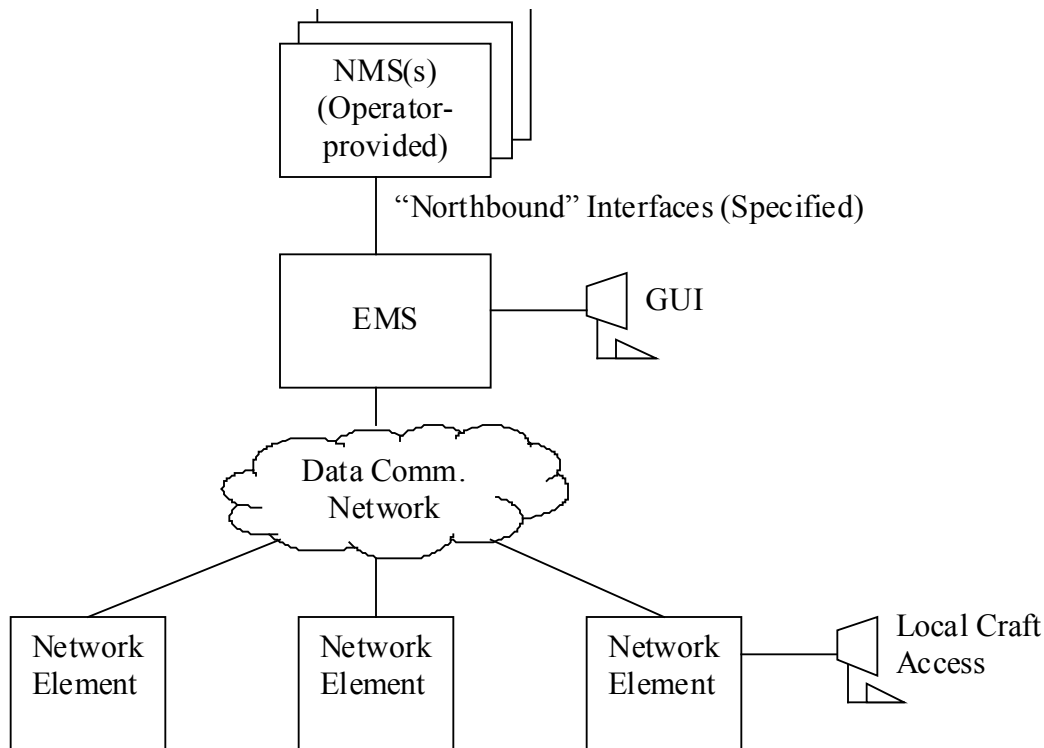


Figure 10
EMS Architecture

EMSs are located in network operations centers, remote from the network elements but connected to them by a data communications network. Since the cost and time required sending a technician to the NE site is burdensome, EMS is a service provider’s primary access for managing NEs. Streamlining network management functions requires the EMS to be integrated with a service provider’s Network Management Systems (NMSs). The requirements are presented with use cases and are divided into management functional areas, e.g. configuration, fault, performance, security, and accounting management.

EMS-NMS Information Model

The EMS-NMS Interface Model defines a protocol-neutral model that allows for consistent definition of the management information required to manage MENs. The protocol neutral model can be used as a basis for defining the

specific EMS-NMS management interface using a well-defined method such as CORBA, IDL, SNMP, JAVA, XML, etc.




Figure 11

The model defines a set of relationships in UML format and managed object descriptions, and follows the network view paradigm defined by the ITU-T (e.g., G.805, G.809, G.85x, and M.3120). Information content of the model makes use of MIBs defined in the IETF for Ethernet and MPLS management.

Ethernet Services OAM

Ethernet services can be offered over multiple types of transport using a variety of tunneling technologies. In all such layered models, it is important to provide basic OAM capabilities in each layer of the hierarchy. Ethernet Services OAM addresses the OAM functionality in the Ethernet Service (ETH) layer, which remains independent of the underlying TRAN layer(s), each of which may have its own OAM capability.

The requirements of OAM functions for the ETH layer focus on monitored parameters e.g. connectivity, delay, delay variation (jitter) and status monitoring. The Ethernet service interface is considered to be the OAM source and termination of ETH layer OAM. In particular, the Ethernet service interface on each device is assumed to have a MAC address that can be used for OAM packet addressing.

Performance Monitoring

Ethernet Performance Monitoring (PM) functionality addresses the need to monitor and collect performance data for various purposes such as monitoring the MEN state and fulfilling the SLS requirements. The required PM counters are defined that are necessary to: a) monitor SLS compliance, and b) monitor the MEN state using the EVC level PM counters and ETH layer PM counters. In addition a framework is provided for collecting the PM counters according to various methods.

5) Test Methods Area Projects

As the various Technical Documents mature, this catalyzes the work in the Test Methods Area. The first of several planned work items in this area is entitled “Requirements and Test Procedures for Network Devices Delivering Ethernet Services”

This document defines the requirements and corresponding test procedures that determine the readiness of a networking device to deliver various Ethernet Services, such as Ethernet Line (E-Line) and Ethernet LAN (E-LAN) services. The requirements defined in this document are based on the MEF Ethernet Services Model, Ethernet Services Definitions and Traffic Management documents described in an earlier section of this white paper.

The document does not define methods for the measurement or monitoring of Service Level Specifications (SLSs) but it does define test procedures, which determine the readiness of a device to deliver various Metro Ethernet Service related SLSs.

The corresponding test definitions provide vendors, carriers, subscribers and laboratories with a common well-known set of criteria to qualify the readiness of a device for service on a Metro Ethernet Network.

Various types of criteria can be used to establish the requirements a device must satisfy to deliver Ethernet services. The document refers to four such criteria: functional, conformance, interoperability and performance.

Functional criteria: A device will have to satisfy certain well-defined functional criteria some of which can be safely taken for granted and others which it may be useful to test. An example of the former would be that a device can be fully configured through a Command Line Interface and, of the latter, that a device correctly filters BPDUs (Bridge Protocol Data Units) when configured to do so.

Conformance criteria: A device must also satisfy certain well-defined conformance criteria. An example of this would be the ability of a device to never forward a frame to its source address.

Interoperability criteria: A further criteria for such devices is interoperability. An example of this would be the ability of a device to participate on the same arbitrary VLAN as the device of another vendor or, more generally, of other vendors.

Performance criteria: A device must also meet certain performance criteria. An example of this would be the ability to successfully forward a certain percentile of frames without loss.

Specification Format and Example

A well-defined template is used to describe the intended tests. This specification directly maps to relevant MEF Services Specifications, and consequently is easily extensible as the MEF work unfolds. A typical example using the test template is shown in Table 1 below

Name	EVC leakage
Test Definition ID	M.6-2
Source	Ethernet Services Model, Phase 1, Section 6
Type	Conformance
Status	Mandatory
Description	A device MUST NOT deliver frames to a UNI which is not in the EVC
Type of DUT	MEN
Object	Determine if a device forwards frames to a UNI which is not in the EVC
Configuration	System test bed
Procedure	Tester offers frames with destination MAC addresses of a DUT which is not in the EVC and monitors if they are forwarded to it
Units	Number of frames
Variables	Unicast, multicast, broadcast, frame lengths, port count
Results	Pass or Fail
Remarks	

Table 1

Anticipated Future Work

The scope of this area will unfold commensurate with the rest of the technical committee progress, and where needed would include test methods for topics such as Services protection and OAM

Conclusion

It should by now be evident that the MEF has clearly targeted the limitations of Ethernet as a carrier class service technology, and is addressing them in a coherent fashion. This is being done by making use of existing technologies wherever possible, and defining new approaches only when necessary.

Disclaimer

This paper reflects work-in-progress within the MEF, and represents a 75% member majority consensus as voted by the 70+ members in the August 2003 Denver Technical meeting. Some technical details may change in due course (by 75% vote) and this paper will be updated as deemed necessary to reflect such changes. The paper does not necessarily represent the views of the authors or their commercial affiliations.

References

“Metro Ethernet Services – A Technical Overview”, <http://www.metroethernetforum.org/metro-ethernet-services.pdf>

Acknowledgements

This paper has drawn upon a large number of contributions presented to the MEF technical committee and the efforts of the respective authors, reviewers and technical writers are acknowledged.

Authors

Mark Whalley	Dinesh Mohan
Agilent Technologies	Nortel Networks
Australia	Canada

Glossary

ALNP – Aggregated Line and node Protection
APCP – Application Protection Constraint Policy
ATM – Asynchronous Transfer Mode
BIP-8
BLSR – Bidirectional Line Switched Ring
CE- Customer Edge (MPLS VPN context) / Customer Equipment (generic context)
CES – Circuit Emulation Services
CIR – Committed Information Rate
EPP – End-to-end Path Protection
EFM – Ethernet in the first Mile
EMS- Element Management System
EVPL – Ethernet Virtual Private Line (Service)
EPLn – Ethernet Private LAN (Service)
IETF – Internet Engineering Task Force
IEEE – Institution of Electrical and Electronics Engineers
IP – Internet Protocol
IWF – Interworking functions
LAN – Local Area Network
LOS – Loss of Signal
LSP – Label Switched Paths
MAC – Media Access Control
MAN – Metropolitan Area Network
MBS – Maximum Burst Size
MEF – Metro Ethernet Forum
MEN – Metro Ethernet Network
MPLS- Multi-Protocol Label Switching
NE – Network Element
OAM – Operations, Administration and Maintenance
OTN – Optical Transport Network
OTUk - Optical Channel Transport Unit of level k
PDH – Plesiochronous Digital Hierarchy
PE – Provider Edge
PIR – Peak Information Rate
QoS – Quality of Service
RDI – Remote Defect Indicator
RPR – Resilient Packet Ring
SDH – Synchronous Digital Hierarchy
SLA – Service Level Agreement
SLS – Service Level Specification
SONET – Synchronous Optical Network
UPSR – Universal Path Switched Ring
UML – Unified Modelling Language
UNI – User-Network Interface
VC – Virtual Channel

VLAN – Virtual LAN
WAN –Wide Area Network