Telco Systems
A BATM Company

Making metro Ethernet services carrier class

A WHITE PAPER FROM TELCO SYSTEMS

Telco Systems
A BATM Company

Contents

There are many compelling reasons for deploying an Ethernet service in the metro access network – not the least of which is its suitability for converging video, voice and data applications over a simple, cost-effective infrastructure. So it is not surprising that service providers are deploying Ethernet to enable new services and create new revenue streams.

In order to achieve carrier class Ethernet service performance, operations, administration and maintenance (OAM) functions must be readily available to ensure QoS and meet provisioned SLAs. Realizing the potential of Ethernet as a service toward the customer and as a transport in the access network requires that we address the key issues of intelligent demarcation in separating the service provider network from the customer and delivering and managing carrier-class Ethernet services to the customer edge.

The challenge

The same services that have long been available in the telephony world are now needed in the Ethernet access network, which is no longer just a metro access technology, but has become a wide-area LAN with converged packet-based voice, data and video services, spanning national and global boundaries. LAN-type management tools are simply not adequate to the task of addressing the security, performance and access control issues across the multiple administrative domains of the wide area and its access networks.

Management requirements for a packet-based network as a whole and an Ethernet access network in particular closely conform to the high uptime and guaranteed service levels of voice networks. So, Ethernet, as the transport of choice for network access, must now play "catch up" in order to meet the exacting performance, reliability and management standards that voice customers have long demanded and received. Managing a stream of packets carrying various types of traffic end-to-end across diverse networks is becoming a challenge. So, a key place to implement management is at that critical demarcation point where the service provider meets the customer edge.

Without comprehensive networking solutions and management of the operational and administrative aspects of both transport and services, Ethernet will not be as easily deployed or as well managed as comparable services in traditional carrier networks. Nor will it be as reliable and cost-effective for service providers and their subscribers.

Providing clear separation – the Ethernet demarcation device

The concept of providing a clear demarcation or separation point that marks the divide between the service provider or carrier network and the end user is not new. Such a separation has long been available as part of the function of the TDM CSU/DSU (channel services unit/data service unit). The CSU and DSU perform separate functions, but are nearly always co-located in the same physical unit, since both are needed to connect to the provider network.

Ethernet demarcation devices (EDD) are physically located at the subscriber site, but are owned by the service provider. These devices can perform the link layer functions like physical connectivity and diagnostics on both sides of the demarcation point. Acting as a Layer 1 Network Interface Device, or NID, the EDD includes performance monitoring, as well as troubleshooting tools like Dying Gasp, remote failure indication, remote loopback and cable integration testing as defined by the IEEE 802.3ah standard. The device can also test the connections to the subscriber's network without dispatching a technician using a "virtual cable test". But, not only does the EDD need to support transport administration – perform loopbacks, collect PM stats, verify the remote device's health and status, it also needs to provide standards- based service administration (IEEE 802.1ag and ITU-T Y.1731). These standards enable the device to be service-aware and manage subscriber traffic flows to assure service levels for different business models and operate in different markets.

For example, in an out-of-franchise configuration, where service providers lease the access network from another provider, there is a non-negotiable requirement to transparently tunnel subscriber traffic over the access provider's network, while still being able to manage the service and access networks. Given that the service domain spans edge-to-edge, or demarc-to-demarc, Ethernet demarcation systems must provide specific OAM tools to allow delivery of carrier-class, revenue-enhancing services. They, therefore, not only must support interoperable OAM for transport, connectivity and service delivery, but also provide the ability to map into different deployment models.

Differentiated carrier-grade services

There are several major areas in which providers seek to differentiate their offerings in terms of the carrier-level services they deliver—ranging from ease of installation and deployment to a full suite of networking features and comprehensive end-to-end management of all the Ethernet services traversing a sophisticated multi-carrier network. Overall, the importance lies not in the supported function, but in how well each function is implemented.

Like the services available on voice networks, Ethernet services need to be measured to ensure that they meet the high availability, reliability and performance standards required for the end-to-end voice, video, data or converged service on which users rely. And the ability to guarantee certain levels and quality of service depends on the provider's ability to measure and assess the capabilities of equipment at all points within the service infrastructure. Tools are available today which enable the provider to define the service level capability across the network, provision the service, and monitor the service performance to determine if the agreed SLA is being met.

Ease of deployment

Nothing hinders the adoption of new technologies more than costly, time-consuming installation procedures. And nothing impacts an ongoing maintenance operation more negatively than the inability to isolate and correct faults quickly and easily without truck rolls.

Support for the IEEE 802.3ah OAM protocol for fault isolation is a key component in the initial network deployment and testing. Adding support for carrier-oriented, domain-specific DHCP (Dynamic Host Configuration Protocol) for auto-discovery and dynamic assignment of IP addresses will significantly shorten configuration times, while reducing the overall complexity of the operation.

Once network elements have been provisioned as part of service creation, it is essential that such services be tested and validated during service turn-up. While external test heads may be used, external test heads disrupt existing services. Network-based test heads cannot ensure that the test/qualification traffic follows the same path as would actual customer traffic.

Embedded test heads on demarcation devices ensure that service validation and assurance is conducted end-to-end and within the actual subscriber data path. These embedded test heads interwork with loopback entities to ensure that not only can assurance be performed over different Ethernet Virtual Connections (EVC), but also for specific customer traffic transported inside those EVCs.

Once validated, services are monitored to ensure that they are performing to their agreements with respect to bandwidth, jitter and frame loss, and that the network continues to provide the underlying performance and networking capabilities.

Edge-to-edge service level verification

With continuing network convergence of voice, video and data applications, both transport-layer and application-layer SLA maintenance is critical to business-class service offerings and user service differentiation.

Edge-to-edge service level verification means that the service provider can provide measurement of conditions in the network including:

- Delay, jitter and bandwidth between network components in order to validate SLA capability of the network prior to initiation of the service offering as well as ongoing verification after the service has been deployed
- Fault detection and isolation to identify the occurrence and location of a problem. Fault detection and isolation ensures that truck rolls are dispatched to the right place for the right reason
- Performance monitoring to create a proactive response before the network experiences problems

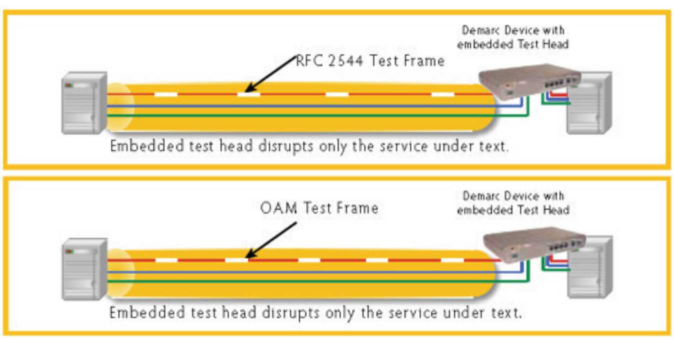

Embedded test head using industry standards (MEF, IEEE and ITU) enables non-disruptive test for frame delay, frame delay variation and throughput

Edge-to-edge service level verification requires that the connectivity layer, the transport layer, and the services layer all work together.

Service providers who own the network end-to-end may be able to guarantee service levels over the network that they own, however, once the service enters someone else's network, such guarantees become more difficult to determine. The goal of edge-to-edge service level verification, therefore, is to be able to verify and maintain SLAs across multiple providers.

An intelligent demarcation device can solve the problem of service- can solve the problem of service level verification across heterogeneous networks by ensuring that the network, transport and services layer are all under common management and can communicate end to end.

Managing services across the multi-provider network

Functions like intelligent services loopback can provide end-to-end service-level verification across multi-provider networks. Logical services loopback (LSL) provides a transparent capability for a network element to loop traffic back to the sender. By monitoring frames sent to a pre-determined MAC address, logical services loopback swaps source and destination MAC addresses in incoming frames and forwards them back towards the sender. Running at line rates – something that prevents gratuitous traffic loss – LSL directs traffic back to a test initiator over an interface, a VLAN, an EVC or at a certain traffic class. Intelligent loopback implemented in the demarcation device will allow the primary provider of the service to monitor for accessibility, CoS agreements, round-trip delay and delay variation.

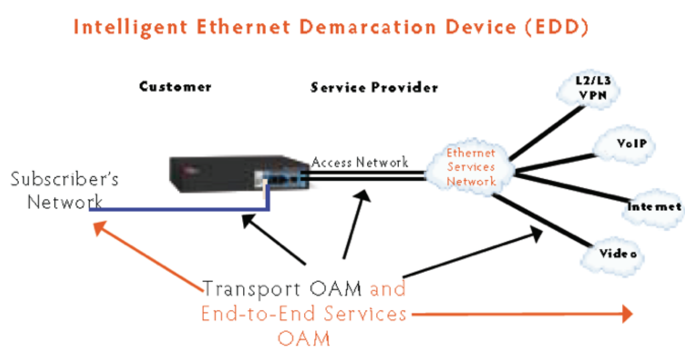

Intelligent EAD performs functions of both the transport OAM (NID) and Services OAM (UNI)

- IEEE and MEF Services Compliant
- Customer-Carrier Networking Interworking
- Traffic Engineering: Policing, COS
- Built-in Management Access and Security Controls
- Extensible Management Options for Network
- Multilayer OAM: Dev ice Connectivity, Link Integrity, etc.
- SLA Management
- Designed for NEBS Level 3 Compliance

LSL also provides the added capability of inserting a unicast MAC address after the original source and destination address have been swapped. This feature allows frames to be sent to a multicast address, thereby allowing configuration fault management capabilities and element discovery.

Driving standards to support intelligent demarcation

Carrier-quality services require sophisticated management tools, and good management requires good measurement. Ensuring adequate service performance levels — and improving on them over time — starts with meaningful measurement and analysis. And it's not meaningful if it's not standardized and measured in the same way regardless of the vendor, the transport or the protocols in play as data traverses the network.

Providers of a service need, first and foremost, a clear view of how that service is performing end-to-end in heterogeneous environment. In order to meet the terms of their service-level agreements, they must have accurate and consistent data on the reliability and throughput/ delay and characteristics of the devices and network connection over which their services run.

The ability to run standardized tests on multiple occasions at different times of day can provide valuable trending information that can be used for capacity planning, improved bandwidth management and writing better SLAs.

Standards organizations like the IEEE and IETF, as well as vendor consortia like the Metro Ethernet Forum (MEF), are working hard to define Ethernet services and the mechanisms to support the converged video, voice and data applications over a simple, unified infrastructure.

Remote monitoring, troubleshooting and remote site management are available today through the support of industry-standard management access protocols like SNMP v1/v2/v3, PING and IEEE 802.3ah – but they only address the link or transport-layer management functions.

As a transport OAM, IEEE 802.3ah provides link management and operates in a point-to-point fashion. It provides for fault detection and isolation and neighbor discovery. Since the 802.3ah OAM protocol is only a transport standard, services management is left to each vendor's discretion for implementation in propriety extensions.

Fortunately, work in the MEF, the IEEE, and the IETF promise to help extend the management interface for the OAM protocol, so that there is interoperability at the management level from different vendors.

As IEEE 802.1ag, ITU-T Y.1731 and the MEF Services OAM standards evolve and converge, we expect one day to see any or all of the following features.

Converging standards

- RFC 2544 foundation
- Addressing
- Discovery
- Connectivity
- End-to-end performance measurement
- Multihop and multipath measurement capability
- Performance monitoring
  - Frame loss
- Jitter
  - Average latency
- A consistent data model
- Topology isolation
- Logical
- Physical
- Service

The test head function as defined in RFC 2544 provides carriers with a certain "comfort level" by defining well-known procedures familiar to carriers – and methodologies that they know and trust. And long accustomed to having remote access to service performance data from their circuit network management tools, they now require the same sophisticated access to management data for their new Ethernet services.

Today's work in RFC 2544-compliant testing provides a sound procedural foundation for further developments and closer coupling with other evolving standards for measuring and managing performance in the delivery of Ethernet-based services. It is important to understand where each standard fits and how each relates to the others.

IEEE 802.1ag Connectivity Fault Management (CFM)

Transport OAM is complemented by connectivity OAM that is being standardized as IEEE 802.1ag. The IEEE 802.1ag Connectivity Fault Management (CFM) is an initiative that describes a methodology to isolate the exact point of failure in multi-carrier networks, enabling the end-to-end (rather than just link-to-link) management of connectivityand services. CFM introduced the concepts of domains and supports autonomy for customers, providers, operators, etc. Multiple domains can be logically-integrated, or each domain can run its own OAM. For example, a provider can run its CFM OAM and isolate a problem to a single operator. The operator can then isolate the problem in its own network via its own CFM OAM. Such a capability is critical to keeping the costs of "virtual" truck rolls to a minimum, since it isolates which network provider in a multi-provider scenario owns the point of failure.

CFM defines messages and protocols across entire networks, and is also responsible for discovery and connectivity activities. In addition to finding maintenance end points and maintenance intermediate points (MEPS and MIPS), this standard defines the link trace and loopback functions.

ITU-T Y.1731 Ethernet OAM for fault and performance management

ITU-T Y.1731 Ethernet OAM for fault and performance management deals with fault management (in a manner similar to that defined in 802.1ag) and layer-2 performance management (using a methodology similar to that described in RFC2544); it also defines mechanisms for testing frame loss and delay. This standard adds timestamps to the test packets defined by IEEE standard.

ITU techniques go beyond the testing defined in 2544, by allowing the measurement of delay, delay variation and frame loss using real user traffic, instead of just using simulated user packets as is the case in 2544-based testing.

The ITU standard and the MEF both use RFC 2544's underlying mechanisms, but not the 2544 test procedures themselves. Message packets are defined by the IEEE and ITU standards and are not addressed by the MEF initiative, which simply requires that the two-way delay be computed.

Metro Ethernet Forum (MEF) services OAM initiative

Metro Ethernet Forum (MEF) Services OAM initiative defines end-to-end OAM for service assurance. (See also MEF 10.1 which describes service performance attributes and MEF 9 and 14 which define testing methodologies for those Ethernet services.) This initiative also defines policies, but not the message packets or the test mechanisms to be used.

Eventually we anticipate seeing the foundation established by the message and protocol definitions in RFC 2544 evolving and being absorbed into the IEEE and ITU standards, which will define those first two requirements, while the MEF Services OAM will address standardized policy definitions to ensure the interoperability that today's proprietary implementations preclude.

This higher level of integration in the measurement and management of Ethernet services and procedural standardization at its core will broaden the appeal of Ethernet as a service-delivering technology and enhance the ability of Ethernet access devices to deliver carrier-quality, multiservice access to the customer edge.

Conclusion

The ability to accurately analyze and manage services more efficiently enables providers to say, with confidence, that these services are truly carrier-class. Standardized tools in the Ethernet access network enables service providers to leverage the ubiquity of Ethernet to provide service scalability and end-to-end service level agreements (SLAs). The ability to measure, manage and support multiple services and multiple service levels will allow providers to differentiate their offerings in a very competitive market.

The Telco Systems Carrier Ethernet Solution

Telco Systems Ethernet demarcation solutions offer a rich management toolset to reduce operational expenses (OPEX) and scale management to large access networks. Multi-level Operations, Administration and Maintenance (OAM) is used to measure and ensure provisioned SLAs, and security controls are embedded to ensure protection against denial-of-service attacks.

Telco Systems' T-Marc™ is a family of cost-effective, fully-managed carrier Ethernet demarcation devices that provide service termination and demarcation over service providers' packet-based networks. As a multi-port customer-located intelligent demarcation device, the T-Marc delivers managed converged services (voice, video and data) over virtual Ethernet in a metro Ethernet network. The T-Marc allows service providers to drop multiple services on separate customer interfaces. Because each service is isolated, providers can troubleshoot an individual service without impacting another.

T-Marc converges up to 4 digital T1/E1 TDM trunks and 10/100/1000 copper or fiber links over Ethernet to deliver both high-speed Ethernet and legacy TDM services which are prioritized over different traffic-engineered paths.

For single port/customer configurations, the Metrobility® Services Line Card provides an intelligent optical Ethernet demarcation point that supports both 100Mbps and 1Gbps connections to the customer site.

These demarcation devices augment existing IP/SNMP-based management frameworks with the IEEE 802.3ah Operations, Administration and Maintenance (OAM) protocol, the IEEE 802.1ag Configuration Fault Management (CFM) protocol, and emerging MEF frameworks to provide proactive health and status updates on network topology and application behavior.

For more information, visit www.telco.com

Telco Systems
A BATM Company

Contact information

International Headquarters
13 HaYetzira St., Yokneam Ilit,
20692, Israel
Tel: +972-4-993-5630
Fax: +972-4-993-7926

North & Latin America
15 Berkshire Rd
Mansfield, MA 02048
Tel: +1-781-255-2120
Fax: +1-781-255-2122

Asia Pacific (APAC)
10 Anson Road,
#17-03 Intl Plaza
Singapore, 079903
Tel: +65 6224 3112
Fax: +65 6220 5848

Europe, Middle East & Africa (EMEA)
Peterstr. 2-4,
52062 Aachen
Tel: +49 241 463 5490
Fax: +49 241 463 5491