

METRO ETHERNET FOR FLEXIBLE, SCALABLE CONNECTIVITY TO DATA CENTERS

INTRODUCTION

Enterprises have been migrating applications from running locally at each business location to running in centralized data centers for some time. Businesses are pursuing this migration because it results in tremendous savings in capital and operational expenses. It also simplifies and more efficiently enables a business to scale their application and storage requirements during changing economic conditions. However, once applications become centralized, the network becomes a critical component to ensure continuous data center connectivity to maintain business continuity.

Ethernet is the technology used for interconnecting individual users, servers and most networking devices on local area networks. Metro Ethernet is one of the fastest growing wide area networking (WAN) technologies to interconnect enterprise locations and provide high speed Internet access. Data centers have used Ethernet to interconnect servers and some types of storage. In the data center, Fibre Channel storage area networks (SAN) are widely deployed and continue to grow. Ethernet is advancing to support Fibre Channel through a technology called Fibre Channel over Ethernet (FCoE). With Ethernet widely deployed in LANs, increasingly being deployed for WANs and evolving to support SANs, isn't it time you considered using Metro Ethernet to connect to your data centers?

This paper discusses the Metro Ethernet attributes that provide more flexible, scalable connectivity between your business locations and your data centers than SONET TDM private lines or IP VPNs. Unlike other WAN technologies, Metro Ethernet has some unique capabilities to allow you to more cost effectively connect multiple locations to your data centers.

BENEFITS OF METRO ETHERNET FOR CONNECTING TO DATA CENTERS

Metro Ethernet is the fastest growing wireline data service throughout the world, according to Vertical Systems Group. The drivers for this growth are highly relevant when using Metro Ethernet for connectivity to data centers.

Bandwidth Flexibility

Metro Ethernet's incremental bandwidth options enable you to purchase the bandwidth you need rather than purchase the amount of bandwidth dictated by the WAN technology. For example, if you need 1Gbps to connect to your data center, with SONET, you need to purchase an OC-48 (2.5Gbps) private line because it provides the closest SONET bandwidth increment to 1Gbps. Similarly, if you require 100Mbps, you need to purchase an OC-3 (155Mbps) private line.

With Metro Ethernet, bandwidth increments are commonly offered in increments as low as 1Mbps. This enables you to best match the Metro Ethernet service bandwidth you need to purchase with the bandwidth connectivity requirements to your data center.

*With Metro Ethernet services,
you only purchase the bandwidth
you need, when you need it*

More Flexible Connectivity Options

Metro Ethernet provides more flexible connectivity than legacy WAN approaches since it can support all types of topologies to connect to your data centers. Metro

Ethernet supports point-to-point, point-to-multipoint (hub and spoke), and multipoint (any-to-any) connectivity. This flexibility enables you to optimize your data center connectivity to your business locations. You can also leverage this flexibility to more quickly and cost-effectively connect additional business locations to your data centers or add additional data centers for business continuity.

Metro Ethernet's flexible connectivity options simplify adding additional data centers for business continuity

In Figure 1, three business locations are connected to a primary data center using an Ethernet Private LAN (EP-LAN) service which provides any-to-any connectivity between all locations. The bandwidth for each location is tailored to meet the site's requirements. Any site's bandwidth can be increased as required. If a new, secondary data center is added to the EP-LAN service, all other sites, including the primary data center, have connectivity to it.




Figure 1: Data Center added to EP-LAN Service

The EP-LAN service serves a dual role providing site-to-site connectivity as well as a backup connection to a secondary data center to maintain business continuity in the event of a disaster or loss of connectivity to the primary data center.

Scalability

Businesses are migrating applications from running on equipment at each business location to running on equipment centralized at data centers. Examples include mail servers, file servers and other application servers.

Furthermore, additional equipment such as firewalls and information security scanners filtering for malware, spam, and other types of Internet threats can be consolidated at centralized Internet access locations. The rationale to consolidate storage, compute, and threat prevention resources is to reduce capital equipment and operational expenses. This approach requires WAN bandwidth to your data center to be scalable to meet your current, planned and sometimes unanticipated bandwidth needs as you grow your business, consolidate IT resources and maintain business continuity.

Metro Ethernet enables you to use a single protocol that remains unchanged as you scale your bandwidth from 1Mbps to 10Gbps today evolving to 100Gbps in the near future. Other data center interconnection technologies such as SONET require a new service deployed, e.g., a SONET OC-3 to an OC-12 private line, as you scale to a higher bandwidth. This results in a service disruption as you decommission the former service and activate the new service in addition to the cost of new equipment to support the different interface.

Metro Ethernet uses the same Ethernet protocol as you scale from 1Mbps to 10Gbps today and to 100Gbps in the near future

With Metro Ethernet, you purchase a sufficiently high speed Ethernet port at your data center to meet your current and anticipated bandwidth requirements, yet only purchase the amount of service bandwidth you need, when you need it. The service bandwidth is ordered via a Committed Information Rate (CIR) which is the amount of bandwidth you purchase up to the Ethernet port speed. At other business locations connecting to the data center, you can purchase a different speed Ethernet port and CIR that best matches each locations needs.

QoS Performance

QoS Performance includes packet latency (frame delay) and packet loss (frame loss) per class of service (CoS). Packet latency and loss can have a significant impact on user satisfaction with application performance or even the usability of certain applications running in a remote data center.

Many legacy software applications were designed to operate in a LAN environment where network latency is insignificant since the applications run on servers located

in the same building, or even same LAN subnet, as its users. Some of these applications are very “chatty” in that they require frequent communications between the client (user) and server application. Such applications either perform poorly or are unusable when running in a data center connected over a WAN, which can have a ten to fifty times higher packet latency than that of a LAN.

Metro Ethernet provides the lower packet latency required by legacy applications designed to operate on a local server on the LAN

Class of Service (CoS) is an important consideration when selecting a packet-centric data service such as Metro Ethernet. Data centers are used for many purposes such as data storage for backup and remotely running applications on virtual servers. Data backup for archival purposes may not require low latency. However, transactional applications running in the remote data center over a WAN may require low latency to function properly or achieve acceptable performance for the user.

Metro Ethernet services provide service level objectives (SLOs) in a service level agreement (SLA). The SLOs are the key performance metrics for a service that are captured in the SLA. The SLA is the business agreement defining the service terms and conditions for which the Ethernet service provider will meet the specified SLOs. Example SLOs are average or maximum latency across a maximum distance (since latency is distance sensitive) and packet loss for services in a metropolitan area or nationwide. SLOs are also provided for different classes of service (CoS) for a given Metro Ethernet service.

For example, you purchase a 100Mbps Ethernet service with two classes of service to connect your business location to your data center. Refer to Figure 2.




Figure 2: Two CoS Ethernet Service Example

You assign 30 Mbps for the Gold CoS with a <15ms latency and <0.001% packet loss for your latency and loss sensitive applications and 70Mbps is for the Silver CoS with a <50ms latency and <1% packet loss for applications which have little to no sensitivity to packet latency or packet loss. Refer to Figure 3.

SLO Metric	Gold CoS	Silver CoS
Packet Latency	< 15ms	< 50ms
Packet Loss	< 0.001%	< 1%

Figure 3: SLOs for example 2 CoS Ethernet Service

Network Availability

As applications are migrated to run in remote data centers instead of locally on the in-building LAN, network availability becomes a critical factor in order to maintain normal business operations. If the WAN connection to your data center experiences an outage, you no longer have access to your remote applications. Therefore, you need to carefully evaluate the performance metrics related to Metro Ethernet network availability such as:

- Availability (or Service Availability)
- Mean Time to Respond
- Mean Time to Restore

Availability is typically defined as the percentage of time that a service is operational, measured over a 30 day period. A service may be deemed non-operational depending on the parameter being measured, e.g., packet loss, which would be specified in the SLA. For example, a service would be considered non-operational if there were a circuit failure caused by a fiber cut. This is considered a hard failure because the service is non-operational until the fiber is repaired, unless there is a redundant fiber connection.

The Availability metric also includes other types of failures that may be intermittent such as the service’s packet loss or packet latency exceeding the maximum service level objective over a brief time period. The Availability metric is clearly an important factor in WAN connections to data centers. Availability can be significantly improved by ensuring that the Metro Ethernet network is self healing, i.e., it has multiple levels of redundancy built-in to keep the network operational as faults occur. The Availability metric is critical to understand when using Metro Ethernet to connect to data centers.

Network Availability becomes a critical Metro Ethernet performance metric as you migrate your applications to remote data centers

The Mean Time to Respond can be defined as the average time the Metro Ethernet service provider has to respond to troubleshoot a problem after it is either identified or reported to the network operations center. Mean Time to Restore can be defined as the average time to restore service to an operational condition as defined by the Availability metric in the SLA. These two metrics are important and interrelated because a small Mean Time to Restore would be negatively impacted by a long Mean Time to Respond and vice versa. These two metrics in conjunction with the Availability metric ultimately determine the Metro Ethernet network availability which again, is a critical metric as you move your applications to remote data centers using a Metro Ethernet service.

IP Version Transparency

Metro Ethernet does not require you to coordinate IP addressing and routing information with your Ethernet service provider. You can use existing IP addressing to your data centers with the freedom to use any IP addresses or routing protocols as you add additional data centers or business locations. Also, Metro Ethernet services transparently support IPv6-based applications alongside IPv4.

SUMMARY

Enterprises are migrating applications to centralized data centers to achieve significant capital and operational expense savings while enabling efficient scaling of their organizations' requirements during changing economic conditions. Ethernet enables you to standardize on a common technology for LAN, WAN and a converged data center. Metro Ethernet provides flexible, scalable bandwidth and network connectivity options with the QoS performance and IP transparency you need to connect to your data centers. As you migrate your applications to a remote data center, the WAN becomes a critical component and Metro Ethernet has become the technology of choice to meet these data center connectivity requirements.

REFERENCES AND RESOURCES

- Understanding Business Ethernet Services, Comcast white paper
- Ethernet Services for Multi-Site Connectivity, Comcast white paper
- Metro Ethernet Services – A Technical Overview, Ralph Santitoro, Metro Ethernet Forum

ABOUT COMCAST BUSINESS ETHERNET

Comcast offers a complete range of MEF certified business Ethernet services including Ethernet Private Line, Ethernet Virtual Private Line, Ethernet Network Service (MEF E-LAN compliant) and Ethernet Dedicated Internet. Each service is offered with a 10 Mbps, 100 Mbps, 1 Gbps, or 10 Gbps Ethernet port in customer-selectable bandwidth increments ranging from 1 Mbps to 1 Gbps. For more information or to request a consultation about Comcast's Business Ethernet Services, please visit <http://business.comcast.com/ethernet>.

