AccessEtherLinX

Differentiated Optical Ethernet Services

Optical Ethernet Services

Table of Contents

Introduction

This technical overview will discuss design considerations for Ethernet-based access networks, and the feature requirements for customer premises equipment (CPE).

While end-to-end Ethernet networks provide operators with significant economic and operational benefits, service providers usually require line provisioning, remote monitoring and troubleshooting capabilities which traditional routers and CDU/CSU have, while typical Ethernet-based CPE may not.

AccessEtherLinX from IMC Networks is a full-featured yet cost-effective optical access device that provides a remotely managed demarcation point at the customer's network edge. In addition to SNMP management, *AccessEtherLinX* supports rate limiting, full range of tagged VLANs and packet prioritization.

In the following we will discuss how *AccessEtherLinX* and its unique features help service providers to offer differentiated Transparent LAN and other Metro Ethernet services to business customers, while minimizing the initial capital investment and cost of operating the network.

Fiber Access Requirements

The bandwidth requirements for most businesses have already exceeded the level of xDSL / T1 circuits. At the same time, Optical Ethernet technology has emerged as a more cost-effective alternative to the copper-based access technologies. Legacy TDM-based circuits and access equipment were designed and optimized for carrying voice traffic but are not efficient in handling data, not to mention the "triple play" of services that combines voice, video and data traffic over the same circuit.

Providing Optical Ethernet access has many advantages both to the customer and the service provider. In addition to Ethernet's unlimited scalability, services are delivered more efficiently when the network is end-to-end Ethernet packet based. On the physical layer (OSI 1), using single-mode fiber optic cabling will future proof the "last mile" access network indefinitely. Fiber is the only media that will support anything you throw at it, from today's bandwidth-hungry applications to tomorrow's technology that has not been even invented yet.

The optical access unit must provide an adequate level of Operations, Administration and Maintenance (OAM) capabilities that will enable service providers to guarantee the quality and the consistency of optical services. At the minimum this means a Simple Network Management Protocol (SNMP) agent resides in the access unit, enabling network operators to proactively monitor equipment in the field, and quickly react to any potential issues before the customer even notices a problem with the service. Remote management allows the service provider to receive real-time device and link status information, traf`fic statistics, enable/disable services, etc. Furthermore, remote management and configuration capabilities are essential in reducing the cost of operating the network by minimizing the need to send a technician to the customer site after the initial set-up.

In order to deliver the profitable "triple play" of services, channels of information sharing the same fiber circuit to the subscriber(s) need to be separated and differentiated so that the appropriate traffic security and Quality of Service (QoS) rules may be applied throughout the network. With *AccessEtherLinX,* a single device can provide multiple services, with various priorities, to a single customer. This differentiation of services will increase the service provider's revenues and profits: Rather than just selling commodity bandwidth, establish service levels that justifiably charge more for additional bandwidth, and prioritizing traffic types such as Voice over IP and video conferencing that are extremely sensitive to latency.

Last, when building optical access networks, most of the capital cost incurs at the point-of-presence where the Layer 3 / 4 switching routers are deployed, leaving little budget for the Customer Premises Equipment (CPE). For service providers to realize a quick return on the initial investment, the CPE needs to be available at a reasonably low cost.

AccessEtherLinX Features

Designed to eliminate the need for additional edge devices, *AccessEtherLinX* is an all-in-one optical access device. The compact, cost-effective unit offers service providers with a wealth of features that will help lower the initial capital investment and the long-term cost of operating the network. *AccessEtherLinX* features line provisioning, rate limiting, traffic separation and differentiation, as well as remote management, monitoring and configuration capabilities. In the following, we will highlight these features in more detail.

Optical Network Interface

The most basic function of an Optical Ethernet access device is terminating the optical fiber and presenting a standard twisted pair interface to the customer network. When deployed at the customer's network edge, *AccessEtherLinX* provides the customer with a familiar 10/100 Ethernet RJ-45 access point to the MAN/WAN fiber optic network. The *AccessEtherLinX* fiber uplink port features a standard ST or SC duplex or a single-strand fiber optic connector so installation is straightforward, and customers do not need to worry about the technical issues such as fiber wavelengths, modality or optical power budgets.

As the optical network interface and demarcation point, *AccessEtherLinX* also provides both the service provider and the customer with a clearly defined point where the customer network and technical responsibility ends and the service provider network begins.

Remote Configuration, Management and Monitoring

As mentioned before, traditional Ethernet-based CPE does not provide the adequate level of management and troubleshooting features service providers are used to with the legacy TDM-based equipment. *AccessEtherLinX* features an SNMP management agent for remote configuration, monitoring of the status and activity on both the twisted pair ports and the fiber uplink, as well as for proactive fault management.

Providers can perform the initial setup of the unit and future upgrades via:

A) IMC Networks' GUI-based *iView*^{2} SNMP application,

B) Telnet, or

C) local serial connection

In addition to static settings, the real-time operational information available via SNMP includes link status, total bytes transmitted and received at each port, and any link-level errors seen on either copper or fiber ports.

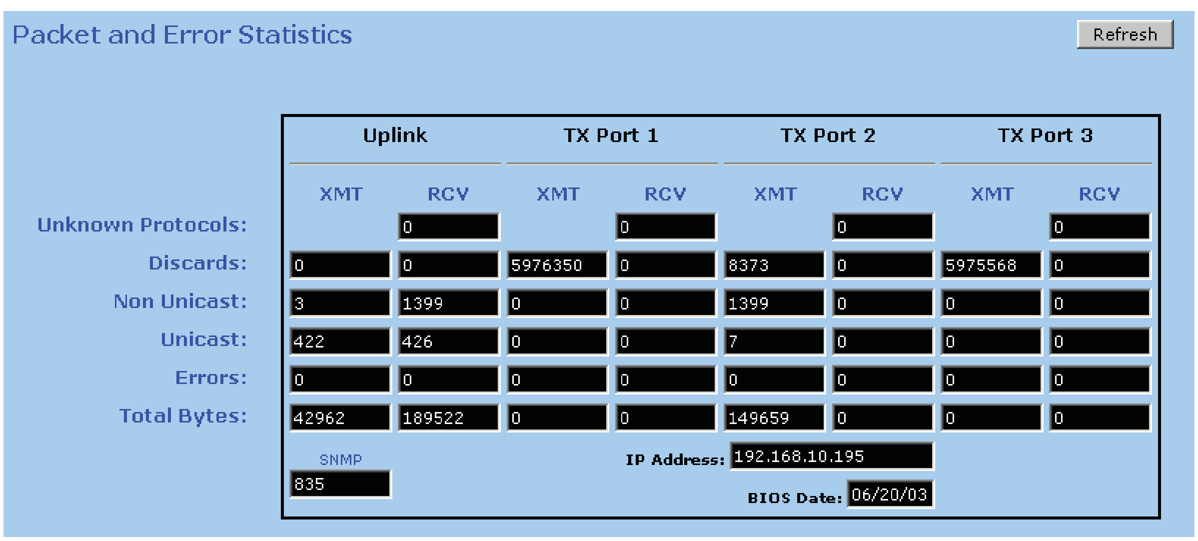

Packet and Error Statistics

Refresh

	Uplink		TX Port 1		TX Port 2		TX Port 3	
	XMT	RCV	XMT	RCV	XMT	RCV	XMT	RCV
Unknown Protocols:		0		0		0		0
Discards:	0	0	5976350	0	8373	0	5975568	0
Non Unicast:	3	1399	0	0	1399	0	0	0
Unicast:	422	426	0	0	7	0	0	0
Errors:	0	0	0	0	0	0	0	0
Total Bytes:	42962	189522	0	0	149659	0	0	0

SNMP: 835

IP Address: 192.168.10.195

BIOS Date: 06/20/03

Remote management is essential in guaranteeing the consistency and integrity of Metro Ethernet services. Equally important is the cost savings network operators are able to realize by minimizing the need for customer site visits after the initial device setup. Via SNMP or Telnet, network operators can control all the functions of *AccessEtherLinX*, including customer rate limiting.

Rate Limiting/Bandwidth Control

AccessEtherLinX provides the ability to control the available customer bandwidth independently on the upstream (copper to fiber) and downstream (fiber to copper). Bandwidth limitations in both directions are individually configurable on each of the copper ports, and globally on the fiber port.

Let's discuss this unique rate limiting feature in more detail. What is rate limiting? It's simply a "speed" or bandwidth limitation applied to an Ethernet connection so that all traffic, regardless of protocol, is subjected to the bandwidth set on a link. Rate limitation is normally expressed as a "port-to-port" function, i.e. you can limit the bandwidth available for transmission of traffic from port A to B, as well as from B to A. This is called bi-directional rate limiting. The speed on a link is normally expressed in bits per second (bps), or Kilobits per second (Kbps). In practice, the service provider can set rate limits in any of the following ways:

Uplink Rate Limiting

Separate receive/transmit bandwidths applied to the uplink port of the *AccessEtherLinX*. Applying bandwidth limits on the uplink will affect all traffic passing through the unit uplink to the downlinks and vice versa. The bandwidth may be symmetrical e.g. 1Mbps receive, 1Mbps transmit; or, it may be asymmetrical e.g. 1Mbps receive, 256Kbps transmit. Finally it may be unidirectional e.g. unlimited receive*, 256Kbps transmit.

* In practice, unlimited is Ethernet wire speed; in the case of *AccessEtherLinX* this could be 10Mbps or 100Mbps

Independent Downlink Rate Limiting

Each downlink can have its own unique bandwidth applied to it. Different devices and/or networks attached to these downlinks can then have individual bandwidths available to them. As in the case of the uplink, bandwidths applied may be bi-directional, asymmetrical or unidirectional.

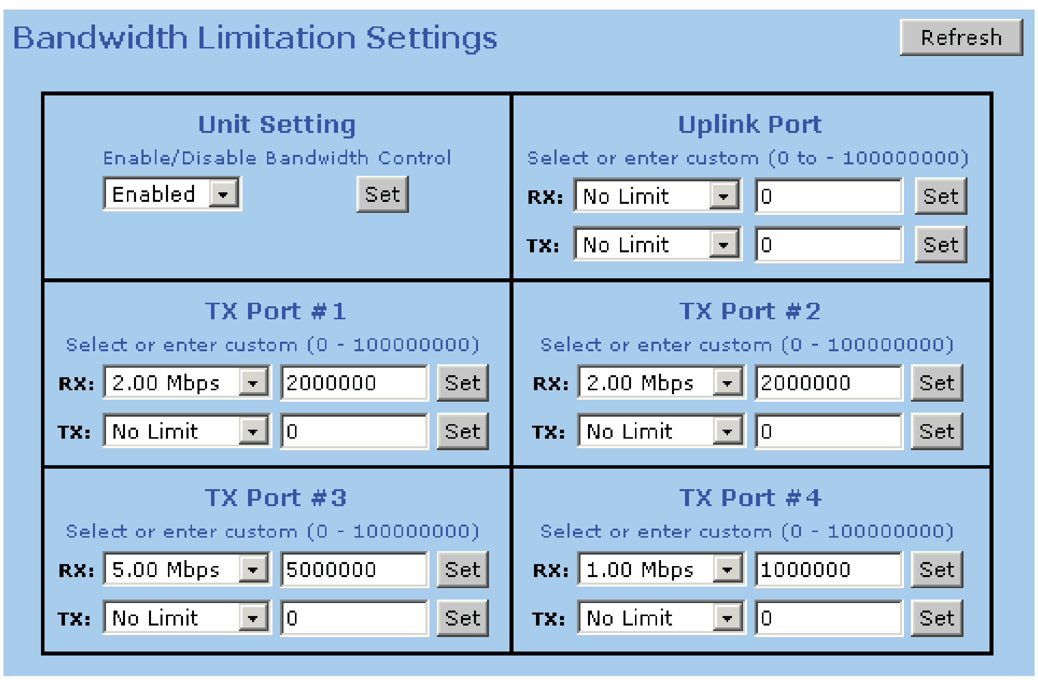

How Bandwidth Control Works

Within the *AccessEtherLinX* if the bandwidth limitation is enabled, a hard coded algorithm measures the number of bits passing through the port. This measurement is done every second. At the start of the measurement cycle the number of bits transmitted (or received), is set to zero, the *AccessEtherLinX* then starts to count the number of bits passing through the port (this is done individually for both the transmit and receive queues of the port). If the number of bits passed reaches the configured bandwidth limit (e.g. 256,000 bits), then the transmitter (or receiver), is paused until the one second measurement cycle is complete. After the one second is up, the counters are set to zero and the transmit/receive queues can start again.

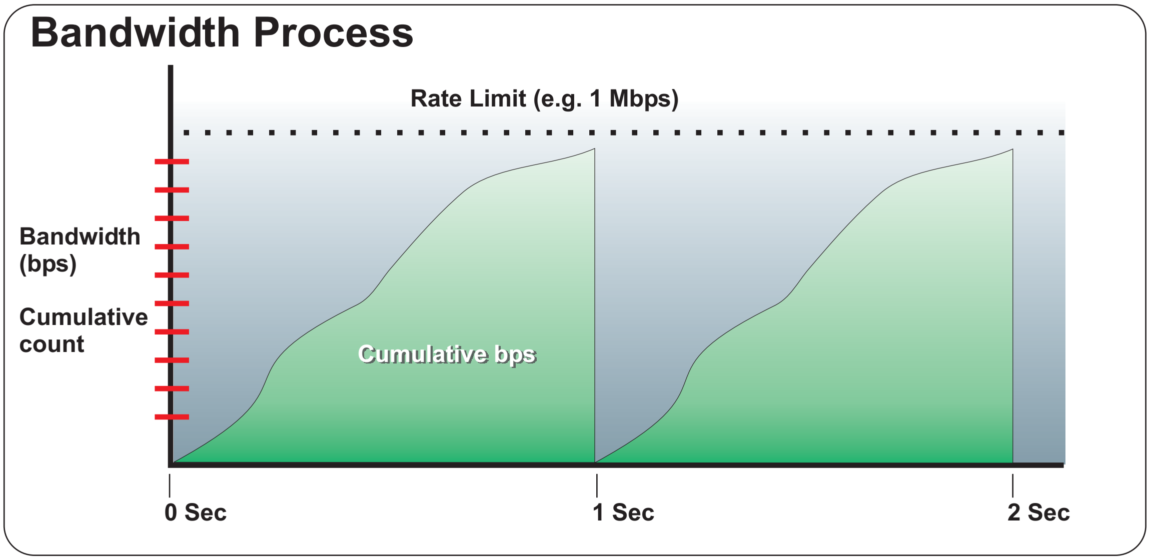

Since *AccessEtherLinX* is a five port Ethernet device, with each port having it's own separate transmit/receive, up to ten queues could be monitored independently, with zero impact on the functionality of the *AccessEtherLinX*.

Setting Bandwidth

Each port has an independent transmit (TX) and receive (RX) bandwidth. If you enable bandwidth control functionality, you can set the limit from a pull down option box or type in manually (note, this would be in bits per second e.g. 1Mbps would be represented by 1,000,000). *AccessEtherLinX* rate limits in 32Kbps increments. If you manually set a rate limit which is not a multiple of 32Kbps, the rate will be rounded either up or down to the nearest rate which is a multiple of 32Kbps. Please refer to the section, "Manually Setting Bandwith" in the Appendix for examples.

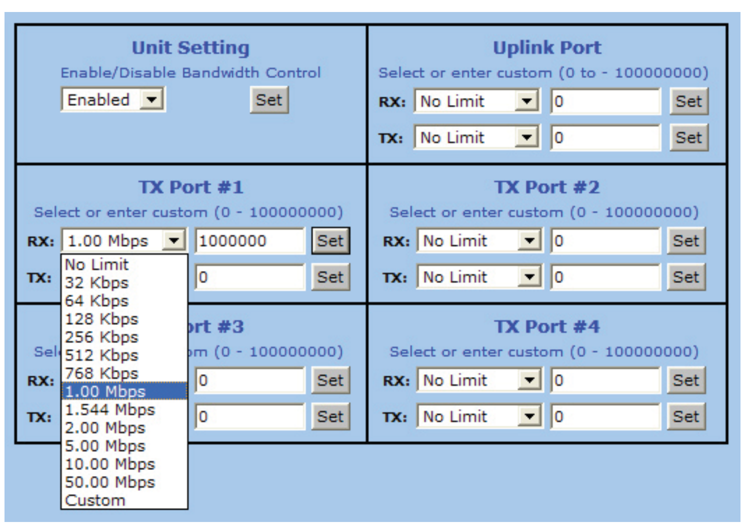

Measuring Bandwidth

If you try to measure the bandwidth that has been set you may well end up with a lower figure e.g. the transmit bandwidth is set to 128Kbps, but measuring the bandwidth via an FTP file transfer results in a transfer speed of only 108Kbps. The reason for this difference is simple: the *AccessEtherLinX* counts all bits passing through— this includes the Ethernet preamble; the MAC layer (layer 2) information; and lastly, the CRC of the Ethernet frame. The actual FTP component of the Ethernet frame may well be only 70% of the entire frame size, hence the difference between what's been set and what can be observed using for example, an FTP file transfer.

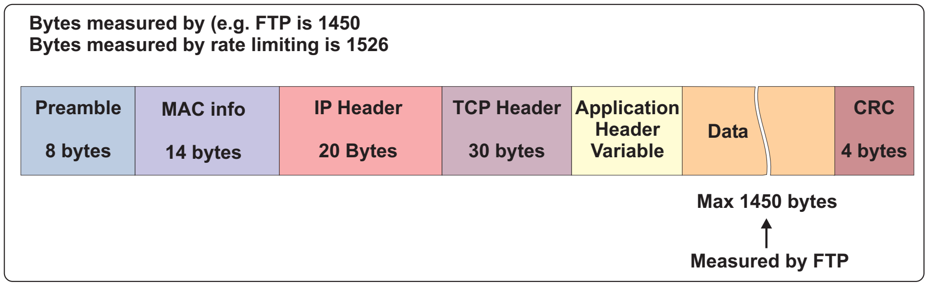

Flow Control

In the bandwidth description above it was mentioned that if the bandwidth limit (expressed in bps) of a specific port is exceeded within the one second time window, then the *AccessEtherLinX* will stop receiving on that port. The question is what happens to Ethernet frames trying to access that port? Before reaching the limit (in most typical cases 8Kbytes before the configured limit), the *AccessEtherLinX* will trigger flow control.

To be specific, the IEEE802.3x flow control standard allows an 802.3x compliant device (e.g. the *AccessEtherLinX*) to issue an XOFF pause frame. If the device attached to the *AccessEtherLinX* port that issued the XOFF frame is also IEEE802.3x compliant, it will stop transmitting frames on that port. Once the one second window is completed an XON flow control frame will be issued telling the other device to recommence transmission.

802.3x flow control is available in the vast majority of Ethernet switches on the market today. Establishing if 802.3x is available is normally done by the device on boot up as part of the auto-negotiation process that also establishes line speed (10 or 100Mbps) and duplex (HDX or FDX).

In addition to supporting 802.3x auto-negotiation, *AccessEtherLinX* can also be configured to force 802.3x flow control ON. You can enable flow control functionality for the unit, then independently configure settings for each port (Based on Auto-Negotiation or Forced ON; if based on Auto-Negotiation, you can configure to Advertise or Don't Advertise). This ability to enable/Force ON 802.3x is of great importance— it will ensure that the devices attached to the *AccessEtherLinX* (if they're 802.3x complaint as well) will enable flow control.

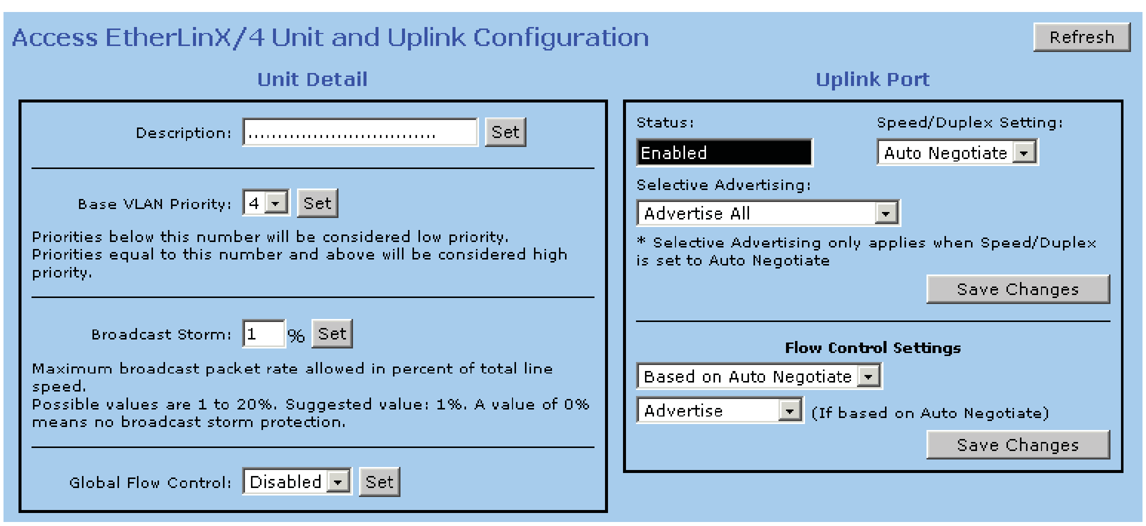

Broadcast traffic

Broadcast frames are used mainly to identify other devices on a network. Broadcast frames received on a downlink port of an *AccessEtherLinX* are passed to the uplink. In the opposite situation, when a broadcast appears at the uplink, the frame is passed to all the downlinks. Broadcasts are important— they allow devices to obtain the MAC ID of an unknown device they're trying to connect with. However, too many broadcasts can radically slow down a network. Broadcast frames unfortunately are too common; some devices (normally switches or routers), if configured incorrectly, may well create broadcast storms. *AccessEtherLinX* has the ability to filter broadcast packets; the unit can be configured so that if broadcasts are equal to or higher than a certain percentage of all traffic within a specific time period, then the *AccessEtherLinX* will simply drop (or ignore) the excess broadcast frames.

Customer Traffic Separation and Security

Customers connecting to Ethernet-based a Metropolitan Area Network want their traffic and sensitive company data to be separated and secured while it travels through the service provider network. *AccessEtherLinX* supports IEEE standards 802.1Q and 802.1p which specify a means for "tagging" Ethernet frames with 12 bits assigned to identify a Virtual LAN (VLAN) and three bits for a priority designation (see Appendix). Network equipment that is VLAN-aware may use this information in a variety of ways, such as directing data from/to specific end users or establishing differentiated quality of service levels based on customer or traffic type.

AccessEtherLinX supports all 4,096 VLAN IDs and eight 802.1p priority levels, with each port having an independent tag and priority setting. VLAN tagging has a number of practical uses. For instance, assign a VLAN tag to distinguish data coming from each downlink port, or to direct traffic from the uplink fiber port to each of the downlink ports.

Only packets with the appropriate VLAN ID are transmitted from the fiber port to each copper port; upstream data from each copper port goes only to the fiber port (unless specifically administered otherwise and is appropriately tagged by *AccessEtherLinX)*. The tags may be used upstream in the network as data flows are aggregated either locally or at a servicing point. Note that you can also assign a unique VLAN ID to SNMP management traffic intended for *AccessEtherLinX* internally; the tag differentiates SNMP traffic on the provider's backbone.

First, the service provider may use the tags to differentiate classes of data from a single customer without requiring a VLAN switch or other piece of equipment beyond the demarcation point to be controlled by the service provider. For example, one downlink port could be used for basic data service and be tagged with one VLAN ID, while a second is connected to a VoIP server using another VLAN ID. If the service provider is using the priority bits across the network core, these can also be independently assigned at the source, *AccessEtherLinX*, so that various end-to-end QoS paths can be implemented. Within *AccessEtherLinX*, two queues are maintained for ingress and egress so that higher priority traffic (the 802.1p level which is user-defined) can be provided with a better QoS.

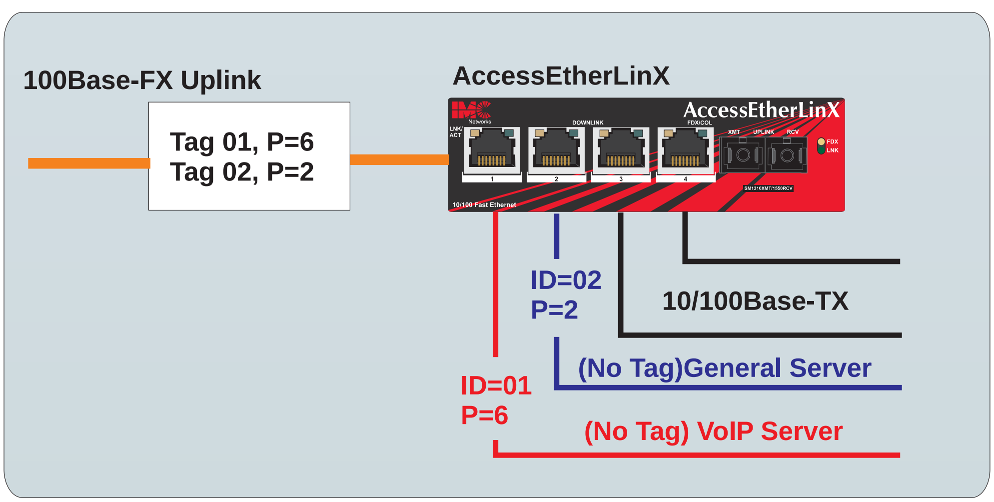

Second, the service provider may wish to use VLAN tags to segregate traffic from multiple services/devices at the same premise or multiple customers in the same location. Potentially, four subscribers may be connected to one *AccessEtherLinX* with each receiving a unique tag, so that all data remains secure while the hardware and fiber are shared.

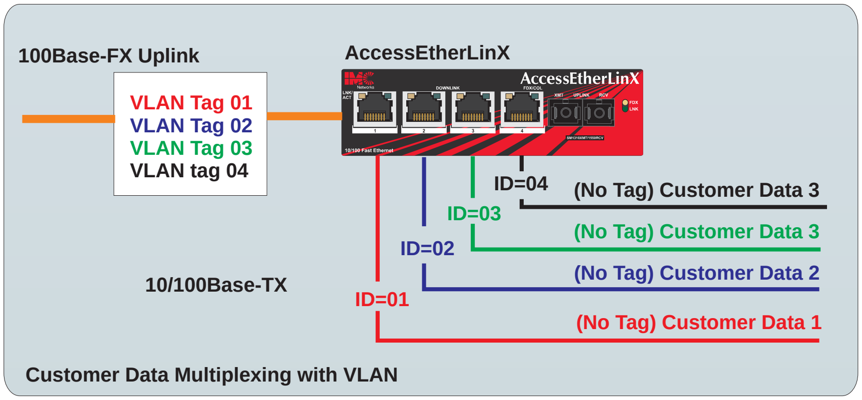

Customer Data Multiplexing with VLAN

Network Protocol Independence

AccessEtherLinX advanced features are not limited to IP traffic. Bandwidth limitation and VLAN tagging is carried out on the MAC layer, so that transparent LAN services supporting network protocols such as IPX, AppleTalk, APPN, NetBEUI, etc. can be provided. If it can be carried on Ethernet, *AccessEtherLinX* can tag/untag and rate limit it. If the customer maintains different physical networks for different protocols, these can each be connected to one of the four downlink ports and have independent VLANs / priorities and bandwidth limitations.

Multicast Pruning

Switches that are not IP Multicast aware will propagate that type of traffic out to all connected ports. By "snooping" on the IGMP packets that control IP Multicast, *AccessEtherLinX* limits transmittal of this traffic to only those ports for which it is intended. For example, if customer A is receiving TV Channel 2 as a streaming video multicast, that traffic would be sent only to the appropriate downlink port, not all four ports. If by chance customer B was also on that multicast, *AccessEtherLinX* would propagate the packets involved to that port as well. This would not be the case if a standard switch was used instead of *AccessEtherLinX*, resulting in a flood of unwanted traffic on those ports not enrolled in the multicast.

Summary

Service providers seeking to implement fiber optic access to subscribers have a number of technology alternatives available. When weighing the choices, consider both the cost and revenue sides of the equation. Ethernet has become the dominant technology used in enterprise networks, pushing aside contenders and pretenders from ARCnet to token ring to ATM.

Direct optical Ethernet access to customers can leverage the production volumes and knowledge base attained in the last 30 years of development. In one inexpensive unit, *AccessEtherLinX* provides the optical interface and maintenance functions required to deliver Ethernet over fiber to the subscriber at both low initial cost and low ongoing cost of ownership.

In order to maximize revenue, service providers need to be able to sell more than a wide-open pipe of commodity bandwidth. *AccessEtherLinX* has advanced features such as rate limiting and VLAN/priority support that allow for differentiated traffic at multiple price points. *AccessEtherLinX* is an optical access device that provides high performance at low cost in order to maximize provider profitability.

Contact IMC Networks Fiber Consulting Service for more information or assistance in configuring the right *AccessEtherLinX* model for your network.

Appendix

VLAN Tag Details

In order to use VLAN tags on the MAC level, a place for the tag in the Ethernet frame needed to be made. After discussion and testing, the IEEE 802.3ac committee determined that extending the frame from a maximum of 1518 bytes (by convention the Ethernet frame length does not include preamble or start frame delimiter) to 1522 bytes. A legal but previously reserved Length/Type of 8100 hexadecimal is used to designate the presence of a VLAN tag. Two bytes of 802.1Q Tag Control Information follow, and then the normal Length/Type field (see illustration.) The TCI field contains the three 802.1P priority bits, a Canonical Field Indicator that is zero for Ethernet, and 12 bits for the VLAN ID.

A zero VLAN ID indicates that VLANs are not being used, but that the priority field is meaningful. While the usage of the priority levels is not formally defined, it is recommended that 7 be reserved for network-critical traffic such as routing table updates; 5 and 6 be used for delay-sensitive traffic such as voice and video; 1 through 4 for varying levels of miscellaneous traffic; and 0 is best-effort default traffic. Not all equipment manufacturers use all eight levels of priority, and those who do may use each level differently so careful consideration of the total network architecture is required to ensure consistent quality of service from end-point to end-point.

In addition, it is important to ensure that any Ethernet equipment receiving a tagged frame is 802.1Q aware, as older gear will see the 8100H field as a length that does not match the data and will discard the frame as having a bad FCS.

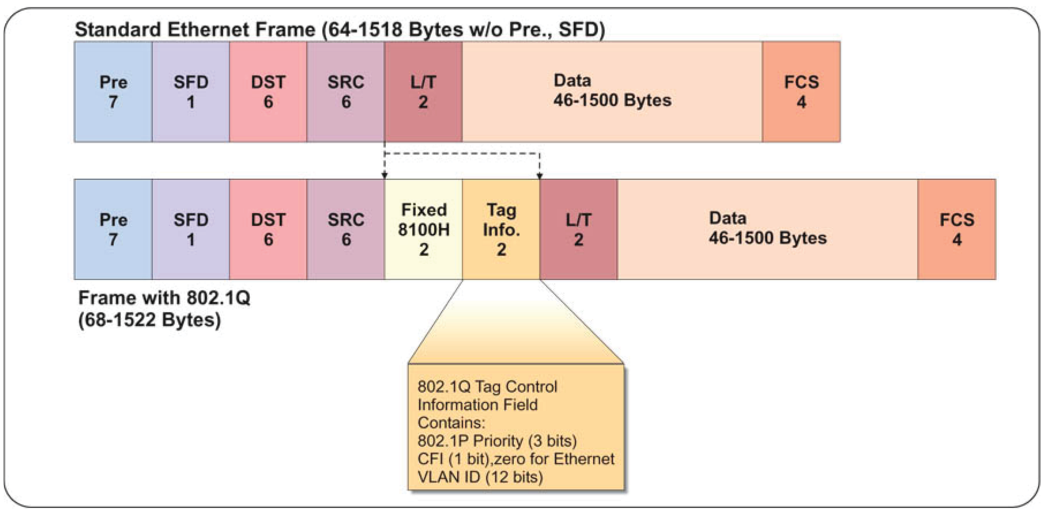

Manually Setting Bandwidth

Preset rate limits are in 32Kpbs (32 * 1024= 32,768) increments. If you manually set a user definable rate limit which is NOT a multiple of 32Kpbs, the rate limit will either be rounded down or up depending on whether the rate is set above or below the mid point (16 * 1024=16,384) between two increments.

For example—

If you set the user defined rate to 44,000: 44,000 - 32,768 = 11,232 : the rate will round DOWN to 32,768.

If you set 50,000 as the rate: 50,000 - 32,768 = 17,232 : the rate will round UP to 65,536.

.

About IMC Networks

IMC Networks is a leading ISO 9001 certified manufacturer of optical networking and LAN/WAN bandwidth management solutions for enterprise, telecommunications and service provider applications. The company provides the industry's widest variety of copper-to-fiber media converters, fiber mode converters, as well as optical repeaters and wavelength division multiplexers. In addition to physical layer products, IMC Networks offers remotely managed Customer Premises Equipment and Layer 3 and Layer 4 bandwidth control and packet classification solutions.

Fiber Consulting Services

IMC Networks' Fiber Consulting Services (FCS) assists network managers and system integrators with the design and development of fiber-based networks. Consulting services are free of charge. Please contact us at fcs@imcnetworks.com or by calling 800-624-1070 in the USA or +1-949-465-3000 outside of USA

Networks

Making Your NetWork Better

www.imcnetworks.com

IMC Networks
Headquarters
19772 Pauling
Foothill Ranch, CA 92610
TEL: 949-465-3000
FAX: 949-465-3020
sales@imcnetworks.com

IMC Networks
Eastern US/Latin America
28050 U.S. Hwy. 19 North, Suite 306
Clearwater, FL 33761
TEL: 727-797-0300
FAX: 727-797-0331
latinsales@imcnetworks.com

IMC Networks
Europe
Herseltsesteenweg 268
B-3200 Aarschot | Belgium
TEL: +32-16-550880
FAX: +32-16-550888
eurosales@imcnetworks.com

Document # WP-2001
0905