Building a Bigger Pipe:
Inverse Multiplexing for Transparent Ethernet Bridging over Bonded T1/E1s

PATTON Electronics Co.

Network Access & Connectivity

Copyright

Printed in the USA.

Executive Summary

This paper proposes a solution for extending the availability of transparent high-speed Ethernet by re-using the existing T1/E1 TDM network. By employing inverse multiplexing technology with support for jumbo Ethernet frames, service providers can use bonded T1/E1 circuits to transparently backhaul Ethernet from remote cell sites and IPDSLAMs. This paper covers the rationale, protocols and key features that an Inverse Multiplexer (Inverse Mux) must support in order to close the current capabilities gaps in carrier and provider networks.

Introduction

Ethernet is finally breaking out into the Wide Area Network (WAN). In order to support Ethernet based networks, service providers are aggressively investing billions of dollars to install fiber and create Metro Area Networks (MANs).

However, since WAN Ethernet technology is still maturing, Ethernet interfaces are not yet ubiquitous. Even in the most developed countries, fiber MANs reach less than 12% of businesses with more than 20 employees. This data indicates there is a huge capabilities gap in provider networks all around the world.

Solution Overview

Service providers with access to wholesale T1/E1 circuits can use inverse multiplexing equipment to bond multiple T1/E1 circuits for transparent Ethernet backhaul. An Inverse Mux transmits a data stream from a high-speed link (Ethernet) over a single high-speed communications channel comprised of multiple lower-speed circuits (Bonded T1/E1 TDM circuits).

The feature set in most inverse multiplexing equipment is designed for delivering service directly to individual subscribers. Yet an Inverse Mux with the right feature set can support transparent Ethernet backhaul from remote access concentration sites (typically cell sites and IPDSLAMs). Since these remote sites serve multiple users, using Inverse Mux equipment for Ethernet backhaul accelerates subscriber acquisition and leads to a faster capture of market share.

Why T1/E1 Circuits?

Why use T1/E1 circuits for Ethernet backhaul? We answer this question by citing five crucial points:

- T1/E1 circuits are available and installed in volume. Their widespread accessibility makes it easy to backhaul Ethernet from almost anywhere in a network without delaying to install infrastructure. By following this business case, service providers can expand services into areas where the demand for Ethernet service is high, but next-generation infrastructure (fiber) does not yet exist.
- Because T1/E1 circuits are ubiquitous, Ethernet can be terminated just about anywhere an access concentration point exists. An inverse mux can close the infrastructure gap by terminating traffic in such locations as points of presence (PoP), central offices (CO), and network peering points.
- Because the T1/E1 circuits have already been installed, as well as the core networks that tie them together, the capital expense has already been incurred. The service provider does not have to raise large sums of additional money.
- TDM circuits have a developed management plane that provides both near and far end statistics and end-to-end alarm reporting. OAM becomes a non-issue.
- Because TDM circuit testing, maintenance and management have been around for many years, service provider personnel possess core competencies that perpetuate its sound operation

Overview: Bonding T1/E1 Circuits

Bonding lower-speed T1/E1 circuits into a higher-speed communications channel is not a new invention. It has

been around for many years. Bonding has been used before in TDM applications to backhaul traffic from remote terminals (RTs). The most important technologies available to bond T1/E1 circuits for data transmission include MLPPP, ATM/IMA, and GFP. The following section covers the relevant characteristics of each.

ATM IMA

For networks heavily invested in ATM, Inverse Multiplexing over ATM (IMA) is a viable approach for bonding multiple T1/E1 circuits into a single higher bandwidth channel. However, because of its simplicity and low cost, Ethernet has become the clear winner in the transport technology wars. As a result ATM IMA is no longer as popular as it once was. Furthermore, no one wants to pay the ATM cell tax. Breaking up Ethernet frames for encapsulation within ATM cells introduces huge protocol overhead, causing operators to lose an average 20% of the network bandwidth available for data payloads.

GFP/VCAT/LCAS

A suite of three separate protocols, GFP, VCAT, and LCAS, comprise another method for transporting Ethernet over bonded T1/E1 circuits. Generic Frame Procedure (GFP) encapsulates Ethernet prior to transmitting the frame over the T1/E1 circuits. Virtual Concatenation (VCAT) handles the actual bonding of the T1/E1 circuits. Link Capacity Adjustment Scheme (LCAS) provides dynamic circuit removal and restoral, similar to the LQM features of ML-PPP (see next section). These protocols show great promise but are relatively new and not yet widely deployed. For providers that need to build out their bandwidth quickly, such a relatively embryonic solution raises some serious interoperability and testability concerns.

Multi-Link PPP

Originally issued in 1994 as an Internet Engineering Task Force (IETF) draft, the Multi-Link Point to Point Protocol has enjoyed close to 14 years of deployment. As an underlying technology solution for Ethernet backhaul over bonded T1/E1 circuits, the following eight points argue that ML-PPP is the clear choice.

- **Stability**. ML-PPP is stable. The protocol has been around so long that there are few unknowns and many known workarounds.
- **Interoperability**. ML-PPP is proven to be interoperable. Without interoperability providers would get locked into a single vendor solution.
- **Testability.** More test equipment vendors support ML-PPP than any other multi-link protocol.
- Low overhead. The average ML-PPP protocol overhead can be as low as 3%. Low overhead means more bandwidth for actual customer traffic.
- **Automated configuration.** ML-PPP uses the Link Control Protocol (LCP) to provide automated configuration of endpoints.
- **Scalable.** More bandwidth can be added by increasing the number of T1 or E1 circuits that comprise the bonded channel, enabling a pay as you go approach to increasing Ethernet bandwidth.
- **Load balancing** is native to ML-PPP. By fragmenting the Ethernet frames and distributing them in equal pieces over the links, ML-PPP minimizes end-to-end latency when transmitting Ethernet over the TDM network.
- **Self-healing.** ML-PPP employs Link Quality Management (LQM) to give the broader channel a self-healing quality. LQM automatically detects a

failed T1 or E1 circuit and avoids sending Ethernet traffic over any failed links. Once a failed circuit is restored and stable, the T1/E1 circuits are automatically re-enabled for Ethernet data.

Challenges Posed by the Evolution of Ethernet

Before delving into the details of transmitting Ethernet over bonded T1/E1 circuits, we first need to look at certain challenges that stem from the evolution of Ethernet technology. Ethernet is going through a transformation as it grows into a WAN-grade protocol. A major focus has been making Ethernet capable of seamlessly interconnecting millions of devices end-to-end. To improve Ethernet's scalability the strategy to date has been to add layers of tags and encapsulations. As a result, the size of a standard Ethernet frame has grown well beyond the 1500 byte maximums, which legacy Ethernet equipment supports. The expanded Ethernet frame not only poses a challenge for legacy equipment, but it also complicates equipment configuration. Let's recap the tagging and encapsulation technologies in use today:

- **Virtual LANS (VLANs).** Service providers commonly use VLANs to isolate and prioritize Ethernet traffic. By appending a tag to the Ethernet frame, VLANs identify the virtual network to which a subscriber belongs. The tag isolates a subscriber's traffic from computers, printers and other devices that belong to different subscribers. The VLAN tag also includes 3 priority bits that provide QoS, which can be used as a basis for Service Level Agreements (SLA)–residential, access only, or multi-media traffic, for example.
- **Q-in-Q or VLAN stacking.** Q-in-Q takes the VLAN approach a step further. It adds a second VLAN tag to the Ethernet frame. Service providers use this second tag to aggregate traffic from multiple end users sharing an access concentrator (DSLAM, etc). VLAN stacked Ethernet frames

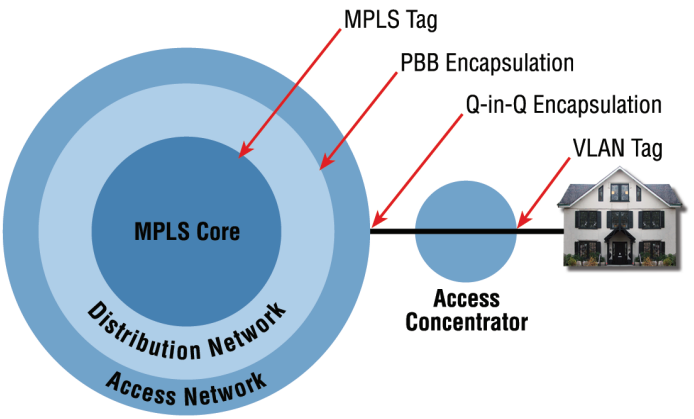

streamline the backhaul of Ethernet through the access network by simplifying the switching decisions at the network access layer.

- **Provider Backbone Bridging (PBB) or MAC-in-MAC.** PBB encapsulates Ethernet with an additional Ethernet MAC address. It basically stuffs VLAN tagged Ethernet into another Ethernet frame. PBB encapsulated Ethernet limits the number of Ethernet MAC addressees that carrier equipment must learn at the distribution layer of their networks. This additional encapsulation further improves scalability.
- **Multiprotocol Label Switching (MPLS).** MPLS is a transport layer technology. It adds a tag to Ethernet traffic so that core networks can focus on switching instead of IP routing. MPLS is the technology of choice in the network core because of its simplicity and speed, its support for Quality of

Service (QoS), and its ability to support Virtual Private Networks (VPNs).

Together, these developments have created the technical challenge of how best to backhaul Ethernet—with its ever-increasing frame size–over a fixed point-to-point connection.

Transparent Ethernet Backhaul

Now that we understand how and why the Ethernet frame became so large, we can deal with the question of how best to transport it. Certain approaches advocate having the network equipment terminate and interpret the tags and encapsulations before forwarding it. This white paper, in contrast, clearly advocates Ethernet transparency. For backhaul from remote locations frame forwarding is a point-to-point exercise. The most efficient and cleanest approach for this application is to forward jumbo Ethernet frames transparently.

Routers with multiple T1/E1 interfaces tend to be cumbersome to configure, and create an overly complex solution to a simple problem. On the other hand, an Inverse Mux that supports jumbo Ethernet frames can simplify the network configuration by transparently passing tagged and re-encapsulated Ethernet frames.

Such a simplified network configuration expedites the process of deploying Ethernet backhaul from remote access-concentration sites. Other approaches introduce increased complexity in the network design, which slows the deployment of Ethernet services. Furthermore, such complexity slows fault resolution and restoration of service, insidiously impacting customer satisfaction and retention.

When it supports jumbo Ethernet frames, an inverse mux works well as an Ethernet backhaul solution for many different traffic types of scenarios. With jumbo frame support, a channel comprised of bonded T1s or E1s can transparently carry IPDSLAM Ethernet traffic (which is typically VLAN and Q-in-Q encapsulated), PBB and MPLS tagged traffic, and even traffic from non-standard Cisco Inter-Switch Links (ISL).

QoS

Without QoS, transparent Ethernet backhaul is inadequate for addressing the bandwidth gap in today's existing infrastructure. Modern networks rely on QoS and traffic prioritization to support real-time and multimedia traffic. QoS reduces latency and jitter for real-time traffic and enables such applications as VoIP and IPTV. Without QoS, service providers would need to overprovision their network.

The only real QoS is hard QoS. Hard QoS involves more than just setting and responding to priority bits. Hard QoS requires queuing, shaping and policing the traffic as it ingresses and egresses through the network equipment. To properly support transparent Ethernet backhaul, hard QoS must be supported according to multiple criteria.

In access environments where Ethernet is VLAN or Q-in-Q tagged, the inverse mux must first prioritize ingress traffic based on VLAN priority bits, queue it appropriately, and then service the egress queue deterministically. The inverse mux must also be able to tag egress traffic while appropriately setting the VLAN priority bits.

When transparently backhauling PBB encapsulated, MPLS tagged or Cisco ISL encapsulated traffic, the inverse mux must also support robust filtering capabili-

ties to enable prioritization based on various criteria. MAC address filtering, queuing and prioritization are essential. MAC filtering also provides call admission control (CAC), allowing the service provider to control the egress and ingress of traffic through the network.

Conclusions

Ethernet has finally made it to the WAN and is emerging as the service provider technology of choice. Until the transformation to a pure Ethernet network is complete, service providers can use inverse multiplexers to leverage the existing T1/E1 TDM network to deliver Ethernet bandwidth. With support for the key features discussed above, service providers can use inverse mux technology to rapidly expand their footprint into areas where next-generation infrastructure is not yet in place.

Not only have Ethernet standards changed, todays offered services now include real-time VoIP and IPTV. Because of these changes, an Inverse Mux used for Ethernet backhaul (rather than access service provisioning) must support a powerful and unique feature set. Core requirements include the following capabilities:

- **ML-PPP**, a field-hardened technology, leverages years of technology and innovation to efficiently transport Ethernet with proven testability, interoperability, and reliability.
- **Jumbo Ethernet frame support** makes Ethernet backhaul simple and transparent regardless of the number of layers, tags and encapsulations embedded in the Ethernet frame. Alternative (routed) solutions complicate network configurations and slow service provisioning.
- **Traffic filtering** is a must have feature, enabling providers to offer application-level SLAs and provide support for CAC. Because they fail to account for PBB, MPLS and ISL, approaches that prioritize based on VLAN tags fall short.

A Viable Solution

One commercial implementation of the technology discussed above is available from Patton Electronics. The IPLink™ Model 2888 Multi-Megabit Inverse Mux is a Transparent Ethernet bridge with two (2) Gigabit Ethernet ports and either two (2) or four (4) T1/E1 ports. Patton's Inverse Mux transparently forwards jumbo Ethernet frames over bonded T1/E1 circuits.

Complete with Layer 2/3 filtering, Layer 2/3 traffic shaping and Active Layer 2/3 QoS, the Model 2888 offers carriers an immediate transparent bridged Ethernet backhaul solution for rapid deployment of broadband services.

PATTON Electronics Co.

7622 Rickenbacker Drive Gaithersburg, MD 20879 USA
Phone **+1-301-975-1007** • Fax +1-301-869-9293
URL **www.patton.com** • E-mail **marketing@patton.com**
Document 07MD-VO-IND-ENET-WP