

On the Vision of Implementing A Truly Native Ethernet-Based Global Multi-Service Infrastructure

Mohamed A. Ali and Antonios Hadjiantonis

Dept. of Electrical Engineering, City College of the City University of New York

{ali,antonios}@ccny.cuny.edu

Haidar Chamas, William Bjorkman, and Stuart Elby

Verizon Communications, White Plains, NY

{haidar.a.chamas,william.a.bjorkman,stuart.d.elby}@verizon.com

EXECUTIVE SUMMARY

This work proposes and devises a novel native Ethernet-based networking architecture and switching paradigms to implement a truly end-to-end Optical Ethernet infrastructure seamlessly stretching from enterprise LAN to Metro to Global. To this end, the City College Next-Generation Data-Centric Networking Lab along with Verizon Communications scientists and engineers are leading extensive research aimed at defining the networking technology and architecture for scaling Metro Ethernet networks into a Global multi-service infrastructure. The proposed Optical Ethernet networking architecture is a true two-layer model, realizing the significant goal of Ethernet-over-WDM, where native Ethernet frames are mapped directly over WDM. It offers significant advantages over existing Layer-2 and MPLS solutions in that it divorces the Ethernet from legacy transport mechanisms like SONET/SDH. The key for realizing such an ambitious initiative rests on resolving three key issues: 1) How to replace the legacy layer-3 switching (routing) and hierarchal IP addressing scheme with layer-2 switching and non-flat hierarchal MAC/VLAN-based addressing scheme. 2) How to provide comprehensive end-to-end Operations, Administration, Maintenance, and Provisioning (OAM&P) in a unified Ethernet-Optical environment. 3) How to fully integrate layer-2 Ethernet control plane functionality with that of the optical transport layer (layer-1)

The primary rationale behind our vision is decreased cost, complexity and increased scalability in the network. This is how Ethernet won the LAN years ago; it was not necessarily the best technology, it was the most cost-effective and easiest to implement. Simplifying network design and reducing costs by utilizing Ethernet as an end-to-end LAN/MAN/WAN protocol is the cornerstone for Ethernet's success in the MAN and WAN. The outcome of this research will bridge the gap that exists between Optical Networking and Ethernet research communities. This will generate new interest in the research community to expand on the proposed work and further explore this important area.

1. Introduction

The tremendous acceptance of Gigabit Ethernet in most Local Area Networks (LANs) has created pressure on carriers and service providers (SPs) to offer native Ethernet services at gigabit rates in the metropolitan area networks (MANs) environment. Enterprises traffic is rapidly transcending existing capabilities, becoming more complex with a growing percentage of this traffic being time sensitive. Recent advances in Ethernet and optical networking technologies, such as the 10 Gigabit Ethernet standards and long reach optics have enabled a whole new generation of Ethernet networking, pushing Ethernet deeper into the MAN and WAN. Further technology innovations are allowing for Ethernet services to be transported over a variety of media, including SONET/SDH, WDM and Multiple Protocol Label Switching (MPLS) systems, which effectively extend the reach (point-to-point) nationally and globally.

Despite the promise of Ethernet and the potential of becoming the universal telecom service, a number of challenges remain before Ethernet technology can truly support carrier-class Ethernet services over a global scale. Specifically, the vision of implementing a pure layer-2 global Ethernet infrastructure still faces several key technical hurdles, including scalability and end-to-end QoS guarantees. Ethernet's major scalability bottleneck is the use of a spanning tree (ST) to route traffic. A single ST utilizes at most N-1 links in a network of N nodes and allows only one loop-free path, which results in limited resources utilization, uneven load distribution, and potential bottlenecks. Ethernet's path establishment via ST means non-optimal path exists, which could introduce packet loss, jitter, and delay. A second issue is the flat addressing structure and lack of routing hierarchy, which can lead to very large routing tables. The limited VLAN tag space (the IEEE 801.Q standard defines an address space of only 4096 available tags) still poses major scalability bottleneck.

2. Motivations and Proposed Work

Enterprise data traffic is nearly all Ethernet. But once it leaves the corporate LAN (or campuses) and heads onto the wide area, it's translated into some other protocol, only to be translated back into Ethernet once it reaches its destination. What's lost in translation, in this instance, is time and money. These conversions are inefficient and expensive, requiring specialized software on both carrier and customer switches. They're also unnecessary—if native Ethernet can be transported end to end. In addition, Ethernet framing preserves VLAN IDs, QoS tags, and virtually every other packet-level control function available at the MAC layer. Since well over 90 percent of all data traffic originates and ends on an Ethernet LAN, the envisioned data-centric next generation networking infrastructure must have the capability of transporting native Ethernet frames across any segment of the network. Thus, transporting native Ethernet frames end to end from the access network through the metro and core networks to another access network is the most cost effective, simple, and efficient solution.

It is the main objective of the proposed work to scale metro Ethernet networks into a global multi-services infrastructure. Specifically, this work proposes a truly native end-to-end layer-2 MAC frame-based Optical Ethernet infrastructure seamlessly stretching from enterprise LAN to Metro to Global. We show that by combining the simplicity and cost effectiveness of Ethernet technology with the ultimate intelligence of WDM-based optical transport layer, Optical Ethernet (Ethernet-over-WDM) could evolve as a next generation networking paradigm providing a seamless global transport infrastructure for end-to-end transmission of native Ethernet frames.

Today's notion of supporting "IP directly over WDM" (IP/MPLS-over-WDM interconnection models) is little more than cleverly disguised marketing; i.e., IP-over-WDM is almost invariably IP packets mapped into SONET/SDH, coupled with SONET/SDH-based point-to-point DWDM systems. The proposed "Ethernet-over-WDM" model is truly a two-layer model where native Ethernet frames are mapped directly over WDM. It offers advantages over existing Layer-2 and MPLS solutions in that it divorces the Ethernet from legacy transport mechanisms like SONET/SDH and other layer-2 protocols. The key for realizing such a very high-risk initiative rests on devising innovative networking and switching solutions to replace the legacy layer-3 switching (routing) and hierarchal IP addressing scheme (but above we state that flat addressing is an issue with scalability – so why I are you here proposing a non-hierarchical addressing scheme?) with layer-2 switching and non-flat hierarchal MAC/VLAN-based addressing scheme?

To implement the proposed ambitious vision of a global multi-services Ethernet infrastructure, several key critical issues need to be thoroughly examined and addressed including: 1) How to totally eliminate the reliance on ST/RST/MST routing and redundancy functionality? 2) How to reliably transport native Ethernet frames that have no overhead capability to perform network OAM&P across the WAN? 3) How to integrate Ethernet control plane functionality with that of the optical transport layer? 4) How to provide comprehensive Operations, Administration, Maintenance,

and Provisioning (OAM&P) in a unified Ethernet-Optical environment? 5) How to devise a novel global layer-2 MAC and/or VLAN ID-address structure and space that is unique, hierarchal and scalable with a source and destination addresses?

The main characteristics of the proposed Ethernet-over-WDM model are:

1. Conventional Ethernet MAC frames and/or jumbo Ethernet frames must be transported natively (translation into some other protocol is not allowed) end to end from the access network through the metro and core networks to another access network.
2. Only pure layer-2, switching at the packet/frame granularity, is allowed throughout the entire network including access, MAN and WAN.
3. Unlike layer-2 MPLS VPNs (point-to-point and multipoint Virtual Private LAN Services (VPLS)), where a full mesh of static label switched paths (LSPs) must be set up between all L2 VPN sites (I think we do not use static LSPs – are they dynamic via LDP?), the proposed optical Ethernet is dynamically reconfigurable network that supports real-time additions/deletions of all customer connections (EVCs). It can also support a fully automated optical networking service (layer 1) at any bandwidth granularity.
4. Supports an IP/GMPLS-based unified control plane that offers a tighter integration between layer-1 (optical transport layer) and layer-2 (Ethernet layer), leading to the collapse of the two layers into a single integrated layer managed and traffic engineered in a unified manner. The unified control plane supports real-time provisioning and restoration of both full lambda and EVCs by running a single instance of an integrated routing and signaling protocols (use of ST, RST, and MST routing are totally eliminated).
5. Native Ethernet frames are routed across the MAN/WAN using only layer-2 addressing scheme (MAC and/or VLAN ID).

3. Implementation Strategy

It is important to emphasize from the outset that there are strong analogies between the IP-over-WDM interconnection models and the proposed Ethernet-over-WDM model. Anyone who has followed the development of IP-over-WDM interconnection models (the overlay and peer models) throughout 1990s can easily observe that most of the initial problems encountered in the development process were mainly due the optical network-Internet (IP) gap. A partitioning between the optical networking and the IP/MPLS research communities caused this gap. It took the industry, the standards bodies, and the two research communities nearly ten years of extensive continuous hard work and collaboration to narrow this gap. Likewise, we strongly believe that the main problem that will hinder the viability of implementing the vision of a global optical Ethernet infrastructure is the wide gap that exists between the optical networking research communities and the Ethernet communities. It is our expectations that the proposed research program would be an important starting point to bring the two communities together, leading eventually to bridging this gap and the realization of the proposed vision.

Now we have the opportunity to reapply a lot of this technology, suitably modified to Ethernet technology. Our strategy is first to take full advantage of the knowledge and developments gained during the past ten years course of developing the IP-over-WDM interconnection models. The key to a successful strategy rests on taking the best features from both the overlay and peer models while avoiding their limitations. Specifically, it is imperative, when devising the proposed integrated L2-L1 control plane that will manage both GigE and optical switches, to avoid the major limitations of the peer model's integrated control plane, particularly the scalability problem. Previous attempts to address the practical feasibility of implementing the peer model and its integrated control plane failed in large part due to the complexity of the model and the edge router scalability problem.