Carrier-Grade Ethernet for Core Networks

Andreas Kirstädter, Claus Gruber, Johannes Riedl, Thomas Bauschert
Siemens Networks, Otto-Hahn-Ring 6, D-81730 München, Germany
andreas.kirstaedter@siemens.com

Abstract: This paper considers the functionality and standards required to enable carrier-grade core networks based on Ethernet-over-WDM. Possible Ethernet backbone network architectures will be discussed followed by an evaluation of the CAPEX and OPEX performance.

OCIS Codes: (060.4250) Networks; (999.9999) Ethernet over WDM, Economic analysis and business cases

1 Introduction

Backbone networks represent the top of the carriers' network hierarchy connecting networks of different cities, regions, countries, or continents. The complexity of these technologies imposes substantial financial efforts on network operators, both in the area of Capital Expenditures (CAPEX) and Operational Expenditures (OPEX).

Ethernet claims to be a possible enabler of cost-efficient networks, as it is characterized by simplicity, flexibility, interoperability, and low cost. While Ethernet is traditionally a Local Area Network technology, continuous developments already enabled its deployment in Metropolitan Area Networks. Recent research and standardization efforts aim at speeding up Ethernet to 100 Gbit/s and at resolving scalability issues, thus supplying Ethernet with carrier-grade features for core networks.

2 Carrier-Grade Ethernet-Based Core Networks

2.1 Carrier-Grade Requirements

In order to be suited for core networks, Ethernet needs carrier-grade performance and functionality. It has to offer and implement the required Quality-of-Service (QoS) and has to enable traffic engineering to fine-tune the network flows. Furthermore, it has to provide fast and efficient resilience mechanisms to recover from link and network element failures and has to enable various Operation, Administration, and Maintenance (OAM) features for the configuration and monitoring of the network. Last but not least, it has to provide secure network operation. Additionally, a high degree of scalability is needed for handling different traffic types and for user separation inside the network. This scalability in terms of address space, maximum transmission speed, and maximum transmission distance becomes an important issue for the next Ethernet generation. E.g., multi-layer operation and optimization can only be used if facilitated by reasonable values of the maximum transmission distance.

At a closer look, it becomes visible that many of these required features are currently implemented repeatedly at different network layers. E.g., resilience mechanisms are found in the WDM layer and in an intermediate Sonet/SDH layer as well as in the packet layers above them. A cost-efficient network and protocol architecture therefore has to evaluate these functional redundancies between the layers very carefully.

2.2 Forwarding Technology and Scalability

The necessary scalability requires new approaches to packet switching and forwarding within meshed end-to-end Ethernet networks. Traditionally, within Ethernet networks the Spanning Tree Protocol (STP) calculates a single tree structure based on configurable IDs of switches, configurable port weights, and priorities to connect any switch with each other. Although loop-less forwarding is guaranteed with this mechanism, STP provides only one path between two locations and a MAC address learning of any equipment is performed at the switches.

However, in the case of combining large networks and adding hundreds of customer networks with an Ethernet-based core network the number of MAC addresses will grow rapidly. Thus, scalability can no longer be provided with current layer-2 approaches and a separation of networks or an additional hierarchy between them has to be introduced to allow a scalable forwarding of data.

Also, the use of a single tree structure providing only a single path between two locations prevents the use of efficient traffic engineering and resilience mechanisms. Thus, several connection-oriented forwarding techniques for carrier-grade Ethernet transport networks are currently under discussion at standardization bodies: VLAN Cross-Connect (VLAN-XC), Provider Backbone Transport (PBT), and Transport Multi-Protocol Label Switching (T-MPLS) [1].

3 Multi-Layer Operation and Optimization

Another important aspect in the area of scalability is the maximum transmission distance of Ethernet signals. Multilayer network grooming approaches are very attractive for the purpose of reducing unnecessary packet processing in intermediate nodes [2] as transit traffic is allowed to bypass intermediate nodes. Traffic between two network edge nodes can either be transported transparently in the optical domain or can be converted to the electrical domain to allow electrical grooming along the path. The effort spent on extending the signal reach of Ethernet signals is rewarded by equipment savings. Figure 1 illustrates the possible port-count savings in an Ethernet core-network where optical grooming is applied up to the maximum transmission distance avoiding unnecessary electrical processing of transit traffic. These results were derived for an Ethernet network following the topology of a generic German backbone that is also used during the CAPEX and OPEX analyses below [3].

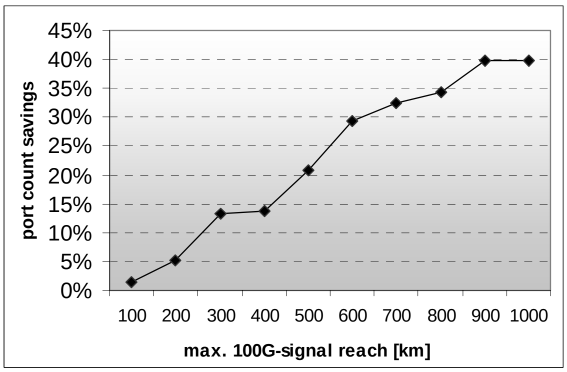

Figure 1. Port-count savings in grooming-enabled Ethernet networks.

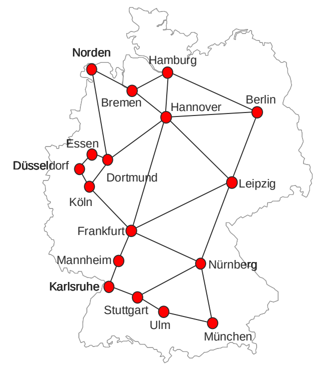

Reference network topology.

Next to IP services, VPN business services transfer increasing traffic and generate high revenues for network providers. In particular, Ethernet services (E-Line and E-LAN) are evolving. Today these layer-2 services are commonly transported via IP/MPLS tunnels. However, the complex functionalities and protocols of the IP layer are often not required to transport these pure layer-2 services. Native End-to-end Ethernet structures will arise where Ethernet business services will be transported on pure layer-2 infrastructures without the need of complex data transformations and changes in the functional layer structure. In continuation, also a merger between Ethernet-based packet transport and IP networking appears on the horizon further reducing superfluous functional redundancies between the single packet protocol layers – just in the sense of the current clean-slate thinking for the future Internet.

4 CAPEX and OPEX Performance

In order to calculate the total CAPEX of specific network architectures future traffic loads, network device counts, and network device prices have to be estimated. For the physical reference topology shown above and for the years 2009 - 2012 typical traffic matrices were extrapolated assuming a homogenous traffic growth rate of 40% per annum. A shortest-path routing algorithm was then applied to determine the single link loads from which the number of switches, routers, and line card ports were finally obtained depending on the network architecture. Future equipment prices were extrapolated following a careful analysis of market data and price development of the last four years (for details on this approach see [3]). The following *generic* network architectures were considered:

(a) IP/PoS-over-WDM: The backbone consists of Label Edge Routers (LERs) and Label Switch Routers (LSRs) that are all equipped with Packet-over-Sonet (PoS) interfaces. Sonet/SDH is only used for transporting the IP packets node to node - directly over WDM. A 1+1 protection scheme is applied.

(b) IP/PoS-over-SDH-over-WDM: Similar to (a), this network scenario considers a backbone where LERs are located at the ingress and egress points of the backbone. However, the traffic is groomed via SDH add-drop-multiplexers and cross-connects inside the core. A 1+1 protection scheme is applied.

(c) IP/PoS-over-OXC-over-WDM: This case is similar to the previous architecture. However, grooming is done by optical cross-connects (OXCs). The range of the optical signal is assumed to be large enough to enable end-to-end optical grooming. Thus, again no LSRs are required inside the network. A 1+1 protection scheme is applied.

(d) IP/MPLS-over-Ethernet-over-WDM: LERs at the edge and MPLS-enabled Ethernet switches in the core of the backbone (MPLS-Ethernet architecture). 1:1 protection scheme used: I.e. all capacity over-provisioned by 100%.

(e) Ethernet-over-WDM: The considered and the connecting (core) networks are native Ethernet networks with Ethernet switches both at edge and core. A reduced number of LERs are deployed to handle a small share of traffic that requires IP routing (share assumed to be 30%). However, Ethernet traffic does not have to traverse LERs at the ingress and egress points of the backbone. 1:1 protection is applied.

(f) Ethernet-over-WDM with service-level protection: The network architecture is identical to (e) except that only premium traffic (share set to 30%) is protected against failure.

If we split the architecture-dependent CAPEX components according to the different layers it first becomes visible that a pure PoS-over-WDM network (a) creates relatively high cost as only expensive PoS router-interfaces are used. The high prices of PoS interfaces generally lead to a high cost component of LERs for all SDH-related infrastructures. PoS-over-SDH architectures (b) prove to be much cheaper as the SDH network employs SDH switches and interfaces instead of expensive LSR equipment for core switching. A PoS-over-OXC network (c) has an even better CAPEX performance due to lower switch and optical transceiver prices of OXC hardware.

Without exception, the three Ethernet-over-WDM scenarios perform better than the previously mentioned architectures. In the MPLS-Ethernet case (d), still a considerable amount of CAPEX is related to LERs and their interfaces. A native 100 Gbit/s Ethernet over WDM network (e) enables higher savings in the LER category. Applying a service-level differentiated protection scheme (f), the CAPEX can be reduced even further.

The OPEX were evaluated via a process-oriented approach [4]. While generally OPEX include a lot more processes, the repair process was selected in this study since the impact of 100 Gbit/s Ethernet gets most visibly in this area. For each network architecture described above, the OPEX per year were evaluated via first determining the total number and type of equipment. By using availability figures [5] the average repair time for a given backbone architecture was estimated. The related costs were derived by multiplying the total repair time with the average salary of a field or point-of-presence technician. As a general result, 100 Gbit/s Ethernet networks are more economical due to the reduced device count (less switches and line cards). A service-level protection scheme further reduces the required network transport capacity, the network element count, and thus the related OPEX.

5 Related Research Fields

Ethernet transmission at speeds of 100 Gbit/s over long distances is very desirable in terms of architecture-related network cost. The transmission of high speed data rates above 100 Gbit/s is well understood (e.g. [6]) although second degree (slope) chromatic dispersion has to be exactly compensated, birefringence effects become grave, and the signal-to-noise ratio of 100 Gbit/s signals is generally lower as fewer photons are transmitted per optical impulse. Additionally, recent trials demonstrated the ability to process electronically the required bit rates of 107Gbit/s [7]. The major problem still is to find efficient optic-electrical and especially electro-optical conversion techniques for these high speeds. Pure electrical solutions are preferable to handle the data at the transmitter and receiver since OTDM techniques still seem to be too complex and difficult to implement in commercial products.

6 Conclusions

In the past, Ethernet evolved from LAN into Metro areas covering speeds from 10 Mbit/s up to 10 Gbit/s. Next-generation Ethernet with transmission-speeds of 100 Gbit/s will facilitate cost-efficient Ethernet transport. As soon as carrier-grade issues like scalability, network resilience, QoS, and OAM of Ethernet-based core network architectures are solved we might see a complete Ethernet-over-WDM core-network infrastructure with resolved redundancies between the single layers.

7 Acknowledgements

The author wishes to thank A. Schmid-Egger for the input to the CAPEX and OPEX analysis. Special thanks go to the whole E3-project team at Siemens Networks.

8 References

1 A. Kirstädter, C. Gruber, J. Riedl, T. Bauschert, “Carrier-Grade Ethernet for Packet Core Networks“, Proc. APOC2006, Gwangju, Korea, Sept. 3-7

2 M. Scheffel, R. G. Prinz, C. G. Gruber, A. Autenrieth, D. A. Schupke, “Optimal Routing and Grooming for Multilayer Networks with Transponders and Muxponders“, IEEE Globecom, San Francisco, USA, 2006.

3 A. Schmid-Egger, A. Kirstädter, “Ethernet in Core Networks: A Technical and Economical Analysis”, Proc. HPSR2006, Poznan, Poland, June 2006.

4 S. Pasqualini, A. Kirstädter, A. Iselt, R. Chahine, S. Verbrugge, D. Colle, M. Pickavet, P. Demeester, ”Influence of GMPLS on Network Providers’ Operational Expenditures: A Quantitative Study”, IEEE Communications Magazine, Vol. 43, No. 7, pp. 28-34, July 2005.

5 NOBEL Project, “Availability Model and OPEX related parameters”, to be published.

6 J.P. Turkiewicz, E. Tangdiongga, G. Lehmann, H. Rohde, W. Schairer, Y.R. Zhou, E.S.R. Sikora, A. Lord, D.B. Payne, G.D. Khoe, and H. de Waardt, “160 Gb/s OTDM Networking using Deployed Fiber”, J. Lightwave Technol. (invited), Volume 23, Number 1, January 2005.

7 C. Schubert, R. H. Derksen, M. Möller, R. Ludwig, C.-J. Weiske, J. Lutz, S. Ferber, C. Schmidt-Langhorst “107 Gbit/s Transmission Using An Integrated ETDM Receiver“, Proc. ECOC2006, Cannes, France, Sept. 24-18, 2006.