GMPLS-controlled Ethernet segment protection

Wouter Tavernier, Didier Colle, Mario Pickavet, Piet Demeester
Ghent University – IBBT, Department of Information Technology,
Gaston Crommenlaan 8, 9050 Gent, Belgium
{wouter.tavernier, didier.colle, mario.pickavet, piet.demeester}@intec.ugent.be

Abstract

Efforts of standardization bodies illustrate that Ethernet is conquering carrier environments. Resulting technologies intend to overcome Ethernet bridging limitations. Validated through emulation, we show the gain of segment protection in connection-oriented Ethernet.

Introduction

For decades Ethernet is and has been known as the technology dominating the LAN environment. Its low-cost, plug-and-play characteristics and extensive backward compatibility result in the high interest it is now receiving for usage in access and metro networks. Initiatives such as IEEE PBB(-TE), ITU-T T-MPLS and IETF GMPLS extensions to support these illustrate that Ethernet indeed is conquering the carrier market. The main driver for these developments can be found in the lack of scalability, resource efficiency, manageability and robust resilience in standard Ethernet bridging.

Rapid Spanning Tree Protocol-based recovery

The plug-and-play character of connectionless Ethernet as we know it is mainly driven by the automatic learning and flooding behaviour of the technology. To avoid that nodes endlessly flood frames the topology has to be loop-free. This is enforced by restricting the physical topology to a logical tree through the (R)STP protocol. Not further mentioning the issue of resource efficiency this brings about, its (re-)convergence characteristics are subject to the well-known count-to-infinity problem and prove to behave badly in larger ring topologies as deeply investigated in [1]. The first can be easily understood resulting from the fact that (R)STP is essentially a distance-vector routing protocol. The second aspect is a result of the agreement/proposal mechanism used by RSTP: control frames will travel half way around the ring to reach bridge X which has a (blocked) alternate port which leads to the root in the other direction. Bridge X will answer back to inform the root neighbor that the root is still alive. The obvious result of this is that the larger the ring, the more is suffered from the needed propagation time of the handshake (see figure 2).

GMPLS-controlled Ethernet

Becoming a carrier technology, Ethernet meets the GMPLS protocol suite as an interesting control possibility. With its increasing PHY speeds, optical Ethernet has now become a very interesting alternative at the physical layer. Therefore the IETF ccamp working group started some efforts for supporting new connection-oriented Ethernet data planes to be controlled by, and fitting in the hierarchy of GMPLS. Two newly evolving connection-oriented Ethernet data planes are PBB-TE and Ethernet Label Switching

(ELS). Both disable (R)STP, MAC learning/flooding and allow an IP-based control plane such as GMPLS (incl. OSPF-TE and RSVP-TE) to enable traffic engineering, advanced protection and OAM. This paper uses the ELS data plane to evaluate GMPLS-controlled protection switching. ELS forwarding uses the translation capability of IEEE 802.1ad-based switches to allow for connections based on link local S-VID (Service VLAN ID) switched frames. This way a frame going through an ELS connection, undergoes a series of push and pop operations on the SVID field similar as in MPLS. The resulting connection is called a Ethernet Label Switched Path (ELSP). Segment Protection of ELSPs now becomes possible as in [5]. This way, any segment between two points in a network can be protected by a precomputed backup segment, provided that a cycle in between both points is available. Fault detection mechanisms at various layers then allow triggering a fast protection switchover. Detection can be performed at the hardware layer, but can also be done at a higher levels, using for example Bidirectional Forwarding Detection ([6]) which will be used in these tests.

Figure 1 - Emulation platform node architecture

Emulation architecture

Few or no equipment is currently available for benchmarking these evolving technologies. Therefore we have set up a platform based on open-source software running on Linux. We emulated the ELS data plane using the Click Modular Router Platform ([2]). Click was extended such as to enable connection-based forwarding using SVID label swapping and to perform BFD fault detection. The control plane software used for tests is based on VLSR software of the Dragon project [3]. We extended this GMPLS package with protection signaling and

notification as defined in [5] and the signaling of ELSPs. The node architecture is shown in figure 1.

Results

A performance comparison of RSTP-based recovery opposed to GMPLS-controlled segment recovery was made in the emulation environment. Fault detection heavily influences overall recovery time. Therefore we distinguish between hardware detection (loss of signal detection) and software based detection. These are available in evaluated technologies, but here we only used the software detection. The software failure detection mechanism of RSTP needs 3 times the expiration of the hello interval, which results in a detection time of 6 seconds. Because this would put the figure out of scale, this is not included into the results. Obviously this is already a major drawback of RSTP-based detection (and thus recovery) compared to BFD-based detection, where we measured detection times of tens of milliseconds.

Figure 2 - RSTP vs. GMPLS-protection in a 4-node ring

For the performance comparison we started with a 4-node ring network in which we break a link and evaluate reconvergence time of discussed technologies. In a second step, the effect of larger ring topologies was evaluated. GMPLS-controlled segment/path protection is done by dividing the ring into two halves, using one half as a protection segment for the other.

In both the 4-node network and larger ring topologies, GMPLS-controlled segment protection proved to behave stable regarding time performance. This is a logical consequence given the fact that both failure detection and switch-over steps are independent of the topology. This is illustrated by the following time diagram.

Timeline (ms)	Event
0	Failure
25	BFD Failure notification (5ms TX/RX multiplier 5)
26	Starting triggering protection switch.
27	Send notify to branch/merge
35	Receive notify ack
35	Start fast-switchover
38	Switch-over finished

However, adding an extra switch into the ring does influence the recovery time of RSTP, resulting into an additional recovery time of 3 ms for RSTP. As expected, this has no effect on switch-over time of GMPLS-controlled segment protection.

The resulting time performance of both techniques is shown in Figure 3. Here we see that for very small networks (less than 7 nodes including failure detection) both RSTP and GMPLS-controlled protection perform very well, meaning sub 50 ms performance, which is typically called carrier-grade performance. From the moment larger topologies are used, RSTP time performance linearly decreases, while GMPLS-controlled protection keeps performing at carrier grade level.

Future development and benchmarking is planned to evaluate the gain of traffic engineering, OAM and manageability characteristics of GMPLS-controlled Ethernet compared to competing technologies in the access and metro network market.

Figure 3 - Reconvergence times on rings with x nodes

Conclusions

Ethernet resilience mechanisms as we know them in LAN environments are not carrier grade. We verified this both in literature and in emulation. We presented technologies and mechanisms under construction to overcome this and verified the value of segment protection in custom made emulation platform.

References

- [1] K. Elmeleegy, "On Count-to-Infinity Induced Forwarding Loops in Ethernet Networks", INFOCOM 2006., pp. 1-13.
- [2] E. Kohler, "The Click modular router", ACM Transactions on Computer Systems, 1999, Volume 18 number 3, p236-297
- [3] Dragon Virtual Label Switch Router (VLSR), <http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/VLSRL>.
- [4] Berger, "IETF RFC 4873 - GMPLS Segment Recovery",
- [5] D. Katz, IETF draft "Bidirectional Forwarding Detection", draft-ietf-bfd-base-08.txt

KEIO UNIVERSITY

GLOBAL COE PROGRAM

**High-Level Global Cooperation
for Leading-Edge Platform on Access Spaces**

2nd GENT Univ. and Keio Univ. Global COE joint workshop 2008

Sep. 26th-27th 2008 Gent University, Gent, Belgium