ETH zürich

Prof. Laurent Vanbever
Networked Systems Group

Route Convergence Optimization in the SWITCH Network

Master thesis proposal

The backbone network operated by SWITCH connects most of the Swiss higher educational institutions, including the ETH Zurich. It consists of about 3000 km of fiberoptic cables and on the order of 100 routers. It uses OSPF [4] and BGP [5] to route traffic between customer networks and the global Internet, carrying around 300k prefixes in its BGP table.

The majority of external traffic is exchanged with neighboring autonomous systems via the switching fabrics of Internet exchange points or direct links (private network interconnections, PNI). Some of those peering sessions, especially those with "route servers" at prominent exchange points like AMS-IX and DE-CIX, are associated with a large number of prefixes.

In case of a failure on one of these peering sessions, the internal BGP mesh [2] is flooded with update messages causing a massive amount of updates to the routing and forwarding information bases throughout the backbone network until the systems have converged to the new routing topology. In some cases, this process was observed to have taken up to several minutes to complete, leading to network outages for the customer networks.

The goal of this thesis is to analyze the dynamics of such events and propose changes to the router configurations to minimize route convergence time and thus impact on customer traffic. The work can be divided into for major parts:

a) Understand the current routing setup in the SWITCH network (OSPF, BGP, IPv4/IPv6)

b) Analyze the effect of large-scale routing changes in the internal BGP system

c) Research the mechanisms for faster route convergence available on the routing platforms in use at SWITCH (e.g. next-hop tracking [1] and hierarchical FIBs [3])

d) Propose a set of changes to the routing protocol configurations to reduce the impact of routing changes on customer traffic

Requirements

- Good understanding of BGP and OSPF (IPv4 and IPv6)

Contact

- Tibor Schneider, sctibor@ethz.ch
- Prof. Laurent Vanbever, lvanbever@ethz.ch
- Simon Leinen, simon.leinen@switch.ch
- Alexander Gall, gall@switch.ch

References

[1] A. Ala, M. Essaaidi, et al. Fast convergence mechanisms and features deployment within operator backbone infrastructures. In *2009 Mediterrannean Microwave Symposium (MMS)*, pages 1–5. IEEE, 2009.

[2] T. Bates, E. Chen, and R. Chandra. Bgp route reflection: An alternative to full mesh internal bgp (ibgp). RFC 4456, 2006.

[3] C. Filsfils. Bgp convergence in much less than a second. *Presentation at NANOG*, 40, 2007.

[4] J. Moy. Ospf version 2. RFC 2328, 1998.

[5] Y. Rekhter, T. Li, S. Hares, et al. A border gateway protocol 4 (bgp-4). RFC 4271, 1994.