On Understanding OSPF Convergence Dynamics in Presence of Multiple Failures

Dan Zhao[1], Xiaofeng Hu[1] and Chunqing Wu[1]

[1] School of computer, National University of Defense Technology
Changsha, Hunan, China
danzhao.nudt@gmail.com, {xfhu,wuchungqing}@nudt.edu.cn

Abstract. Open Shortest Path First (OSPF) is a popular link state routing protocol used in Internet infrastructure. It usually takes several tens of seconds for OSPF network to recover from a failure. The convergence time is delayed mainly by failure detection and routing calculation scheduling. In this paper we analyze OSPF convergence behavior in presence of multiple failures, where the interactions between failure detection and routing calculation scheduling could generate complicated dynamics during convergence process. We also present experimental study to understand the impact of multiple failures on convergence. The results demonstrate that multiple failures have a greater chance to delay the convergence. This suggests that operators should take it into account while configuring OSPF network.

Keywords: OSPF, link state protocol, multiple failures, convergence

1 Introduction

Open Shortest Path First (OSPF) [1] is a successful link state protocol widely used in intra-domain ISP networks. In OSPF network, every router establishes adjacency with its connected counterparts and describes the connection status using Link State Advertisement (LSA). Once the topology changes, routers utilize specific mechanism, typically Hello protocol prescribed in OSPF, to detect the failure and generate new LSAs. After the synchronization of LSAs throughout the network by flooding, routers are capable of calculating the correct routing table for packet forwarding.

Multiple timers, such as ***HelloInterval***, ***spfDelay*** and ***spfHold***, are implemented to limit protocol overhead. However, the reaction and interaction of timers are likely to increase convergence dynamics, especially in presence of multiple failures. In recent years, there has initiated a research trend in measuring and analyzing the impact of multiple and regional failures [3], [4]. In this paper we aim to understand OSPF convergence behavior in presence of multiple failures, and demonstrate its intrinsic cause.

2 Impact of Multiple Failures on Convergence

In this section we will illustrate the details of OSPF convergence behavior in presence of multiple failures.

2.1 Asynchronous Detection of Multiple Failures

The adjacency between each pair of routers is independent from each other. More precisely, it depends on the timing of hello exchange on specific interface. Therefore, the failure detection time (*FDT*) of the two neighboring routers differs (denoted as ΔFDT) because they connect to different failed routers and have asynchronous failure detection behavior. Given that hello packet is sent every ***HelloInterval***, the most time variation can't exceed ***HelloInterval*** range. Thus it would be only delayed by ***spfDelay*** and the SPF algorithm is executed for only once.

2.2 Scheduling Routing Calculation

When routers have receives LSAs, routing calculations are scheduled. The calculation is normally delayed by ***spfDelay*** and ***spfHold***, depending on the timing of last scheduling. When $\Delta FDT_{max} \leqslant$ ***spfDelay***, it means all LSAs would arrive at a particular router within [0, ***spfDelay***] range (assuming propagation delay is negligible). However, if $\Delta FDT_{max} >$ ***spfDelay***, there must be some LSAs arrive at a router after the first routing calculation. Therefore ***spfHold*** can take effect, and the convergence is further delayed. Considering that multiple failures greatly change the network topology that requires reconstruction of many end-to-end paths, delayed convergence definitely slows down the reconstruction and largely impact the service quality of applications.

3 Experimental Study

We perform experiments on an emulation system named *CORE* [5] that uses *OpenVZ* [7] virtualization technology. Each virtualized node in *CORE* runs *Quagga* [8]. The ***spfDelay*** is set to default value of 5s, and the initial and maximum value of ***spfHold*** is 1s and 10s respectively. We use real ISP backbone topology, AS3967, reported from Rocketfuel [6] and monitor the convergence time to see how convergence is delayed in presence of multiple failures.

In order to clarify the delay by routing calculations, we refer to convergence time here as the duration from the moment that the failure is initially detected to the time the last router has updated its routing and forwarding table. The result of convergence time is shown in figure 1.

When multiple failures happen, more delay is introduced into convergence process. When ***HelloInterval*** is set to 10s, only less than 30% of experiments converge less than 6s where there is only single run of routing calculation delayed by ***spfDelay***. For those convergence takes time in [6, 7] range, a considerable number of routers experience 2 routing calculation delayed by ***spfDelay*** and ***spfHold*** which is 1s at the beginning. We can observe that convergence time of some scenarios locate in [8, 9] range. This is because routing calculation is scheduled for 3 times, and the last delay of ***spfHold*** has increased to 2s. When convergence time exceeds 11s, there are two situations. In some circumstances, there are 4 successive running of routing calculations and ***spfHold*** is adjusted to 3s. Thus the delay adds up to 11s. Another possibility is that routing calculation is scheduled twice, but the second is scheduled beyond ***spfHold*** after the fist calculation. Therefore the second convergence is delayed for 5s as well, and the additional time comes from the detection time variation. In our experiments, there is more than 50% convergence that takes more than 10s to stabilize the network after multiple failures.

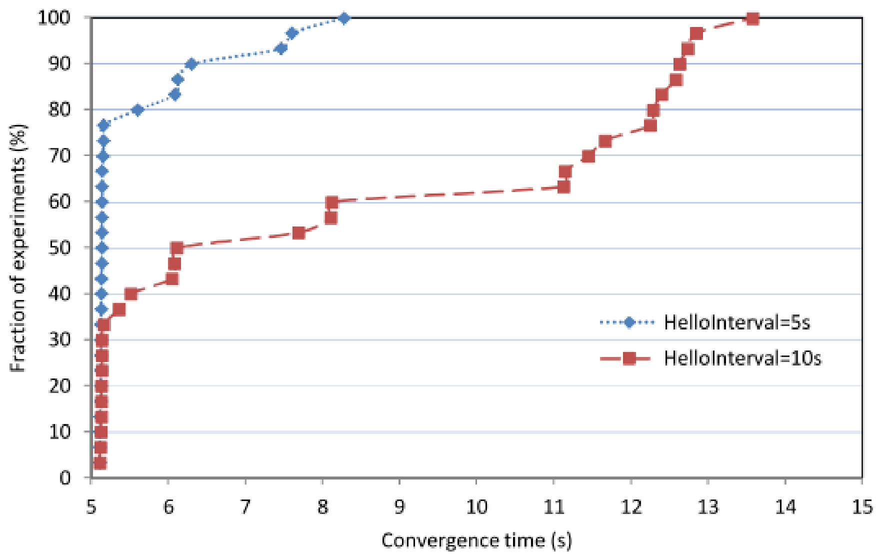

Fig. 1. Convergence time in presence of multiple router failures

4 Conclusion

In this paper we formally analyze OSPF convergence dynamics in presence of multiple failures. According to our analysis, the convergence is mainly and largely delayed by protocol timers. To our point of view, the cause of such dynamics mainly lies in that the detection of multiple failures is asynchronous. Reducing detection timers could help alleviating the impact, whereas smaller timers may overreact to subtle network change and enlarge network instability. The results of our experimental study also demonstrate that convergence is greatly delayed in presence of multiple failures. This suggests that network operator should carefully configure OSPF protocol taking into account their dependency and possible network failure scenarios, aiming that network could gain faster convergence while keeping the processing overhead in considerable level. Furthermore, we only refer to concurrent multiple failures in this paper. We believe that the failures which have cascading behavior can result in tremendous impact than normal multiple failures, and this is the issue that we hope to address in the future.

Acknowledgments. The work described in this paper is supported by the NSFC under Grant No.61103189 and No.61070199, Program for Changjiang Scholars and Innovative Research Team in University (No.IRT 1012), Program for Science and Technology Innovative Research Team in Higher Educational Institutions of Hunan Province: "network technology", and Hunan Province Natural Science Foundation of China (11JJ7003).

References

1. J. Moy, OSPF version 2, Internet Engineering Task Force, Request For Comments(Standards Track) RFC 2328, April 1998.
2. P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, Achieving sub-second IGP convergence in large IP networks, Computer Commun. Rev., vol. 35, no. 3, 35-47 (2005)
3. Sebastian Neumayer and Eytan Modiano, Network Reliability With Geographically Correlated Failures, in Proc. of IEEE INFOCOM, 2010
4. Sebastian Neumayer, Gil Zussman, Reuven Cohen, and Eytan Modiano, Assessing the Impact of Geographically Correlated Network Failures, in Proc. of IEEE MILCOM, 2008
5. J. Ahrenholz, Comparison of CORE Network Emulation Platforms, In Proc. of IEEE MILCOM Conference, pp. 864-869, 2010
6. N. Spring, R. Mahajan and D. Wetheral, Measuring ISP topologies with RocketFuel,In Proc. of ACM SIGCOMM, pp. 133-145, 2002
7. OpenVZ Linux Containers, http://wiki.openvz.org/
8. Quagga Routing Suite. http://www.quagga.net