An evaluation of IP-based Fast Reroute Techniques*

Pierre Francois, Olivier Bonaventure
Dept. CSE, Université catholique de Louvain (UCL)
Belgium
francois@info.ucl.ac.be, bonaventure@info.ucl.ac.be

Categories and Subject Descriptors: [C.2.2]: Routing Protocols

General Terms: Reliability, Performance, Algorithms.

Keywords: IP, Routing, Fast-Reroute.

1. INTRODUCTION

Today, IP-based networks are used to carry all types of traffic, from the traditional best-effort Internet access to traffic with much more stringent requirements such as real-time voice or video services and Virtual Private Networks. Some of those services have strong requirements in terms of restoration time in case of failure. When a link or a router fails in an IP network, the routers adjacent to the failing ressource must react by distributing new routing information to allow each router of the network to update its routing table. A realistic estimate of the convergence time of a tuned intradomain routing protocol in a large network is a few hundred of milliseconds [1].

For some mission critical services like voice or video over IP, achieving a restoration time in the order of a few tens of milliseconds after a failure is important [2]. In this paper, we first present several techniques that can be used to achieve such a short restoration time. While most of the work on fast restoration has focussed on MPLS-based solutions [2], recent work indicate that fast restoration techniques can be developed also for pure IP networks. Recently, the RTGWG working group of the IETF started to work actively on this problem and several fast reroute techniques are being discussed. However, as of today, no detailed evaluation of the various proposed IP-based fast reroute techniques has been published.

The goal of this short paper is to firstly provide a brief overview of fast restoration techniques suitable for pure IP networks, in section 2. Then, in section 3, we evaluate by simulation how many links can be protected by each technique in large ISP networks based on their actual topology. This coverage is an important issue as some techniques cannot protect all links from failures.

*This work was supported by Cisco Systems within the ICI project. Any opinions, findings, and conclusions or recommandations expressed in this paper are those of the authors and do not necessarily reflect the views of Cisco Systems.

CoNEXT'05, October 24–27, 2005, Toulouse, France.
ACM 1-59593-035-3/05/0010.

2. IP FAST REROUTE TECHNIQUES

The first technique proposed, implemented and deployed to quickly reroute IP packets when a link fails is to use MPLS's label swapping forwarding [2]. In IP networks that are not using MPLS to forward IP packets, it is possible to use MPLS only to provide protection. The MPLS protection LSP can be established by the protection router by using RSVP-TE. If the network is bi-connected, then those MPLS LSP can be used to protect any single link failure. Thus the coverage of this technique is 100%. Its main drawback is that it requires to enable RSVP-TE in the network even if it is not utilized to forward packets when the network is stable.

The first IP-based protection technique being considered within the IETF is the utilisation of loop-free alternates [3, 4]. If a router I is using a link $I \rightarrow J$ to reach destination d, then a loop-free alternate is a direct neighbour, say router N, of router I if N reaches destination d without using link $I \rightarrow J$. When the link $I \rightarrow J$ fails, router I can send the packets towards d to N instead of J and those packets will reach d. Formally, a loop-free alternate for destination d at router N is defined in [3] as a router N such that $Cost(N \rightarrow \ldots d) < Cost(N \rightarrow \ldots I) + Cost(I \rightarrow \ldots d)$. Since routers use shortest path routing, an equivalent condition is that $(I \rightarrow J) \notin ShortestPaths(I, d)$. Although loop-free alternates are defined on a per destination basis in [3], we argue that from an implementation viewpoint, this is not the best solution. Assume that 1000 destination prefixes are reached by router I via link $I \rightarrow J$ and that all those destinations are protectable with a loop-free alternate. When router I detects a failure of link $I \rightarrow J$, it should be able to update its FIB to point each of those 1000 entries to their respective loop-free alternates. Given the time required to update a FIB entry [1], this could be above the 50 millisecond budget. A better approach is to consider only the loop-free alternates that are able to protect *all* destinations that are currently reached via the link to be protected. Formally, a neighbour N will be a valid downstream node to protect link $I \rightarrow J$ if $(I \rightarrow J) \notin SPT(N)$. We will show in section 3 that even by using this more constraining condition, it is possible to protect a large fraction of the links in real ISP networks.

A closer look at ISP topologies showed that when there is no loop-free alternate to fully protect a link, there is often a router two hops away that does not utilize the link to be protected. This motivated the introduction of U-turns in [5]. A neighbour U of a router I can act as a U-turn to protect link $I \rightarrow J$ if one of its neighbours, say router

R, does not utilize link $I \rightarrow J$ inside its SPT. To serve as a U-turn alternate, router U must be able to support two types of forwarding. When the network is stable, router U uses its normal FIB to forward packets. For the packets affected by the failure that are *u-turned* by router I, router U must detect that these are affected packets and forward them directly to the alternate router, router R without using its normal FIB. Compared with the loop-free alternates, the main drawback of the U-turns is that they require a cooperation among the neighbours and some modifications to the interfaces.

The loop-free and U-turn alternates discussed in the previous section are not sufficient to provide a full coverage in large networks. This coverage can be improved by using IP tunnels as proposed in [4]. These tunnelling schemes can be used to create virtual links between routers. While in the past packet encapsulation and decapsulation were performed by the central CPU with a limited performance, interfaces on current high-end routers are now able to encapsulate and decapsulate tunnelled packets at wire speed. To protect the directed link $I \rightarrow J$, router I needs to find a router N that is reachable without using the link to be protected and that is also able to forward packets to any destination without using link $I \rightarrow J$.

A method to find such a the tunnel endpoint was proposed in [4]. To protect link $I \rightarrow J$, router I must compute the intersection of the set of routers that it reaches without using the link and the set of routers that does not use the link. If the set contains several routers, then a criteria must be defined to select the best one. If the set is empty, then no protection tunnel can be established to protect this link. When it is not possible to find a valid tunnel endpoint to protect a link, a protection tunnel may still be used by forcing one potential tunnel endpoint to forward all the packets received via the tunnel to a particular neighbour, instead of using its FIB. We will see that using this technique allows to protect more links in some topologies.

A last protection technique was proposed recently in [6]. This solution can be considered as an extension of the protection tunnels described earlier, but it requires a cooperation among all the routers of the network. Intuitively, the idea of this solution is that to protect link $I \rightarrow J$, router I should be able to send the affected packets inside a tunnel towards a special address of router J : J_I. This address is a special *not via* address. Its semantics is that all routers of the network must have computed their FIB such that they *never* use link $I \rightarrow J$ to forward packets towards destination J_I.

3. COVERAGE OF THE IP FAST REROUTE TECHNIQUES

A potential issue with the IP-based fast reroute techniques is that several mechanisms may be required to fully protect all links in networks. All protection techniques are not equivalent in terms of coverage and complexity. As a complete protection is required, we implemented a simulator to test the simpler techniques first and only try to use the more complex techniques when the simple techniques do not suffice.

To evaluate the network coverage of the IP-based fast reroute techniques, we considered five very different ISP topologies. The first one is Abilene, a research network deployed over the continental US. It is composed of 11 routers and 14 (28 directed) links. The second one is GEANT, a pan-European research network, composed of 36 (72 directed) links. ISP1 is a commercial network covering a European country. The core of this network is composed of 190 directed links (64 directed links are backup links) and 50 routers. ISP2 is a also commercial network in a European country. The core of this network is composed of 11 routers and 26 links. ISP3 is a Tier-1 ISP whose core is composed of 83 routers and 286 directed links. Due to the setting of the IGP weights, 21 directed links do not carry traffic and one link is only used in one direction. In this network, the setting of the IGP weights was tuned to meet some specific traffic requirements.

Table 1 summarises the coverage of the IP-based fast recovery techniques in the studied network topologies. It clearly shows that by combining loop-free alternates, U-turns and protection tunnels, it is possible to protect all links in real ISP topologies. The values describe the percentage of links that can be protected by combining the first protection techniques. For example, in GEANT all links are protected by using LFA, U-turns and protection tunnels, while in Abilene protection tunnels with directed forwarding are required in addition to the techniques used in GEANT. The *not-via* addresses were not necessary to protect unicast IP traffic in the topologies that we considered.

Network	Links	LFA	U-turns	Tunnel	Directed Tunnel	Notvia
Abilene	28	42%	85%	92%	100%	-
GEANT	72	66%	93%	100%	-	-
ISP1	114	54%	71%	71%	100%	-
ISP2	26	15%	42%	100%	-	-
ISP3	265	65%	95%	96%	100%	-

Table 1: Combined coverage of loop-free alternates, protection tunnels and not via address

4. CONCLUSION AND FURTHER WORK

In this paper, we showed by simulation that loop-free alternates combined with U-turns are sufficient to protect between 40 and 90% of the directed links in the studied networks. Furthermore, adding protection tunnels to those two basic techniques was sufficient to achieve a full coverage. We are planning to study the impact of the different protection techniques on the traffic by considering the traffic matrix of the studied networks.

5. REFERENCES

[1] P. Francois, C. Filsfils, O. Bonaventure, and J. Evans. Achieving Sub-Second IGP Convergence in Large IP Networks. ACM SIGCOMM Computer Communication Review, July 2005.

[2] J.-P. Vasseur et al. *Network Recovery: Protection and Restoration of Optical, SONET-SDH, and MPLS*. Morgan Kaufmann, 2004.

[3] A. Atlas, et al. Basic specification for IP fast-reroute : loop-free alternate. Internet draft, draft-ietf-rtgwg-ipfrr-spec-base-01.txt, work in progress, September 2004.

[4] S. Bryant, C. Filsfils, S. Previdi, and M. Shand. IP Fast Reroute using Tunnels. Internet draft, draft-bryant-ipfrr-tunnels-01.txt, work in progress, Oct 2004.

[5] A. Atlas. U-turn alternates for IP/LDP Local Protection. Internet draft, draft-atlas-ip-local-protect-uturn-00.txt, work in progress, November 2004.

[6] S. Bryant and M. Shand. IP Fast Reroute using Notvia Addresses. Internet draft, draft-bryant-shand-ipfrr-notvia-addresses-00.txt, work in progress, March 2005.