Muliple Routing Configurations Demonstrator*

IP Fast Reroute in Practice

Tarik Cicic
Simula Research Laboratory
PB. 134, Lysaker
Norway
tarikc@simula.no

Ole Kristoffer Apeland
Simula Innovation
PB. 134, Lysaker
Norway
oleap@simula.no

ABSTRACT

We demonstrate IP fast reroute using a recently proposed method called Multiple Routing Configurations. The demonstrator comprises a small network of Linux routers, where the effect of network failures can be observed on real-time applications with and without IP fast reroute in effect.

1. INTRODUCTION

Routing protocols are known to converge slowly after topology changes or link or router failures. Intra-domain protocols in common use, such as OSPF and IS-IS, may need seconds to exchange new link state information and converge to a common view of the network. In this period, the lack of valid network routes will affect the transport services. A disruption of a link or a router in central parts of a network has the potential to affect thousands of phone conversations or TCP connections, with obvious adverse effects.

Recent research has shown that a large portion of the link and router failures have a temporary scope, lasting for a short period of up to a few minutes before becoming operative again. It has been recognized that starting the intra-domain reroute process may not be the best reaction to such events, as the reroute process would cause network-wide changes and possibly instabilities, only to be restarted when the link or node becomes operative again. It is a much better strategy to restrain from sending the link state updates until it has become clear, or at least overwhelmingly probable, that the change is permanent. This however not always possible, and waiting in the defect routing state prolongates the transport service interrupt.

Recently, there has been much attention in the research community and the IETF on resolving these problems using *proactive, local* IP recovery. In this approach the backup paths are prepared proactively (apriori), and should be able to recover any link or router failure in the network. Local (on the point of failure) reaction eliminates the network-wide signaling and thus realizes rapid reaction and minimal communication interrupt. The proactive, local IP recovery is proposed used as the first aid in the period between a failure is detected and the regular rerouting process is finished. If a waiting period is incorporated before the reconvergence starts, transient failures can be handled without any notification to the rest of the network.

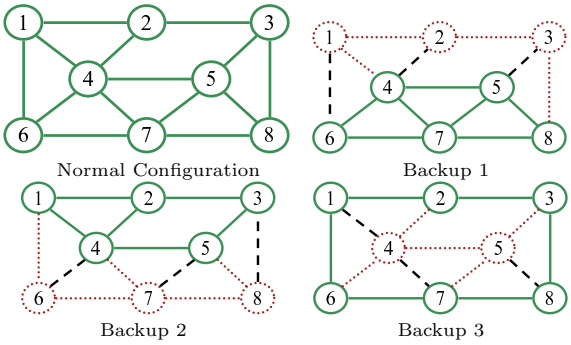

Figure 1: An example network topology with three backup configurations. The dotted links and nodes are isolated from any traffic forwarding, while the dashed links are restricted to only forward the traffic to their attached isolated node.

Multiple Routing Configurations (MRC, [1]) represents one of the attractive recent fast IP recovery proposals. MRC guarantees recovery from any single failure, and starts the recovery immediately after the failure is detected. The main idea of MRC is to use the network graph and the associated link costs to produce a small set of backup network configurations. The link costs in these backup configurations are manipulated so that for each link and node failure, the node that detects the failure can forward the incoming packets around the failed component toward the destination. This is achieved by isolating some links from traffic forwarding using infinite link weight, and by restricting the traffic on selected other links using a high, finite link weight (Fig. 1). The backup configurations can be used in conjunction with the standard shortest-path routing protocols such as OSPF.

2. DEMO DESCRIPTION

MRC has been extensively studied using simulations and formal modeling. In order to better understand how it can be embedded into a real routing system, we have built a

*This work is supported by the Norwegian Research Council FORNY programme Nr. 179820.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
MC'07, November 26-30, 2007, Newport Beach, CA Copyright 2007 ACM
ISBN 978-1-59593-935-7/07/11...$5.00

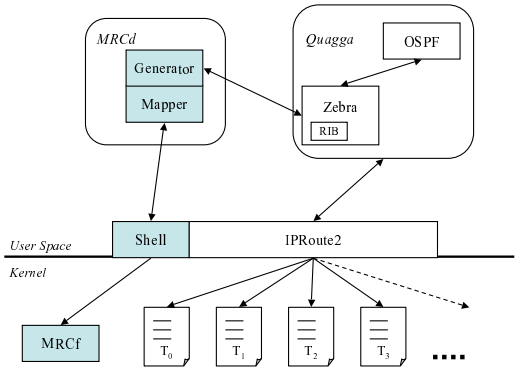

Figure 2: Network node architecture overview. The shaded boxes implement the MRC functionality.

prototype based on the Quagga routing suite and the native Linux forwarding.

2.1 Architecture and Implementation

In network nodes, the demonstrator extends the Quagga and Linux routing and forwarding with three additional components: a backup configuration generator, a routing information mapping processor, and an extended forwarding engine.

The backup configuration generator reads the network topology from Quagga's OSPF deamon, and produces the configurations that satisfy the formal requirements of the MRC model. This process is distributed on each node.

The routing information mapping processor uses the backup configurations and the existing routing information in form of IP route prefixes and the default next hops to populate the backup forwarding tables. One such table is constructed per backup configuration.

It is the task of the forwarding engine to encounter the packets destinated to a failed outgoing interface and to determine which backup configuration and operative outgoing interface the packet is going to use. As MRC requires, packets forwarded in a backup configuration have to be tagged as belonging to that configuration, so that the routers further down the forwarding path can use the appropriate backup forwarding table.

In our implementation, all routers have the same structure, depicted in Fig. 2. The routers are implemented on the Linux platform and use kernel version 2.6.22. Linux provides the basic packet forwarding functionality, and in fact supports up to 256 parallel forwarding tables accessible through user routing policies. This provides a good basis for our MRC implementation. Each node runs OSPF through use of the Quagga open source routing suite version 0.99.6 [2].

The backup configurations are generated on each node, and the generator is co-located with the routing information mapper in the MRC deamon (MRCd). MRCd retrieves the routing information from Quagga, which controls the default routing as normal (table T_0 in Fig. 2).

We implemented a simple shell extension to IPRoute2 component of the built-in Linux routing system. The shell is used to configure the MRC forwarding (MRCf) and the backup forwarding tables $T_1, T_2, \ldots$. The shell is also used for debugging purposes.

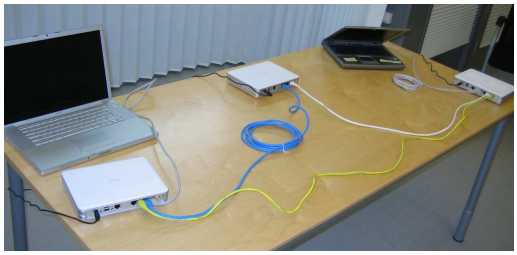

Figure 3: Demonstrator configuration. The shortest path (yellow cable) can be disconnected without affecting the video transmission between the laptops.

2.2 The Setup

The demonstrator comprises three ASUS WL-500g routers running OpenWrt Linux distribution, connected through their Ethernet interfaces into a triangular topology. Two Apple Powerbook laptops with video camera are used as data terminals, running Skype video or the traceroute application (Fig. 3).

When the OSPF routing converges, the shortest (single-hop) path between the terminals is used. MRC fast recovery can be turned on and off in the network. When MRC is on, removal of the link used for data forwarding does not affect the video transmission. If MRC is not used, link removal leads to OSPF reconvergence, which causes a temporary freezing of the video transmission.

3. PRACTICAL INFORMATION

3.1 Requirements

The demonstrator needs approximately $2m^2$ of the table space and an AC power supply.

3.2 Expected Impact

IP fast reroute is being standardized within the IETF and is expected to be available on the next generation routing equipment from major manufacturers. Enabling instant recovery of data streams in case of failure, it represents a key component for the support of the real-time applications in the Internet.

This project demonstrates a state-of-the-art IP fast recovery implementation, which we believe is relevant for ACM Middleware audience.

4. REFERENCES

[1] A. Kvalbein, A. F. Hansen, T. Čičić, S. Gjessing, and O. Lysne. Fast IP network recovery using multiple routing configurations. In *Proceedings of IEEE INFOCOM*, April 2006.

[2] Quagga routing software suite. Online, 2007. http://www.quagga.net/.