

Characterizing network events and their impact on routing

Amelie Medem Kuate Renata Teixeira Michael Meulle
Universite Pierre et Marie Curie and CNRS FT R&D

1. INTRODUCTION

We call *network events* incidents that disturb the normal behavior of one or more elements of an IP network. Routers, network interface cards, and IP links can fail or malfunction for many reasons. For example, operators may need to reboot a router for a software upgrade, an interface card may crash, and IP links may be overloaded because of a denial-of-service attack. Any of these network events can impact customer's traffic (packets can be lost or delayed, and, in extreme cases, customers may lose connectivity to parts of the network). When customers complain, network operators need to intervene to diagnose and, hopefully, fix the problem. In this work, we characterize network events according to their *causes* by using data collected from the Virtual Private Network (VPN) backbone of a large European provider. The European VPN network do not connect to Internet, but interconnects sites of over ten thousand enterprise networks.

Many studies so far characterize failures using routing-message logs [3, 2, 1]. This implies that their analysis shows how network events manifest in the control plane, but not the root-cause of routing instabilities. The only exception is the early work of Labovitz et al. [3], which studied the causes of routing instability of an Internet Service Provider (ISP) using trouble-ticket logs. Our work complements theirs by studying all network events in trouble-ticket logs, not only the ones that impact routing.

In this work, we first characterize network events from trouble-ticket logs of a large European VPN provider, and then study their impact on intra-domain routing messages. We start from the trouble tickets to capture all events that require operator intervention. These events are the most costly for service providers and the ones that have the most severe impact for customers. Thus, understanding the nature of network events is a

necessary first step to reduce operational cost and to improve the robustness of the network for customers. This work makes three main contributions: A taxonomy of network events according to their causes, which can be used by operators to catalog trouble reports in a more systematic manner; a methodology to correlate network events from trouble-ticket logs with IS-IS messages, which is important to understand the impact of events in the control plane; and the insights from the characterization of three years of network events. The following sections of this abstract present each of these contributions.

2. TAXONOMY OF NETWORK EVENTS

We propose a taxonomy of network events from the perspective of a given network, called the *network of interest*. This taxonomy labels events according to a hierarchy of four levels. The classification of network events was inspired from data recorded on trouble tickets. The first level distinguishes between events that happen inside the network of interest (*internal* events) and events that happen in neighbor networks (*external* events). This distinction is important because internal events are under the control of network operators, who can fix them or minimize their impact on routing. However, operators can not directly fix external events, but they can find means to limit their consequences on customers' traffic.

The second level separates events that are the results of maintenance activities (*planned events*) and those that are completely unexpected (*unplanned events*). Planned events are known in advance and operators can prepare the network beforehand and warn customers. It can take time for operators to detect an unplanned event and schedule the reparation activity.

The third level of the taxonomy distinguishes between events caused by failures on *network elements* and events from failures on the *environment*. Environment refers to all equipments outside the network, such as power supplies and fans. We separate environmental events into: *human errors*, *power outage* and *heating*. Network elements can be *IP* or *lower-layer* equipments. We further distinguish between events due to *hardware* or *software* problems.

3. MEASUREMENT METHODOLOGY

Our methodology is organized in three steps:

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.

CoNEXT'07, December 10-13, 2007, New York, NY, U.S.A.

Copyright 2007 ACM 978-1-59593-770-4/ 07/ 0012 ...\$5.00.

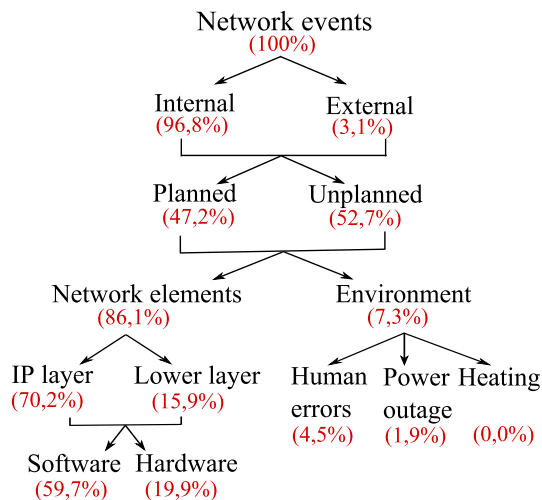


Figure 1: Classes of network events in trouble tickets of the European VPN network

- Parsing of trouble tickets:** We first process three years of trouble tickets of the network of interest to obtain events, and classify them according to our taxonomy. Network operators log all complaints of customers in trouble tickets. These logs contain all interactions between customers and operators from the occurrence of an event until its resolution. There are many challenges on the data analysis. First, the datasets are imprecise, because the information on failures is logged by hand by operators. Second, there is no standard way of describing trouble tickets. To deal with these challenges, we have to automatically and manually process the trouble tickets.
- Process IS-IS messages:** We continue the work with the process of eight months of IS-IS routing messages to extract information on the status of IP links. When an IP link status changes (for example, link up or down), all routers adjacent to the link flood the network with routing messages to report the event. We extract the start and end-time of all IP link disconnections in the network.
- Correlate network events with IS-IS messages:** Finally, using time and location information, we correlate classified network events issued from trouble tickets with IP link state changes. Location information issued from IS-IS messages and routers' configuration help identify all IP links and neighboring IP links of routers.

4. CHARACTERIZATION RESULTS

This section discusses our preliminary results of the analysis of network events in the European VPN network. Figure 1 presents the taxonomy annotated with

the percent of all the events in the three-year period studied fall in each category.

During three years of analysis, the European VPN network only experiences a few number of external events (3.1%). This can be related to the number of neighbors of the network (only two in our case).

Planned events are nearly half of all events. Thus, using carefully developed maintenance procedure can solve a large part of the problem.

Most of network events are related to network elements. IP level events represent the majority of network events (70% of events). Hardware events only account for 19.9%, whereas software problems are responsible for 59.7%. Software events do not always represent failure or defect. In fact, 68% of software events are part of maintenance activities, which mainly consist of installation, upgrade or downgrade of operating systems.

Previous work point out misconfiguration as a significant cause of network events, but we only find few instances of misconfiguration. They are included in human-error events, which only account for 4% of all events. One explanation might be that Internet configuration is more complex or change more often than in VPNs, which would make misconfigurations more frequent in the Internet.

Our correlation with IS-IS shows that 46,6% of the network events do not impact IS-IS. We find that this depends on routers' configuration and sometimes on the nature of events. Studies that focus on IS-IS messages cannot capture these events, even though they are important enough to initiate a trouble ticket. The analysis of router configurations reveals that 12,1% of these events happen in routers without IS-IS configuration.

5. PERSPECTIVES

Future work will proceed in three directions. First, we intend to gather trouble tickets from other networks. Automatically generated trouble tickets will be helpful, since the recorded information is more accurate. Second, we aim to analyze the impact of network events on inter-domain routing messages as well. Third, we intend to progress towards designing automatic tools that will help network operators to quickly fix disruptions inside their network.

6. REFERENCES

- [1] Y. Huang, N. Feamster, A. Lakhina, and J. Xu, Diagnosing network disruptionss with Network-wide analysis, in *SIGMETRICS*, JUN, 2007.
- [2] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah and C. Diot, Characterization of Failures in an IP Backbone, in *INFOCOM*, 2004.
- [3] G. Labovitz, A. Ahuja, and F. Jahanian, Experimental study of the Internet Stability and Backbone failures, in *FTCS*, JUN, 1999.