

Resilience Characteristics of the Internet Backbone Routing Infrastructure

Craig Labovitz, Roger Wattenhofer, Srinivasan Venkatachary
 Microsoft Research
 {labovit, cheenu, rogerwa}@microsoft.com

Abha Ahuja
 Merit Network, Inc.
 ahuja@merit.edu

I. INTRODUCTION

It is widely believed that the Internet is a highly fault-tolerant, survivable network. In particular, the Internet is attributed with the ability to route packets around faults quickly, in a matter of seconds. However, in this position paper we provide empirical data from the experimental injection and measurement of several hundred thousand inter-domain routing faults that shows the time required for Internet backbone routing protocols to re-route around failures is actually several orders of magnitude longer, sometimes taking more than 30 minutes. We show that these fail-over delays stem from systemic problems with the design and implementation of the current Internet inter-domain routing infrastructure.

Further, we explore large-scale vulnerabilities in Internet backbone routing. We describe means by which a single malicious or misconfigured router can cause routes to large parts of the Internet to fail. We argue that the lack of effective routing fail-over and lack of security associated with routing protocols seriously undermines the dependability of the Internet routing infrastructure in the presence of failure or malicious attacks.

The Internet's sustained exponential growth and the continued emergence of new and varied network applications provides testament to the scalability of the backbone infrastructure and protocols. The original TCP/IP decision to place network intelligence and state almost exclusively on end-nodes has enabled a diverse progeny of applications ranging from multimedia to collaborative learning. This scalability, however, comes at a price. Since its commercial inception in 1995, the Internet has lagged behind the public switched telephone network (PSTN) in availability, reliability and quality of service (QoS). This relative lack of reliability stems in part from the absence of intermediate backbone state and synchronization between routers. Despite the remarkable tolerance demon-

strated by end-users for failures and delays in today's predominant network applications, including email and web browsing, the relative lack of Internet backbone reliability poses a significant challenge for emerging transaction-oriented and interactive applications such as Internet telephony, online trading and laboratories.

Although recent advances in the IETF's Differentiated Services working group promise to improve the performance of application-level services within some networks, across the wide-area Internet these QoS algorithms are usually predicated on the existence of a stable underlying forwarding infrastructure. In this paper, we present results from our research and related research efforts that show that the dependability and failure repair properties of the Internet routing infrastructure leave much to be desired.

II. ROUTE AVAILABILITY AND FAILURE

We first looked at the availability of inter-domain routes. We define the *availability* of a given default-free route from a provider as the period of time that a path to the network destination, or a less specific prefix, was present in the provider's routing table. We include less specific prefixes in our definition as provider's regularly aggregate multiple more specific network addresses into a single supernet advertisement.

The graphs in Figure 1 show the cumulative percentage of time default-free routes were available from each provider during our ten month study [4]. The horizontal axis shows the percent time available; the vertical shows the cumulative percentage of routes with such availability. Both graphs only include routes available for more than 60 percent of the time during our study. The two graphs in Figure 1 represent the same data, but Figure 1(b) provides an expanded view of route availability above 99.9 percent.

A recent study [7] found that the PSTN averaged an availability rate better than 99.999 percent during a one year period. From the graph in Figure 1(b), we see that

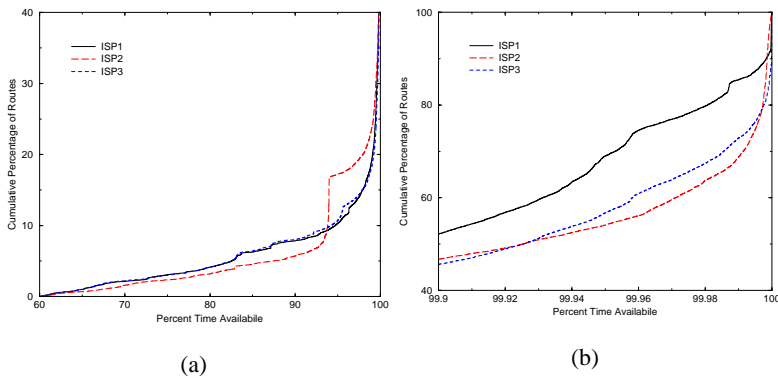


Fig. 1. Cumulative distribution of the route availability of three service providers.

the majority of Internet routes (65 percent) from all three providers exhibited an order of magnitude less availability.

In addition to availability, we examine the rate of failure and fail-over in inter-domain paths. We define an inter-domain route *failure* as the loss of a previously available routing table path to a given network, or a less specific, prefix destination. A *fail-over* of a route represents a change in the inter-domain path reachability of that route.

The two graphs in Figure 2 show the cumulative distribution of the mean number of days between route failures (a), and route fail-over (b) for routes from ISP1, ISP2 and ISP3. The horizontal axes represent the number of ISP routes that exhibit a specific mean-time to failure/fail-over or less; the vertical axes show the cumulative proportion of the ISP's routing table entries for all such events. Examining the graph in Figure 2(a), we see that the majority of routes (greater than 50 percent) from all three providers exhibit a mean-time to failure of fifteen days or more. By the end of thirty days, the majority (75 percent) of routes from all three providers had failed at least once. The distribution graphs for ISP1, ISP2 and ISP2 share a similar curve, with ISP1 exhibiting a slightly lower cumulative MTTF curve starting at ten days.

III. FAULT OCCURRENCE CHARACTERISTICS OF INTERNET ROUTING

We briefly describe some Internet outages which directly, or indirectly, impacted a majority of Internet backbone paths. We provide the following summaries as anecdotal evidence of the sources of major Internet failures.

- April 25, 1997 — A misconfigured router maintained by a small Virginia service provider injected an incorrect routing map into the global Internet. This map indicated that the Virginia company's network provided optimal connectivity to all Internet destinations. Internet providers that accepted this map automatically diverted all of their traf-

fic to the Virginia provider. The resulting network congestion, instability, and overload of Internet router table memory effectively shut down most of the major Internet backbones for up to two hours. Incorrect published contact information for operations staff, and lack of procedures for inter-provider coordination exacerbated the problem [1].

- November 8, 1998 — A malformed routing control message stemming from a software fault triggered an interoperability problem between core Internet backbone routers manufactured by different vendors. This problem led to a persistent, pathological oscillation and failure in the communication between most Internet core backbone routers. As a result, Internet end-users experienced wide-spread loss of network connectivity, and increased packet loss and latency. The majority of backbone providers resolved the outage within several hours after adding filters which removed the malformed control message [2].

Overall, both Internet and telephony outages stem from a wide range of sources, including faults in the underlying telecommunication switching system, and the higher level software and hardware components. Like Pradhan [3], we are interested in estimating the reliability of Internet backbone paths at specified probability and duration thresholds such as the mean number of events per year, and the mean time spent in events. The significant findings of our work include:

- The Internet backbone infrastructure exhibit significantly less availability and a lower mean-time to failure than the Public Switched Telephone Network (PSTN).
- The majority of Internet backbone paths exhibit a mean-time to failure of 25 days or less, and a mean-time to repair of twenty minutes or less. Internet backbones are rerouted (either due to failure or policy changes) on the average of once every three days or less.
- Routing instability inside of an autonomous network does not exhibit the same daily and weekly cyclic trends

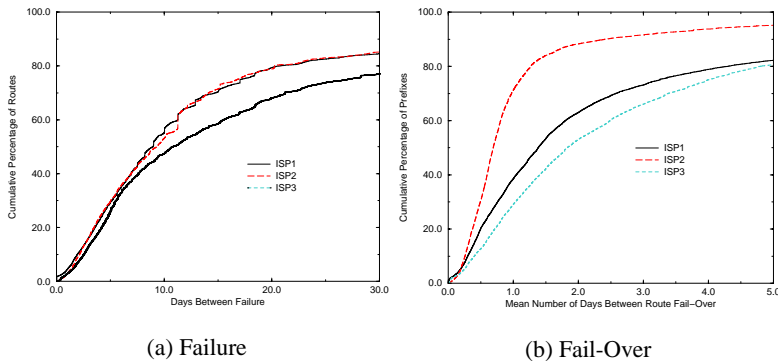


Fig. 2. Cumulative distribution of the mean-time to failure and mean-time to fail-over for default-free routes from three ISPs.

as previously reported for routing between Inter provider backbones, suggesting that most inter-provider path failures stem from congestion collapse.

- A small fraction of network paths in the Internet contribute disproportionately to the number of long-term outages and backbone unavailability.

IV. FAULT REPAIR CHARACTERISTICS OF INTERNET ROUTING

While the Internet backbone routing protocol, BGP, is believed to be resilient to faults and converge on new routes very quickly, our measurements showed that [5] the time to repair in the case of a fault is actually in the order of minutes, sometimes taking up to 30 minutes. We also showed that n BGP routers connected in a complete graph may potentially explore $n!$ routes, or all possible paths of all possible lengths between each router after a fault.

In this section, we present both empirical observations and analysis of the data collected by our fault injection experiments. We first provide several examples illustrating the repair process.

Our measurement methodology consists of deliberately injecting faults into the Internet at specific points and then observing the repair events caused from different points across the Internet. Our fault injection apparatus consists of probe machines maintaining geographically and topologically diverse BGP peering sessions with more than 40 commercial Internet providers.

As we only injected routing information for addresses assigned to our research effort, these faults did not impact routing for commodity ISP traffic with the exception of the addition of some minimal level of extra routing control traffic.

During the six months of our study, we analyzed the routing topologies between more than 400 pairs of Internet providers. We graph only three representative topologies

in Figure 3 for clarity. We note that all of the other monitored topologies in our study exhibited related behaviors.

In each diagram, we label the *steady-state path*, or the path normally selected by ISP4 in the absence of a fault. The steady-state paths include IS1-ISP4 in Figure 3(a), ISP2-ISP4 in (b) and ISP3-ISP4 in (c). Similarly, we label backup paths chosen by ISP4 in each diagram with the letter P followed by integers denoting the frequency with which we observed that backup path (i.e. P1..P6 in Figure 3(c)). For clarity, we graph only the most common backup paths observed during our study. In addition to the paths illustrated, ISP4 announced an additional 11 unique paths for $R2$ and 7 additional paths for $R3$ after 23 and 27 percent of the faults, respectively. We note that ISP4 only observed a single backup path for the topology in Figure 3(a).

In [6], we present measured repair behaviors of more than 20 unique routes between more than 400 pairs of Internet service providers. We provide analysis of the impact of specific inter-domain policies and topologies on the speed of routing fail-over. Our major results include:

- The upper bound on repair delay when a route to a destination fails is linearly related to length of the longest possible path between a source and destination.
- On average, larger Internet Service Providers (ISPs) provide faster repair times than smaller providers for a given route.
- Errant paths are frequently explored. These “vagabond” paths likely stem from misconfiguration or software bugs.

Throughout the six months of our study reported in [6], we observed frequent examples of misconfigured, or *vagabond* paths between the majority of the 400 pairs of Internet providers we monitored. For example, between two routers that were co-located at a major ISP peering point called Mae-West, we found that the routes from one to the other went around through a Mediterranean country!

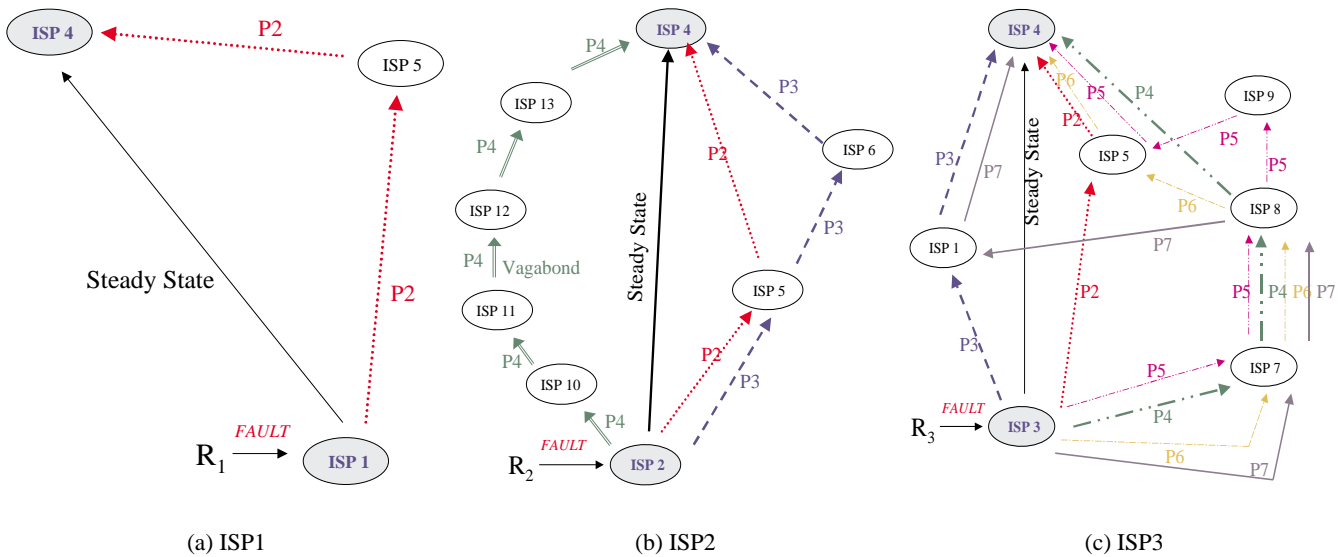


Fig. 3. Example of backup paths explored during process of repair for routes from a Japanese provider to three different ISP routers at the California Mae-West exchange point.

This was later confirmed as due to a misconfiguration.

Currently, BGP is vulnerable to such misconfigurations and can cause major outages. Secure BGP [8] is a proposal to avoid malicious routers from being able to spread false routing information. The Internet Routing Registry [9] attempts to provide a database of authentic origin points associated with prefixes.

As the national and economic infrastructure become increasingly dependent on the global Internet, the availability and scalability of IP-based networks will emerge as among the most significant problems facing the continued evolution of the Internet. This paper has argued that the lack of inter-domain failover due to delayed BGP routing convergence will potentially become one of the key factors contributing to the “gap” between the needs and expectations of today’s data networks. These results suggest a strong need to reevaluate applications and protocols, including emerging QoS and VoIP standards which assume a stable underlying inter-domain forwarding infrastructure and fast IP path restoration.

REFERENCES

- [1] R. Barrett, S. Haar, R. Whitestone, “Routing Snafu Causes Internet Outage,” Interactive Week, April 25, 1997.
- [2] North American Network Operators Group (NANOG) mailing list, <http://www.merit.edu/mail.archives/html/nanog/msg03039.html>.
- [3] D. Pradhan, “Fault Tolerant Computer System Design”, Prentice Hall, New Jersey, 1996.
- [4] C. Labovitz, A. Ahuja, F. Jahanian, “Experimental Study of Internet Stability and Wide-Area Network Failures,” in Proc. International Symposium on Fault-Tolerant Computing, Madison, WI. June 22, 1999.
- [5] C. Labovitz, A. Ahuja, A. Bose and F. Jahanian, “Delayed Internet Routing Convergence,” to appear in Proc. of ACM SIGCOMM. August, 2000.
- [6] C. Labovitz, R. Wattenhofer, S. Venkatachary, A. Ahuja, “The Impact of Internet Policy and Topology on Delayed Routing Convergence”, Microsoft Research Technical Report (MSR-TR-2000-74).
- [7] R. Kuhn, “Sources of Failure in the Public Switched Telephone Network”, IEEE Computer Vol. 30, No. 4. April, 1997.
- [8] Kent, S., Lynn, C., Seo, K., “Secure Border Gateway Protocol (S-BGP)”, in IEEE JSAC. April, 2000.
- [9] Merit Network, Inc., “Internet Routing Registry”. <http://www.radb.net>