IETF Next Steps in Signaling
Internet-Draft
Expires: April 2005

C. Shen
H. Schulzrinne
Columbia U.
S. Lee
Samsung AIT
October 2004

Internet Routing Dynamics and NSIS Related Considerations
draft-shen-nsis-routing-00.txt

Status of this Memo

Copyright Notice

Abstract

This document presents the main results from a recent Internet routing dynamics measurement and discusses their impact on NSIS protocol design. It also provides an evaluation of the simple, low cost packet TTL monitoring route change detection mechanism in the context of different NSIS deployment models.

Table of Contents

1. Introduction

Interaction with IP routing is an important aspect in Next Step In Signaling (NSIS) protocol design [1]. Solving this problem requires a good characterization of today's Internet's routing behavior. In this memo we summarize the main results from a routing measurement experiment conducted between April and August 2004, and discuss their impact to the design of NSIS. The focus of our routing study is route change. We look at the types, duration and likely causes of different route changes observed.

Various route change detection mechanisms have been proposed in [2]. Some of them are simple and low cost, such as packet TTL monitoring. It will be very helpful to know how good such a mechanism will be. We evaluate the effectiveness of packet TTL monitoring in detecting route changes based on our measurement experiments. We also introduce the concept of NSIS-concerned route changes, and look at the TTL monitoring method in the context of different NSIS deployment models.

For details about our measurements and analysis, the readers are referred to our technical report [5].

2. Measuring and Analyzing Internet Routing Dynamics from NSIS

2.1 Measurement Methodology and Data Sets

We collected four data sets in our measurements that involved totally 24 public traceroute servers located in US, Iceland, Netherlands, Australia, Germany, Switzerland, Bulgaria, Sweden, and Thailand.

Data Set I (DS I) contains 12 sites and is collected from April 9, 05:14:02 AM EDT 2004 to April 24, 00:09:34 AM EDT 2004, at an average per-site exponential sampling rate of one traceroute every 15 minutes. This corresponds to an average of 2.75 hours sampling interval for each path. The independent and exponential sampling allows us to approximate the amount of time that the Internet spends in a particular route from the number of times we observe that particular route in our measurements.

Data Set II (DS II) and Data Set III (DS III) contain 24 sites. Each site is exponentially sampled on average every 30 minutes. This corresponds to an average of 11.5 hours sampling interval for each path. Data Set II is take
n from May 22 12:20:11 AM EDT to Jun 13
12:24:19 PM EDT, 2004 and Data Set III is taken from June 14 03:05:23 PM EDT to July 06 10:34:14 AM EDT, 2004.

Data Set IV (DS IV) contains both a ten-minute fixed interval and a two-hour exponential interval measurement conducted between four designated paths that are believed to be showing typical routing characteristics, i.e., they are not paths that display extremely high route fluctuations. DS IV is collected from July 14 to August 3, 2004. The purpose of this data set is to provide additional insights on short scale routing dynamics and also to obtain some understanding about impacts of longer measurement intervals (posed by infrastructure restrictions) on the accuracy of measurements.

2.2 Summary of Findings from an NSIS Perspective

2.2.1 Route Prevalence and Route Persistence

We identified the most frequently observed route, or dominant route, for each path, and computed the prevalence of each dominant route, which is the number of times the dominant route is sampled divided by the total number of samples for that path. We also studied how long the path is likely to stay in one route, or route persistence, and obtained the duration distribution of long-lived routes. The results confirmed some of the conclusions from earlier measurements [3][4], i.e., Internet paths are strongly dominated by a single route, but significant site to site variation exists. The strong dominance of a single route is good for NSIS. To cope with significant site to site variations, it will be helpful for NSIS to employ an adaptive approach in dealing with routing dynamics. For example, the NSIS protocol should be more aggressive in detecting route changes on a path that is particularly instable, by using a shorter path refresh interval, or advanced route change detection mechanisms.

2.2.2 Different Types of Route Changes

Our statistics of route changes shows that route changes occur over a wide range of time scales, ranging from seconds to days, and over different network scales, ranging from intra-AS changes to inter-AS changes. The majority of route changes are found to be involving no change of total number of hops (i.e., TTL-invisible). However, a large proportion of these TTL-invisible route changes are caused only by a small set of (mostly local) router pairs belonging to the same service provider or host network, which are probably doing route splitting and load balancing. The router administrators have the best information on how the routers are configured and therefore are the best persons to take care of these route changes.

Route splitting is a cause to extremely short period route changes (called route fluttering [3]), at the scale of seconds or minutes. It is controlled by a special option of an IPv4 router and can be turned off [5]. Load balancing usually leads to short or medium

scale route changes and are seen in routers within both host and service provider networks. Extra mechanisms need to be deployed in these networks to allow proper functioning of NSIS in the presence of load balancing. These may include special route monitoring mechanisms to quickly detect the route change and notify related NSIS modules about the change, or establishing redundant state information in load balancing routers. The latter approach tends to cost more resources but gives a more prompt response time.

2.2.3 Accuracy of Measuring Path Characteristics

Our analysis of DS IV showed that the accuracy of measuring the path characteristics with a longer measurement interval depends heavily on the stability of the path being studied. Using a ten-minute fixed interval and a two-hour exponential interval measurement, we have seen examples where both measurements capture the same number of routes, AS-paths and their changes; we have also seen examples where the two-hour interval measurement missed about half the number of unique routes, AS-paths and their changes. (During our processing of other data sets, we used several techniques to remove outliers that showed great instability in order to counter this effect.) This issue is essentially the same as the site to site variation problem mentioned in 3.2.1 and calls for adaptive NSIS mechanisms to characterize or interact with path dynamics. An example will be an NSIS refresh mechanism that starts with a short interval value and gradually increases that interval depending on the actual routing characteristics it observes.

2.2.4 Impact of Multi-Homing

Multi-homing is often seen in our measurements and causes AS level route changes as well as asymmetric routing. A particular example we encountered in DS IV is a site that uses one ISP as the incoming ISP and another as the outgoing ISP. We've also seen the same site using one ISP as its primary outgoing ISP in May and June and then switch to another one as primary outgoing ISP in July. The result is a dramatic change of route characteristics, from a very stable route to a very fluctuating route. Moreover, the site still occasionally uses the old AS-path for a while, causing short term AS level route changes. If the multi-homed site has control over such behaviors, it should deploy appropriate mechanisms to notify related NSIS entities about the switch to allow a fast NSIS recovery.

3. NSIS-Concerned Route Changes and NSIS Deployment Models

We have mentioned that the majority of route changes are local and

TTL-invisible so they are best tackled by the router administrators, or by keeping NSIS in mind when turning on related routing options. The remaining of the route changes, many of which non-trivial in that they could affect more than one hops in different geometrical locations, tend to have bigger impact on the applications using NSIS protocol and are more concerned by the NSIS community. It will be very nice if many of these changes can be detected by simple mechanisms like packet TTL monitoring, without special routing monitoring mechanisms installed by service provider or host network administrators. We will look at what our experiments show about the effectiveness of the TTL monitoring method. But before that, we introduce the concept of NSIS-concerned route changes and NSIS deployment models.

From a network perspective, generic route changes can be classified into inter-AS and intra-AS route changes depending on whether the route change affects the AS set in the path. Intra-AS route changes may further be divided into ingress-point, egress-point and mid-point route change, depending on the location of the routers in the AS where the route change occurred.

NSIS needs to deal with all generic route changes only in a full NSIS network where all nodes are NSIS Entities (NEs). A far more likely network scenario will be a mixed deployment of NEs and normal routers. In this case, only a subset of all generic route changes will involve change of NEs. These are NSIS-concerned route changes that should be dealt with by NSIS.

To better understand NSIS-concerned route changes, we need to make assumptions about the actual NSIS deployment models. We list below possible NSIS deployment models in a mixed environment of NEs and normal routers, together with the correspondin
g NSIS-concerned route
changes in each model.

- AS model: in the AS model each AS deploys a central NE that is responsible for NSIS related signaling for this AS. The NSIS-concerned route changes in AS model include all route changes that involve AS changes, i.e., inter-AS route changes.

- Entry model: in the entry model the ingress routers of each AS are NEs. NSIS-concerned route changes in this model include inter-AS and intra-AS ingress point route changes.

- Border model: in the border model both ingress and egress routers of each AS are NEs. NSIS specific route changes in this model include inter-AS, intra-AS ingress point and intra-AS egress point changes.

- Edge model: in the edge model the access routers of the source and destination sites are NEs. NSIS-concerned route changes in this model include inter-AS route change involving the first or last AS, intra-AS ingress point route change in the first AS as well as intra-AS egress point route change in the last AS.

- Generic model: we define the generic model as all other mixed deployments that cannot be clearly mapped to any of the above four categories. It might be a combination of the above models plus more NEs in certain parts of the network. There is no straightforward way to define an NSIS-concerned route change in this case other than to observe whether a change of NEs is involved once a route change has occurred.

4. Evaluation of Packet TTL Monitoring Based Route Change Detection

The data in this section is abstracted from our routing measurements (Section 3). We started by assuming a full NE deployment model where all route changes count. Table 1 shows the proportion of TTL visible route changes and AS path changes in our data sets I, II and III.

Item description	DS I	DS II	DS III
TTL-visible route changes	38%	25%	23%
AS level route changes	18%	8%	8%
TTL-visible AS level changes	77%	83%	88%

Table 1: TTL-visible route changes and AS changes

The percentage of overall TTL-visible route changes does not seem very promising. However, even in an all NE network, the route changes that concern us most are always non-trivial ones. If we focus on those AS level changes, we can see that TTL does a fairly good job in detecting roughly 4 out of 5 such changes.

We also examined the effectiveness of TTL in detecting NSIS-concerned route changes in four mixed NSIS deployment models as shown in Table 2. Clearly the TTL detection mechanism works better in these mixed models than in the all NE model. This is because NSIS-concerned route changes in all these four models are usually non-trivial changes, where TTL monitoring tends to be more effective. The table also shows that the successful TTL detection ratio tends to be higher when the deployment model is sparser in terms of NEs and thus fewer route changes become NSIS-concerned. Therefore, the edge model and

the AS model have higher detection ratio than the entry and border models. It is not possible to check the effectiveness of TTL detection mechanism in a generic mixed model when detailed deployment information is unknown, but the answer is expected to fall between the full NE model and the four mixed models we have studied.

NSIS model	DS I	DS II	DS III
AS model	77%	83%	88%
Entry model	51%	41%	40%
Border model	45%	39%	38%
Edge model	74%	90%	92%

Table 2: Effectiveness of TTL detection in the four mixed models

Overall, the TTL monitoring method appears to be a reasonably good way to find non-trivial route changes, although this mechanism alone is hardly enough to detect all NSIS-concerned route changes.

5. Conclusions and Future Work

We performed a routing measurement on today's Internet and studied the impact of general Internet routing dynamics, especially route changes, on NSIS design. We concluded that different types of route changes require different handling. Frequent yet local route changes caused by route splitting or load balancing are best handled by routing monitoring inside the network; many non-trivial NSIS-concerned route changes may be detected by hosts employing simple packet TTL monitoring mechanism, as we have seen in typical NSIS deployment models. We are currently also investigating other route change detection methods to see when and how they are likely to be useful.

6. Security Considerations

There are no explicit security issues associated with current version of this document.

7. Acknowledgements

We would like to thank Jongho Bang for their comments and support in this work.We would also like to thank all traceroute server sites that participated in our measurement.

8 Informative References

[1] Hancock, R., "Next Steps in Signaling: Framework", draft-ietf-nsis-fw-06 (work in progress), July 2004.

[2] Schulzrinne, H., "GIMPS: General Internet Messaging Protocol for Signaling", draft-ietf-nsis-ntlp-04 (work in progress), October 2004.

[3] Paxson, V., "Measurements and Analysis of End-to-End Internet Dynamics", PhD thesis, University of California, Berkeley, 1997.

[4] Zhang, Y., "Characterizing End-to-End Internet Performance", PhD thesis, Cornell University, 2001.

[5] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[6] Shen, C. and H. Schulzrinne, "Internet Routing Dynamics and NSIS Related Considerations", technical report, Columbia University, October 2004.

Authors' Addresses

Charles Shen
Columbia University
Department of Computer Science
1214 Amsterdam Avenue, MC 0401
New York, NY 10027
USA

Phone: +1 212 854 5599
EMail: charles@cs.columbia.edu

Henning Schulzrinne
Columbia University
Department of Computer Science
1214 Amsterdam Avenue, MC 0401
New York, NY 10027
USA

Phone: +1 212 939 7004
EMail: schulzrinne@cs.columbia.edu

Sung-Hyuck Lee
SAMSUNG Advanced Institute of Technology
San 14-1, Nongseo-ri, Giheung-eup
Yongin-si, Gyeonggi-do 449-712
KOREA

Phone: +82 31 280 9585
EMail: starsu.lee@samsung.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.