Consolidated Review of

A Comparison of Syslog and IS-IS for Monitoring Link State

1. Strengths:

The paper analyzes accuracy of syslog messages for determining network failures by comparing it against IS-IS routing messages collected from CENIC network. The paper examines the cause of the differences between the syslog and IS-IS link failure data.

The paper is the first attempt to analyze the effectiveness of using syslog messages for detecting network failures. The experiments are conducted in a production network. The paper is easy to read and follow. The analysis is sound and interesting.

The paper analyses why syslog overestimates link failures: the idea syslog fail to accurately capture short failures and link flapping. The paper presents a very interesting result: that despite the inaccuracies in syslog data, the distributions remain the same. The paper explores different explanations for duplicate UP and Down messages in syslog message, highlights the most probable explanation and presents an effective solution to this problem.

2. Weaknesses

Some parts of the paper are not as clear as they should be (see more detailed comments below).

This work expands a proposal by the authors to use common data sources to analyze failures. The practical implications of the work are limited IMO.

3. Comments

This reviewer likes the premise of this short paper. Namely, verify the assumptions that syslogs and email reports accurately describe network downtimes by comparing it with ISIS routing data. The findings that often times certain spurious messages can bias the accuracy is both interesting as well as significant and provides good insights regarding the reliability of various datasets.

The paper focuses solely on IS-IS and syslog. It is not immediately clear how this technique lends it self to other protocols such as BGP or other IGP protocols such as OSPF. Extending this study to one more IGP would greatly strengthen the arguments made within the document.

I would not call “dedicated tracing of routing protocol state” a “significant instrumentation” in today’s networks. Assuming network failure analysis is more challenging and interesting in larger networks, the marginal cost of collecting routing protocol state should be relatively low. Given the (admitted) limitations of the scheme proposed in this paper, a much stronger motivation is needed. The results presented in the paper are interesting and I find the paper an interesting exercise. However, based on the argument in previous paragraph I think the applications might be limited in practice.

I enjoyed reading this paper and believe this makes a valuable contribution.

Table 1: Shouldn't you list router configuration files as one of the data sources used in the study? Section 3.3, second sentence: "to do" => "due to".

Section 3.3, last sentence of first para: Insert "and" before "specific diagnostic message".

Section 3.4: Did you try values other than ten seconds for matching syslog events with IS-IS events? Section 4.1: For each DOWN and UP event, you will get two IS-IS messages -- one from each side of the link. Why didn't you match each syslog message to one of these messages?

Section 4.1: Intuitively I understand what link flapping is, but for the purpose of the analysis, can you explain how you defined link flapping?

Section 4.2: The clause "it is possible that a syslog-based link downtime..." does not parse.

Section 4.2: How was manual verification performed for the 25 syslog failures lasting more than 24 hours? Did you talk to the operators? Or looked at some other data source?

Table 4: What's your rationale for presenting separate numbers for core and CPE failures? I feel it just overloads the reader with too many numbers without any significant new insights.

Fig. 1(c): minor nit-pick: in the legend, can you swap the order of 'IS-IS CPE' and 'Syslog CPE'?

Section 4.3: I felt a little lost in this section. You might want to summarize the numbers in a table. Section 4.4, first para: 'applified' => 'amplified'.

Section 4.4, last para: Your analysis accounts for 82 + 99 of the 399 events that were reported by IS-IS but not syslog. Do you have explanation for the remaining events?

4. Summary from PC Discussion

The PC thought the study was well done. The paper could be better motivated since it isn't clear that collecting IS-IS data is a very large burden for ISPs. The authors could, for example, point out other work that has used syslog to make failure inferences (e.g., "Troubleshooting Chronic Conditions in Large IP Networks").

5. Authors’ Response

The authors would like to thank our shepherd and anonymous reviews for their feedback. We are confident that our approach to evaluating the accuracy of syslog will work for networks that use a different link-state protocols, such as OSPF. Performing such an analysis using distance-vector protocol data would to require a different approach and likely multiple measurement points.

Some reviewers believe that we are overstating the difficulty of operationally collecting routing protocol updates. Indeed, from a technical point of view, configuring and deploying an IGP listener, as we have done, is not difficult. However an active listener such as ours carries the risk of disrupting network operations in the event of a catastrophic malfunction or compromise. Our own deployment proposal was met with initial resistance from network operations staff, requiring discussions and thorough testing before data collection could begin. Collecting routing protocol updates is simply not standard practice. Practically speaking, Syslog data is much easier to obtain, and in some cases, may be the only data available. Thus, in many cases, network analysis is conducted using Syslog data. We believe that our study, highlighting the discrepancies between the two types of data, provides a useful comparison between these two data sources.