Control-Plane-Triggered Traffic Performance Degradation Detection

Mingwei Zhang, Ph.D.

Background

Network Diagnostic Tool (NDT) is a network performance tool that measures the characteristics of a TCP connection under heavy load. The various RTT values and other performance metrics defined in NDT results (e.g. defined in ndt7 protocol) can reveal the network performance of an NDT client toward the defined NDT servers. By establishing normal performance profiles for clients (or client-server pairs), we can learn about the expected network performance as well as incidences of performance degradation.

Internet traffic performance degradation, traffic slowdown, or complete outages, happen for different reasons. One particular group of traffic degradation that we are interested in exploring is the control-plane-triggered degradation. Inter-domain routing changes that happened on the Internet can hurt the performance of the Internet users observed via NDT performance metrics degradation. There is a wide range of control-plane behaviors that can potentially affect the performance of data-plane traffic. Here is a list of such behaviors:

- **non-optimal routes**: when a router is misconfigured to change to use a non-optimal route, such as choosing a long AS path toward the traffic destination, the data-plane traffic suffers from the long RTT and thus overall degradation of performance accordingly;
- **traffic hijacking**: when malicious autonomous systems start announcing prefixes and hijacks the traffic, they can choose to forward the hijacked traffic back to the original destination, thus causing non-optimal paths to be selected for traffic;
- **traffic blackout**: malicious players can also black-out hijacked traffic, thus causing complete data-plane traffic outages;
- **route leaks**: accidental route-leak could potentially taint a large portion of the Internet with a less-optimal path and also cause bottlenecks.

Knowing if and which control-plane behavior triggers a data-plane performance degradation allows network operators to better improve performance and potentially prevent future incidences.

Apart from NDT performance metrics, Measurement Lab also conducts traceroutes from the NDT servers back to clients. The traceroute results include IP-level path information that can be mapped to AS-level path information and correlate to corresponding control-plane messages.

Together with historical path information, we can better correlate performance degradation events with traffic path changes from both control-plane and data-plane.

Proposed Tasks

In this proposal, we propose a multi-phase study that involves the following tasks:

A. Develop network performance degradation detection model using NDT data.

B. Traceroute path correlation with detected network degradation events.

C. AS-level path correlation with detected network degradation events.

D. Root-cause inference model using historical IP-level and AS-level paths information correlated to events.

Task A mainly aims to develop a data pipeline to understand the sequence of NDT data and convert the data into discrete anomalous events with basic machine learning models. The usage of machine learning models is on the application of existing models instead of developing novel models.

Task B mainly involves parsing and understanding the traceroutes triggered alongside the NDT tests, and extracting IP-level paths from the servers to the clients. The paths during normal performance time ranges will provide a topological baseline and then can be used to derive differences against paths collected during anomalies.

Tasks C aims to develop code that collects and correlates relevant BGP messages around detected events. The BGP messages and the corresponding AS-level path changes will provide insights into the control-plane behaviors during the detected anomalies.

Task D finalizes the study by compiling detected degradation events, historical and on-event IP-level and AS-level paths, and related BGP messages together and trying to depict a holistic picture of all the changes that happened before and during anomalous events.

Timeline

Time-line for the total five-month period:

- Month 1: Develop a model to mark and detect traffic performance anomalies using NDT data sets. Conduct short-term data tests.
- Month 2: Collect anomaly detection results using long-term NDT data.
- Month 2-3: Develop a data collection pipeline to collect and correct control-plane information with detected data-plane anomalies.
- Month 4: Conduct inference study to detect potential control-plane-triggered traffic degradation incidences, and collect evidence for further study.
- Month 5: Collect long-term inference results and write reports for the discovery.