TopHat: The OneLab topology information component

Thomas Bourgeau[1], Timur Friedman[1], Benoit Donnet[2]
[1]LIP6-CNRS and UPMC Paris Universitas, France [2] Université catholique de Louvain, Belgium
{thomas.bourgeau,timur.friedman}@lip6.fr benoit.donnet@uclouvain.be

***Abstract*— Today, PlanetLab [1] provides limited facilities to make applications aware of the underlying network. The European Commission funded OneLab project [2] aims at building a sophisticated monitoring system allowing to get a view of the network topology and allowing applications to query this view. In order to deepen the current PlanetLab monitoring capabilities, we are proposing the addition of an always-running measurement service providing to any application a dedicated service that replies to specific queries on network topology.**

Index Terms: Network topology, Measurement

I. Introduction

The aim of the OneLab topology information component is to provide a common measurement service that continuously performs distributed measurements at the IP interface level; this component is called TopHat. The idea is to provide an open and public topology information service, providing both data from both the present and the past.

As one of the most well-known tracing system, CAIDA's *archipelago* [3], we will probe the Internet from multiple vantage points in order to study the dynamic of the network based on real time view of the network topology. We will progressively deploy our probing engine up to larger and larger numbers of monitors around the world using the PlanetLab test bed and by federating with other existing topology monitoring systems.

As traceroutes emanating from a large number of monitors and converging on selected targets can easily appear to be a DDoS attack, TopHat employs known techniques, such as the *Doubletree* [4] algorithm, to avoid excessive redundancy in probing.

TopHat also conducts reactive probing. It uses AS-level measurements to analyse the *BGP communities* attribute of a given AS in order to reveal a topology change. This information is used to launch a targeted active measurement where the topology has changed.

TopHat is designed as a service that distributed applications can remotely query to obtain information about the underlying network, such as path information (interfaces traversed, delays, loss rates) or various statistics on the topology as a whole. In addition, this probing engine will be able to perform on-demand measurements and so should be able to replace the traceroute tool for PlanetLab users.

This paper is organized as follows: Sec. 2 provides an overview of our topology information component; Sec.3 is dedicated to the measurement subcomponents; Sec. 4 explains TopHat methods and, finally, Sec. 5 concludes this paper and discusses further potential directions.

II. Topology information component

In this section we present an overview of the OneLab topology information component, TopHat, summarizing its architecture and usage. On one hand, probing monitors must be able to efficiently and easily store the data collected. On the other hand, the information must be available for easy retrieval for research purposes.

TopHat provides an API that serves as an interface between an active measurement tool, an AS-Level tool, a topological database and any application that wishes to send queries to determine the characteristics of the measured topology.

The topological data can be obtained by existing active measurement tools or through new tools developed for this purpose. TopHat is used by any application that is willing to send queries to determine the characteristics of the paths, either measured or predicted, between itself and a list of destinations.

There are two different entities defined to interact with TopHat:

- The *Agent* refers to any probing engine that is allowed to add topological information into the TopHat topological database. It can be either an active measurement or an AS-Level measurement tool.
- The *Client* refers to an application that queries TopHat for obtaining any information about the network topology.

Any authenticated entity (Client or Agent) will have an access to the API methods through the web interface, this feature will allow a user to have a direct access to the topology information stored in the TopHat database. TopHat should be accessed remotely via XML-RPC over HTTPS. To apply for a TopHat account, a web interface was produced (*https://www.top-hat.info*).

III. Measurement subcomponent

In this section we describe the measurement subcomponents that are used to gather either AS or IP level information in the TopHat topological database.

A. Active measurement subcomponent

This subcomponent aims at continuously probing the network in order to obtain topology information at the interface level. The probe engine is based on Doubletree [4], a distributed and cooperative Internet topology discovery algorithm that allows tracing monitors to share information through distributed hash tables (DHTs). This architecture presents advantages in terms of scalability, robustness, and flexibility.

B. AS-Level subcomponent

This subcomponent maintains information about the mapping from IP address to AS numbers (and inversely). This mapping is performed using information collected by the iPlane [5] and the cymru project [6]. We will provide a specific BGP attribute, the BGP communities attribute that can be used to trigger a targeted topology discovery process as a change in the BGP communities attribute of a given AS can reveal a topology change. Thus, a traceroute is triggered based on a set of rules. Currently, 4 rules are defined: "*modification of the BGP communities of a given route*", "*modification of the AS_PATH of a given route*", "*withdrawn of a given route*", and "*addition of an unknown route*". Standard values for BGP communities are described in RFC 1997 [7].

IV. TopHat methods

The main purpose of the TopHat implementation is to allow applications to have easy access to topological data.

TopHat defines methods to gather topological information in its database emanating from either the active measurement or the AS-Level subcomponent. It also allows any application to retrieve topological information through the use of methods that can be implemented in a large set of development languages.

Our current implementation of the query interface exposes a database-like view of path properties between every pair of end-hosts measured in the Internet. Any query to TopHat involves an SQL-like query on this view; TopHat does not compute a priori the entire table for every source-destination pair; instead it derives necessary table entries on-demand.

The query interface exported by TopHat must be carefully designed to enable a diverse range of applications. To reach this goal we use a standard method call XML-RPC to let application access the topology information. This common library can be used by several development languages, such as *Java, C, perl, or python*. A client can develop any kind of application that can retrieve or fetch data from the API if the XML-RPC library is imported to the designed program.

Moreover, we would like to let a user download a complete annotated view of the Internet map seen by TopHat and to add the possibility to plot the topological data regarding the geographical region where they have been probed. This view should be stored during a long time period to allow the research community to study the network dynamics.

V. Conclusion

In this paper, we focused on the Topology Information Component that is part of the OneLab project. In particular, we described how this component allows any application to query TopHat to obtain topology information such as the path from a source to a destination, or various statistics on the topology. We also describe our measurement tools that provide large-scale IP information and trigger topology changes at the AS-level.

TopHat is currently limited to topology information but as an ongoing project, we will continue to enrich TopHat functionalities to provide a better and extended monitoring service on new heterogeneous environment deployed on PlanetLab.

References

[1] "Planetlab." [Online]. Available: https://www.planet-lab.org/.

[2] "Onelab." [Online]. Available: http://www.one-lab.org/

[3] "Caida: Archipelago measurement infrastructure," February 2008. [Online]. Available: http://www.caida.org/projects/ark/

[4] B. Donnet, P. Raoult, T. Friedman, and M. Crovella, *Deployment of an Algorithm for Large-Scale Topology Discovery*. IEEE Journal on Selected Areas in Communications, Sampling the Internet: Techniques and Applications, 24(12), 2006.

[5] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iplane: An information plane for distributed service," in *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, November 2006.

[6] D. Deitrich, E. Lancaster, D. Munson, R. Thomas, R. Connolly, J. Martin, N. Long, and S. Gill, "Cymru project." [Online]. Available: http://www.cymru.com/

[7] R. Chandra and P. Traina, *BGP Communities Attribute*, Internet Engineering Task Force, RFC 1997, August 1996.