

Network Transparentization in Alibaba

Dennis Cai, Jie Cao, Weiguo Chen, Hongqiang Harry Liu, Xin Wu, Ming Zhang
{d.cai,jie.caojie,weiguo.cwg,hongqiang.liu,x.wu,m.zhang}@alibaba-inc.com
Alibaba Group

1 BACKGROUND

Traditional networks are nearly black-boxes to operators. In these networks, network devices only provide coarse-grain information such as SNMP or NetFlow counters, output of command lines or system logs. Such information typically contains aggregated data defined by standard protocols or device vendors, being agnostic of the specific needs of network operators. Moreover, both of the data and control plane behaviors are entirely controlled by device vendors, leaving little room for network operators to customize features or logics inside the devices to obtain fine-grained information.

The invisibility of fine grain and real time network information makes it difficult to improve the reliability and performance of a network. For instance, one prevalent problem in network operation is packet loss. Without the ability to exactly trace individual packets on the data plane, it is hard to localize the loss point. For another example, it is common that a network device's control plane enters an incorrect status causing routing policy violations. Without accurate status information of the routing protocols, operators have to collect data from a large number of devices to perform inference, which is time-consuming and inaccurate. Moreover, fine-grain monitoring usually requires putting new logics into data and/or control planes, which can take months working with vendors without programmability.

To solve the preceding problems, an essential trend in both industry and academia is making network devices "white-boxes" by enabling programmability on both data and control planes. On the data plane, network forwarding chips open programming interfaces to define customized packet forwarding behaviors and meta-data store inside the chips themselves. On the control plane, network device firmware is becoming more inclusive to customized software. For example, most of the prominent vendors we know are moving their switch operating system to Linux running on x86, allowing network operators to run their own protocol stacks or customized agents on their switches. This trend has provided tremendous opportunities to network operators to entirely transparentize their networks and build a new generation of network management systems.

2 DEMO DESCRIPTION

This demo presents Alibaba's recent effort to transparentize its network with newly available programmability on network devices and its new network management framework based on the network transparency that has been deployed inside Alibaba's global scale production networks.

The demo includes four major parts:

- First, it illustrates the deployments of programmable data and control plane, and the overall system architecture of network management on top of transparent networks. This part will be shown in a poster.
- Second, it presents the fine-grain and flexible visualization of the overall real time network status of Alibaba global networks and the ability to diagnose particular traffic problems in real time. This part will be shown in a screen with continuously played video.
- Third, it shows the ability to quickly deploy new switch software agents with Docker to quickly collect specific information and perform failure mitigation locally in real time. This part is a live demonstration. We will build a small testbed with a switch and one or two servers.
- Fourth, it introduces several pragmatic problems that were solved merely by network transparentization and some experiences in deploying programmable data and control planes and running transparent networks. This part will be in the poster.

3 DEMO REQUIREMENTS

The entire demo needs the following environment support:

- A table with at least 2 meters \times 1 meter area;
- At least three power sockets;
- A 27-inch screen (we can bring ourselves' if needed);
- A support (or a wall) and enough ground area for a poster;
- A seat for at least one person;

The demo will be operated in turns by Dennis Cai, Jie Cao, Hongqiang (Harry) Liu, Xin Wu and Ming Zhang from Alibaba Group.