

Improving VPP for Path Tracing

Implementation of the draft draft-filsfils-spring-path-tracing-02 into Vector Packet Processing (VPP)

Student



Julian Klaiber

Introduction: Today's networks, especially service provider networks, are becoming more complex and significant. Increasing amounts of data have to be processed. It is only in the last few years that networks have been continuously developed not only at the hardware level but also at the protocol level. The emergence of protocols such as segment routing via MPLS or IPv6 shows the trend toward more intelligent networks. New protocols are being developed or existing ones enhanced to meet this growth and the need for more intelligence in the network. A new technology that is used in this work is path tracing. This project aims to bring the new technology path tracing into VPP. Cisco invented both technologies; this work was done in collaboration with them.

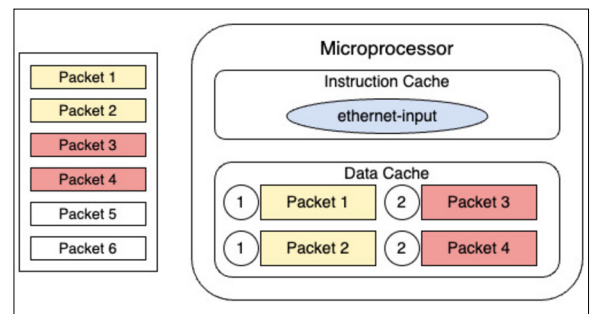
Approach / Technology: In typical networks, so-called ECMP (Equal Cost Multi-Path) routes can be encountered. These ECMP routes are critical since they provide higher route availability and load-balancing to the destination. Even if they are a good thing, they bring much more complexity to the network. The goal was to provide a solution that allows tracing the packets from source to destination and detecting routes that are not going the optimal path to the destination. Path tracing provides a record of the packet path as a sequence of interface IDs to detect these routes. It also allows a complete reconstruction of the packet delivery path. Path tracing works by injecting a source node's probes into the network. These probes are then routed to the sink node, the destination. Each node in the path is a so-called midpoint and adds information to the probe. When the packet arrives at the sink, it will be forwarded to a regional collector and then processed by a pipeline to further process and investigate them.

VPP, on the other hand, is a fast and scalable dataplane implementation. It provides a graph of instructions that can be applied to packets. The packets are processed in a vector that travels through the graph. So on each graph node, the underlying instruction set is applied to the packets in the vector. An instruction set in our case is the path tracing behavior (source, midpoint, or sink). The different behaviors must be applied after looking up the IPv6 destination address before sending it out of the interface. A new node must be introduced into the graph to implement path tracing into VPP. The new node will be inserted in the current graph of instructions as soon as a so-called path tracing interface is configured on the host. The defined behaviors are applied when a packet in the vector is a so-called path tracing probe. That means that the current timestamp, interface id, and interface ingress/egress load are added to the packet before sending it out of the interface.

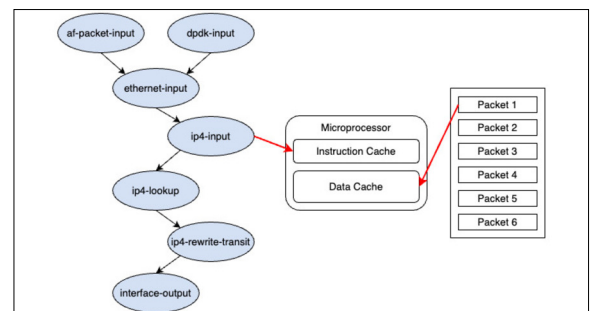
Conclusion: Path tracing is a new technology that will

help make today's networks more intelligent and robust. This technology will allow us to get a better and deeper insight into a network and hopefully allows us to make it more reliable and ready for the future. With the implementation of path tracing into VPP, the first step is testing and building applications for this new technology. VPP allows for a fast and easy deployment of a path-tracing network and, therefore, a much easier test environment to develop and research further in this area.

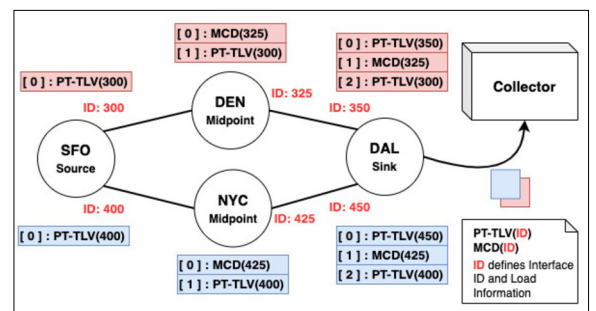
Microprocessor and Data Cache
Own presentation



Vector and Graph Interaction
Own presentation



Path Tracing in Network
Own presentation



Advisor

Prof. Laurent Metzger

Subject Area

Software and Systems,
Computer Science

Project Partner

Cisco EMEA