

Monitoring and Troubleshooting OpenFlow Slices with an Open Source Implementation of IEEE 802.1ag

May 2011
Version 1.3
Ronald van der Pol

SARA

The IEEE 802.1ag Ethernet OAM standard

The IEEE 802.1ag standard offers three kinds of messages:

- Loopback Message and Loopback Reply (LBM/LBR)
- Link Trace Message and Link Trace Reply (LTM/LTR)
- Continuity Check Messages (CCM)

Loopback messages can be compared with the L3 IP ping program. But instead of pinging to IPv4 or IPv6 addresses, LBM enables pinging to Ethernet MAC addresses.

Link trace messages can be compared with the L3 IP traceroute program. But instead of tracerouting to IPv4 or IPv6 addresses, LTM enables tracing to Ethernet MAC addresses.

Continuity check messages are periodic hello messages. These are used to detect loss of connectivity.

All these messages are sent as Ethernet frames with ethertype 0x8902. The IEEE 802.1ag protocols are usually implemented as part of an interface of an Ethernet switch or a router. But not much equipment supports this protocol yet. Fortunately, the protocol can also be used in environments where not all equipment support 802.1ag. Because the protocol uses normal Ethernet frames, switches that do not support 802.1ag will forward the frames as any other Ethernet frame.

Open source implementation of IEEE 802.1ag

SARA has implemented the IEEE 802.1ag protocol as a set of open source utilities that can be run on servers. Currently CentOS, FreeBSD, Linux, and MacOSX are supported. But they can also be run on Arista Networks Ethernet switches by installing the utilities on the embedded Linux server inside the Arista Networks switch. The utilities make it possible to send LBM, LTM and CCM PDUs from the server and also let the server respond to LBM, LTM and CCM PDUs. In this way end users can troubleshoot connectivity problems in

large Ethernet VPNs by themselves without needing the help of network managers of the switches in between. Each Ethernet switch in the network that has IEEE 802.1ag enabled can be probed by the end user by sending LBM or LTM PDUs to that switch.

The open source IEEE 802.1ag utilities are currently tested by SARA and partners on Ethernet based circuits between Amsterdam and Chicago and beyond. We will soon start testing with OpenFlow switches connected to these Ethernet VPNs.

Nature of the experiments and description of approach

The idea of our experiment is to use the IEEE 802.1ag utilities as an external daemon connected to an OpenFlow switch. It will use the OpenFlow feature of sending certain frames to a controller. Ethernet frames with Ethertype 0x8902 will be sent to the external 802.1ag daemon connected to the switch. This external daemon will respond to LBM, LTM and CCM PDUs on behalf of the OpenFlow switch. By doing this each OpenFlow switch can support Ethernet OAM features and it makes troubleshooting of large multi-domain Ethernet based VPNs much easier.

We propose to set this up in the SC11 sandbox and connect some circuits to Chicago and Amsterdam (SARA) to demonstrate how IEEE 802.1ag can be used for monitoring and troubleshooting such VPNs/slices and how OpenFlow can be useful in such a setup. We will also investigate if we can setup a link to the GENI GEC12 meeting, which is currently scheduled to take place in Kansas City during the same week as SC11.

Description of the network requirements and vendor collaborations

We ask for a couple of strategically placed OpenFlow switches in the sandbox. There should be a flow entry for Ethertype 0x8902 to be sent to the external controller. In the next months we will investigate how we can integrate our IEEE 802.1ag utilities with an OpenFlow switch more closely, e.g. by deciding whether or not the switch should forward the 802.1ag PDUs, just like a "normal" 802.1ag capable switch would do.

Desired outcomes and takeaways

We would like to learn if IEEE 802.1ag is a useful monitoring and debugging tool in a large multi-domain Ethernet and how integration with OpenFlow can help.

End users can use the 802.1ag utilities to send LBM and TLM PDUs to the 802.1ag enabled OpenFlow switches. The 802.1ag enabled OpenFlow switches will also send periodic CCM PDUs. A website will be setup that will show a live reachability matrix of all 802.1ag enabled switches.

Relevance to the HPC community

OpenFlow based networks are expected to play a role in both HPC datacenters and VPNs/slices used by Scientific Virtual Organisations. Monitoring and troubleshooting such networks, especially in the case where multiple domains are involved, is important for the success of this new networking paradigm. We think the IEEE 802.1ag protocol is a powerful tool for end-users and network managers to monitor and troubleshoot these complex virtual networks.

References

1. IEEE Std 802.1ag-2007
<http://standards.ieee.org/getieee802/download/802.1ag-2007.pdf>
2. Open Source Implementation of the IEEE 802.1ag Protocol
<http://nrg.sara.nl/dot1ag-utils>