

OpenFlow support with HP Procurve 5400zl series

ProCurve is a division of HP that offers various classes of enterprise switches. HP-Labs is the central research organisation of HP. A three-way partnership between Stanford, HP Labs and the ProCurve division enabled the implementation of Openflow as an experimental feature on ProCurve's network hardware.

HP-Labs implemented OpenFlow in the firmware for the HP Procurve 5400zl series, which represents a class of enterprise switches. New code was added to interface the OpenFlow code with various parts of the switch, such as sending and receiving packets, enumerating ports and VLANs and setting up the TCAM. A user interface was added to monitor and set up the OpenFlow module, this user interface was tightly integrated into the existing the CLI and SNMP interface. The OpenFlow test suite from the reference implementation was modified to test OpenFlow running on the ProCurve switch from a Linux PC.

The OpenFlow implementation manages both a hardware path and a software path, because not all OpenFlow rules can be instantiated in hardware. The hardware path is managed using a TCAM, an OpenFlow rules is translated into a TCAM entry that capture packets of that flow and process packets according to the specified actions. A special TCAM entry captures all packets not matched by those other TCAM entries and send them to software. The software path uses the full set of OpenFlow rules to process and forward packets, if no rule match a packet it is sent to the OpenFlow controller.

Any new feature that is added to existing network hardware takes significant effort since there may be architectural limitations for the task at hand. OpenFlow defines a flow as a 10 tuple pattern, but this is more bits than common TCAM implementations can match and therefore the ProCurve switch cannot match the L2 header in hardware. The result is that when instantiating a OpenFlow rule in hardware, MAC source and MAC destination are ignored, and EtherType has to be IP. The remaining 7 tuples are matched properly either as exact match or wildcards. Non IP rules are kept in software. The hardware can only drop or forward to a single port, therefore any rule that specifies multiple ports or header rewrite is kept in software. The hardware can support around 1500 rules per line card, any additional rule is kept in software.