Huawei Technologies
SDN Showcase at SDN and OpenFlow World Congress 2013

Introduction

While OpenFlow already penetrated data centers, service providers are now turning to OpenFlow to ease network operations and provisioning. In preparation for one of the biggest European conferences focusing on SDN and OpenFlow, Huawei Technologies commissioned EANTC to validate Huawei's OpenFlow solution targeted at service providers. This report explains the use cases, test cases and results we collected during the test execution in EANTC's lab in Berlin, Germany.

Huawei built a service provider access network in EANTC's lab based on a mix of Huawei OpenFlow switches and legacy switches that do not support OpenFlow. The test network was managed by a pool of Huawei OpenFlow controllers. EANTC used Ixia's IxNetwork testers to emulate subscriber traffic.

Test Setup

The test bed network was composed of four Huawei SN640 OpenFlow 1.3 switches controlled by Huawei's OpenFlow 1.3 controller cluster, called Smart OpenFlow Controller (SOX). Huawei also brought two AR1600 routers that did not support OpenFlow to represent legacy networks. One of these routers was configured as part of the OpenFlow domain with Autonomous System (AS) 100, while the other router was configured as part of a legacy domain in AS 200. All devices were interconnected using Gigabit Ethernet.

Each of the OpenFlow switches established a single OpenFlow channel to the OpenFlow controller cluster using an out-of-band management network. The controller automatically discovered and displayed the complete network topology including the non-openFlow routers. The controller negotiated and learned the proper version for each OpenFlow device.

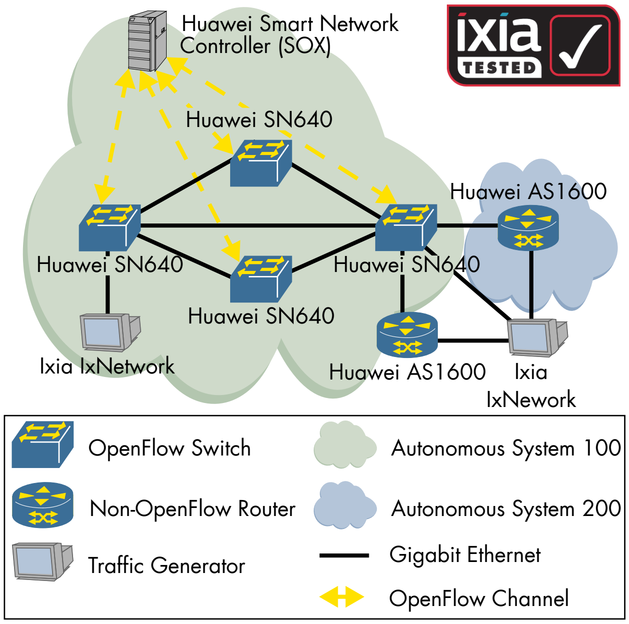

Figure 1: Test Bed Network Topology

Results

Interworking With Legacy Devices

Unless a greenfield deployment is an option, service providers are likely to expect OpenFlow devices to be installed in the network gradually. It is therefore obvious that OpenFlow-based networks have to interwork with legacy network components. This interworking could be part of the same administrative domain or a different domain.

Test Results Highlights

- ➔ **High-availability including OpenFlow controller resiliency**
- ➔ **Traffic management with network-wide Quality of Service and multi-path forwarding**
- ➔ **Self service traffic management**

In this first test we looked into how common and well trusted protocols such as IP/MPLS and External BGP interwork between OpenFlow and non-OpenFlow devices. The first step involved the OpenFlow Controller discovering the complete network topology. The automatic topology discovery algorithm utilized Link Layer Discovery Protocol (LLDP) messages to discover the network topology. Since LLDP was enabled on the legacy devices these and the links connected to the OpenFlow switches were also discovered by the Controller.

We then investigated IP/MPLS interworking between the OpenFlow switches and the Non-OpenFlow devices. Huawei configured Resource Reservation Protocol – Traffic Engineering (RSVP-TE) to be used as Label Switched Path (LSP) signalling protocol. Huawei used Open Shortest Path First – Traffic Engineering (OSPF-TE) as link state routing protocol to distribute the link state and traffic engineering information between nodes in the same autonomous system. All those tasks were performed by the controller. Both the OpenFlow and the non-OpenFlow switches did not need to be provisioned by hand at all.

We verified that OSPF and RSVP-TE sessions were established between the Non-OpenFlow switch in AS 100 and the SOX via the controller's GUI and devices' CLI interface. While MPLS service and user traffic was running, we did not observe any packet loss. The OpenFlow switch that was connected to the Non-OpenFlow switch had installed flow entries to push and pop MPLS header.

Next on the verification list was the use of eBGP to inter-connect AS 100 (the OpenFlow domain) and AS 200 (the non-OpenFlow domain). One OpenFlow switch was configured as Autonomous System Border router (ASBR) for AS100 and another router was configured as an ASBR for AS 200. Once the controller discovered the neighbor, it installed a flow entry for the control traffic and established the BGP sessions toward the Non-OpenFlow switch in AS 200. After the routing information was exchanged between AS 100 and AS 200, we verified that IPv4 prefixes advertised from Non-OpenFlow switch in AS 200 were installed in SOX's Routing Information Base (RIB) and vise versa. We did not observe any packet loss for the user traffic.

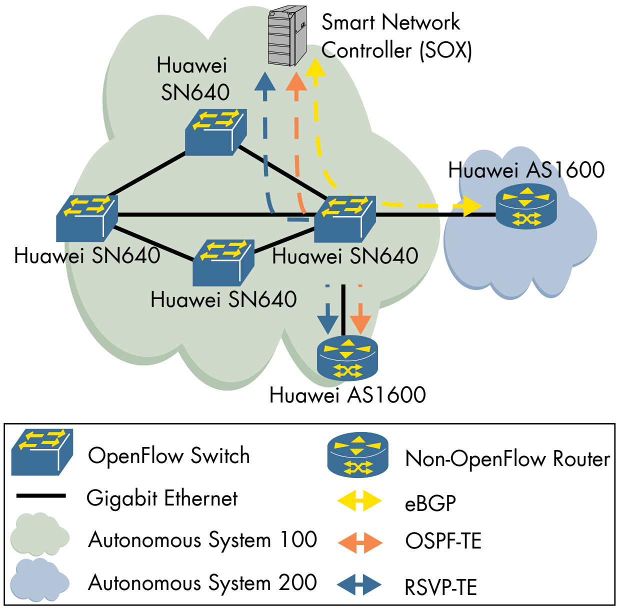

Figure 2: Interworking With Legacy Devices

Multi-Path Forwarding

Network resource optimization is one of the big challenges that traditional networks are facing these days. Various methods, such as Equal Cost Multipath (ECMP), are employed in the networks to better utilize the network resources. ECMP means that multiple paths to the same destination are used to provide load balancing.

According to OIF, centralized network control allows more granular network control and optimization than distributed network control. The flow-based OpenFlow control model allows network administrators to apply policies at a very granular level, including session, user, device, and application levels.

In order to test these capabilities of OpenFlow multipath forwarding, we sent IPv4 user traffic for three different network services, distinguished by DSCP, IP source and IP destination address. For each network service we applied a different bandwidth profile as show in table 1. I

Traffic Class	Service A	Service B	Service C
High	20 Mbit/s	40 Mbit/s	100 Mbit/s
Medium	30 Mbit/s	60 Mbit/s	150 Mbit/s
Low	50 Mbit/s	100 Mbit/s	250 Mbit/s

Table 1: Bandwidth Profiles

The user traffic was equally load balanced between two equal cost paths based on the flow level. The direct link between the left and the right OpenFlow switch was set with a higher link metric, therefore both indirect links were used. We did not observe any packet loss or re-ordered packets for all services and Class of Service (CoS).

Link and Node Resiliency

In current service provider networks, each network layer has its own resiliency mechanism resulting in more expensive networking equipment, higher operational and management costs. We asked Huawei if OpenFlow can help provide a uniform and reliable resiliency mechanism while at the same time provide carrier-grade resiliency?

To answer this question we looked into the resiliency options implemented by Huawei's OpenFlow-based solution. We emulated the traditional service interruption scenarios such as link and node failure and measured the service interruption time.

We tested the protection approach utilizing the OpenFlow 1.3-defined fast-failure bucket group type. This mechanism enables the OpenFlow switch to change the forwarding path without requiring round trip communication to the controller. This method significantly reduced the failure reaction and recovery time. The results from both resiliency categories are shown in the following figure:

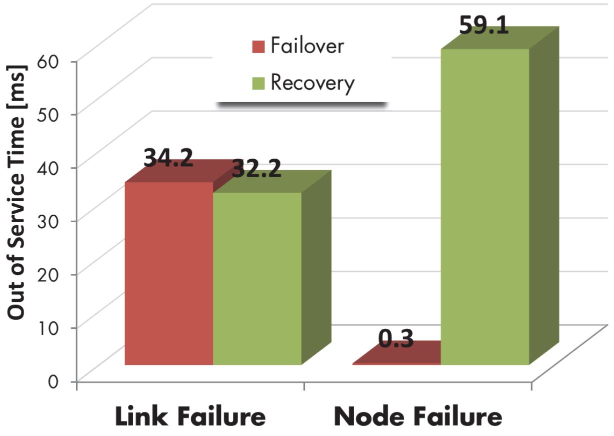

Figure 3: Service Interruption Time per CoS

Controller Resiliency

Service provider networks need to handle a potentially large amount of user flows and traffic. If all user flows are controller through a single device, that device represents a single point of failure as well as, perhaps, a performance bottleneck. Therefore, the Open Networking Foundation (ONF) introduced a multi-controller feature in the OpenFlow 1.2 standard with the main goal to avoid a single point of failure. In this architecture each of the OpenFlow switches establishes an OpenFlow channel to all OpenFlow controllers in the domain.

Huawei chose to implement their OpenFlow controller as a cluster composed of multiple OpenFlow controllers. All of the controllers in the cluster can be viewed as one single logical controller. Huawei explained that having multiple controllers in one cluster provides scalability and reliability as a big benefit to the service provider.

The number of controllers in the cluster could dynamically be adjusted to manage the load based on real time observation of the network state. We tested the Huawei's SOX controller cluster ability to react to controller failure and dynamic load balancing.

In our test we used one physical machine running several independent controller processes. Initially, Huawei configured the controller cluster with two OpenFlow controller instances. Both instances were configured to run as *equals* in the cluster. We verified that the *packet-in* (data packets that are sent to the controller) load was equally balanced between both instances. Each of the controller instance was configured to handle at most 10,000 *packet-in* packets/second.

Once we increased the *packet-in* load to 32,000 packet/second, the number of controller instances increased automatically to 4 to handle the load. We verified that each of the controller instance was handling 8,000 *packet-in* packets per second without loss. While the traffic load was running, we disabled one of the controller instances. The controller cluster detected the failure of one of the instances and instantiated a new controller instance. The *packet-in* load was again distributed equally between the controller instances.

Rate Limiting

Current Quality of Service (QoS) deployments in service provider networks typically handle customer traffic based on 8 bits classification available in IP headers. This means that at most, service providers can distinguish between 8 different classes of service. The SDN approach allows to classify based on flow information. In our test we classified the packets based on DSCP, IP source and destination information.

We tested automated rate limiting per service and CoS using *per-flow* meters introduced in OpenFlow 1.3. These per-flow meters provide measurement and the control at the flow level.

We used a service with bandwidth profile (*Service C)* as detailed in Table 1. We verified that the service was provisioned according to the bandwidth profile by sending user traffic at Committed Information Rate (CIR) for all classes of service. The SOX' GUI showed that the flow entries were installed in both OpenFlow switches.

As a second step we increased the traffic rate for the High traffic class to 200 Mbit/s (twice the CIR) and observed that half of the traffic was remarked to the Low traffic class. After increasing the Low traffic class rate to 500 Mbit/s, we observed that this traffic class was rate–limited to 250 Mbit/s.

On-Demand Elastic Quality of Service

Huawei explained that in contrast to traditional static QoS implementations, SDN provides mechanisms to allocate network resources in an elastic way, determined by individual user profiles and application requirements to ensure an optimal user experience.

In this QoS-focused test we looked into the elasticity of services – on-demand modification of Committed Information Rate (CIR) attribute per Class of Service. We verified that Huawei's SOX controller can provide an interface to the customers/applications and change the CIR service attributes on demand or automatically. For this test we used the same service as used in the previous test - *Service C*.

In our first scenario we verified that the CIR can automatically be changed when specific traffic rate threshold was reached. In addition, we verified that the bandwidth of the traffic flow was reduced by 10% automatically when the quota for the total traffic volume of 10 GByte for Low CoS was exceeded, and increased back at CIR when the customer provisioned additional data volume. This kind of quota implementation helps the service provide to keep control about their sold service plans.

Summary

Service providers could be assured SDN solutions are being created to provide high availability, QoS and self service traffic management. Our test results show that industry standard 50 ms failure recovery could be reliably supported by Huawei's OpenFlow based networks in case of link, node and controller failure.

Virtualization, mobility, and the need for network resource monetization and optimization place significant demands on the network— we verified that those demands can be handled by Huawei's SDN network architecture and solution components.

Many service providers operate today a packet based converged network for data, voice, and video traffic. Those existing networks can provide differentiated QoS levels for different applications, however, the provisioning of those resources is highly manual and static. Because of the static nature, the network can not dynamically adapt to changing traffic, application, and user demands.

Huawei's SDN solution provides a new, dynamic network architecture that transforms traditional network backbones into service-delivery platforms.

About EANTC

The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier Ethernet, Triple Play, and IP applications.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/