EANTC

Multi-Vendor NETCONF/YANG-Based SDN Management Interoperability Test

White Paper

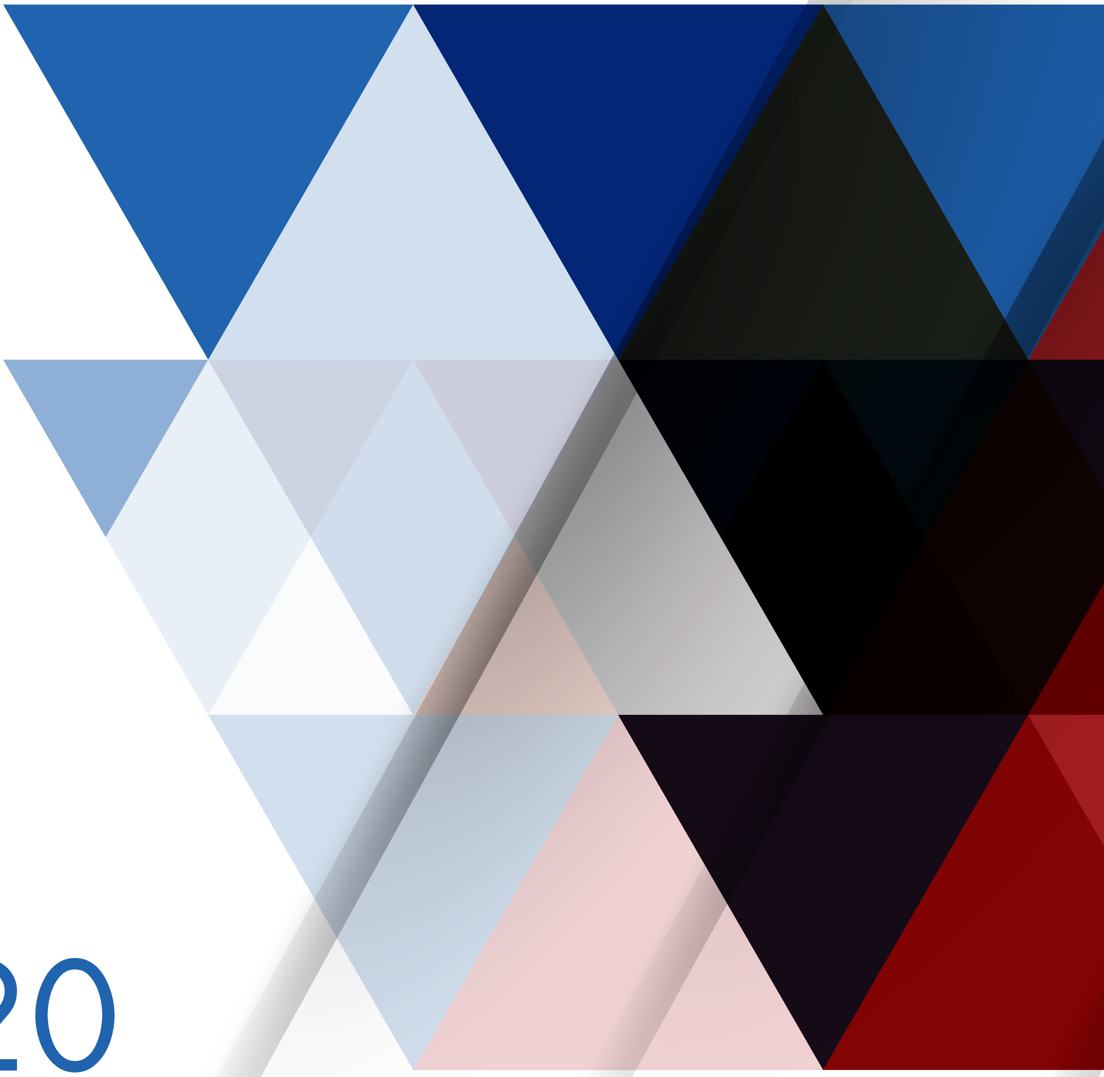

2020

Table of Contents

Editor's Note

With great pleasure, we are publishing this report of our first multi-vendor interoperability test exclusively dedicated to network configuration management and performance monitoring. The event was focused around NETCONF, a protocol to manage network device configurations in a standardized way. NETCONF is increasingly implemented by manufacturers, specifically due to the demand by network operators for centralized orchestrator solutions and automated provisioning. The IETF, the OpenConfig initiative, and other standards bodies have worked avidly to create a wide range of standardized YANG models complementing the NETCONF protocol. Multi-vendor management interoperability unfolds its full power with such standardized YANG models. On many occasions, manufacturers and operators have vowed to enable full mix-and-match combinations of orchestrators/controllers and network devices.

To support this significant initiative at EANTC, we decided to split off the area of management testing from our annual MPLS SDN interoperability test events, dedicating much more attention and time to a new series of NETCONF/YANG management interoperability events. We invited more than 100 vendors; more than 40 were interested and participated in the first preparation call. Subsequently, the more details of the test coverage were discussed, the fewer companies were ready to participate.

As it turned out, most vendors are still in the early stages of NETCONF/YANG implementations. The good news is that many promised that their implementations will be ready in the next few months when we will continue our series of NETCONF interop events.

Our first event in November proved to be very successful and productive, thanks to outstanding support from strong Cisco, Huawei, and Nokia teams. All vendors participated with orchestrators, and two of them additionally with routers. Initially, we covered the basics of device configuration, followed by more advanced L3VPN network service configurations. Finally, we evaluated multi-vendor telemetry through gRPC collectors.

All tests were carried out remotely during the Covid-19 pandemic. We would like to thank all vendor teams for their flexible and efficient collaboration across 16 time zones. Centralized compute resources were made available by EANTC in collaboration with VMware, who provided a complete VMware Integrated OpenStack (VIO) cluster and support.

The support of standardized YANG models will be the big debate in 2021: Which models to standardize on—IETF (more liked by manufacturers, although it is difficult to generalize) or OpenConfig (more looked at by service providers)? In our event, we looked at both worlds with mixed results. The OpenConfig models turned out to be slightly less interoperable; with YANG models, the devil is in the details. The IETF models benefit from the "IETF YANG Doctors" providing quality assurance before models are released. In any case, orchestrators often had to revert to proprietary YANG models supported by the router vendors during this event.

We faced a number of rather straightforward interoperability issues, reconfirming that vendors did not test with each other extensively before and that our event was one of the first of its kind. Some issues found and resolved related to standards interpretation and general design rules of YANG models. EANTC hopes to contribute to a joint understanding of NETCONF/YANG implementation with our test events, gradually increasing the multi-vendor interoperability in this important area of innovation.

Introduction

Today's telecommunication networks are becoming ever more complex. Enterprise customers increasingly demand broadband connectivity for private, public, and hybrid cloud services. The transition to 5G creates new advanced service and network scale requirements. All of these innovations have a vital and growing impact on business agility and efficiency on telecom operators. Advanced end-to-end network orchestration is required to manage the transport networks providing a wide range of VPN services with customer-specific traffic engineering. These end-to-end infrastructures typically span multi-vendor router solutions, which are challenging to configure and maintain efficiently. Automation of provisioning and monitoring tasks at the network layer is essential to the operators' success in the future.

The IETF (Internet Engineering Task Force) has standardized NETCONF (Network Configuration Protocol) and YANG models to simplify the configuration management and open up a disaggregation of management components (controllers, orchestrators) from routers. With network automation, operators can efficiently create, monitor, and deliver an array of new network services in a scalable infrastructure.

Executive Summary

We designed the interoperability test cases to focus purely on NETCONF and transport network-specific YANG models. The test cases helped us to investigate the management plane between the controllers and the network devices. We focused on use cases for integrating NETCONF applications to address key topics such as 5G stations, related clocking management, Layer 2 and Layer 3 VPN service provision in a multi-vendor environment, covering their required monitoring, configuration, and automation functions.

The event covered eleven test groups to evaluate Cisco, Huawei, and Nokia solutions. In total, we evaluated 39 multi-vendor test group combinations. We carried out functional tests across four categories: Device functions and configurations, service provisioning, monitoring, data export, and Precision Time Protocol (PTP) features.

The Device Functions and configurations area includes the necessary test scenarios for configuring and monitoring interface-specific parameters, common system properties on a network device, MPLS, and Segment Routing based parameters.

The Service Provisioning section covers the creation of Layer 2 EVPN service, L3VPN service, and configuring Bidirectional Forwarding Detection (BFD). The Monitoring and Data Export includes telemetry streaming and alarm management. The fourth section covers the configuration of the PTP profile. Initially, we planned a dedicated test area for microwave equipment, but we did not test the area due to the lack of participants with microwave solutions.

All tests have successfully demonstrated the management of IP implementations test. However, one of the controllers faced an issue when recovering the modified data in the router since the router exposed multiple views of the same information to the controller via multiple YANG models. The controller selected a YANG model to roll back the interface configuration. In another case, one of the controllers faced an issue with the "Merge" operation to revert the configuration. Instead of the “Merge” operation, the controller used the "Replace" operation. During the test of one combination, a proxy between one of the controllers and the routers had to be set up to change the namespaces to the XML file's accepted form. This was used in all the following tests for this particular combination.

In the management of interface – OpenConfig test, the controllers successfully configured the interface parameters and modified the parameters using the OpenConfig YANG model. Once the controller faced the issue with changing the MTU size with the OpenConfig model. The participated test combinations were successfully demonstrated the System Management test by configuring the NTP, DNS resolver, and RADIUS. However, system identification was not demonstrated with the same YANG model. In the Retrieve Interface Frame Sizes Distribution test, the controllers successfully retrieved the statistic for the number of packets received on the interface via the YANG model. However, the good and bad packets were not differentiated on the statistic. The test case Multiprotocol Label Switching was performed with four different combinations. In the test, controllers managed to configure the OSPF, BGP, and MPLS parameters via the YANG model. However, a controller faced a compatibility issue with their GUI to use the web browser. The test cases MPLS with Segment Routing, Layer 2 EVPN service, L3VPN service, Bi-Directional Forwarding Detection Telemetry streaming from the devices, and the participated test combination were successfully demonstrated. In the Precision Time Protocol test, the routers supported some basic features. The router did not support the PTP profile.

From the testing perspective, the NETCONF interoperability event was very successful, and most of the issues were fixed during the test execution. Prevalent issues relate to adapting to vendor-specific YANG models.

Remote Collaboration Aspects

Following our successful mixed on-site and remote MPLS SDN interoperability test event in March 2020, and our series of fully remote interoperability and performance test programs in Network Function Virtualization (NFV) since 2016, we decided to take a step for SDN interop testing we had planned for a long time. Covid-19 changed the environment, encouraging us to go this way for multi-vendor testing as well. The remote collaboration had many aspects and layers:

The Test Setup: EANTC had the responsibility of setting up the test environment to be used during the Hot Staging, which contains providing the connectivity between three labs laying on three different continents (Europe, North America, and Asia). EANTC installed the virtualized Nokia NSP Controller components and virtualized Cisco NSO components at EANTC's lab. The Huawei iMaster NCE-IP controller was located at Huawei labs in China. All routers under test were located in the vendors' labs, specifically at Cisco and Huawei labs.

Connectivity: As part of the test setup, EANTC created and configured IPsec tunnels between the EANTC's firewall and the test sites. The connections were verified one week before the Hot Staging with the cooperation of the participated vendors.

In collaboration with VMware, EANTC provided a VIO platform (VMware Integrated OpenStack) to host Cisco and Nokia's controllers/orchestrators. The platform has sufficient resources to host much more large-scale workloads.

Remote Access: EANTC created remote access accounts for all participants to reach their VMs in EANTC's lab to configure their devices and run the tests.

Wiki: A Wiki collaboration space was used to plan the test and document all results.

Digital Meeting Rooms: EANTC provided the participants with links for digital meeting rooms to meet together, discuss, and run the tests.

Adapted Work Time: The collective work between three continents requires the technical abilities and the ability to adapt the work time and routines to reach the work goals. All participants were brilliantly able to adapt to one work-time period, which allowed this wonderful collaboration experience.

Devices Under Test

Vendor Name	Device Name	Device Role	Software/ Firmware Version
Cisco	Network Services Orchestrator (NSO)	Controller	NSO 5.4.0.2
Cisco	N9K-C9316D-GX	PE Router	Version 9.3(5)
Cisco	N3K-C36180YC-R	PE Router	Version 9.3(5)
Huawei	iMaster NCE-IP	Controller	V100R020C00
Huawei	NetEngine 8000 M8	PE Router	NetEngine 8000 V800R013C00SPC003T
Huawei	NetEngine 8000 X4	PE Router	NetEngine 8000 V800R013C00SPC003T
Huawei	NetEngine 8000 M14	PE Router	NetEngine 8000 V800R013C00SPC003T
Nokia	Network Services Platform (NSP)	Controller	20.11 (beta)

Table 1: Software and Hardware Details

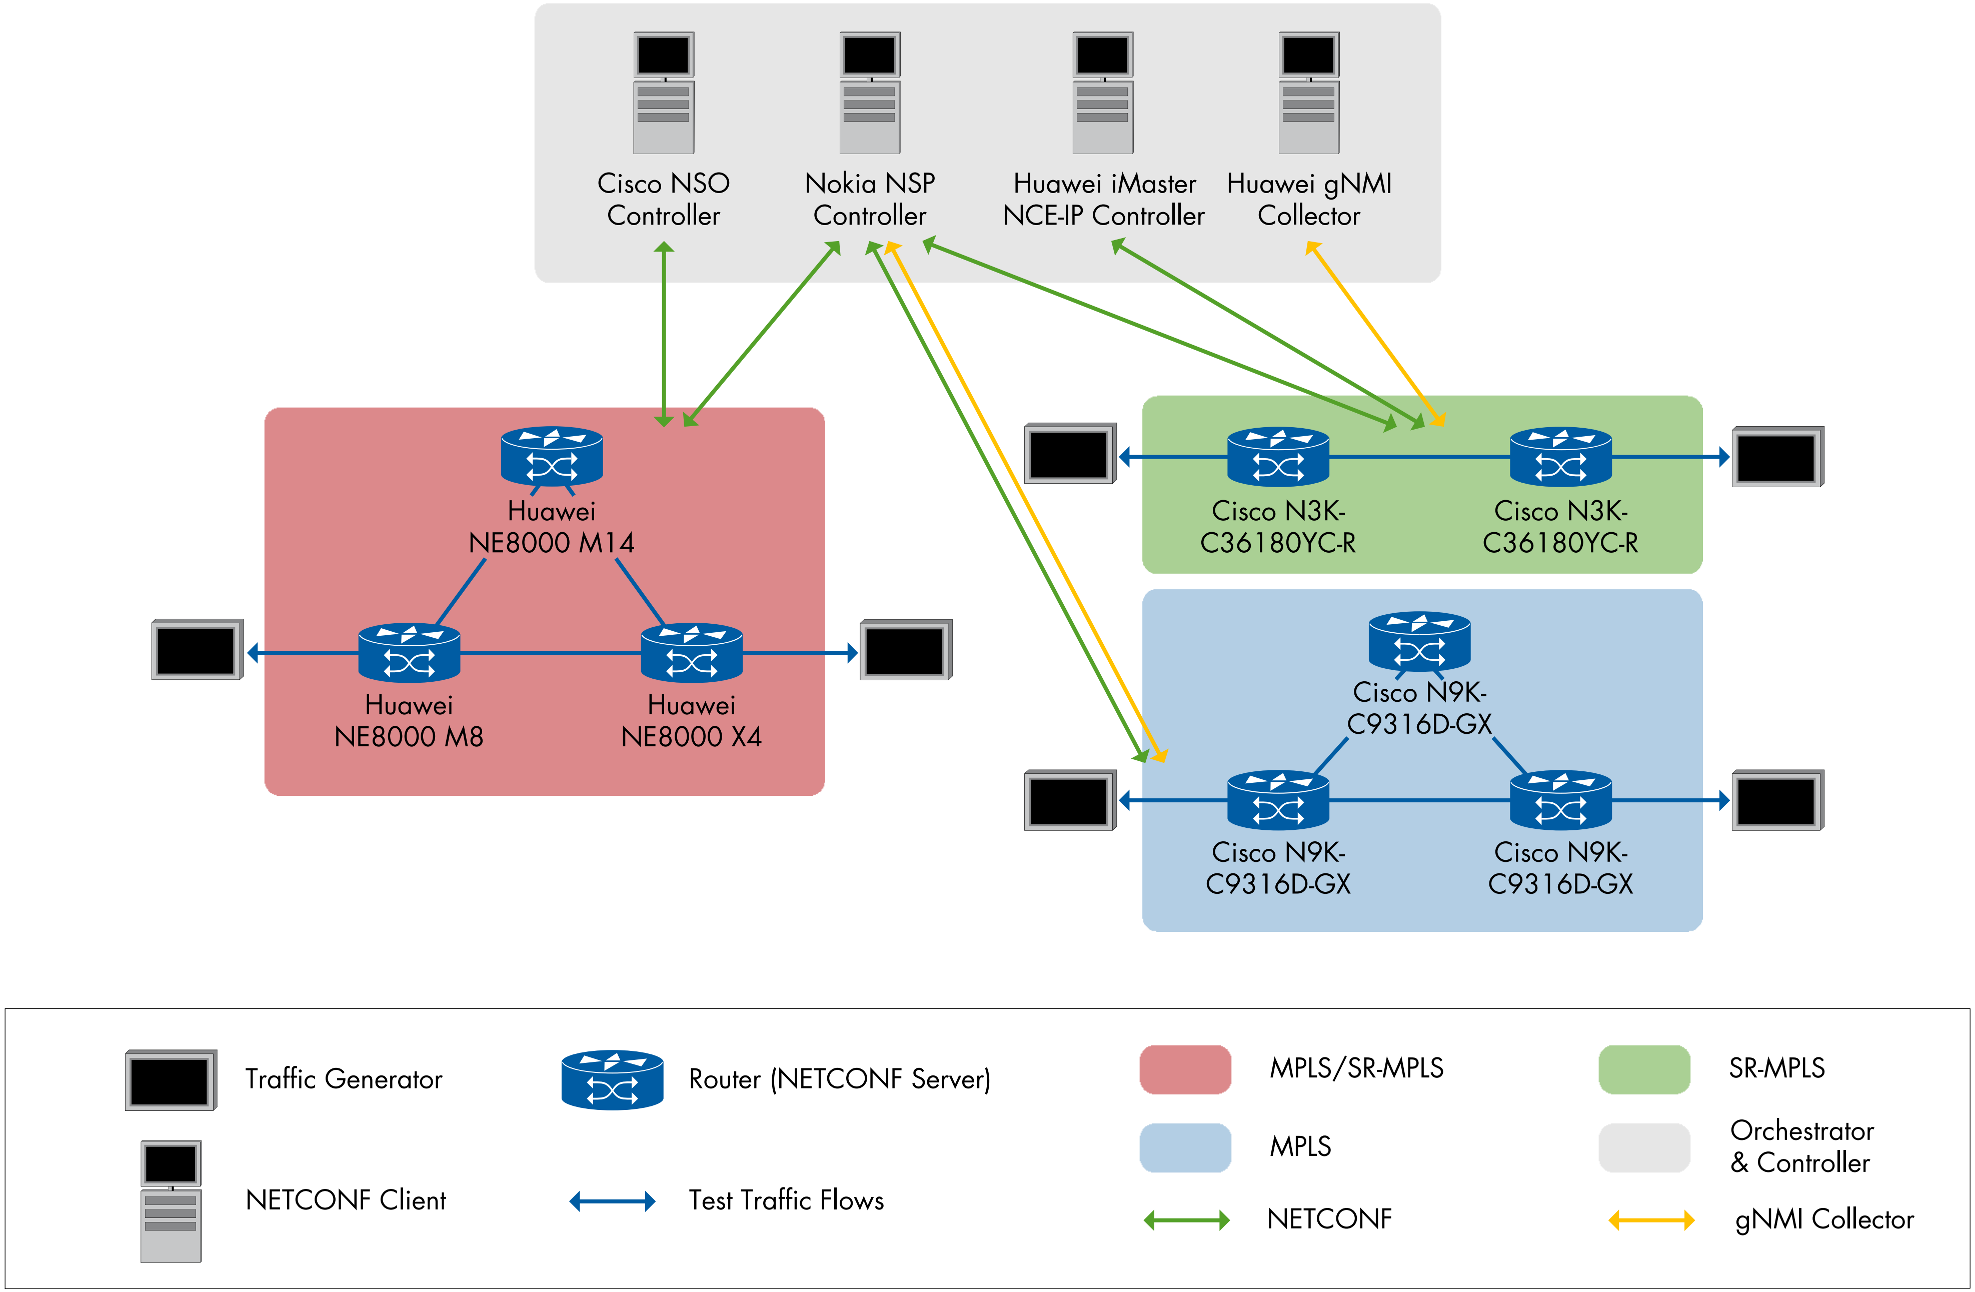
Cisco NSO Controller
Nokia NSP Controller
Huawei iMaster NCE-IP Controller
Huawei gNMI Collector
Huawei NE8000 M14
Huawei NE8000 M8
Huawei NE8000 X4
Cisco N3K-C36180YC-R
Cisco N3K-C36180YC-R
Cisco N9K-C9316D-GX
Cisco N9K-C9316D-GX
Cisco N9K-C9316D-GX
Traffic Generator
Router (NETCONF Server)
MPLS/SR-MPLS
SR-MPLS
NETCONF Client
Test Traffic Flows
MPLS
Orchestrator & Controller
NETCONF
gNMI Collector

Management of IP Implementations

The configuration of an IP interface is the most basic part of VPN provisioning. But in a multi-vendor environment, this configuration can be severe having the IT engineer remember the command level of details on every vendor's device and the dependencies under this parameter, such as IP address and MTU size. Especially the manual operation workload cannot be ignored in a large scale or on-demand business model. In particular, this configuration is the foundation of the network service, and correctness is the key to determining whether the VPN can run normally. NETCONF acts as an automation and network template, becoming the key to repetition and scale.

We verified the SDN controller's ability to add a Layer 3 interface via NETCONF on the PE routers using the YANG models. We performed four test combinations from all different vendor devices. To start a DUT pair's test (between controller and router), we performed the following steps. We confirmed that a NETCONF session has been established on each device. Then, we started to edit the configuration via NETCONF from the controller and observed the router's configuration updates via CLI. Finally, we pinged and sent traffic through the IP address that has been configured. We repeated the above steps for each test pair.

We observed that a NETCONF session has been successfully established with the PE router over SSH from the controller. Both parties also sent HELLO messages to advertise capabilities. The following Figure shows an example output taken from one of the test pairs.

The established session allowed the controller to perform one additional step: To use the GET-SCHEMA operation to retrieve the router's YANG model. This step is optional if the YANG model hasn't been released online or available before testing.

From here, the configuration of NETCONF began. Its actual configuration work started from a configuration data storage. As defined in RFC 6241, the configuration data storage is a complete configuration data set that allows network devices to operate normally. The controller obtained this configuration from the router under test via get-config, referred to as running configuration. The running configuration datastore holds the complete configuration currently active on the network device. Only one configuration datastore of this type exists on the device, and it is always present. We opened the YANG model previously retrieved and added the data of interface parameters to be configured, like IPv4 address, MTU size, and interface description, into the YANG model. The controller pushed the modified model to the router via NETCONF edit-config.

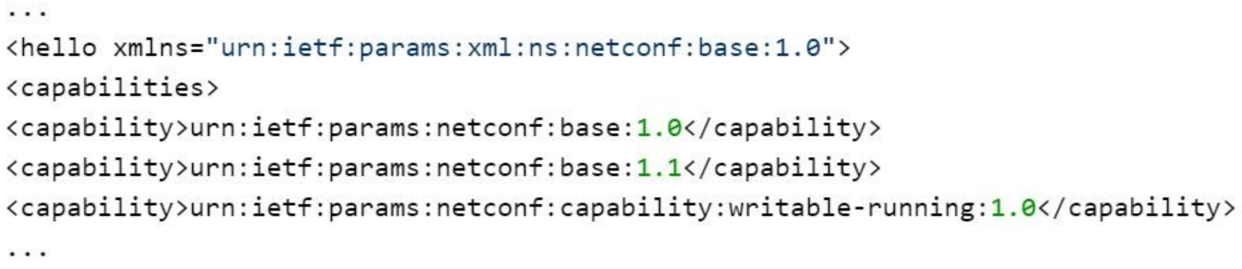

```
...
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<capabilities>
<capability>urn:ietf:params:netconf:base:1.0</capability>
<capability>urn:ietf:params:netconf:base:1.1</capability>
<capability>urn:ietf:params:netconf:capability:writable-running:1.0</capability>
...
```

Figure 1: Writable-Running Capability

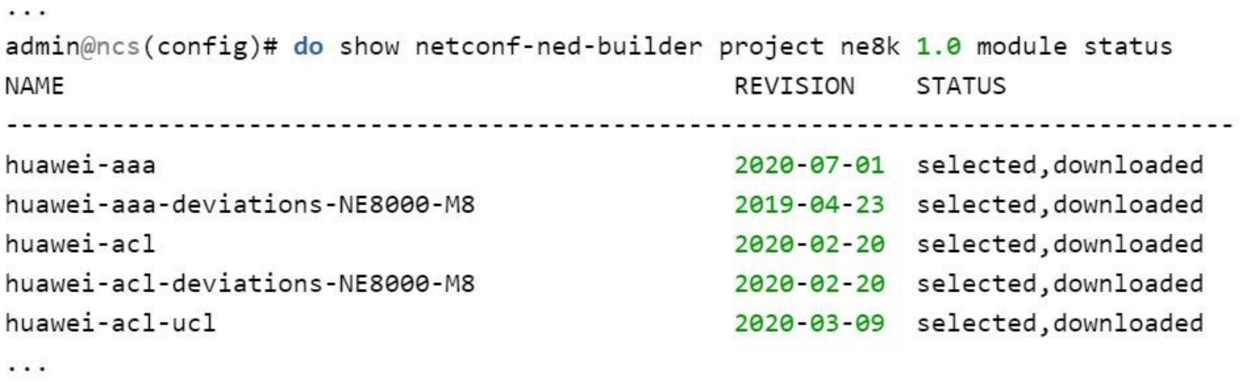

```
...
admin@ncs(config)# do show netconf-ned-builder project ne8k 1.0 module status
NAME                                                REVISION     STATUS
-------------------------------------------------------------------------------
huawei-aaa                                          2020-07-01   selected,downloaded
huawei-aaa-deviations-NE8000-M8                     2019-04-23   selected,downloaded
huawei-acl                                          2020-02-20   selected,downloaded
huawei-acl-deviations-NE8000-M8                     2020-02-20   selected,downloaded
huawei-acl-ucl                                      2020-03-09   selected,downloaded
...
```

Figure 2: Example of Retrieved YANG Models

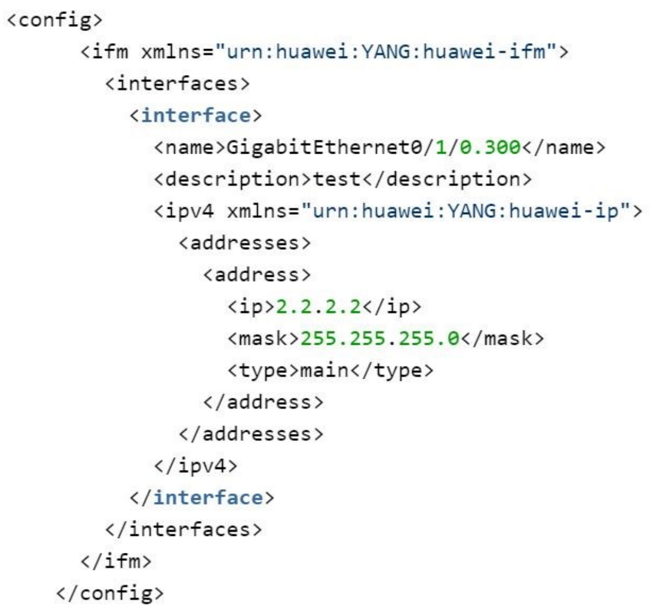

```
<config>
      <ifm xmlns="urn:huawei:YANG:huawei-ifm">
        <interfaces>
          <interface>
            <name>GigabitEthernet0/1/0.300</name>
            <description>test</description>
            <ipv4 xmlns="urn:huawei:YANG:huawei-ip">
              <addresses>
                <address>
                  <ip>2.2.2.2</ip>
                  <mask>255.255.255.0</mask>
                  <type>main</type>
                </address>
              </addresses>
            </ipv4>
          </interface>
        </interfaces>
      </ifm>
    </config>
```

Figure 3: Added IP Address in YANG Model

Once the NETCONF edit-config operation was finished, all pings reached the gateway IP address, indicating that the IP address has been successfully configured on the router. We also observed the updates of the MTU size and interface description via CLI in the status counter on the router.

As an optional step, we verified the controller's ability to identify changes in the router's IP configuration. This functionality required the controller to update the running configuration datastore first. However, this function in the current solutions supported the manual operation, so we manually started get-config same as described in the previous steps. Through the obtained configuration, the controller successfully identified the changed IP interface.

As the final step, we verified that the controller updated the IP interface configuration. The controller pushed the YANG model on the router the same as described before. We observed that the IP configuration has been successfully updated on the router, and the previously configured information has been deleted.

In the test, we observed the following issues. However, it did not have any impact on the tests since the engineers quickly fixed it. The configuration redeployment failed on one of the controllers since the controller has two views of the same modified information through the native model and the IETF model. The change was reported twice, once using the native YANG model of the router and once using the IETF models. The IETF model has been used to rollback the interface's configurations. In one case, the controller used the operation "Replace" to change the manually configured IP address in the last step, using the "Merge" operation resulted in two IP addresses on the interface. During the testing of one combination, the PE routers replied with error messages regarding the XML files' namespaces. The controller engineers set up a proxy between the controller and the routers to change the namespaces to the accepted form. This was used in all the following tests for these particular combinations.

All combinations were tested successfully. EANTC observed the change of the basic configurations, like IP address with the Subnet Mask, the interface (Up, Down) status, and setting the MTU size. The following Table shows the test combinations and the supported YANG models.

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Huawei Proprietary Model / Re-enforce Manually changed settings with the IETF model
Huawei iMaster NCE-IP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	OpenConfig: To read the parameters from the routers Cisco Proprietary Model: To configure the parameter
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	IETF YANG Model: IETF Interfaces

Table 2: Management of IP Implementations - Successful Combinations

Management of Interfaces—OpenConfig

Network management describes the process of configuring, monitoring, and troubleshooting the networks, which have multiple layers of applications, protocols, logical and physical links.

The management of large networks requires the controller and the network nodes to support multiple model types. This motivated EANTC to test the interoperability of using the NETCONF protocol to push the management configurations of the network devices interfaces through the OpenConfig model besides the IETF or the proprietary models used for other tests. We verified the ability to use the OpenConfig model to manage the interface configuration of the PE routers.

After the successful NETCONF session establishment between the controller and the routers, the controller retrieved the routers' current running configuration. The participating controllers added several parameters like interface name, IPv4 address, IP prefix length, MTU size, and a loopback interface in the retrieved model. Then, controllers committed the configuration to be applied on the routers. We successfully verified the configuration changes on the routers. The controllers also modified the configuration, like rename the interface and changed the interface's status via the YANG model. In one of the combinations, the controller couldn't change the MTU size with the OpenConfig model.

The Table below shows the test combinations and the supported YANG models.

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	OpenConfig
Huawei iMaster NCE-IP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	OpenConfig
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	OpenConfig
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	OpenConfig

Table 3: Management of Interfaces—OpenConfig - Successful Combinations

System Management

Devices managed by NETCONF and perhaps other mechanisms have common properties that need to be configured and monitored automatically.

This test verified the YANG models used for system identification, time-of-day management, DNS resolver, and RADIUS configuration.

The test topology includes one controller and two routers. After the successful NETCONF session establishment, the controller synchronized the current running configuration from the routers and added the retrieved model parameters for the configuration of the NTP server, DNS resolver, RADIUS server.

The configured parameters were successfully pushed from the controller to the routers by committing the configuration. We verified the configuration successfully in the routers. As the last step, the controller managed to delete or rollback the configuration.

We faced YANG model compatibility issues. Some of the models had a slightly different structure from the models we used before. The controllers got the XML models, adapted them, and built the proper drivers to be sent to the routers again. The system identification part, like retrieving the software version of the routers, was not demonstrated in the test. The following Table shows the test combinations and the supported YANG models.

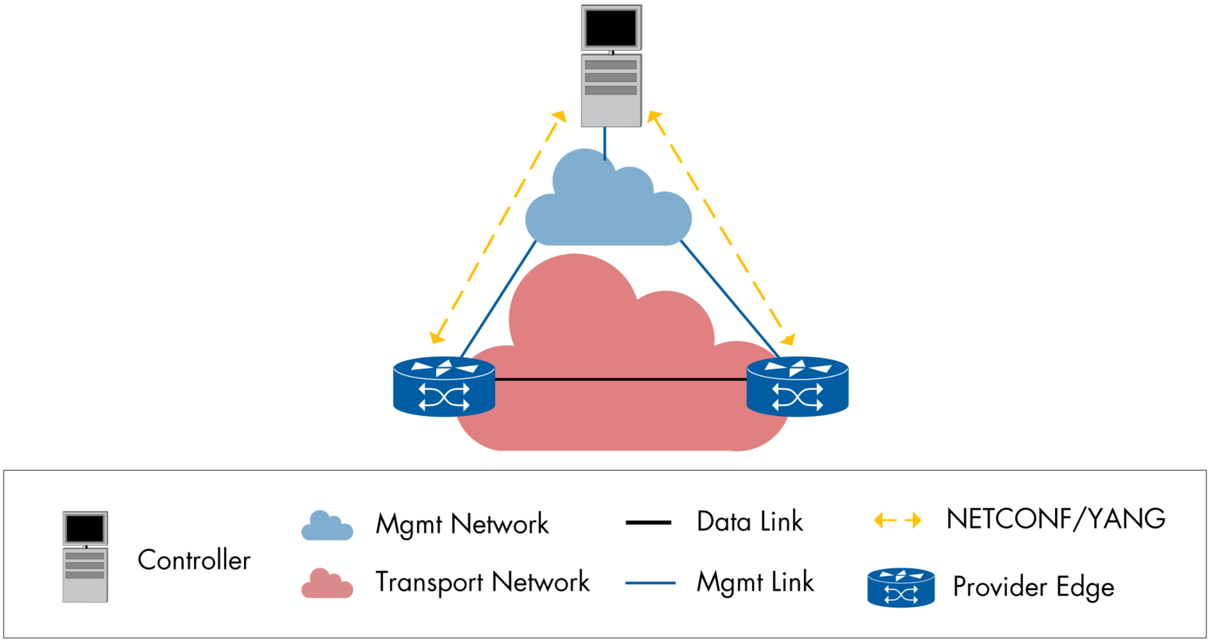

Figure 4: System Management - Test Setup

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Huawei Proprietary Model
Huawei iMaster NCE-IP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	Cisco Proprietary Model
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	IETF YANG Model IETF Interfaces

Table 4: System Management - Successful Combinations

Retrieve Interface Frame Sizes Distribution

From the network administrative perspective, the packet size distribution across the large network is an important factor. In this test, we verified the interface Ethernet frame sizes distribution statistics between two PE nodes by retrieving via the NETCONF client.

As shown in the topology below, one controller and two PE routers were set up for this test. The NETCONF session was established between the controller and the routers. Once the session was successfully established, the controller could retrieve the interface statistic for the frame distribution via the RPC log message.

The following packet sizes were displayed on the controller with the statistic of the number of Rx packets, the number of Tx packets, and the total number of packets.

- Pkts64Octets
- Pkts65to127Octets
- Pkts128to255Octets
- Pkts256to511Octets
- Pkts512to1023Octets
- Pkts1024to1518Octets
- Pkts1519to1548Octets

The statistic for the number of packets received on the interface can be retrieved via the YANG model. However, the good and bad packets were not differentiated on the statistic.

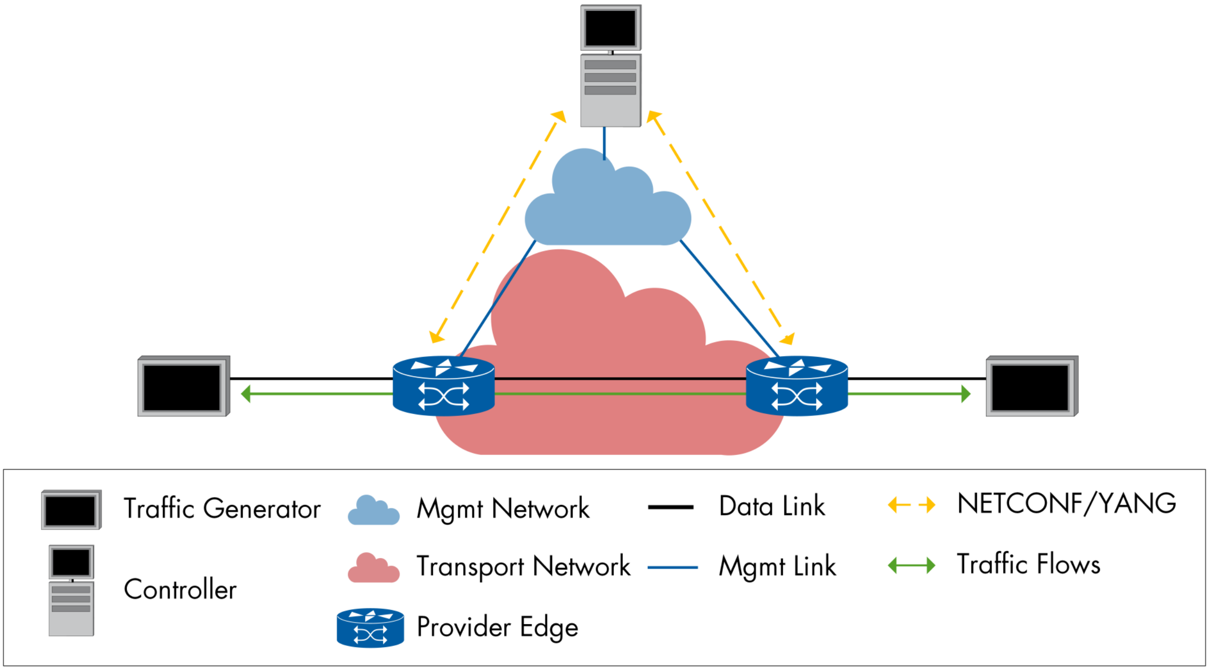

Figure 5: Retrieve Interface Frame Sizes Distribution - Test Setup

Controller	Router 1	Router 2	YANG Model
Huawei iMaster NCE-IP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	Cisco Proprietary Model
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model

Table 5: Retrieve Interface Frame Sizes Distribution - Successful Combinations

Multiprotocol Label Switching

The configuration of the services and the affected equipment is among the largest cost-drivers in provider networks. In that sense, the NETCONF and YANG model helps the service providers automatically configure the services like Multiprotocol Label Switching (MPLS). This test verified a YANG model that can configure and manage the Label Distribution Protocol (LSP) specific parameters for IGP-congruent LSPs.

For the test setup, NETCONF sessions were established between a controller and two PE routers. Before configuring the MPLS, the controller pushed the parameters to configure the Open Shortest Path First (OSPF) and the Border Gateway Protocol (BGP) in the routers using the native YANG model. After the successful establishment, the controller retrieved the current running configuration from the router and added three templates to add OSPF, BGP, MPLS, and RTP configuration parameters. For the OSPF configuration, the interface name and the router ID was added to the template. BGP configuration template was added with the parameters static AS number, Router ID, peer AS number, and peer router ID. As the MPLS-LDP configuration parameter, interface name, the router ID, hold time were configure. After pushing the three templates, we verified the configured parameters on the routers, and the traffic flow was successful between the PE routers. As the last step, the controller managed to delete or rollback the configuration. In the test, we observed the following issues. However, it did not have any impact on the tests since the issues were fixed quickly.

During one test combination, the PE Routers had to add static routes towards the traffic generator. The controller was able to add but not delete static routes, so we had to use IS-IS to work around the issue. In one case, the controller had a software bug during the test, which prevented pushing the configurations to one particular vendor's PE routers. The GUI of the controller had a compatibility issue with one web browser.

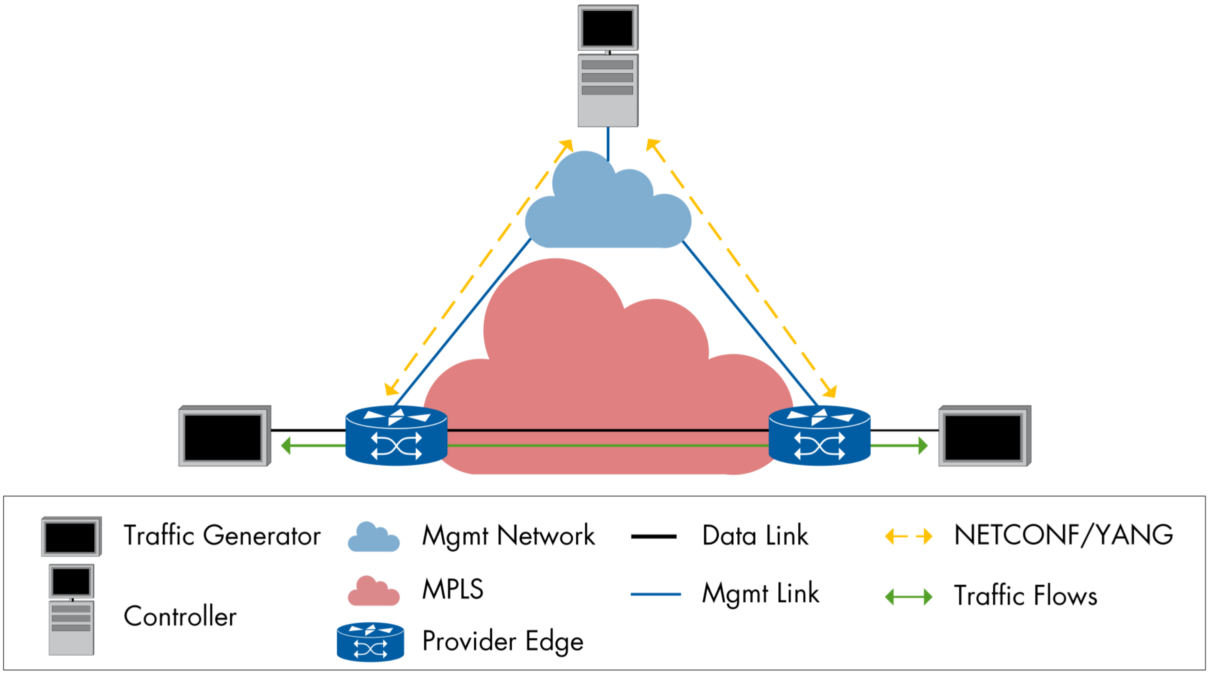

Figure 6: Multiprotocol Label Switching - Test Setup

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Huawei Proprietary Model
Huawei iMaster NCE-IP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	Cisco Proprietary Model
Nokia NSP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	Huawei Proprietary Model

Table 6: Multiprotocol Label Switching - Successful Combinations

MPLS with Segment Routing

NETCONF and YANG allow the service providers to easily configure the Segment Routing (SR) on top of the MPLS data plane. This test verified a YANG model that can configure and manage the SR on the MPLS data plane.

In this test, three routers were used to configure SR with the native YANG model. After the successful NETCONF session establishment, the controller retrieved the routers' running configuration and added the parameters like router ID, interface names, neighbor router IDs, ISIS name, ISIS net, and segment ID. After adding the parameters, the controller pushed the parameter with the XML payload to the routers.

We successfully verified the configured parameters on the routers. Traffic Generators were used to verify the connectivity and the transportation of the packets between the PE routers. EANTC also required packet traces to verify the MPLS tags were added to the packets. As the last step, the controller deleted or rolled back the configurations successfully from the routers. All test combinations were successful.

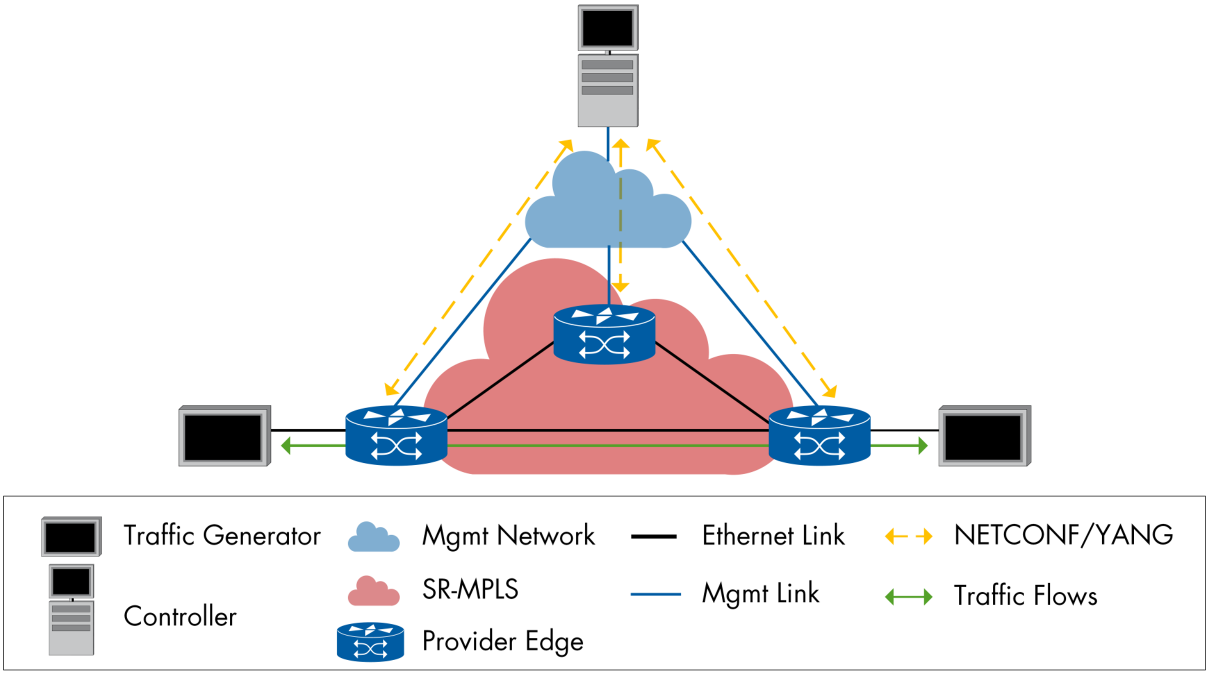

Figure 7: MPLS with Segment Routing - Test Setup

Controller	Router 1	Router 2	Router 3	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	-	Huawei Proprietary Model
Huawei iMaster NCE-IP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	-	Huawei Proprietary Model

Table 7: MPLS with Segment Routing - Successful Combinations

Layer 2 EVPN Service Creation

In a campus network with multiple buildings connected via WAN, the EVPN-MPLS Layer 2 VPN technology facilities equipment needs Layer 2 adjacency across buildings. NETCONF and YANG help, in this case, to configure the Layer 2 EVPN on top of the MPLS or SR-MPLS underlay network. This test verified that the controller could configure and manage the Layer 2 EVPN on top of the MPLS and SR-MPLS underlay network.

The test topology consists of one controller and two routers. After the successful NETCONF session establishment between the controller and the router, the controller retrieved the current running configuration from the router and added the parameters like device name, VLAN ID, IP address and subnet mask, service name, VRF name, auto-evi, AS number, and MPLS encapsulation for EVPN in the payload of the controller. After adding the parameter, the controller pushed the parameter successfully to the routers. We verified the parameters for the Layer 2 EVPN on the routers, and the EVPN service was up and running on both routers. The traffic flow was successful between the PE routers. The controller successfully deleted or rolled back the Layer 2 EVPN service from the routers as the last step.

In this test, a controller used a service YANG model in the northbound interface to configure the parameter. The configuration parameters are converted from the service model to the router's proprietary model to push the parameters into the routers.

During the test execution, we observed some minor issues. In one of the test combinations, the EVPN service was down on a PE router after the service creation. The issue was resolved by restarting the router.

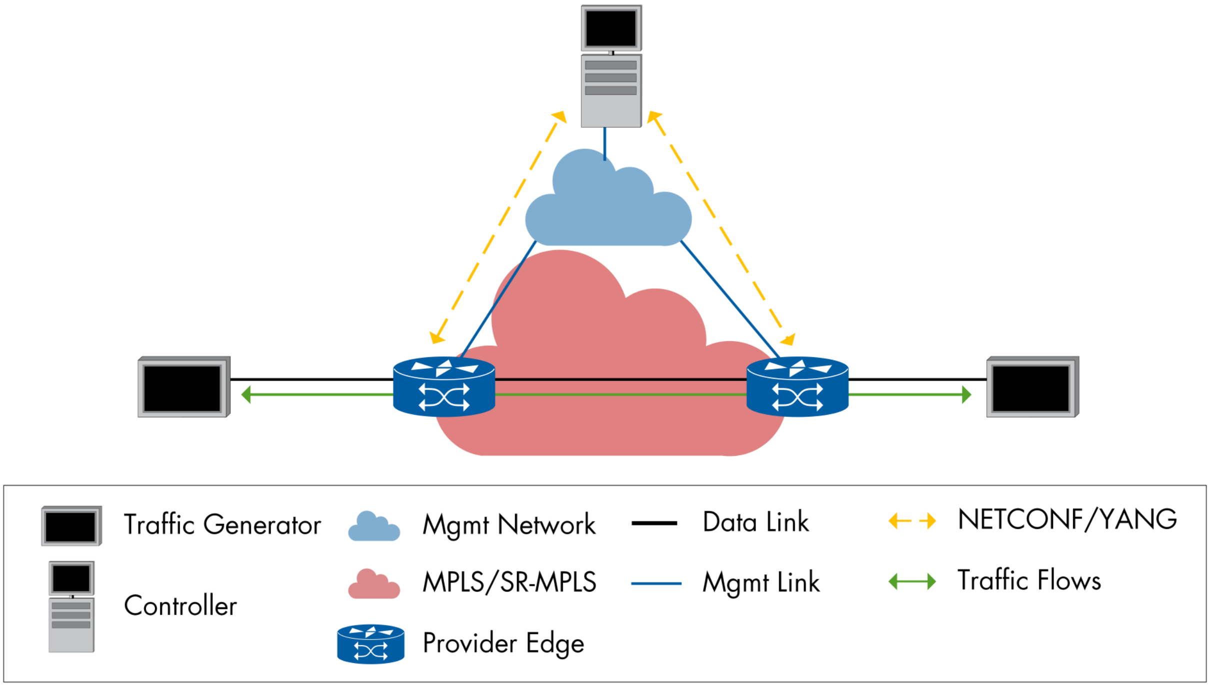

Figure 8: Layer 2 EVPN Service Creation - Test Setup

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Huawei Proprietary Model
Huawei iMaster NCE-IP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Northbound: Huawei proprietary model Southbound: Cisco proprietary model
Nokia NSP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	Huawei Proprietary Model

Table 8: Layer 2 EVPN Service Creation - Successful Combinations

Layer 3 VPN Service Creation

The Layer 3 VPN Service Creation (L3VPN) technology is used to route the VRF aware VPN packets over the service provider backbones like MPLS or SR-MPLS. NETCONF and YANG model helps the service providers to create and modify the L3VPN service simply in an automated way. This test case verified a YANG model that can configure and manage the L3VPN service on top of the MPLS or SR-MPLS.

For the test, NETCONF sessions were established between a controller and two PE routers. Before starting the test, the underlay configuration like IGP and BGP was configured via the previous test case controller. In this test, a native YANG model was used to configure the L3VPN service on the routers. After the successful NETCONF session establishment, the controller retrieved the running configuration from the routers (PE1 and PE2) and added the parameters like interface name, VLAN ID, IP address, VRF name, access ID, Router ID, AS number, device name, router target import, and router target export in the retrieved model. The controller successfully pushed the parameters to the routers. We confirmed that the added parameters belong to L3VPN on the router and the service status was up on the router. As the last step, the controller successfully deleted or rolled back the L3VPN service from the routers. All combinations of these tests were successful. Two controllers in the test used the RFC8299 as a service YANG model in the northbound interface to update the controller's data store with the L3VPN parameters. After the successful update, the controller converted the data model to the native YANG data model and successfully pushed it to the router via the southbound interface. The L3VPN configuration was based on IPv4 and IPv6 addresses.

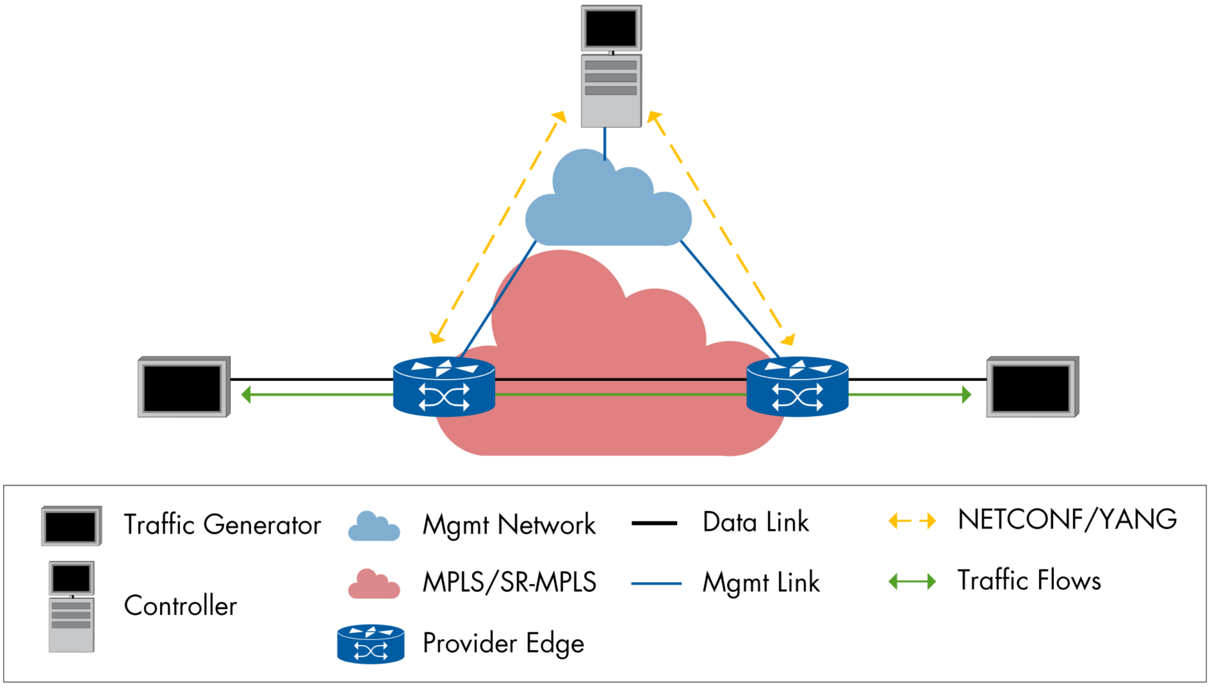

Figure 9: L3VPN Service Creation - Test Setup

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Northbound: IETF L3VPN RFC8299 Southbound: Huawei Proprietary Model
Huawei iMaster NCE-IP	Cisco N3K-C36180YC-R	Cisco N3K-C36180YC-R	Northbound: IETF L3VPN RFC8299 Southbound: Cisco Proprietary Model
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	Huawei Proprietary Model

Table 9: Layer 2 EVPN Service Creation - Successful Combinations

Bi-Directional Forwarding Detection

Bi-Directional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. Besides, it provides a consistent failure detection method for network administrators. In this test, we verified a YANG model to configure the BFD parameters on top of the underlay network.

The test topology consists of three routers and a controller. Before starting the test, the underlay configuration like ISIS and BGP was configured on the routers. As the first step, the controller successfully established the NETCONF session with routers and then retrieved the routers' running configuration.

The controller added the parameters like router name, port ID, peer IP addresses, minimum Rx interval, minimum Tx interval, ISIS name, and BFD detection time in the retrieved model. After adding the parameter, the controller successfully pushed the parameter to the routers. We verified the parameters for the BFD configuration on the routers. As the last step, the controller successfully deleted or rolled back the BFD configuration from the routers. All combinations of these tests were successful. In this test, a controller used a service YANG model in the northbound interface to configure the BFD parameter. The configuration parameters are converted from the service model to the router's proprietary model to push the parameters into the routers.

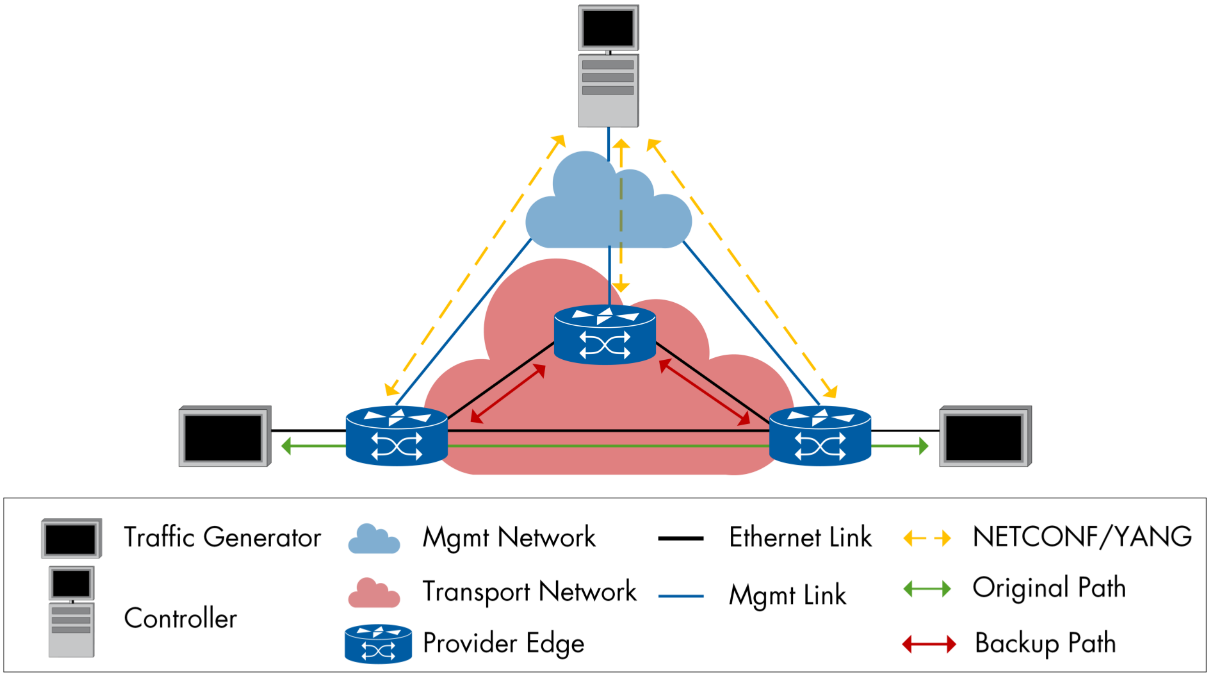

Figure 10: Bi-Directional Forwarding Detection - Test Setup

Controller	Router 1	Router 2	Router 3	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Huawei NetEngine 8000 M14	Huawei Proprietary Model
Huawei iMaster NCE-IP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Northbound: Huawei Proprietary Model Southbound: Cisco Proprietary Model
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	Huawei NetEngine 8000 X4	Huawei Proprietary Model

Table 10: Bi-Directional Forwarding Detection - Successful Combinations

Telemetry Streaming

The Streaming telemetry is the data model to facilitate operational data monitoring with higher efficiency. The network devices work in push mode to send the network operation data to the collector, instead of pull mode comparing with SNMP or CLI. The OpenConfig data models work with different transport protocols, such as NETCONF/RESTCONF and gRPC Network Management Interface (gNMI).

This test verified the Streaming telemetry for interface monitoring using OpenConfig data models with the transport protocol gNMI. The gNMI defines a particular set of gRPC operations such as Capability Request, Get Request, Set Request, and Subscribe Request.

The test topology consists of one controller, one collector, and two PE routers. After the successful NETCONF session establishment between the controller and the routers, the controller retrieved the routers' running configuration and added the parameters to configure the gRPC server, port number, and sensor groups. The gRPC service was configured between the collector and the PE nodes. After the successful connection, the collector requested the subscription to the sensor paths with stream mode as sample or on-demand. In the sample mode, the sample interval was also defined during the subscription. The statistics were successfully streamed from routers to the collector via gNMI. In this test, collectors retrieved the interface statistics, and one of the test combinations displayed the physical inventory of the router via the YANG model. In one of the combinations, a controller supported the gNMI to collect the statistics from the routers. All participating test combinations successfully performed the test.

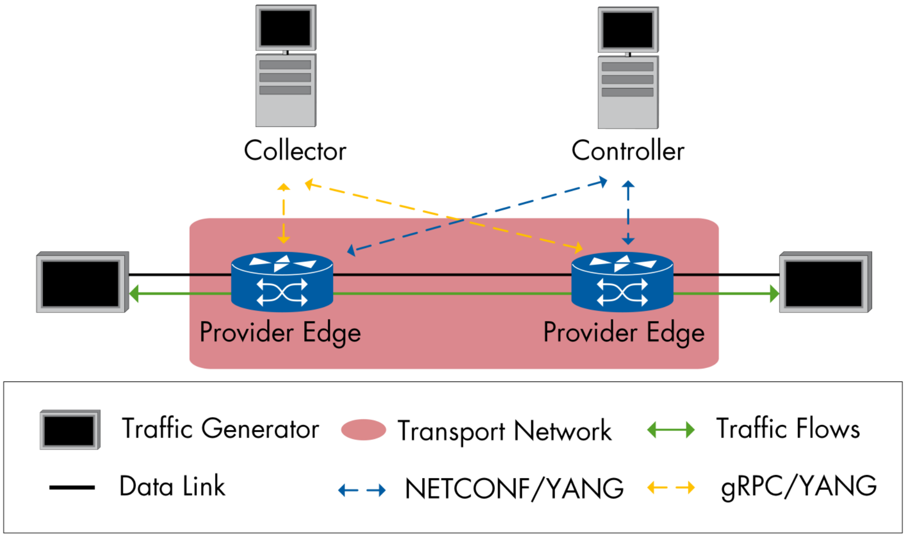

Figure 11: Telemetry Streaming - Test Setup

Controller	Router 1	Router 2	YANG Model
Cisco NSO	Huawei NetEngine 8000 M8	Huawei NetEngine 8000 X4	Dial-In: Apply config: OpenConfig gRPC: Huawei Native
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Dial-In: OpenConfig model
Nokia NSP	Huawei NetEngine 8000 M14	Huawei NetEngine 8000 X4	Dial-Out: OpenConfig model Dial-In: Huawei Native Model

Table 11: Telemetry Streaming - Successful Combinations

Precision Time Protocol

Time synchronization and transferring through the packets network are essential for the industry as the deficient latency applications come to play; managing and monitoring the time and frequency transfer nodes are crucial for the whole network and its services.

This test verified that the controller could set up and read the routers' PTP parameters, which helps simple management and monitoring functions.

At the beginning of the test, the connectivity between the routers was granted by the previous tests. NETCONF sessions were initiated from the controller towards the routers. After the successful session establishment, the controller retrieved the running configuration from the routers and added the parameters in the retrieved model to configure the master and slave clock with the parameters like source IP, interface ID, clock role with master or dynamic, transport mode with multicast or unicast, PTP domain ID.

The controller successfully pushed the configured parameter into the routers. Two routers were used for the test. After the configuration, one router acted as the master clock and the second router acted as the boundary clock. We verified the configured parameters on both routers. In the boundary clock, the PTP clock status was locked, and also it displayed the master clock ID. As the last step, the controller successfully rolled back the PTP configuration from the routers.

One controller deleted the parameters partially since some of the parameters like domain ID and interface ID are not possible to be deleted on the routers. The controller set the default value or zero for non-deletable parameters.

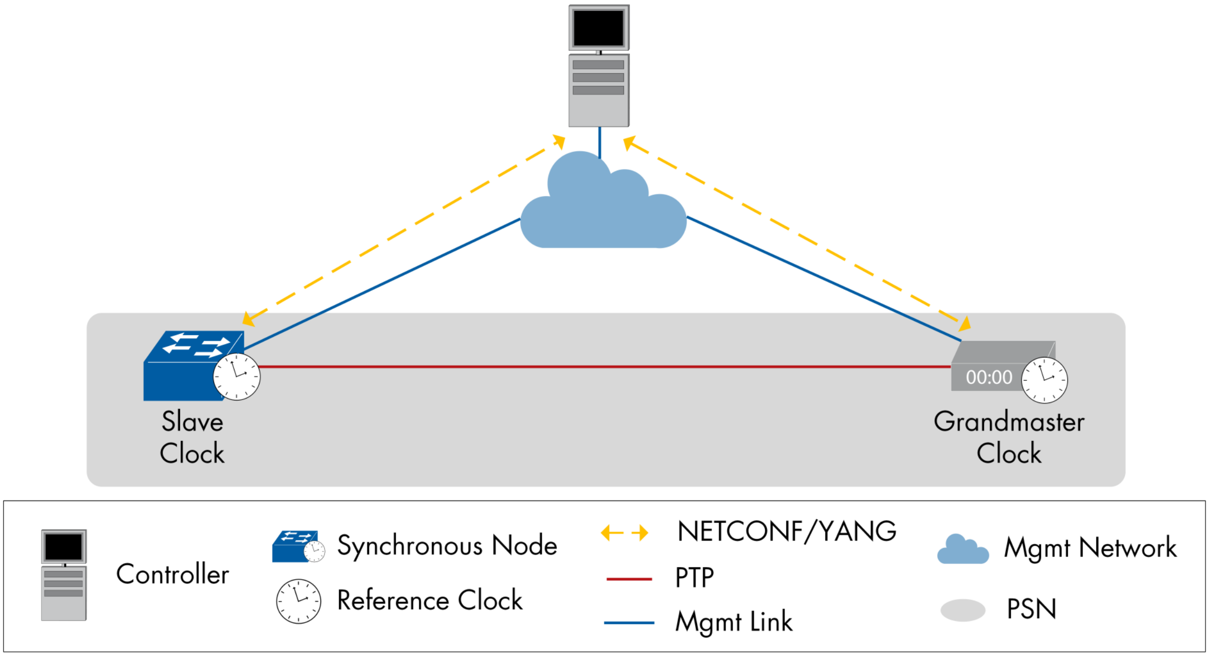

Figure 12: Precision Time Protocol - Test Setup

Controller	Router 1	Router 2	YANG Model
Huawei iMaster NCE-IP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model
Nokia NSP	Cisco N9K-C9316D-GX	Cisco N9K-C9316D-GX	Cisco Proprietary Model

Table 12: Precision Time Protocol - Successful Combinations

YANG Models

In the test, we covered the following YANG models.

Vendor	YANG models
Cisco	Cisco-NX-OS-device.yang
	openconfig-interfaces.yang
Huawei	huawei-mpls.yang
	huawei-network-instance.yang
	huawei-ifm.yang
	ietf-ip.yang
	openconfig-interfaces.yang

Conclusion

In one week of in-depth testing, we successfully performed a wide range of interoperability tests with IETF, OpenConfig, and vendor proprietary YANG models completing a total of 39 multi-vendor test combinations. In general, the NETCONF clients supported IETF models and OpenConfig models. Some adaption was needed to fully support the proprietary YANG models of the router vendors. In one case, we were able to use northbound models. Several interoperability issues occurred between the NETCONF clients and servers as expected in events like this. Three vendor teams solved over 95% of the issues; the other issues need further troubleshooting.

In the Device Functions and Configurations area, we successfully verified configuration and monitor interface-specific parameters, common system properties on a network device, MPLS, MPLS with Segment Routing based parameters, Management of Interfaces using OpenConfig, and Retrieve Interface Frame Sizes Distribution.

The Service Provisioning test covered Layer 2 EVPN service creation, L3VPN service creation, and Bidirectional Forwarding Detection. Layer 2 EVPN and L3VPN service creation were demonstrated with MPLS and SR-MPLS data plane.

In the Telemetry streaming using OpenConfig, the controllers successfully pushed the YANG model to the routers, and the collectors successfully subscribed to the sensor groups in the routers to collect the interface statistic. The controllers were able to push the configuration of the PTP master and slave ports and successfully locked PTP between the DUTs.

While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein. All brand names and logos mentioned here are registered trademarks of their respective companies in the United States and other countries.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, https://www.eantc.de/
[v1.1 20210126]