



EANTC Independent Test Report

Juniper Contrail SD-WAN Solution

September 2018



Introduction

Since its inception just a few years ago, SD-WAN has become a go-to solution for Service Providers and Enterprises that require an affordable and flexible Wide Area Network (WAN) connectivity solution. Additionally, Enterprises and Service Providers are looking for SD-WAN solutions to help them simplify and automate the WAN. Key capabilities of SD-WAN that they desire include Zero-Touch-Provisioning, centralized management, automated and simplified through the use of software defined network (SDN) technology and graphical user interface (GUI) tools, integrated security, application policy based routing and advanced monitoring, analytics and reporting tools, to name a few.

At EANTC, we are well aware that there is no single, uniform SD-WAN standard. The number of different types of SD-WAN solutions in the market continues to increase as more vendors enter the market, each bringing their own strengths and capabilities based on their area of focus as a network equipment vendor: routing and switching, cloud software, WAN acceleration, and security.

Test Highlights

- ➔ Zero Touch Provisioning supported for sites using CPE/uCPE via GUI or REST API
- ➔ CSO managed 10,021 simulated SD-WAN spoke sites
- ➔ Failover implemented for active-passive dual Hub
- ➔ Advanced UTM security features, content and web filtering, verified
- ➔ Application Forwarding verified for non-SLA based application flows
- ➔ Verified failover mechanisms triggered by application-level SLA violations along with APBR and AppQoS capabilities
- ➔ Application visibility and SLA performance monitoring performed as expected

Report Overview

Juniper Networks® commissioned EANTC to conduct an independent review of its SD-WAN solution, Contrail SD-WAN, in order to help the industry better understand the breadth and depth of its capabilities. Contrail SD-WAN is comprised of an SD-WAN controller and customer premise equipment (CPE) endpoint devices. Juniper uses its Contrail Service

Orchestration (CSO) product as its SD-WAN controller. It leverages its broad portfolio of CPE devices included its Universal CPE platform, the NFX Series Network Services Platform, SRX Series Services Gateway, and the vSRX, a virtual machine version of the SRX. EANTC conducted the tests at the Juniper Networks headquarters facility in Sunnyvale, California in August 2018.

Juniper requested that we verify a specific set of test cases, which focused on the breadth and depth of its SD-WAN solution features and capabilities. Therefore, our tests showcase the combination of Juniper's vast experience as one of the leading routing and security vendors and their long-standing investment in telco cloud infrastructure and management software.

Specifically, Juniper asked us to confirm the following aspects of its SD-WAN solution: application-based traffic steering, which was tested by forcing an SLA violation, Zero Touch Provisioning using both the GUI and API to provision Juniper's two different CPE offerings: SRX Series and NFX Series Universal CPE devices, High Availability and resiliency, multi-tenancy, network segmentation, RBAC, highly-customizable traffic and log monitoring, which allows you to present network data for management-level reports or to troubleshoot a system issue. Also, site scalability of the controller was put to the test along with many additional features.

The EANTC test cases were all successfully completed and verified, confirming Juniper's claims that their Contrail SD-WAN solution is a secure, scalable, and customizable solution with many add-on features, and it is ready to take on today's and future Enterprise and Service Provider requirements.

Hardware and Software

Hardware Type	Software Version
NFX250	15.1X53-D491.1
SRX340	15.1X49-D144
SRX4100	15.1x49-D144.1
vSRX	15.1X49-D144.1
Contrail Service Orchestration (CSO)	4.0.1

Juniper Networks Contrail SD-WAN can be configured in both a Hub and Spoke and Full Mesh architecture. In the Hub and Spoke architecture, MX Series routers, the SRX4100, and SRX1500 can be used as hub gateways. In the case of our test the SRX4100 was exclusively used for the Hub sites.




Figure 1: Physical Test Bed

The SRX4100 was chosen over the SRX1500 because multi-tenancy requires higher capacity. The CPE devices we used included the NFX250, which is one of Juniper Networks Universal CPE families, and the SRX340. As a Universal CPE device the NFX Series has the ability to load VNF images. For our test cases, we used the vSRX as the VNF. For the tests, we generated traffic that simulated a number of different applications. For each application we assigned a percentage as to how much of the overall traffic its data would represent, along with its flow size. Each application was then mapped to a predefined traffic type profile (QoS). This data can be found in Table 1.

App	%	Bytes/Flow	Type
Skype Call	20	645	Voice-Video
SMTP	20	369267	Premium Internet
Youtube	13	2059712	High Priority Video
Bittorrent	15	520848	Internet
Facebook	20	394024	Internet
HTTP Enterprise	12	10398	Hosted AV

Table 1: Traffic Profile Specification

Test Bed Creation

The testing started with a physical topology as shown in Figure 1. This scenario represents a small SD-WAN customer. We proceeded from scratch, installing Contrail Service Orchestration (CSO) and then setting it to support multiple tenants.

CSO Installation

To manage this SD-WAN test bed, CSO, the Contrail SD-WAN Controller, was installed on a bare metal x86_64 server running Ubuntu 14.04.5 LTS.

First, the CSO downloader and Installer Applications was retrieved from Juniper Network's website and installed on a MAC OS laptop. Information about the server and the hypervisor that would host CSO was provided. Hypervisor options include KVM or ESXi. In our case Juniper selected KVM (Kernel-based Virtual Machine) version 1.2.2 as the hypervisor. The downloader then checked to confirm that the minimum requirements for the installer were met. A dedicated Virtual Machine (VM) was created for the installation and CSO package was automatically downloaded.

Once the download process is complete, the user is automatically redirected to the CSO installer page, where they log into the CSO Installer VM. The user then chooses the deployment size (small, medium, large), which relates to the desired scale of the SD-WAN deployment and defines the resource required to install the solution. The user then provides some

additional, basic provisioning information, such as whether or not CSO will be reachable directly or sit behind NAT, the IP address range for the VMs and the NTP server details. The VMs included:

- Infrastructure components (central + regional)
- Microservices (central + regional)
- Virtual Route Reflector (vRR)
- Load Balancer




Figure 2: CSO Automatic Installation in Progress

Figure 2 shows the installation in progress.

CSO leverages a distributed architecture, which means that it has the option to separate infrastructure and micro-services into various regional components and one central component. In the lab, we used a simplified architecture which contained both in the same physical server. After CSO finished its automatic installation, we were able to login to the Administration Portal and see an empty dashboard with its getting started steps, as shown in Figure 3. This dashboard was later observed to be customizable through the addition of modular widgets.




Figure 3: CSO Administration Portal With Customizable Dashboard

CSO can also be installed in a High Availability mode in AWS. Juniper showed us their site onboarding workflows for this deployment option. However, we did not perform the actual site onboarding. Juniper Networks also showed us an active deployment of CSO 4.0.1 in the AWS cloud. For the rest of this

report we used our CSO 4.0.1 instance, which we installed on a bare-metal server.

Service Provider POP Onboard and Activation

In this step we focused on the onboarding and activation for a hub device in a Service Provider POP. Juniper Networks Contrail SD-WAN supports a wide variety of their MX Series and SRX Series devices as SD-WAN gateways. We used a SRX4100 as our hub device. The MX Series was not used in our testing.

A POP was created. Once the POP was identified a hub was created using with the device template "SRX as SDWAN Hub" and the necessary hub parameters were enabled, including number of WAN links, their types, and the IP address for each link type. The device serial number and passcode were used for activation. Once the hub provisioning was completed, CSO automatically changed the status or state of the hub from detected to provisioned, which provided verification that the task had been completed. The forwarding plane was tested in other sections of this report. Juniper indicated the device management and control plane are secured using an IPSec tunnel.

Tenant Creation, Multi Tenancy and Network Segmentation

In this section, we added a tenant. The required information for a tenant includes the user information and topology deployment type, which in this case was full-mesh SD-WAN. For the purposes of our testing, we validated multiple deployment types, including hub & spoke and full-mesh. This was done by adding additional tenants, some of which were configured for hub & spoke and others of which were configured for full-mesh.

CSO provides an administration portal for the managed service provider, which also acts as a customer portal for each tenant. It is important to note that CSO provides one portal, which provides access to all end users based upon their login credentials. The segmentation or restriction of access per user is managed by the RBAC capabilities of CSO. The administration portal and its RBAC capabilities can also be used by Enterprise IT network managers to provide access to various department heads and IT engineers in their organization.

Contrail SD-WAN supports multi-tenancy in the hub; spoke sites are customer specific. We observed a shared hub with two remote sites, each site was owned by a different customer, running separate application traffic simultaneously. The tenant traffic segmentation was achieved by using various VRF instances, along with different BGP communities and VRF import/export policies, which was automated by CSO as part of the provisioning process. To add

another level of route and traffic separation, different departments (e.g. Retail, Finance, guest Wi-Fi, etc.) can be configured for the customer. During our testing, network segmentation, when activated under tenant properties, created a separate VRF per department along with a default VPN for all unallocated LAN interfaces. This allows further traffic segmentation, increasing security and allowing all other benefits from VRF, like using the same IP LAN space.

SD-WAN Zero-Touch-Provisioning

Cost and complexity are expected when deploying a large number of devices across many physical sites. In order to reduce this issues, Zero-Touch Provisioning (ZTP) was introduced.

Initially, the test bed topology only had the required physical connectivity configured. All the CPE devices had their original stock configurations in order to demonstrate that the activation performed was truly zero-touch (ZT). In this section we went through the following steps sequentially:

- Zero-Touch-Provisioning for CPE devices (NFX and SRX) via CSO GUI
- Zero-Touch-Provisioning for Dual CPE with Multi-homing via CSO GUI
- Zero-Touch-Provisioning via REST API

We also configured the following security parameters: SSL Settings, VPN Authentication, and Overlay Tunnel (GREoIPSec) Encryption (default is AES-256-GCM).




Figure 4: Zero Touch Activation Flow Steps

NFX/SRX as SD-WAN CPE Onboard and ZT Activation

The stock configuration of Juniper Networks CPE devices contain a reference to a Juniper managed redirect server, also-known-as a "phone-home server." Prior to initial boot up and configuration the device serial number for each unique CPE device is mapped

to the desired instance or cluster of CSO SD-WAN controllers that will manage the devices. This mapping was created and stored in the redirect server, and a certificate was used for CSO authentication.

Site creation and configuration were performed in the CSO via its GUI, there is also the option to use REST API which is discussed later. This step requires information like site name, address and region, template (e.g. NFX as SD-WAN CPE), and the amount and type of WAN links desired and LAN segments information. WAN links can be MPLS or Internet, and each CPE can have up to a maximum of 4.

While not tested, CSO also automates license activation for the CPE devices it supports as part of the provisioning process. In our case, licensed activation was performed via CSO after the device was provisioned.

The NFX Series is a Universal CPE device, as such it has the ability to host and run multiple VNFs at the same time. In our test case, we only tested two VNFs at the same time. For the VNFs to be loaded on the NFX, they first must be uploaded into CSO, so that CSO can then deploy the VNFs as part of the NFX provisioning process.

Application signature package was also performed to be able to monitor traffic to the application level. This license and application signature management is performed by CSO. Detailed steps to achieve Zero-Touch-Activation are shown in Figure 4.

After site creation and configuration, activation was performed by selecting the devices, clicking activate and entering the activation code.

In this section, only the dual-CPE with multi-homing is described, the other two are presented in the following sections. The details of their activation are not discussed in this report, but the process was performed successfully as part of the tests.

Dual CPE with Multi-Homing

Two NFX250 devices were used as dual CPE in an active-active fashion. Both had initial stock configurations. Hub redundancy was active-passive. The steps required to onboard, configure and activate the site were virtually the same as in the prior sections, with the difference being the used template, i.e. Dual NFX250 as SD-WAN CPEs, and the configuration for primary/secondary hub selection. The topology is shown in Figure 5. We were able to certify that both CPEs were forwarding traffic and only one of the two hubs was active. Fail-over functionality results are described in another section.

Zero Touch via REST API

The topology shown in the previous section was selected to be activated via Zero-Touch Provisioning (ZTP). Devices had their respective stock configurations. For site onboarding/activation via API the following steps must be followed:

- Generate a token in the CSO based on the tenant login credentials, to be used for authentication
- Site creation
- Site configuration
- Site activation

Secure API calls were used to onboard and activate the device via POSTMAN. Each step required one single API call, only the site activation required two (one for each CPE). After finishing the steps, the devices were shown activated in CSO in the same way as before via the GUI.




Figure 5: Dual-CPE Multi-Homing Topology

CSO Site Scalability

SD-WAN and scalability go hand-in-hand as enterprises need a simpler way to manage and optimize WAN services for their high number of campus and branch locations. For service providers, looking to deliver SD-WAN as a Service, they need to be able to deliver this scale for their Enterprise customers. In this test case we focused on the scalability of Contrail SD-WAN, more specifically we focused on the number of sites that could be provisioned and managed by Contrail SD-WAN.

The test bed used in this case was different from the one specified for the other scenarios. It consisted of the Contrail SD-WAN controller and a Juniper in-house simulator. The simulator was created to emulate vSRX devices for both hub and spoke devices by responding to the controller with data captured from the device.

The CSO infrastructure was built using 7 physical servers (each with 48 vCPUs and 256 GB Memory) running a total of 36 VMs:

- 1 Installer
- 6 Virtual Route Reflectors (vRR)
- 3 Contrail Analytics
- 2 Southbound load balancer
- 6 Infrastructure (3 central and 3 regional)
- 6 Micro service (3 central and 3 regional)
- 6 Load balancer (3 central and 3 regional)
- 6 Elk stack (3 central and 3 regional)

The simulated connectivity between vSRX devices was hub and spoke. Each tenant was assigned 5 hubs, and each of the hubs were assigned a broad range of endpoint CPE devices, between 50-500, to best represent a real deployment scenario. Juniper Networks had configured this test in advance, including the CSO infrastructure and simulator. When we started our test, CSO had already provisioned 9780 sites.

The goal was to add 250 additional spokes distributed across the 5 hubs, all under 1 tenant, in order to surpass the goal of 10000. For this, the simulator was configured to add 30 concurrent at a time. This number was chosen by Juniper, because it is the maximum rate at which its simulator can effectively add sites. After creating the tenant and the hubs, the spoke addition started, each following the next steps:

1. Create
2. Configure
3. Activate

All controller servers and VMs were added to a new and dedicated Nagios, which is an open-source systems, network, and infrastructure monitoring tool. It was used to record the CPU/RAM utilization before and after the CPE addition phase and to observe the values during it.

While spokes were being added, the highest utilization in both CPU and RAM was observed in a few of all servers and VMs. It was just under 50% while the sites were added. Before and after this test, the CPU was down to 10% or lower. An extra spoke was created in the end to record the messages between CSO and CPEs during all stages. After the provision, each spoke had an open SSH connection with the controller monitoring aliveness, but without GREoIPSec overlay for the communication, as would be the normal case with non-simulated CPEs.

The final result showed 10,021 CPEs provisioned in CSO resources under tenant devices. A snippet is shown in Figure 6.




Figure 6: 10021 Provisioned CPEs in CSO

Application-based Traffic Steering

Traffic engineering is essential for cost-effective and reliable network operation. This section focuses on the aspects involved in making this work in Juniper Networks Contrail SD-WAN solution, plus features related to path selection. As an overview, the features described in this section are:

- Non-SLA Based Application Forwarding Capabilities
- SLA Violation Automatic Detection
- SLA Based Application Traffic Management

Some of the features interact with each other or rely on one another. The topology used is shown in Figure 8 under subsection "SLA Based Application Traffic Management - APBR enabled".

Non-SLA Based Application Forwarding Capabilities

Each CPE can have up to four WAN links configured. Each of these can have one or more non-exclusive modes, including:

- Default Forwarding
- Local Internet Breakout (LBO)
- Backup Link

The default forwarding option for Contrail SD-WAN is to load balance non-SLA based application traffic flows across all available WAN links. The customer administrator has the ability to assign one or more of the WAN links as preferred. If this is the case, then the

non-SLA application flows will be forwarded across the preferred links. We verified the usage of default and preferred link forwarding options. In one test case we verified the default forwarding option and saw load balancing across all WAN links. In another test case we saw the one configured preferred link, with all non-SLA based traffic flowing across it.

In the case of the Local Internet Breakout (LBO), a LBO link is identified, which allows all internet application traffic to be directly routed onto the internet service. We observed internet traffic from a site being routed directly onto the internet service, while testing this use case.

The backup link functionality allows a WAN link to stay without traffic until all other WAN links connected to the site are out of service. We were able to observe a WAN link configured as a backup without any traffic until all the other WAN links were unavailable. At this point, we saw all traffic switch onto the backup link. Additionally, once the other WAN links became available, the traffic switched back to those links and the backup link went back to being unused.

SLA Based Application Traffic Management - AppQoE enabled

The Service Level Agreement (SLA) fulfillment allows applications to run smoothly and businesses to achieve their goals. Juniper Networks' Contrail SD-WAN supports SLA profiles at the application level via SD-WAN policies. The SLA profile describes the thresholds and the SD-WAN policy enforces this to a specific application or group of applications. Contrail SD-WAN provides Application Quality of Experience (AppQoE) and Advanced Policy-Based Routing (APBR) support for SLA based routing, this section will focus on AppQoE. APBR will be covered in the next section.

For AppQoE, the SD-WAN solution accurately measures the application SLA across multiple WAN links. The application traffic is then dynamically mapped to a path among the available WAN interfaces that best meets the application SLA requirement. In order to measure the application SLA performance, the AppQoE service measures the application's packet-loss, RTT/Latency and Jitter to score the application's performance and thereby select the best possible link for that application. In addition to measuring the application SLA on an active link (also referred to as passive probing), there is a need to measure the SLA of these applications by sending active (synthetic) probes on all the available links. These real-time metrics are used to score an application over a particular link and make decisions with respect to which WAN path should be selected for the particular application.

Two SLA profiles were configured to be used for all tests, and are shown in Table 2.

Parameter	P2P SLA Profile	HPVideo SLA Profile
Application	Bittorrent	Youtube
Traffic type (Qos) profile	Internet	High priority Video
RTT	300ms	250ms
Session sampling	100%	100%
SLA violation counts	5	5
Sampling period	10	10

Table 2: Configured SLA Profiles for Applications

Contrail SD-WAN SLA profiles are used by SD-WAN policy intents for traffic management. SD-WAN policies help optimized utilization of WAN links and the efficient distribution of traffic. Every tenant has an SD-WAN policy and intents are then created in the SD-WAN policy. Policy intents include the following parameters: Source, Destination, SLA profile, and Intent name. These policies can be assigned to specific sites or specific departments within an organization, which can then be automatically mapped. The AppQoE and APBR features both use intent to identify the best possible path for the application.

Two policies were configured, one for Bittorrent and one for Youtube, with their respective SLA profiles.




Figure 7: Application Link Switch on SLA Violation

In the tests we saw the application taking the link assigned by the application SLA policy, which was verified in the CSO SD-WAN analytics information. Then an SLA policy violation was introduced for the application. We then were effectively able to see Bittorrent traffic switch to a different WAN link once

it's current one had a latency greater than 300ms, same case with Youtube and a latency of over 250ms. Latency was introduced using the Ubuntu 14.04.5 TLS as impairment tools connected inline. In Figure 7 we are able to see how the Application SLA Performance Monitoring from CSO displayed the RTT introduced for the Bittorrent case.

SLA Based Application Traffic Management - APBR enabled

As mentioned above, this test verified the Advanced Policy Based Routing capabilities within Contrail SD-WAN. When using the APBR capability, a specified routing instance is set for the application's traffic as the first path. All subsequent packets of a session use that same routing path. In order to ensure routing happens dynamically, the Controller runs RPM based probes to evaluate the quality of other possible links and reconfigure the route for the particular application in APBR if the Service Level Agreement (SLA) of a particular link is not met. This test verifies the Link Affinity capabilities of Contrail SD-WAN.

We configured the preferred WAN type for specific applications with select constraints. In regards to WAN type selection for application traffic, the default order of preference is the following:

1. WAN type preference
2. SLA Link Metric (RTT, loss, jitter, throughput) which may match any condition or all
3. Bandwidth and cost (lower is preferred)

An additional feature is the Failover/Fallback. The failover part refers to being able to switch the traffic, when an application SLA is no longer being met, to a different WAN link than the configured preferred path. The fallback allows the traffic to return to the preferred WAN link when it meets the preferred SLA parameters again.

The application traffic of interest for this case is Youtube. The preferred path is MPLS. Failover/Fallback was activated. Metrics were added to Figure 8 only for this scenario, values were the same for both CPEs and are below as Cost/Subscriber-Bandwidth:

- WAN#0: 1000/1000
- WAN#1: 800/800
- WAN#2: 600/600
- WAN#3: 400/400

Performed steps for this test were the following:

- Youtube application was initially flowing through WAN1. Reason: MPLS, SLA link metric and lower cost.
- SLA violation was forced on WAN1 by increasing latency to 260ms. This caused Youtube traffic to switch to WAN0. Reason: MPLS Link SLA metric was violated on WAN1




Figure 8: Logical Topology Used for Application-based Traffic Steering Scenarios

- SLA violation was caused on WAN0 by increasing latency to 260ms. This caused Youtube traffic to switch to WAN3. Reason: WAN3 was the next best available link from a SLA policy and cost metric perspective. Contrail SD-WAN chose the internet path because Failover/Fallback was activated. If Failover/Fallback was not activated, traffic would not switch to Internet
- SLA was restored on WAN1 by removing impairment. Due to Failover/Fallback being activated, the traffic switched back from WAN3 Internet link to WAN1 MPLS link. This last step is shown in Figure 9.




Figure 9: Traffic Fallback Event Specification

High Availability and Resiliency

In this section, the focus was on the high availability, redundancy and resiliency capabilities of Contrail-SD-WAN. The topology used is shown in Figure 5 in the “Dual CPE with Multi-homing” provisioning section. This scenario allowed us to recreate a failure in the primary hub or in one of the cluster spokes. Initially, traffic was distributed between the two spoke devices and all of it was flowing through the Primary Hub (Hub#1).

Traffic used for this test was a combination of HTTP and DNS and it was not steered to any particular WAN link. In order to simulate the failure, the devices were physically shut down.

Primary Hub Failure

The secondary hub had no traffic flow sessions. After a primary hub failure, we observed application traffic failover to the secondary hub. The traffic graph for DNS and HTTP is shown in Figure 10. We measured failover times for both TCP (HTTP) and UDP (DNS) based applications. The failover time for TCP applications was 3 seconds and UDP application was for 18 seconds, respectively. It is important to note that failover times will vary based on application type and rate.




Figure 10: Primary Hub Failure HTTP Throughput and DNS Total Successful Queries

After we certified that the traffic was switched to the secondary hub, the primary hub was switched on again. As our original configuration was active-passive for the hubs. After the primary hub became fully functional, traffic was observed to fallback to it and the secondary became passive again, and was left again with zero sessions. The fallback time was sub-seconds for both applications. A traffic graph for the failover event can also be observed, from CSO perspective, in Figure 11.



Figure 11: CSO Site Management Traffic Graph for Primary Hub Failure

CPE link failure was tested and verified in the “Application-based Traffic Steering” section of the report.

In all types of failures/recoveries, alarms and application link switch events were observed. Alarm logs indicated overlay tunnels up/down events, cluster switch events, and WAN link status down/up. Application link switch events displayed information about a specific application:

- SLA violation time and profile
- Link switch time
- Spoke site and Hub
- Past/Current overlay tunnel
- Department

Monitoring & Analytics

We were able to observe various aspects of Juniper Contrail SD-WAN monitoring features and customizations. A few are subtly mentioned in other sections and the rest are described next, among them are:

- Device events
- SD-WAN and Security Report Creation
- Infrastructure monitoring
- Application SLA Performance Monitoring
- Application Visibility
- Application Link Switch Events (end of section “High Availability and Resiliency”)
- Alarms & Alerts (end of section “High Availability and Resiliency”)

Device Events

Application session monitoring tracks the following data from each flow and allows for a summary or detailed view: application, hostname, service name, nested application, source and destination zone, Protocol ID, reason, and NAT session information.

Security events allow filtering device events into different categories including firewall, web filtering, content filtering, antispam, antivirus, IPS, and IPSec VPNs.

Report Creation

Contrail SD-WAN provides a library of predefined templates, and customization is also possible. The end user, in this case, could be the tenant administrator/Enterprise network administrator or the service provider who can collect reports for their tenant. The SP administrator can collect reports for any tenants. Once identified or created the desired data/reports they would like to receive can be e-mailed for distribution across the organization. These reports can be created at the tenant level or per site.

We observed the creation of a new customized security report that was sent as a .pdf to an e-mail recipient including different combinations of information like Top Applications/Users/Sites by bandwidth/Sessions/Risk/meeting SLA. Another level of classification is also permitted. We also created SD-WAN report as part of the test, and these are the features that were reported: Device and Security Events, App Visibility and App SLA Performance Monitoring also have the possibility of scheduled reports via e-mail.

CSO Infra Monitoring

To have monitoring and analytics tools to monitor the devices, infrastructure, including the SD-WAN controller itself, Icinga version 2.4.1, an open-source tool, was customized by Juniper. This tool monitors the Infrastructure and microservices VMs (host aliveness, CPU usage, disk IO, disk usage, load average, memory usage, paging stats, etc.) by logging in to the server every minute (by default) via NETCONF and executing the correspondent commands, SRX Series Next Generation Firewalls are monitored the same way. For NFX Series Universal CPE devices streaming telemetry is used.

Application Visibility

This feature is one of the many options for Application-level monitoring. We observed that the AppVisibility capability within Contrail SD-WAN was able to display a number of application details on the site management page: top talkers (applications) for a site, per WAN link, Application SLA performance, BW/Session utilization graph, etc... We were able to observe during test cases using the traffic profile shown in Table 1 from section “Hardware and Software”. An example is shown in Figure 12.

Different templates and layouts for graphs/widgets and data are options to the user.




Figure 12: Customizable Application Visibility Widget

Figure 13 shows other possible widgets added to the customizable dashboard.




Figure 13: Customer Widget Dashboard

Additional Functionality

Besides all features and highlights discussed, other relevant functionality was also tested in the lab, which will be described in the following subsections, among them:

- Unified Threat Management (UTM)
- Design Portal for Custom Service/Service Chaining
- Device RMA
- Role-Based Access Control (RBAC)
- CSO Troubleshooting

Unified Threat Management (UTM)

Contrail SD-WAN's endpoints are all fully integrated with Juniper SRX NGFW software. Therefore, Contrail SD-WAN can be used to provision NGFW policies using the same provisioning interfaces used for SD-WAN policy configuration. Additionally, Contrail SD-WAN can also be used to provide advanced security features such as SRX Unified Threat Management (UTM). Offering the ability to provide multiple security features and the SD-WAN capabilities from one portal simplify management of the WAN.

In this case, we tested web and content filtering for an SD-WAN site. For the web filter, customized URLs (e.g. www.yahoo.com) were specified to be blocked. And for the content filter, file with .exe and .zip extensions

downloaded via HTTP would not be allowed. Both these filters are added to a UTM profile, and this profile is enforced via a firewall policy for the site. All web and content were observed to be permitted before the firewall policy was deployed (via CSO), after this, the specified traffic was blocked successfully. Antivirus and Antispam are other UTM features in CSO, but they were not configured or tested in this case.

We observed firewall policies configuration can be updated and deployed dynamically via CSO GUI effectively blocking/allowing specified traffic after the change.

CSO Troubleshooting

Every time a site is enabled or updated CSO generates specific job logs. This allows the operator to go back into this job log and identify what happened, locate the failure and correct it. The option to share with Juniper Networks is possible. Over time, the database of job logs will naturally continue to grow. Therefore, it is advantageous to have a tool that aids the end user in the search and presentation of data. Juniper Networks has provided Kibana for this purpose.

Kibana is a query based open source tool which is customized and integrated into CSO. It comes with the CSO package. Needless to say by reading the prior sentence, CSO is in part based on Docker and Kubernetes. We observed how all log information was filtered at different levels with tailored charts in order to assist troubleshooting. Kibana version 5.6.8 was used with this CSO 4.0.1. In Figure 14 we are able to see a chart that shows the incoming request per microservice, this chart was obtained after finishing the "CSO Site Scalability".




Figure 14: Kibana Overview - Registered Events

Custom Service Enablement for SD-WAN Site

In this case, we tested the CSO configuration designer. In many cases, there are capabilities that an Enterprise IT manager or Service Provider may want to expose as part of the solution. Instead of waiting for these features to be added to the solution, as long as they are available in Junos, these Junos based features can be exposed or added to the configuration process using customer-defined templates.

The CSO configuration designer allows the end user to create a graphical user interface for Junos capabilities that are accessible through the CLI, but not readily available in Contrail SD-WAN. In Configuration Designer, you can manually type a working configuration or copy and paste an existing golden configuration from your device. You can also use your own data model to configure your template. Once created, the templates are listed on a Design page, where you can review them at a glance. You can also modify the parameters and values of your templates as needed from the Design page. This capability helps the administrator to simplify and automate the configuration and provisioning of an SD-WAN site. These are referred to as Stage 2 templates in Contrail SD-WAN. An overview on how the Configuration Designer works can be seen on Figure 15.




Figure 15: CSO Designer Integrated Workflow

For our test case, we used the configuration designer to create a generic DHCP Relay template, the tool detects variables and those will be user input. Default values can be set for the customer, and the service is published for a specific tenant. This feature works based on jinja2. After using the service via the customer portal we observed the configuration was correctly deployed as per the template and variables specified. Variable customization is seen on Figure 16.

Item	Variable Name	Var Type	Class Type	Key	Required	Default Value	Pattern	Information	Description
1	lan_interface	self	enumeration		✓	---	---	[?a-zA-Z0-9-]{1,64}	Description for Lan Interface
2	dhcp_server	self	string		✓	---	---	---	Description for DHCP Server
3	lan_vrf	self	string		✓	---	---	---	Description for Local VRF
4	tenant	self	string		✓	---	---	---	Description for Tenant

Figure 16: Customizing Variables in the Configuration Designer

Device RMA

This setup showed that Return Material Authorizations (RMAs) can be managed from the CSO customer portal. The RMA was initiated from CSO and it backed up all the configurations from the site. The device was physically replaced, then the RMA was marked as granted. After this, CSO re-provisioned the site using the serial number/activation code required. The device replaced was an SRX340.

Role Based Access Control (RBAC)

The RBAC feature controls which system users can view, read, write, and execute within the Administrative and Customer portals. Administrators can provide granular control over GUI objects within each navigation menu, restricting users to the views and/or capabilities specific to their role. The RBAC capability also includes the OpCo feature, which enables global administrators to define a single service across multiple regions while allowing regional administrators to manage their own local customers. In this scenario, global service providers give OpCo administrators access to a centrally deployed Contrail Service Orchestration instance, along with the local resources they need, enabling them to offer SD-WAN services that meet local regulatory requirements.

This feature can be used within enterprises to provide hierarchical access to capabilities at different levels or to allow or restrict access to specific capabilities across departments. Predefined roles were tested to work only with their specific permissions. Specific roles tested were: service provider admin and operator, tenant admin, operator, and site configuration.

Service Chaining with 3rd Party VNF

The Network Service Designer tool within Contrail SD-WAN enables service managers and administrators to intuitively define a customized service catalog through a simple wizard. This capability is a key enabler for Enterprise organizations and Service Providers to consolidate branch devices onto a single platform to simplify and automate the WAN edge.

Service chaining allows new services to be instantiated as software-only, running on commodity hardware. We used the CSO Network Service Designer portal to create the service with the following data, based on a request:

- VNF image: ubuntu-fw
- LAN or WAN side: LAN




Figure 17: Design for Service Chaining with 3rd Party VNF

In Figure 17 we can see the published network service design. The service provider has to upload the third party VNF image mentioned into the Administration Portal/Resources/Images. In this case Ubuntu 16.04.2 LTS was uploaded into CSO. This was done via REST API. This could also be done by the GUI, as tested and documented earlier in this report for the vSRX.

The service provider can then allocate the service to the tenant via the Administration Portal. Finally, the customer can drag and drop the network service to the desired CPE in the site management page, add the IP address and the routes and start the service. After the service was deployed, we logged in to the Ubuntu and performed ping/traceroute destined to the Internet by pinging the Google DNS IP address of 8.8.8.8 to verify if it was added in the uCPE. We also logged into the vSRX to confirm that the service chain was completed on the NFX Universal CPE, and finally we logged into the SD-WAN hub, to check the traffic flow information in the Hub itself, to make sure the traffic was flowing across the network. No firewall capabilities were tested.

Conclusion

A high-level overview of the testing process is given in this section along with important results.


In order to demonstrate the provisioning process, the initial test bed was only physically connected and devices were zeroized. CSO installation was performed, tenants and hubs were provisioned to allow later ZTP for different CPE devices and templates. After provisioning test bed devices, different traffic steering capabilities were configured and tested. Most important were SLA violation detection in a link and SD-WAN policies to perform Link Affinity with APBR and AppQoS. Path preferences for each application were set and SLA profiles with thresholds to cause automatic switchover of application traffic when violated.

Failover test of active-passive Hub was also performed to demonstrate high availability and resiliency. Custom service creation/activation using a designer portal was used to show service chaining and with Juniper and third party VNFs, and the ability to create new service features with stage 2 configuration templates. Customizable monitoring and troubleshooting portals handling a high amount of events, also allowing data presentation in many different personalized ways to fulfill specific user needs.

A separate test bed was used to perform site scalability in the controller side, resulting in 10k emulated spoke sites provisioned in CSO.

In summary, the EANTC test cases successfully verified Juniper's claims of the Contrail SD-WAN being a secure, scalable and customizable solution with many add-on features corresponding to today's and future service provider requirements.

About EANTC



EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies.

Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. Customers include leading network equipment manufacturers, tier 1 service providers, large enterprises and governments worldwide. EANTC's Proof of Concept, acceptance tests and network audits cover established and next-generation fixed and mobile network technologies.



This report is copyright © 2018 EANTC AG.
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.com, <http://www.eantc.com/>
[v1.6 20181002]