EANTC Test Report
Huawei CloudMetro
Performance and Functionality

Introduction

Huawei commissioned EANTC to verify the functionality, performance and security of their CloudMetro solution. The main goal of the solution is the transition of the broadband access infrastructure into a cloud-based architecture with a separation of the control and user plane.

Our tests included following areas:

- Lifecycle management
- Performance and scalability
- Functionality
- Resiliency
- Security

Executive Summary

Through a series of tests, EANTC verified various functional and performance claims for Huawei's CloudMetro platform

- Automated deployment of Network Cloud Engine software and full lifecycle management of VNFs in accordance with MANO architecture.
- Automatic scale-up and scale-down procedures for the Network Cloud Engine components to adjust for performance demand
- High subscriber session setup rate
- High number of simultaneously managed subscriber sessions.
- Authentication of PPPoE and IPoE subscribers via RADIUS backend.
- Simultaneous support of PPPoE and IPoE subscribers on the same physical interface
- Application of per-subscriber bandwidth profile upon session establishment, or modification on the fly.

Test Highlights

- **Demonstrated lifecycle management, up- and down-scaling of control plane modules**
- **Demonstrated separation of the user and control plane**
- **Up to 1 million subscribers per Network Cloud Engine (default QoS profile and shared IP pool) - verified with bidirectional test traffic.**
- **2,000 PPPoE sessions/s and 1,700 IPoE sessions/s with RADIUS authentication**
- **2,500 PPPoE sessions/s and 2,600 IPoE sessions/s without authentication**
- **Simulated faults of control plane components cause no interruption in current user sessions.**

- Recovery of Network Cloud Engine components from various faults through redundant setup of components.
- Stable operation under ongoing DDoS attacks

CloudMetro Architecture

Challenges of the Conventional Broadband

In the conventional deployment of the broadband services, the provider would assign a number of hardware Broadband Network Gateways (BNG) to serve a specific geographic location. The resource allocation in this case is based on the estimated number of subscribers and their estimated bandwidth utilization. However, even with careful planning, a rigid resource distribution faces a fundamental problem of uneven utilization. Even if the total traffic generated by the subscribers lies within capacity limits of the network, some nodes

may become oversubscribed, both in terms of bandwidth and user sessions.

Furthermore, such rigid deployment of resources prevents the operator from easily deploying or modifying services. Management of IP ranges, VLANs, quality of service parameters, fault detection becomes difficult with the growing number of deployed access devices, locations and subscriber accounts. Hindered by the increased management efforts and inflexible allocation of resources, the providers are missing the opportunity to rapidly deploy new, innovative services, time-limited offerings and flexible broadband plans for the customers with different demands.

Huawei's CloudMetro Solution

With the CloudMetro solution, Huawei seeks the possibility to make the deployment and management of a large network more flexible. The key element of CloudMetro is the separation of the control and user plane in a software defined networking (SDN) and network function virtualization (NFV) architecture. The intention is to handle the subscriber session signaling and additional services in a cloud-based, virtualized controller, while the forwarding of user data is handled by the Multi-Service Edge devices.

The CloudMetro solution relies on standard technologies to provide maximum interoperability in any environment. The central part of the CloudMetro solution is the Network Cloud Engine (NCE) component, which is the implementation of the control plane for the broadband access.

On the southbound interface, the NCE is managing Multi-Service Edge devices (MSE), in our test setup we used Huawei ME60-X8, NE40E-X2-M8A and NE40E-X2-M16A. According to Huawei, any devices from their NE40E Universal Service Router[1] and ME60 Multiservice Control Gateway[2] lineups can be used for this purpose. The MSEs can be configured via OpenFlow protocol to forward and handle specific traffic flows. They use a VXLAN tunnel to transport necessary data from the subscriber session to the NCE during the session establishment process. The NCE itself is a set of virtual network functions (VNF) running in a standard MANO (NFV Management and Orchestration) environment.

In the conventional broadband architecture, a Broadband Network Gateway (BNG, or BRAS in the former terminology) serves as the termination point of the subscriber sessions and as their interface to the Internet. BNG takes the responsibility to signal and maintain the sessions, and to relay subscriber authentication to the AAA backend. Broadband user traffic is typically carried either within PPP encapsulation, in which case the BNG is responsible for assigning the IP configuration to the clients via PPPoE protocol, or as plain

1. X16A, X8A, X3A, X16, X8, X3, X2-M16A, X2-M8A, X2-M16, X2-M8, M2F, M2H and X1-M4
2. X16A, X8A, X16, X8, X3, X2-M8A and X2-M16A

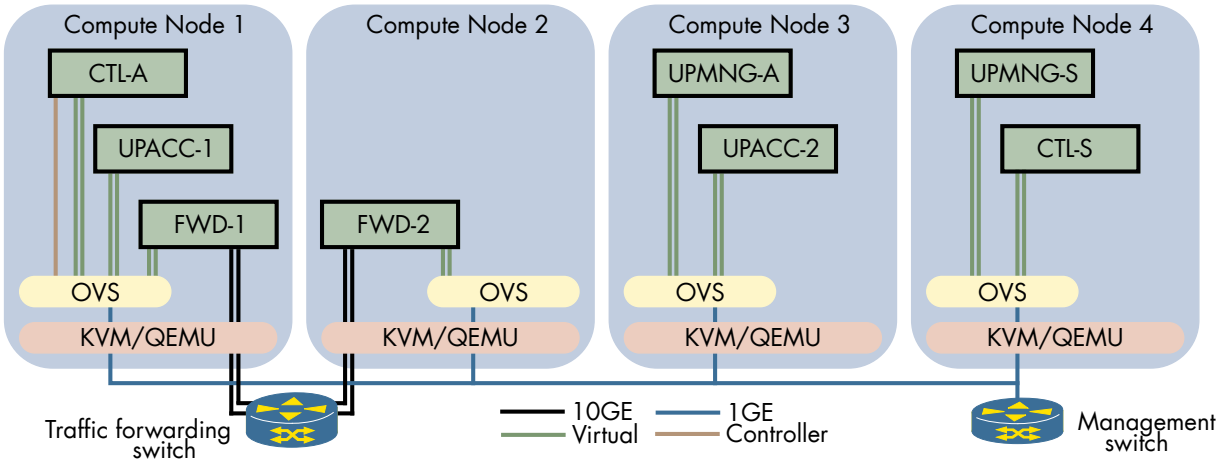

Figure 1: Network Cloud Engine

IP, in which case the BNG provides a DHCP service for the connecting clients. While for the most part BNG will simply forward user traffic, it requires sufficient intelligence for maintaining the control plane as well, performing authentication, managing IP ranges etc.

In the CloudMetro concept, these functions will be moved to a central NCE element which will be capable of managing multiple geographic locations. The centralization and virtualization of this component allows for flexible scaling of the control plane performance. During the session establishment process, MSE will forward the control-plane related user packets to the NCE using VXLAN encapsulation.

Once the session is successfully established, NCE issues flow definitions to the MSEs to forward traffic of the user session to the core router. The MSEs are also responsible for decapsulating user traffic in case of PPP. MSE therefore serves as the endpoint of the PPP session and all forwards the traffic toward network core as plain IP.

Hardware Under Test

Network Cloud Engine (NCE)				
Node Type	**vCPU**	**Memory**	**HugePage**	**Software**
2x CTL	2	16GiB	Not configured	HUAWEI VNE9000 Patch Version: V100R003SPH001
2x FWD	22	16GiB	1GiB	
2x UPMNG	6	32GiB	Not configured	
2x UPACC	14	16GiB	1GiB	
Compute Nodes	Number of nodes: 3 **Hardware:** HUAWEI RH2288Hv3 V100R003 **CPU:** 2x Intel(R) Xeon(R) CPU E5-2690 v3-12Core @ 2.60GHz **RAM:** 8x DDR4 RDIMM-16GB-288pin-0.9ns-2133000KHz-1.2V-ECC-2Rank(1G*4bit) **NIC:** 6x PCIE 82599ES 10-Gigabit SFI/SFP+ Network Connection			OS: SUSE, kernel version 3.0.93-0.8 hypervisor: v2.3.0, Open vSwitch v2.3.2 SR-IOV with DPDK v16.04 libvirtd (libvirt) 1.2.17
Multi-Service Edge devices (MSE)				
Hardware type				**Software**
HUAWEI ME60-X8				Version 8.120 (ME60 V100R003C10B180) Patch Version: V100R003SPH001
HUAWEI NE40E-X2-M16A				Version 8.120 (NE40E V100R003C10B200 Patch Version: V100R003SPH001
HUAWEI NE40E-X2-M8A				Version 8.120 (NE40E V100R003C10B200)

CloudMetro: Test Setup

For our functional and performance tests, we deployed a CloudMetro configuration on a set of 4 compute modules. The deployment, shown schematically in the diagram below included the necessary CloudMetro components in a redundant active/standby configuration, as well as a number of RADIUS servers for authentication of the user sessions.

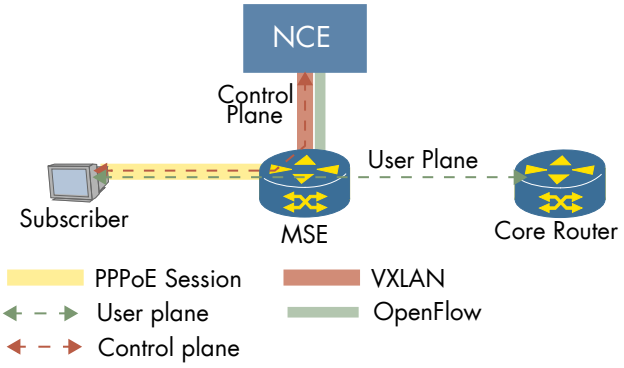

Figure 2: Logical Topology

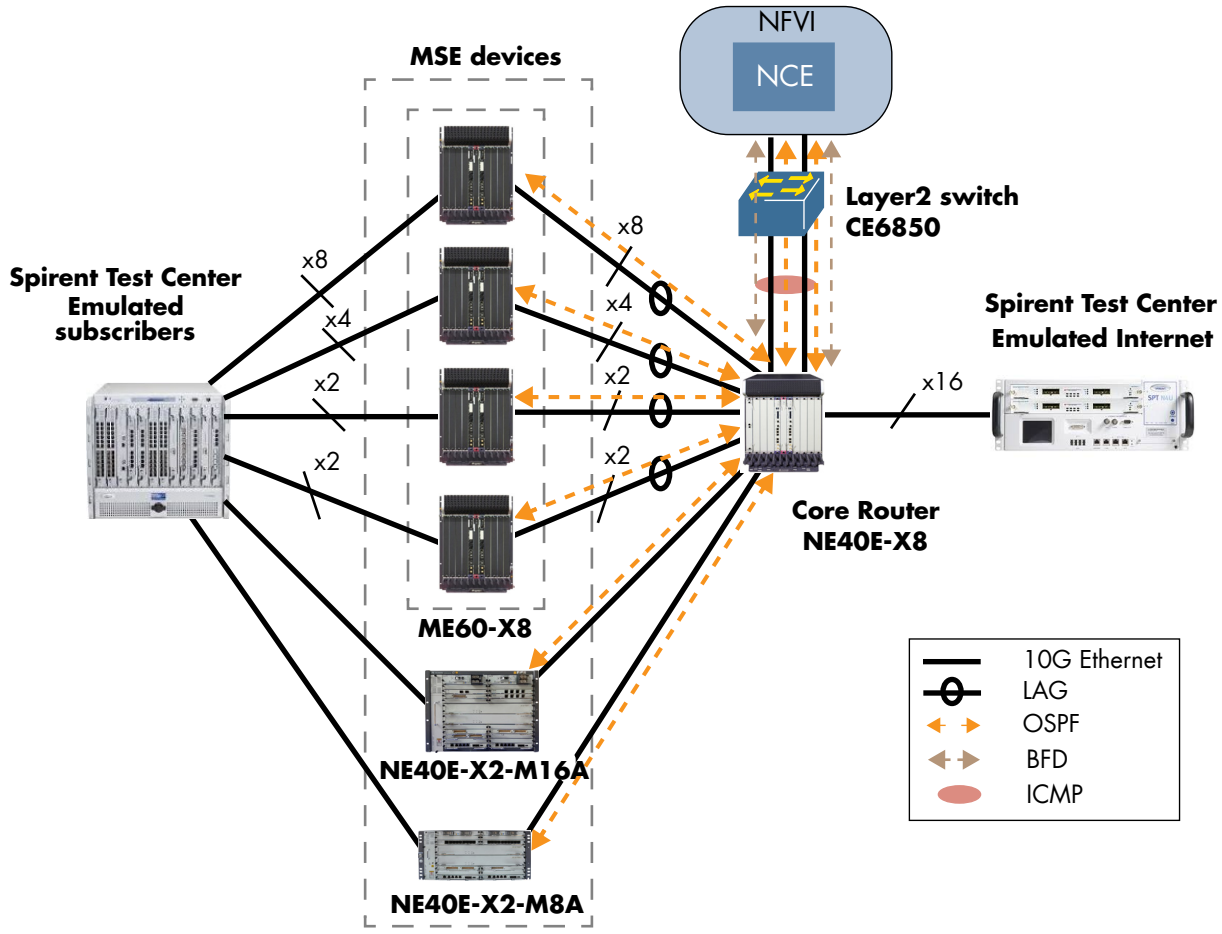

Figure 3: Physical Test Setup

In order to verify the correct function and traffic forwarding of the CloudMetro components, we deployed an external traffic generator TestCenter N4U and N11U from Spirent connected to the MSE devices on the subscriber side and to the core router on the network side. By simulating PPP or DHCP clients on the subscriber side we verified the correct session establishment process and by sending test traffic for each session we verified the correct traffic forwarding.

For the test, Huawei deployed four ME60-X8 devices as MSE, equipped for different number of network interfaces. The intention was to create a realistic example of capacity distribution across multiple locations. In total, we connected 16 10GE ports on the subscriber and 16 on the network side.

Life-Cycle Management and Functionality

Test Highlights

- **Successful demonstration of life-cycle management using FusionSphere Openstack CSM and OM as VIM[a] and CloudOpera CSM as VNF-M[b]**

a: Virtual Infrastructure Manager; b: VNF Manager

The deployment of the Network Cloud Engine consists of multiple VMs carrying the various components and spread across several compute node in order to provide hardware redundancy. The compute nodes use KVM/QEMU as the virtualization platform, and OpenStack Juno for the orchestration of the components.

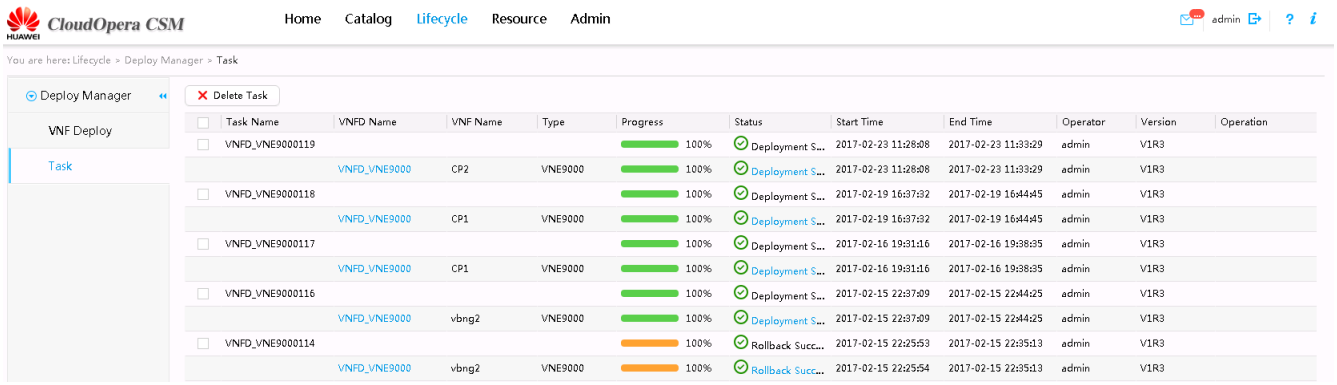

Figure 4: Huawei CloudOpera CSM

The UPMNG (User Plane Management) is responsible for managing OpenFlow connections to a group of up to 8 MSE devices. When new MSE units are added or removed, NCE is able to automatically spawn additional or delete excess UPMNG instances to maintain the desired performance. UPMNG instances are running in an active/standby configuration on different compute nodes to provide resilience against faults.

We verified the scale-in and scale-out functionality for UPMNG modules by modifying the number of MSE devices managed by a single instance. Using command line interface on the NCE, we configured additional MSEs and observed that additional UPMNG instances were automatically spawned to accommodate the number. Similarly, we observed that the excess UPMNG instances were automatically destroyed when the number of MSEs fell below threshold.

UPACC (User Plane Access) components store the information on PPPoE and IPoE subscribers and also provide automatic scale-up functionality to adjust for the performance demand, based on the number of line cards configured on the MSE devices. We verified the functionality by adding or removing line cards from the configuration and observing how additional UPACC instances are automatically spawned or removed

The CTL (Control) and FWD (Forwarding) components are responsible for managing the MSE devices on low level. The two redundant (active/standby) CTLs perform management and configuration tasks on the MSEs, while the FWDs are maintaining the flow state for each subscriber. In our tests, we created a pool of 2 FWD instances on the NCE to provide sufficient capacity for 1 million subscribers. The FWDs provide redundancy, however all instances operate in active state. In case of an FWD fault, the load will be spread to other instances until a new instance is created or rebooted.

We initiated the deployment process of CloudMetro on a cluster of 4 compute nodes managed with OpenStack. The components of NCE are deployed in the Virtual Private Cloud (VPC) that in turn can be spread over multiple Virtual Data Center (VDC).

Huawei FusionSphere OpenStack provides two virtual infrastructure manager (VIM) modules, The CPS module is responsible for mapping of the physical network interfaces, while OM is responsible for creating logical network connectivity inside the cloud. We used these components first to set up our network interfaces in the pass-through mode and then to create a VPC for the NCE components.

As next step, we used VNF manager (VNF-M) CloudOpera CSM to deploy all components of the NCE. The necessary VNFs images, configuration parameters and the network interconnections are defined in a single VNFD template file. The instantiation process ran fully automated and took approximately 5 minutes to complete. The only manual intervention was necessary at the end of the deployment to switch NCE to the control/user plane separation mode from the default "normal" mode.

In a separate test case, we demonstrated the process of adding MSEs to the configuration. For this purpose, we configured the MSE interface and established OSPF, OpenFlow, NETCONF and VXLAN sessions between it and NCE. We demonstrated the identical process for three different

Huawei products: ME60-X8, NE40E-X2-M8A and NE40E-X2-M16A devices that can be utilized as MSE devices.

After verifying that all components are active and properly configured, we verified the functionality of NCE by establishing subscriber sessions and sending test traffic.

Subscriber Session Management

After deploying and configuring the CloudMetro platform, we reviewed and tested the process of subscriber session establishment.

Within CloudMetro architecture, only the session establishment process is handled by the NCE. The MSE, receiving a PADI (PPPoE Active Discovery Initiation) or a DHCP request from a new subscriber forwards the packets transparently to the NCE using a permanently maintained VXLAN tunnel.

A successful session authentication by NCE results in a SDN flow configuration being sent to the MSE via OpenFlow protocol. From now on, the subscriber traffic is forwarded by the MSE directly to the network and does not consume processing resources on the NCE.

In case of PPP, MSE also performs the PPP decapsulation. Furthermore, it is the responsibility of the MSE subscriber-facing interface to maintain the PPP session by requesting LCP Echo Requests to the clients.

In our tests we reviewed and analyzed the process of session establishment for PPPoE subscribers. We verified that in all established sessions, the clients received the IP addresses in accordance with the configured address pools. We also monitored the OpenFlow communication for selected subscribers to insure correct configuration of the MSEs and successfully demonstrated that the subscriber traffic was completely offloaded to the MSEs.

We demonstrated the ability of the CloudMetro platform to support both PPPoE and IPoE subscribers on the same physical interface. We configured 2 sub interfaces on NCE assigned to PPPoE and IPoE domains respectively, but sharing the same IP pool. We verified the correct function by establishing PPPoE and IPoE subscriber sessions on the same physical interface and sending test traffic.

Performance and Scalability

In our next series of tests we asserted the performance of the CloudMetro platform with one Network Cloud Engine and four ME60-X8 MSEs deployed to support up to 1 million subscribers. We measured the maximum session establishment rate, the total session capacity and the throughput of the solution.

We performed all tests using the Spirent Testcenter traffic generator, connected to the MSE devices to simulate PPPoE or IPoE subscribers, and to the core router to terminate the test traffic flows on the network side. Currently, the platform provides no support for IPv6, all test traffic flows used IPv4, additionally encapsulated in PPPoE for the PPPoE subscribers.

Session Setup Rate

We measured the maximum session establishment rate supported by a single NCE installation in combination with 4x ME60-X8 MSE devices. We emulated up to 800,000 PPPoE clients establishing PPPoE sessions against the network at high rate. We configured 6 second timeout and no session reattempts for the PPP emulation.

During the tests, the subscribers were authenticated against the RADIUS server running on one of the compute nodes of NCE using MSCHAPv2 method. Huawei used DHCP Option 60 to provide authentication for the IPoE subscribers.

In order to achieve the desired performance, Huawei adjusted the session establishment process

Test Highlights

- **Established 1 million PPPoE sessions at 2,500 sessions/s without authentication**
- **Achieved 800,000 PPPoE sessions at 2,000 sessions/s with RADIUS authentication**
- **Established 1 million IPoE sessions at 2,600 sessions/s without authentication**
- **Achieved 800,000 IPoE sessions at 1,700 sessions/s with RADIUS authentication**

to reduce communication between components. The RADIUS was used for subscriber authentication only, no accounting was performed. In addition, NCE did not transmit QoS parameters for the established SDN flows, saving on OpenFlow communication between NCE and MSE.

We achieved successful establishment of PPPoE sessions at the rate of 2,000 sessions/s without session failures observed. Due to limitation of the RADIUS server, we had to maintain this setup rate for the first 800,000 sessions only, the remaining 200,000 sessions were established at a lower rate of 320 sessions/s. During the test, the CPU consumption on the CTL reached approximately 70% and the memory consumption 42%.

After the sessions were established, we sent bidirectional test traffic for all sessions to verify the correct forwarding between core network and the emulated subscribers. We observed only 1 failed session in case of the unauthenticated IPoE subscribers, and 135 with authentication enabled. For PPPoE sessions, we observed no failures.

Session Type	Auth. method	Total sessions	Setup Rate [session/s]
PPPoE	none	1,000,000	2,500
	MSCHAPv2	800,000	2,000
IPoE	none	1,000,000 1 failed	2,600
	DHCP Option 60	800,000 135 failed	1,700

Subscriber Session Capacity

Similarly to the session setup rate test, we emulated a large number of PPPoE and IPoE subscribers establishing sessions against the network. In this case however, QoS information was transmitted to MSEs and the sessions were established at a lower rate.

In our test run we verified the successful establishment of 1 million sessions for both PPPoE and IPoE domains at a rate of 2000 sessions/s with retries enabled. Subsequently, we verified the correct traffic path setup of these sessions by running test traffic between endpoints.

During the test, we retrieved the information about user sessions and traffic rate from NCE. This information is made available through the NETCONF sessions between the NCE and MSEs.

Per-subscriber Bandwidth Profile

We verified the assignment of subscriber's total upstream and downstream bandwidth limits. This information is a part of per-subscriber profile and is delivered by the RADIUS server to the NCE after the subscriber is authenticated. The task of NCE is then to apply this profile to the MSE when configuring the traffic flow for this subscriber.

To do the test, we configured a default traffic profile for the subscribers (up/downstream limits) and applied this profile for the PPPoE domain in NCE. We established 100 PPPoE sessions and verified that a correct profile was applied to each of these subscribers via NCE's CLI and also by measuring allowed bandwidth limits using test traffic.

In another test, we verified that subscriber's profile can be dynamically changed. We configured a different profile, however it was not applied automatically to existing subscribers. We performed the test manually by sending out RADIUS CoA-Requests (Change of Authorization) for affected subscribers to MSE and verified that the new profile was applied on the fly for the established sessions.

Resiliency

CloudMetro incorporates several resiliency concept in order to improve availability of the broadband service in case of various faults that can occur in physical equipment or software. We simulated a number of faults described in the table below, while having active subscriber sessions established and with running test traffic.

We measured the possible service downtime by evaluating the number of lost packets, failures during session establishment or other negative effects. We also evaluated the logs of the devices for alarms raised during the faults.

Node failure

We analyzed the behavior of each of the NCE components by simulating a their failure during operation. For this purpose we established up to 1 million subscriber sessions and while continuously

sending test traffic, forcefully restarted the respective VNFs housing the components via OpenStack command line tools.

In case of CTL and UPMNG, all subscriber sessions continued to function without interruption after restarting the component using nova stop/start command. We did not observe any loss in the test traffic being sent through these sessions.

In case of the UPACC component, we first established 937,500 subscriber sessions (15 of the traffic generator ports) and then started the establishment of the remaining 62,500 sessions at a lower rate. During this process, we simulated the UPACC node failure using nova reboot command. This allowed us to analyze the behavior of the platform not just for already established sessions, but also see how the fault affect new subscribers attempting to connect. We observed a window of approximately 150 seconds during which no new sessions could be established. We validated however that there were no interruptions in the test traffic for the sessions already active, or those established after recovery of UPACC node.

In another test case, we simulated a failure of the FWD component by stopping and restarting its VM. As expected, we observed no traffic loss for the existing sessions, only the establishment of new session was briefly interrupted for nine seconds during shutdown, respective for 20 seconds during restart.

As the next step, we simulated the failure of the VXLAN session by deleting it on one of the MSEs. As expected, the VXLAN failure did not cause issues with the already established sessions, but no new sessions could be established. Session shut down initiated by the subscriber however was possible, as the MSE identified defunct sessions through missing LCP packets - after the configured timeout of three packets, the session was dropped.

Protection against Denial of Service Attacks

We verified NCE's ability to operate under ongoing distributed denial of service (DDoS) attacks. Malicious PPPoE subscribers can flood the BRAS with PPPoE discovery (PADI) requests, consuming resources and possibly disrupting the sessions of other subscribers.

CloudMetro provides a possibility to limit the rate of PADI packets to an acceptable maximum. We configured a maximum allowed rate of 1500 packets per second. After establishing 437,500 parallel PPPoE sessions and constantly sending test traffic to verify their stability, we simulated an attack by sending PADI packets at a rate of 1000 packets/s and verified that these are answered. We confirmed that at the rates above 1500 requests per second, excess requests are dropped, keeping the resource consumption on the NCE to an acceptable limit.

Finally, we attempted to attack NCE with a high rate flow of PADI packets at 5 Gbit/s. The result of this test was inconclusive, as the rate of the attack was already limited by the MSE device.

About EANTC

The European Advanced Networking Test Center (EANTC) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies. Based in Berlin, Germany, 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. http://www.eantc.com

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