EANTC Independent Test Report

Metaswitch’s Clearwater IMS Core

Performance, Scalability, Reliability and Functionality

May 2018

About EANTC

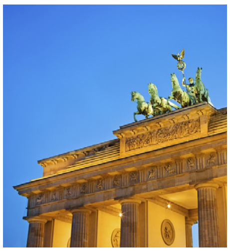

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies.

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.

Table of Contents

Introduction

IP Multimedia Subsystems (IMS) are finally ready to deliver on the promise of ubiquitous VoLTE and converged rich communications services; but to fully realize the potential, network operators must think beyond simple virtualized network functions to fully cloud native IMS components.

EANTC was commissioned by Metaswitch to perform independent tests of their Clearwater Core IP Multimedia Subsystem (IMS) solution.

Test Highlights

- ➔ Verified 20M successful concurrent subscribers capacity at a 5,544 Registrations/s rate
- ➔ Measured 1.7M error-free concurrent calls created at a rate of 2772 Call/s
- ➔ Clearwater Core flawlessly processed 15,000 Calls/s
- ➔ Observed no traffic loss in three node failure and recovery scenarios
- ➔ Automated instantiation (including configuration and verification) of a Clearwater Core deployment for 20M subscribers on Amazon EC2 in 72 minutes

The emergence of Voice over LTE (VoLTE), Voice over WiFi (VoWiFi) and other IP multimedia services place large bets on the IMS core's capabilities. New IoT applications, new services with unknown traffic patterns, or spikes in subscriber sessions can put the network in jeopardy. Service providers need a reliable and scalable solution that can scale, heal, adapt and be 5G ready. On the other hand, the approach of sizing for full scale from day one, or investing in big iron, comes with its own disadvantages, most notably high cost.

Service providers are well aware of the service agility costs. This is one of the main reasons that has driven Network Functions Virtualization (NFV) and cloud-based services to the forefront of any new network design. NFV was designed from its infancy around the idea of resource sharing and optimization.

Metaswitch's powerful and flexible cloud native virtualized IMS (vIMS) core enables network operators to transition fully from circuit-switched telephony infra-structures to all-IP, and to provide globally ubiquitous rich communications services.

Clearwater Core is Metaswitch's implementation of IMS. According to Metaswitch, Clearwater Core has been architected from the ground up for massively scalable deployments within virtualized public, private or hybrid elastic compute clouds.

To prove its ability to scale and its cloud native resiliency, Metaswitch has asked EANTC to test it. The vIMS solution was hosted on Amazon EC2 – a global Infrastructure as a Service (IaaS) platform. This was a convenient option for running high-scale testing. It is not necessarily expected that operators will deploy their solutions in a public cloud environment, and the results should be indicative of what could be achieved in private Telco clouds. For the testing, we sent our team to Metaswitch's offices in the town of Enfield, United Kingdom for one week and the results were impressive. This report will take you through each of the tests we performed and its results.

Component Overview

System Under Test

Clearwater Core is decomposed into a number of different node types, and is deployed with clusters of each node type, scaled as required. Each node type provides a different function as part of the whole. These node types do not map to traditional IMS functional entities, but the combination of them all results in IMS compliant behaviour of the whole.

- SIP Processing Node (SPN): The Clearwater Core node that handles incoming SIP requests SPN is stateless
- Diameter Gateway Node (DGN): The Clearwater Core node that communicates with the AAA elements of the network, such as HSS. DGN is also stateless
- Storage Cluster Node (SCN): Contains subscriber state and transaction timer information

Testing and Orchestration Tools

- Clearwater Core Stress (CWCS): A lightweight traffic generator and analyzer developed by Metaswitch
- Salinas: A library of orchestration scripts used to simplify the lifecycle management of the Clearwater Core deployment created by Metaswitch

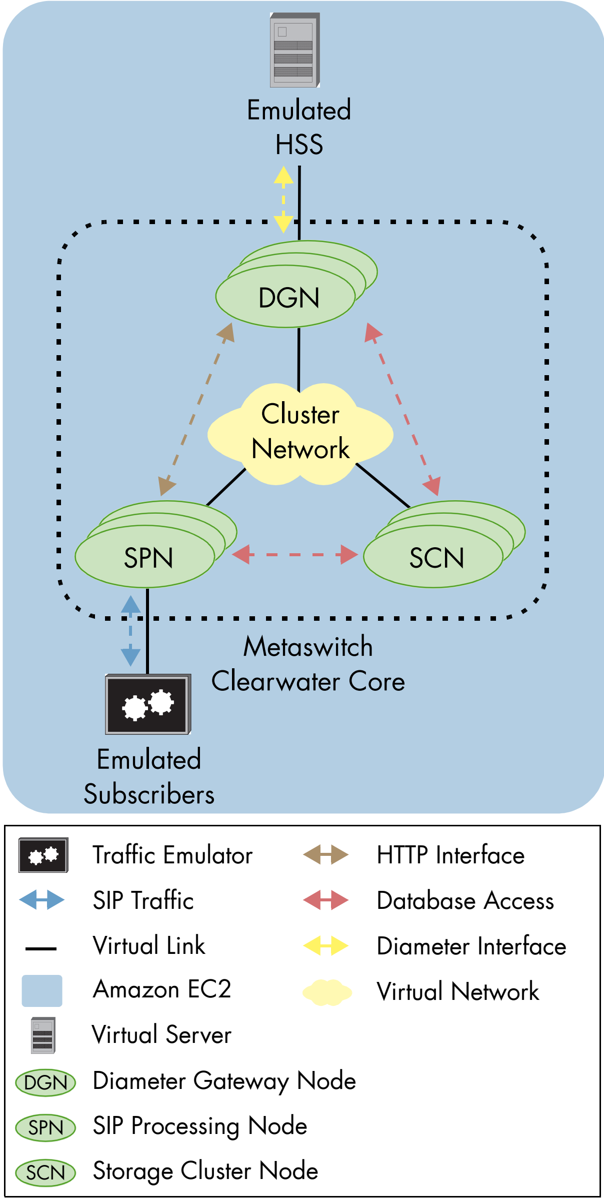

Figure 1: Test Setup

Test Setup

Our tests involved four main components and a traffic generator. The first three components were Clearwater Core's SPN, DGN and SCN instances grouped into three clusters – one cluster per function. The three clusters were directly connected via a single subnet and they exchanged information using Restful API sessions over HTTP. The fourth component was a slightly slimmed down Home Subscriber Server (HSS) function that communicated using Diameter messages with the DGN cluster. The traffic generator was connected to the SPN cluster and they exchanged SIP traffic.

For a more realistic configuration, Metaswitch configured a Breakout Gateway Control Function (BGCF) entry on the SPN instances which pointed traffic back towards the traffic generator. This simulates traffic being directed towards an external serving gateway in a service provider environment and adds an extra lookup task for the SPN for each call it processes.

During the pre-staging, the Metaswitch team introduced us to their home-grown test tools. The traffic generator was hosted on a c5.9xlarge, 36-core machine on Amazon EC2. We evaluated the test tools carefully and found them to perform exceptionally well for the purpose of this test.

We used four differently weighted scenarios to stress the unique functional elements on the Clearwater Core solution. The first scenario focused the pressure on subscription rate and concurrent subscribers. The second scenario stressed SIP call frequency. The third tested concurrent call capacity and the fourth was a uniformly weighted configuration with equal stress on all four performance elements. The latter scenario was used for reliability tests. Detailed numbers are listed in Table 1.

In all the traffic profiles, 50 percent of emulated calls were incoming and the remaining 50 percent were outgoing calls. We emulated VoLTE flavored call signalling. The detailed call flow is depicted in Figure 2.

Traffic Profile	A	B	C	D
Registration Rate [Regs/s]	5,544	3,333	2,772	4,166
Concurrent Subscribers	20M	12M	10M	15M
Call Rate [Invites/s]	-	15K	3,800	2,500
Call Duration [sec]	-	10	1,200	720
Concurrent Calls	-	150K	4.5M	1.6M
Total Calls [BHCA]	-	29M	14M	9M

Table 1: Emulated Signalling Scenarios

During the pre-staging phase, the Metaswitch team prepared a 95-node Clearwater Core deployment. Part of this pre-staging included moving to the new Intel® Xeon® Scalable processors, which are the c5 variants within Amazon EC2. Previous testing by Metaswitch had used the older c4 variants.

We closely monitored CPU and memory utilization levels on the nodes as we ran a selection of tests. Since the nodes reported the combined load of individual CPU cores via SNMP – which sometimes can mask the exact load per core – we manually checked a sample of running instances and verified that load distribution was relatively even across CPU cores. This made us more comfortable with the averaged measurements that the Clearwater Core system provided.

Performance

An IMS system handles signalling and control plane traffic. It is important that cloud IMS deployments can function without error under pressure. In order to quantify Clearwater Core's performance we designed the tests taking into account four factors: subscriber capacity and registration rate, call signalling rate, and concurrent calls. Table 3 contains the detailed resource utilization measured for each of the test scenarios.

Subscriber Capacity

In this test, we emulated a scenario with a high subscriber registration rate towards the SPN cluster and monitored traffic for any signalling failures. Those signalling transactions required Clearwater Core to access its subscriber data at a higher capacity. We also stressed Clearwater Core's concurrent subscriber capacity by maintaining all the sessions that got established by the traffic generator for the duration of the test. See Traffic Profile A in Table 1.

At the end of the test's one hour run time, Clearwater Core successfully registered 20 million subscribers at a rate of 5,544 registrations per second. The test pushed the maximum memory usage on some of the nodes up to 82 percent, while maximum CPU utilization was short of 10 percent.

Call Setup Rate

Increasing call signalling transactions on the SPN cluster should, in theory, increase the load on its computational capacity. The purpose of this test was to measure the SPN cluster performance under a rate of 15,000 call attempts per second (Traffic Profile B in Table 1) and verify the calls' successful establishment and termination. The calls were evenly distributed by the traffic generator amongst the SPNs and comprised of incoming and outgoing calls, equally. We sent the call traffic while maintaining a rate of 3,333 new subscriber registrations per second – to ensure continued operations under load. The test ran for a duration of one hour.

Clearwater Core processed and signalled all of the 29 million calls successfully. We also verified that all 12 million subscription attempts completed without errors. The average call establishment time was 33.8 milliseconds and the maximum was 1.1 seconds. We measured a maximum of 78 percent of CPU load and 49 percent of memory utilization on the SPN nodes.

Concurrent Call Capacity

Metaswitch's design for Clearwater Core separates packet processing from state storage. Of the three node types, only the SCN is stateful – SPNs and DGNs are stateless.

When sessions get established on the SPN cluster, call information is sent towards any of the available SCNs. This test is meant to stress the feature's limitation – maintaining sessions concurrency at a high load.

We used Traffic Profile C in Table 1 for the purpose of this test. Clearwater Core processed all 10 million emulated SIP registrations with a single failure occurrence when memory utilization hit 82 percent on the SCN cluster. The error was for an INVITE request. Maximum CPU utilization was relatively low at 10 percent. We measured an average call establishment time of 17.3 milliseconds and a maximum of 260 milliseconds. Highest CPU utilization stood at 20 percent and memory at 42 percent across the nodes.

High Availability

It is critical that an IMS deployment is reliable during a node or site failure. Metaswitch designed their virtual IMS with redundancy in mind. In order to test their design, we emulated three failure scenarios, one per node type.

Each test run – including failure and recovery – was performed three times. We used a balanced traffic scenario (Traffic Profile D in Table 1) to emulate the network behavior during failure and recovery procedures. Table 3 contains the detailed resource utilization measured for each of the test scenario.

SPN Failure and Recovery

The SPN handles subscriber traffic and represents the first point of entry to the IMS core. Commonly, recovery from a failure in a subscriber facing entity can rely on the network's ability to stop routing traffic towards the failed node. Moreover, in a stateful SIP session signalling use case, it was important that other SPNs continued to handle sessions that had been previously established on the failed node.

To emulate this scenario, we reconfigured the traffic generator in two instances. The first sent traffic towards the SPN which was to fail, and the second towards the rest of the SPN cluster. Upon SPN instance failure, we confirmed that only the test traffic hitting the failed SPN was affected, while all other traffic was processed successfully.

In order to emulate the behaviour of a typical SIP peer, we redirected the SIP requests from the first traffic generator towards the remaining SPN nodes. These requests related to both existing and new SIP sessions. We confirmed that previously established SIP sessions could be terminated successfully, and new sessions could be established as the remaining SPN nodes in the cluster processed all SIP requests without errors.

DGN Failure and Recovery

The DGN handles subscriber authentication and communication towards the HSS. Any failure of this function could have a catastrophic effect on establishing or maintaining user sessions. Metaswitch claims that their DGN is fail safe, so we put that claim under the test.

We emulated network traffic until it maintained a steady state. We then terminated one of the DGNs and kept an eye on emulated traffic. After a 10 minutes period, we recovered that DGN node while maintaining the traffic stress.

The DGN recovered successfully without affecting any of the 9 million calls either whilst it had failed or when it recovered. We observed no CPU or memory spikes on other nodes in the cluster.

SCN Failure and Recovery

As the name indicates, the Storage Cluster Node (SCN) maintains the subscriber information in a database cluster. It uses Memcached to store subscriber information in a redundant fashion across all the SCNs. Memcached is a widely used open-source in-memory database. Metaswitch recommends a minimum of three SCN instances in a cluster for high availability. The Clearwater Core deployment that we tested had five nodes because of the number of subscribers being tested. The nodes interconnected via a single subnet.

Similar to the SPN and DGN, we failed one of the SCNs in the cluster while actively sending user traffic and recovered it after 10 minutes. We verified that the user traffic was not affected and all calls and registrations were processed successfully. No CPU or memory spikes were observed.

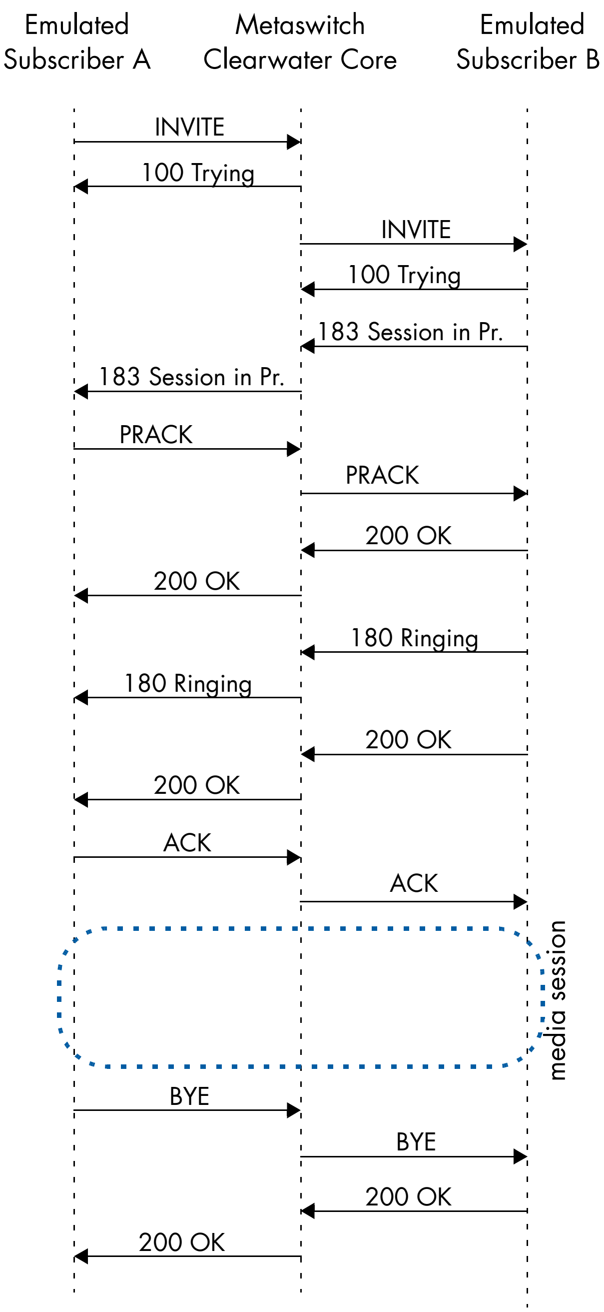

Figure 2: Call Scenario

Virtual IMS Core Orchestration

Cloud orchestration is a major driver of cloud model adoption. The large number of instances make it near impossible for legacy operations to cope. The European Telecommunications Standards Institute (ETSI), for example, has heavily modeled their NFV architecture around management and orchestration – usually referred to as MANO – in order to streamline NFV lifecycle operations.

In large deployments similar to this 95-node setup, a service provider would expect a clean instantiation, configuration and verification of each new instance, followed by scaling out or in, and termination of that instance.

Instantiation and Termination

Automation is key to managing scattered cloud resources. MANO tools often utilize service descriptors to correlate resources throughout the service's lifecycle. The descriptor would include the baseline configuration – such as IP addresses and licenses – while more complex configurations can be handled following the successful instantiation.

In this test, we asked Metaswitch to demonstrate their instantiation tools. Metaswitch explained that their Salinas tool connects to Amazon EC2 resource manager and:

- Updates DNS records to map newly allocated node IP addresses to the cluster configuration files
- Instantiates new nodes inside the clusters
- Pushes the configuration to a single Clearwater Core node (at this point, that node pushes the configuration to other nodes in the cluster– ensuring that no more than one node of each type is out of service at once)
- Validates the instance configuration

In case of termination, Salinas also:

- Terminates instance nodes
- Cleans up DNS records

We verified the above by monitoring instance creation and observing changes via the Amazon EC2 user interface. After instantiation we sent traffic through the newly-created Clearwater Core deployment successfully.

We also confirmed that Salinas terminated all nodes and cleared DNS records after termination by checking the Amazon EC2 user interface.

Scaling Out and Scaling In

One can argue that the biggest advantage of a cloud deployment is the ability to optimize the scale of the application according to demand. One indication of high demand is the resource utilization on any of the cluster nodes. Metaswitch explained that their intention was to make the deployment size as flexible as possible to allow real-time changes on the nodes. Configuration gets pushed from a single node towards other nodes, Metaswitch added. We executed this test using Metaswitch's Salinas tool.

The test started with a smaller deployment flavor of 40 SPNs, 37 DGNs and 5 SCNs as a baseline. The intention was to verify that the Clearwater Core can scale out under overload conditions, given that it had considerably less SPNs than the originally planned 53. We sent traffic through the deployment at a rate of 15,000 Calls/s and 3,333 Registrations/s (Traffic Profile B). We measured an average CPU of 82 percent on the SPN cluster as expected. In a production environment, this should be a relevant trigger for scaling out. In our test environment, we manually triggered the instantiation of one SPN node using Salinas and watched the average CPU go down to 75 percent. No calls were affected by the scale out operation.

To ensure that scaling in is also hitless, using Salinas we resized the scaled-out deployment back to its initial size while maintaining traffic emulation at the same rate. The SPN cluster's average CPU was measured at 79 percent. Only one call invitation – out of 33 million – received a 504 error response. However, it is worth mentioning that the scaled-in cluster was running at a higher traffic rate than it is intended to – for the purpose of this test case.

Clearwater Core Virtual IMS Node	Amazon EC2 Instance Flavor	Verified Instance Count	Guest Operating System	Software Version & Build
SIP Processing Node (SPN)	c5.large	53	Ubuntu Linux 14.04	v11.1 (CC-spn-11.1.0-180217-249.00.50)
Diameter Gateway Node (DGN)	c5.large	37	Ubuntu Linux 14.04	v11.1 (CC-dgn-11.1.0-180217-249.00.50)
Storage Cluster Node (SCN)	c5.4xlarge	5	Ubuntu Linux 14.04	v11.1 (CC-scn-11.1.0-180217-249.00.50)

Table 2: Hardware and Software Configurations

Test Scenario	SPN				DGN				SCN			
	CPU Utilization [%]		Memory Utilization [%]		CPU Utilization [%]		Memory Utilization [%]		CPU Utilization [%]		Memory Utilization [%]	
	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max
Subscriber Session Capacity	7	10	26	27	8	10	8	8	5	6	43	82
Call Setup Rate	67	78	31	33	13	16	8	8	5	6	26	49
Concurrent Calls Capacity	18	24	17	18	8	9	7	7	3	4	22	42
SPN Failover and Recovery	15	20	22	23	8	10	8	8	4	5	18	36
DGN Failover and Recovery	13	25	43	44	5	9	9	9	16	25	10	13
SCN Failover and Recovery	10	17	6	8	6	9	8	9	13	25	11	21

Table 3: Resource Utilization per Test Scenario

Further details on the configuration of specific Amazon EC2 machine specifications:

Instance Name	vCPUs	RAM	EBS Bandwidth	Network Bandwidth
c5.large	2	4 GiB	Up to 2.25 Gbps	Up to 10 Gbps
c5.4xlarge	16	32 GiB	2.25 Gbps	Up to 10 Gbps
c5.9xlarge	36	72 GiB	4.5 Gbps	10 Gbps

Table 4: Amazon EC2 Configuration Details

Each vCPU is a hardware hyperthread on a custom 3.0 GHz Intel Xeon Platinum 8000-series processor optimized for Amazon EC2.

Conclusion

Metaswitch Clearwater Core vIMS proved to be reliable in the cases of node failures. Its hitless failover and recovery would make it a dependable asset for network operators. Clearwater Core also navigated through three performance test scenarios with healthy results.

A cloud-based Clearwater Core deployment capable of supporting 20 million subscribers and 15,000 concurrent call attempts could be deployed in less than two hours.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/
[v6 20180515]