

# Cisco Quantum™ Virtualized Packet Core Functionality Test Report

## Introduction

As an independent test lab we are contracted by vendors to verify the existence of a certain function or to measure performance characteristics of a solution. Towards the end of 2013 Cisco approached us with a request to verify that their new solution, the Cisco Quantum™ virtualized Packet Core (QvPC), functions just as the physical incarnation of the StarOS software - the ASR 5000 family.

From a test lab perspective, the ultimate test methodology is the one that has been tested again and again. Since we have tested the same functions in the past (with more than one vendor) we knew that the methodology we use is solid and when Cisco added to the mix a physical device for us to compare the virtual functions with, we were compelled to respond to Cisco's request with a yes.

## Background

In 2010 we tested Cisco's ASR 5000 in a project that was so large it was fondly named MegaTest. The report was published on LightReading.com, receiving a great deal of attention from the industry. In 2010 the focus of our tests was still on functions relating to UMTS (or 3G as mobile operators billed these services) and the ASR 5000 was shown as a Serving GPRS Support Node (SGSN) and Gateway GPRS Support Node (GGSN).

Fast forward almost 4 years and the ASR 5000 has been further developed, the industry is abuzz with Long Term Evolution (LTE or 4G) deployments, and service providers are pushing their suppliers to develop versions of the same functions that could run on general purpose servers. The push for virtualizing network functions has gotten a great deal of interest from vendors and service providers, all working within The ETSI Industry Specification Group for Network Functions Virtualization (ISG NFV).

### Test Highlights

- **Virtual and physical solutions support the same functions and protocols**
- **Hardware and hypervisor independence**
- **Integrated services function as expected**

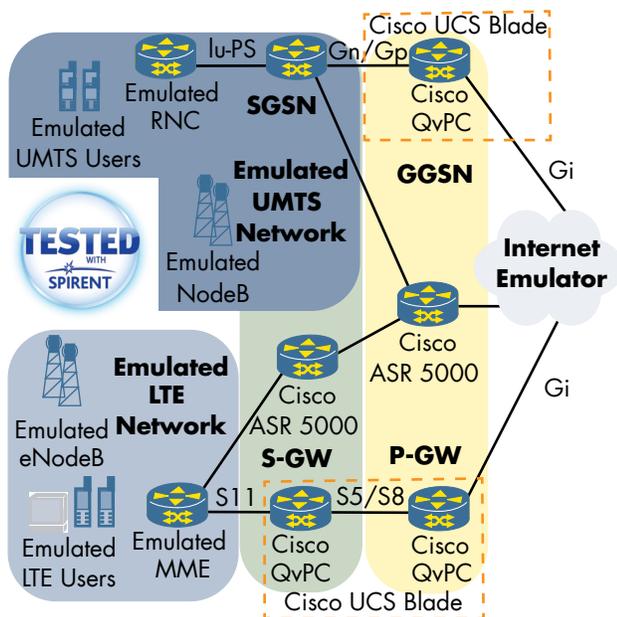
Announcements of virtual implementations of traditionally hardware-bound functions appear with higher frequency, but the evidence or successful deployment stories are rare. For this reason, we are pleased to provide the first independent verification of virtualized mobile core functions.

## Tested Devices

The setup provided by Cisco was actually a lab that is used with a specific tier-1 mobile service provider. The lab was already instrumented with a Spirent Landslide, a tester which we used in past comparable tests, so EANTC's engineers simply took over the lab from Cisco.

In addition to the tester, Cisco provided two UCS B200 M3 Blade Servers on which two hypervisors were residing: KVM's QEMU version 1.4.0 and VMware's ESXi 5.1.0. We selected several test cases to show that regardless of the hypervisor, the same QvPC functionality existed. More on this in the following sections.

Another device that was available for our testing was the ASR 5000. This device, deployed by many mobile service providers and field-proven according to Cisco, was our benchmark. We were told by Cisco that any function that worked on the ASR 5000 was also expected to work in QvPC, running on any of the hypervisors. Figure 1 depicts the logical and physical topology used in the tests.



**Figure 1: Cisco QvPC Logical Test Setup**

## User Interface

Most network engineers spend their work-time at a computer accessing various network elements. Rarely do they physically interact with the system. One can imagine, that if the network engineer does not even sense that the system that he or she are working on is virtual, the vendor implementing the system will have an easier time convincing the engineers to use it.

In order to convince ourselves that the system feels the same as the physical version, we started our verification by taking a copy of the configuration that existed in both ASR 5000 (one serving the P-GW role while the other the S-GW role) and copied them to two QvPC instances. The theory was that once QvPC reboots, the same calls that were successfully supported by the ASR 5000 configuration could be run through a system comprised only of the QvPC.

We quickly realized that in order to do exactly this (i.e. copy the configurations) we would have had to turn off both ASR 5000 since the same IP addresses would exist in the same network. Clearly, this was not going to work. So instead, we asked Cisco to configure the QvPC instances with unique IP addresses as well as create the hardware-relevant configuration, and copy from the ASR 5000 just the configuration that was relevant to the functions: S-GW and P-GW.

Once the configuration was loaded on both QvPC instances, we took the same tester configuration with which we ran LTE subscriber traffic through the ASR 5000, and repeated the same call flows. All calls came up and terminated successfully.

This is a good message to service providers. In essence, with some smart planning, a smooth migration path could be created between the physical and the virtual elements.

With this test we also had a complete virtualized system under test that supported LTE data calls. The next challenge was to verify that the operator will not notice that the system under test was a virtualized one. We figured that the best way to make this point was to execute commands on both virtual and physical systems, and verify if the two systems behave the same.

We collected 15 commands that we would use had we ran a mobile network. Examples included “show resource session,” “show ip interface summary,” “show p-gw service all” and more. We then executed all commands on both the ASR 5000 and QvPC systems and compared the output.

Both systems executed the same commands and provided the same output. It was impossible to tell the difference between the virtual and the physical system in normal operations. This was another good message for mobile service providers looking to transition to virtualized mobile core components: the move is not likely to tax the operational team responsible for the mobile gateways by forcing them to learn a new Command Line Interface (CLI).

## UMTS and LTE Call Flows

Obviously if a mobile service provider decides to replace their physical mobile core elements with a virtualized equivalents, the mobile subscribers should not be aware of the move. A move to virtual mobile core will clearly require performance and scalability verification for the same call flow models that are currently supported, but the scope of our validation was functional, so we tamed our desire to generate high call load and agreed on functional goals for the next tests.

Our call model included 10,000 subscribers, each using two bearers for a total of 20,000 bearers in all test cases.

The subscribers activated and deactivated their calls at a rate of 1,000 per second in both UMTS and LTE call models.

### **UMTS (3G) Call Flow**

While most of the marketing literature these days focuses on Long Term Evolution (LTE), UMTS (or 3G) networks are very much still alive and active. In this verification we aimed to make sure that the QvPC was able to support UMTS calls as a GGSN.

We used the Spirent Landslide to emulate most of the UMTS network components such as the Radio Network Controller (RNC) and the SGSN, as well as the 10,000 subscribers. Our plan was to first validate our call flow against a setup that had the GGSN role implemented in an ASR 5000 and then move the same call flow, without making any changes, to a setup in which the GGSN was the Cisco QvPC. We also asked Cisco to run the QvPC in two different hypervisors: VMware and KVM to make sure that the virtual implementation was hypervisor independent.

We executed the test three times with the same results regardless of which system combination was acting as the GGSN. The call flows obviously were established with the ASR 5000 without a problem and we were pleased to see the same behavior when we ran the Cisco QvPC in its VMware or KVM hypervisor hosts.

### **LTE (4G) Call Flow**

LTE networks are commonly advertised as 4G. From a subscriber point of view 4G means faster connectivity. The architecture of an LTE network was changed significantly from UMTS and with that, the components of the mobile core changed their roles and names. The SGSN function (and then some) is now being done by the Serving Gateway (S-GW) while the GGSN functions are now accomplished by the Packet Data Gateway (P-GW).

In this and the following test setups Cisco provided both S-GW and P-GW functions. We used Spirent Landslide tester to emulate the other elements: eNodeBs, MMEs and the Internet servers. This test actually ran four times: once in an all physical configuration; once with Cisco's QvPC running in VMware hypervisor; once with the system running in KVM hypervisor; and another run in which IPv6 was used.

As we already mentioned, the tests were meant to verify functionality, and were not looking into performance. Still, we set a baseline of traffic for all tests at 350Mbit/s spread across all subscribers. With the same configuration we reached this performance, after all 10,000 subscribers attached to the network, in all but one scenario: the scenario in which the Cisco QvPC was running in the KVM hypervisor. In this scenario, roughly 1/3 of the performance was reached (131Mbit/s).

We investigated the issue and saw no indication that the cause was in the QvPC. We monitored the same CPU utilization in all virtualized instances of the tests (32% – 34% CPU). Based on other testing experience that EANTC gathered last year from virtualized network functions, we knew that optimization of the interface between the hypervisor and the host was often required, so we decided to see if reducing the number of sockets the hypervisor needs to open will help. We used the Landslide tester option to activate "TCP persistent connections," which reduced the number of TCP sessions that the system had to support, and with this we recorded the same performance as in all previous test setups.

At this point we had enough evidence to support Cisco's claim that the virtual solution, regardless of the hypervisor in which it was running, was functioning in the same way as the physical system.

### **Handover Between ASR 5000 and QvPC**

It is entirely plausible that mobile networks will use both physical and virtual instances of network functions for a while to come. We simulated such a condition in which two S-GW instances existed: an ASR 5000 and a QvPC. The point of the test was to verify that when a subscriber initiates a handover, the ASR 5000 running as S-GW could handover all calls to the QvPC functioning as an S-GW.

We monitored if any subscribers or any packets were lost during the handover and found none. The user experience did not change on account of the handover between the physical and the virtual system. We repeated the same test running the QvPC in the KVM hypervisor and saw no issues there either.

## High Availability

The fact that a function is virtualized on a general purpose server, means that the system must take into account the potential that the host CPU could fail. Since Cisco's QvPC is the same code that runs on the ASR 5000-series packet core components, Cisco's solution for high availability should also exist here. Ironically, the solution is called Inter-Chassis Session Recovery (ICSR), a name that makes sense in the physical world, but not when the system does not use chassis.

To verify that indeed the QvPC could use the same mechanism as the ASR 5000 Packet Core, we used two Cisco UCS servers chassis - one was running the primary instances of the System Architecture Evolution Gateway (SAE-GW, the case where both S-GW and P-GW are running in the same system), and the second was running a backup SAE-GW. The idea was to activate our 10,000 subscribers, have them send data, and then pull the plug on the primary Cisco UCS server. We expected that while packets may be lost, no subscriber session will be lost. The test setup is depicted in Figure 2.

We executed the test three times to make sure that we get no outlier results and were pleased to see that the system was behaving in the same manner each time we run the test – no subscriber sessions were lost in any of the test runs, and the out of service time we recorded, 26 seconds, was constant across all three runs.

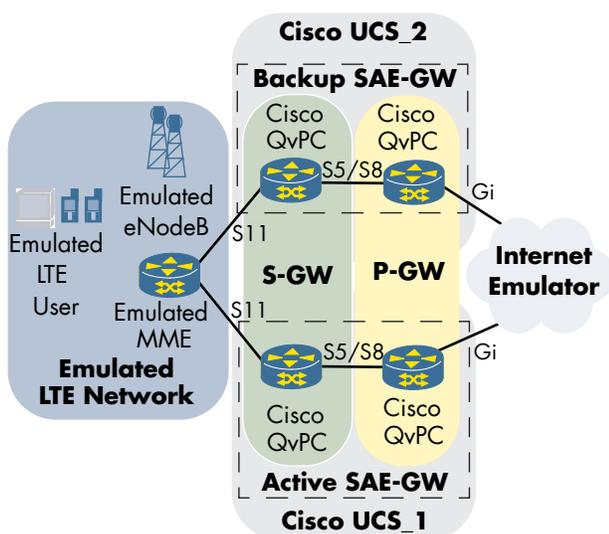


Figure 2: High Availability Test Setup

## Integrated Services

Last, but certainly not least, were three services that Cisco integrated into the QvPC itself: Firewall, Network Address Translation (NAT) and “Enhanced Charging Services” - a deep packet inspection (DPI) implementation

In the NAT test case Cisco configured two mapping rules: one to one and one to many. Both worked without a problem for TCP and UDP traffic that our subscribers were sending.

In the firewall test case, we saw the QvPC block both IP addresses and ports without a hitch and the DPI rules, which we verified by sending full-protocol stack over the subscriber bearers, filtered specific URLs, illegal characters, blocked specific user agent and counted RTP packets. All three test cases met our expectations.

## Summary

As mobile service providers are planning the next round of upgrades to their mobile core, we appreciate that a vendor is already confident enough in their implementation to expose it to independent validation, and to publicize the information to the industry. The Cisco Quantum™ Virtualized Packet Core seems to hold its promise to provide the same look and feel as the Cisco ASR 5000-series hardware-based counterpart. We confirmed Cisco's claims that the Cisco QvPC supports 3G and LTE (4G) calls identical to those run on the ASR 5000 and verified additional functions. It is conceivable that the next logical verification will have to focus on performance, and we are ready.

## About EANTC



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

EANTC AG  
 Salzuffer 14, 10587 Berlin, Germany  
[info@eantc.com](mailto:info@eantc.com), <http://www.eantc.com/>  
 4.0 201.40205, JG