network test

NFV Performance Benchmarking for vCPE

May 2015

OVERTURE

BROCADE

(intel)

SPIRENT
Communications

Integra

Executive Summary

Can Network Function Virtualization (NFV) provide cost savings and high performance at the service provider network edge, particularly when virtual customer premises equipment (vCPE) is involved? That's the question Overture Networks asked Network Test, an independent third-party test lab, to help answer.

Working with Overture and its partners Brocade, Intel, and Spirent Communications, and verified by service provider Integra, Network Test conducted a rigorous set of benchmarks to quantify NFV performance at the service provider network edge.

Among the key aspects of this project:

- The system under test was a pure-play virtualization solution deployed via orchestration onto a low-end commercial off-the-shelf (COTS) platform

- Tests involved measurements of Overture's Ensemble Carrier Ethernet (ECE) VNF by itself and in a service chain with one or more Brocade Vyatta 5600 vRouters

- Integra observed all tests and verified all results

- Overture's ECE switching VNF delivered wire-speed throughput in all tests

- Tests for both Overture and Brocade VNFs used a realistic mix of frame sizes, using a frame length distribution model based on Internet backbone traffic

- Maximum latency was typically measured in microseconds, much lower than the tens-of-milliseconds threshold where many enterprise applications experience degraded performance

This report presents test results along with sidebars explaining the importance of throughput and latency as performance metrics. The test methodology section gives details on key test parameters. Appendix B gives details on all software versions used in this test.
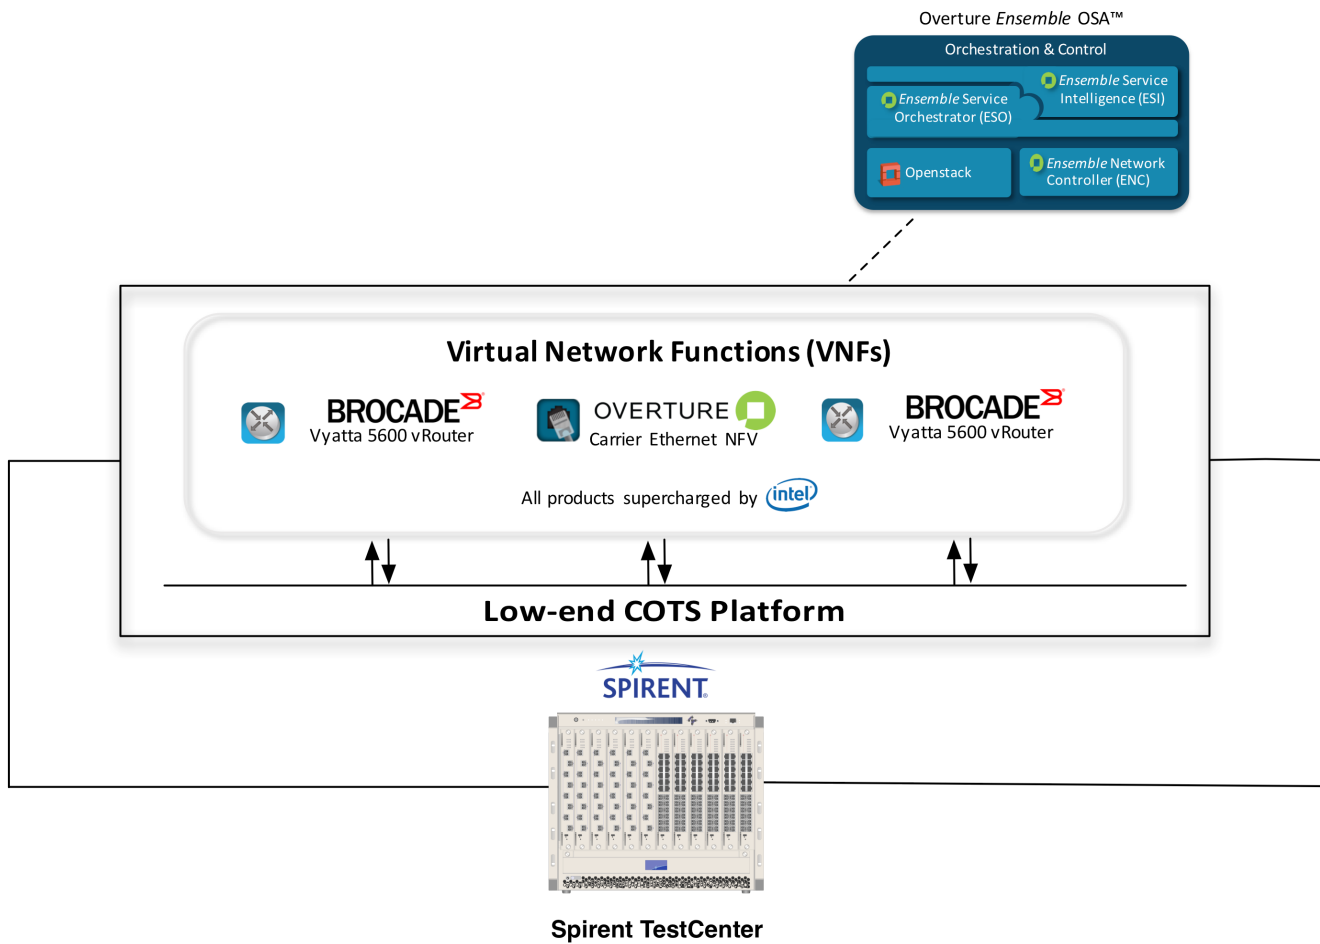
Figure 1: The NFV Benchmarking Test Bed

NFV Benchmark Results

Figure 1 illustrates the test bed, while Table 1 presents results of all performance tests. These NFV benchmarking covered three sets of tests:

- With an Overture ECE running on a COTS platform
- With an Overture ECE and one Brocade Vyatta 5600 vRouter VNF, also on a COTS platform
- With an Overture ECE and two Brocade Vyatta 5600 vRouter VNFs, also on a COTS platform

Frame size	ECE			ECE and 1 Brocade VNF			ECE and 2 Brocade VNFs		
	Throughput (Gbit/s)	% Line Rate	Max latency (usec)	Throughput (Gbit/s)	% Line Rate	Max latency (usec)	Throughput (Gbit/s)	% Line Rate	Max latency (usec)
64	1.52	100%	124	0.69	45%	1,056	0.36	23%	593
128	1.73	100%	143	1.18	68%	931	0.63	37%	704
256	1.86	100%	118	1.86	100%	509	1.17	63%	816
IMIX (362 avg)	1.90	100%	143	1.90	100%	418	1.52	80%	1,501
512	1.92	100%	174	1.92	100%	359	1.87	97%	1,397
1,024	1.96	100%	124	1.96	100%	399	1.96	100%	630
1,280	1.97	100%	120	1.97	100%	458	1.97	100%	825
1,518	1.97	100%	125	1.97	100%	476	1.97	100%	792

Table 1: NFV Performance Benchmark Results

It's important to note that all VNFs used the Data Plane Development Kit (DPDK), a set of libraries and NIC drivers for fast packet processing on Intel x86 processors. DPDK significantly improves performance, even on low-end COTS platforms.

Although not formally part of the test bed, engineers used Overture's Ensemble Service Orchestrator (ESO) system to deploy and manage NFVs.

In each set of tests, engineers configured a Spirent TestCenter traffic generator/analyzer to find the *throughput* rate and to measure *maximum latency* at that rate. (See the "Why Throughput Matters" and "Why Latency Matters" sidebars for brief explanations of why these metrics are such important predictors of application performance.)

Besides using various frame sizes ranging from 64 to 1,518 bytes (the smallest and largest permitted in Ethernet), engineers also tested with "IMIX," a combination of small, medium, and large frames commonly found in Internet backbone traffic.

Benchmarking Overture ECE

In tests with the Overture ECE alone, measured throughput ranged anywhere from 1.52 to 1.97 gigabits per second, depending on frame size (see "Test Methodology" for details). This is wire-speed performance for all test cases, regardless of frame size.

While at first glance throughput appears to increase with frame size, in fact the ECE delivered line-rate performance with all frame sizes. The difference in rates is explained by the interframe gap (IFG) and preamble that accompany every Ethernet frame.

Smaller frames have a substantially higher percentage of overhead from IFG and preamble. For example, even though a pair of interfaces is theoretically capable of carrying 2 Gbit/s of traffic, the actual maximum with 64-byte frames is 1.52 Gbit/s due to IFG and preamble overhead – exactly what the ECE achieved here.

To eliminate any confusion about throughput, Table 1 presents throughput as a rate, and also shows the percentage of line rate for each test result. The key takeaway is that the Overture ECE delivered wire-speed, lossless performance for all frame sizes, including IMIX.

As for delay, Overture ECE introduced maximum latency in the low hundreds of microseconds in all test cases (see "Why Latency Matters" sidebar). Given that enterprise applications may introduce delays in the tens of milliseconds or greater on gigabit networks, these results suggest the Overture ECE is unlikely to have an adverse impact on application performance.

Service Chaining: Adding Brocade VNFs

As noted, these tests also measured performance with one or more Brocade Vyatta router VNFs running on the same COTS platform as the Overture ECE.

Many factors can affect VNF performance, including host system CPU speed, available RAM, disk I/O, and hypervisor scheduling algorithms, among others. Differences in VNF performance may be attributable to these external factors, and not to the VNFs themselves. The tests described here treated the COTS platform as a "black box," and did not involve instrumentation of internal factors that can affect VNF performance.

Table 1 shows the results for tests with one and two Brocade routing VNFs. In tests with one routing VNF, the system delivered wire-speed performance for frames of 256 bytes and larger, including perfect performance with IMIX traffic.

The router's line-rate performance with 256-byte frames is significant for customers with large amounts of voice-over-IP (VoIP) traffic. This frame size is closest to the roughly 220-byte frames generated by the Real-time Transport Protocol (RTP) with G.711 encoding, by far the most popular type used for VoIP.

Maximum latency is higher in all cases with one or more VNFs, but this is expected. Latency is cumulative, meaning the delay added by each networking device contributes to the total. Further, the virtualization architecture dictates that every frame passes through the host server's kernel twice for each VNF, one apiece as it enters and exits the VNF. Thus, latency grows at a nonlinear rate with more VNFs.

Why Throughput Matters

When benchmarking network device performance, "throughput" has a different meaning than the colloquial use of the term. Here, it refers the highest rate a system can forward traffic *with zero frame loss.* In casual usage, throughput refers only to forwarding rate, regardless of whether loss occurs.

This more rigorous definition has important implications for application performance, and is arguably a better predictor of performance in production settings. A layer-2 switch or layer-3 router does not and cannot know whether the applications it carries can tolerate loss. In cases where upper-layer protocols can retransmit lost frames (TCP is by far the most common example), even a single dropped frame can still lead to degraded application performance.

Service provider networks are especially sensitive to dropped frames, since TCP can be very slow to recover from loss in high-bandwidth, high-latency situations. The safest way to ensure high performance for applications is not to drop frames in the first place.

Still, maximum latency is in the hundreds of *microseconds* for most frame sizes, whereas application performance typically are sensitive only to latencies in the dozens of *milliseconds.*

With two Brocade vRouter VNFs running on the same COTS platform in addition to the Overture ECE, throughput is lower, although the system continued to deliver wire-speed performance with frame sizes of 1,024 bytes and larger.

Maximum latency increases with three VNFs present (two Brocade vRouters and one Overture ECE), but again this is expected. In this configuration, each frame passes through the host server's kernel six times (once apiece going into and out of each VNF). Even in the very worst case – 1.5 milliseconds with IMIX traffic – maximum latency is below the threshold where it would have an impact on typical enterprise application traffic.

Test Methodology

The primary objectives of this project were to measure the throughput and latency of Overture Networks' Ensemble Carrier Ethernet (ECE) VNF, both alone and in conjunction with one or more Brocade Vyatta 5600 vRouters, all running on a commercial off-the-shelf (COTS) platform.

RFC 2544 describes an industry-standard methodology for network device throughput and latency measurement. In the context of these tests, "throughput" means the highest rate at which the system under test (SUT) forwards all traffic with zero frame loss. "Latency" means the delay introduced by the SUT. In this case, the tests measured latency of every frame using a first-bit-in, last-bit-out (LIFO) method as specified by RFC 1242 for store-and-forward devices.

The SUT comprised the Overture ECE running on a Supermicro SYS-5108A-FTN4 server equipped with a 2.5-GHz Atom CPU with 8 cores; 16 Gbytes of RAM; and a 4-port gigabit Ethernet NIC. In some tests, one or more Brocade Vyatta 5600 vRouters also ran on the SUT. In all tests, the ECE used 1 of the Atom CPU's 8 cores. Each Brocade vRouter used 2 cores.

Why Latency Matters

These tests measure maximum latency, or the greatest amount of delay introduced by the system under test. Latency is arguably an even more important metric than throughput in determining application performance. After all, forwarding rates may be much lower than the throughput rate in production networks, but a system will always introduce *some* amount of delay for every frame.

Overture specifically requested that this report use maximum latency in describing device delay. This is useful in three ways.

First, maximum latency is a worst-case number. Even under the most stressful possible conditions (forwarding all traffic at line rate with zero loss), the system will never introduce delay longer than the maximum latency. For software and network architects alike, this is a useful upper bound in predicting application performance.

Second, maximum latency helps in understanding device buffering capacity, which is especially important for delay-sensitive voice and video applications.

Finally, maximum latency also helps understand *jitter,* or latency variation, which also can have a significant impact on application performance. Jitter may be zero, near zero, or it may approach maximum latency – but it will never exceed that number.

The SUT used [Data Plane Development Kit (DPDK)](#), a set of libraries and NIC drivers for fast packet processing on Intel x86 processors.

The primary test instrument for this project was the [Spirent TestCenter](#) traffic generator/analyzer from Spirent Communications. Available in physical and virtual formats, Spirent TestCenter is capable of line-rate traffic generation and latency measurement with 10-nanosecond timestamping precision. Test engineers used a Spirent Test Center HyperMetrics SPT-11U chassis and an EDM-2003B gigabit Ethernet module to generate and receive traffic, offering traffic to two gigabit Ethernet physical interfaces on the device under test.

In an RFC 2544 throughput test, the test instrument uses a binary search to determine the throughput rate. Engineers configured the test tool to offer bidirectional streams of traffic to each of two ports on the SUT. The test instrument also can measure latency of every frame. As required by RFC 2544, latency is measured at, and only at, the throughput rate.

The duration for all tests was 60 seconds. Engineers configured the test tool to repeat the throughput test with the seven standard frame sizes given in RFC 2544. These are 64-, 128-, 256-, 512-, 1,024-, 1,280-, and 1,518-byte frames.

These tests also determined throughput and latency for "IMIX" traffic, which uses a combination of frame sizes commonly found in Internet backbone traffic. By default, Spirent TestCenter's IMIX load comprises a 7:4:1 ratio of 64-, 594-, and 1,518-byte frames, respectively. These Ethernet frame sizes in turn reflect 40-, 576-, and 1,500-byte IPv4 packets, as commonly observed on Internet backbones. With the IMIX load, the average frame size is 361.77 bytes.

Engineers repeated the tests with the SUT in three configurations: With Overture Networks' ECE, and with the ECE in concert with one and two Brocade vRouters running on the same Supermicro server. For all tests, engineers measured the throughput rate in gigabits per second, as well as maximum latency.

Conclusion

Network architects are constantly looking for ways to control costs and boost performance, and these have sometimes been conflicting goals. This is especially true at the service provider network edge, where historically providers have had to trade off performance for cost.

As these test results demonstrate, it's now possible deliver both low cost and high performance at the edge of service provider networks. In many test cases, a COTS platform running Overture's ECE switching VNF and Brocade's routing VNFs delivered perfect wire-speed throughput, with zero packet loss. In all test cases, maximum (not average) latency was below the point where it would have an impact on typical enterprise application performance.

Significantly, the Overture ECE switching VNF turned in perfect wire-speed performance in all test cases.

These tests also modeled two key aspects of conditions on live production networks. First, tests used "IMIX," reflecting the most common frame sizes found in enterprise backbone networks. Second, tests with 256-byte frames are useful predictors of performance for VoIP networks, since this is closest to the frame sizes used by VoIP traffic.

These results provide assurance to network professionals looking to deploy low-cost, high performance solutions at the service provider edge. Using technologies such as a low-end COTS platform and DPDK, it's now possible to deploy multiple VNFs at the network edge with little or no tradeoff in performance.

Appendix A: About Network Test

Network Test is an independent third-party test lab and engineering services consultancy. Our core competencies are performance, security, and conformance assessment of networking equipment and live networks. Our clients include equipment manufacturers, large enterprises, service providers, industry consortia, and trade publications.

Appendix B: Software Releases Tested

This appendix describes the software versions used on the test bed. Network Test conducted all benchmarks in April 2015 at Network Test's lab in Westlake Village, California, USA.

Component	Version
Overture Ensemble Carrier Ethernet (ECE) VNF	15.1.0.38
Brocade Vyatta 5600 vRouter	3.5R1
Spirent TestCenter	4.50.5906

Appendix C: Disclaimer

Network Test Inc. has made every attempt to ensure that all test procedures were conducted with the utmost precision and accuracy, but acknowledges that errors do occur. Network Test Inc. shall not be held liable for damages which may result for the use of information contained in this document. All trademarks mentioned in this document are property of their respective owners.

Network Test Inc.

31324 Via Colinas, Suite 113 • Westlake Village, CA 91362-6761 • USA
+1-818-889-0011 • http://networktest.com • info@networktest.com