SDNCentral Labs NFV Performance Test

Brocade Vyatta 5600 vRouter

NFV Routing and Security Performance Benchmark on Mid-Range Cloud Servers

October 2014

BROCADE

sdn central

The Independent Community and #1 Resource for SDN and NFV

Tables of Contents

1 EXECUTIVE SUMMARY

NFV (Network Functions Virtualization) promises to enable rapid innovation by liberating networking functions from proprietary hardware platforms. Core to the success of NFV is understanding the level of performance that can be achieved with software equivalents of hardware appliances. SDNCentral, with the help of its test lab partners, set out to validate the promise of NFV by conducting performance benchmark tests on real-world COTS (commercial off-the-shelf) mid-range server platforms. In this first of a series of reports commissioned by Brocade and conducted by SDNCentral, we will investigate the performance of Brocade Vyatta 5600 vRouter on x86 server platforms commonly found in public and private cloud environments.

The goals of this testing are to establish a baseline of the Brocade Vyatta 5600 vRouter running on a typical x86 COTS server platform, and to validate its L3 forwarding performance, route control-plane scalability, and firewall application performance in a setting relevant to data centers within enterprise, cloud service providers, and telecommunications providers.

Table 1 summarizes the Vyatta 5600 vRouter functions, scalability goals, and key highlights of the validation results on a 2-RU server fitted with dual 10-core Xeon processors (E5-2670v2@2.50GHz).

Validation Categories	Vyatta 5600 NFV Functions	Performance Goals	Finding and Highlights
L3 Forwarding Performance	• Core L3 dual-stack forwarding data plane	• Aggregate 80Gbps, 60 million packets per second L3 forwarding performance on modest COTS computing node • Maintain forwarding performance with IP diversity and up to 2 million flows in forwarding information base (FIB)	 • 70 million packets per second with 64-byte frames • 80Gbps bi-directional traffic with 256- to 1518-byte frame sizes • More than 75Gps of throughput under many scenarios including IMIX frame sizes and IPv4/IPv6 dual-stack • Performance maintained with up to 2 million flows

Route Scalability	• Service aggregation routing	• Routing information base (RIB) scalability up to millions of routes • Across one or many BGP peers	• Scalability up to 8 million routes from a single peer or up to 8 peers
Firewall Application Performance	• Stateful firewall • Stateless firewall (ACLs)	• Stateful connection handling in high-subscriber/user scenarios • High-fidelity application quality of experience (QoE)	**Stateless Firewall** • 200K sessions/sec • 2.5 million concurrent sessions **Stateful Firewall** • 50K sessions/sec • 1 million concurrent connections Application QoE maintained with stateless/ stateful firewall scenarios

Table 1. Summary of Test Results

In the tests conducted by SDNCentral testing lab partners, the Brocade Vyatta 5600 vRouter platform successfully demonstrated dual-stack IPv4 and IPv6 L3 forwarding performance that is comparable to many traditional hardware appliances under the scenarios tested. L3 forwarding for a dual-socket CPU compute node was measured with an impressive aggregate system performance of 70 million packets per second and 80Gbps rates over a wide set of IP conditions.

Furthermore, the validation successfully demonstrated that the Brocade Vyatta 5600 can support key routing and security NFV functions with route learning up to 8 million routes per server (4 million per Vyatta 5600 instance on single dual-socket compute node), as well as stateful connection handling of 1 million concurrent connections while maintaining a rate of 50K new sessions/second.

The validation results demonstrate that the Vyatta 5600 can deliver NFV cloud service aggregation routing and stateless/stateful firewall security functions at a high-level of performance on real-world servers commonly found within enterprise data centers, cloud service providers, and telecommunication carrier data centers.

2 VALIDATION PRECURSORS

2.1 NFV Overview

NFV is an industry term used to denote the capability of delivering network functions through virtualized compute infrastructure.

These network functions span capabilities that traditionally have been delivered using hardware-based networking products. Routing, security, and other network functions like load balancing can be implemented using virtual machines in single-VM or multi-VM service chained topologies. Service providers have keen interest in leveraging commercial off-the-shelf (COTS) computing hardware coupled with NFV software because the solutions can reduce capital expenditures for dedicated networking hardware. Carriers also have incentives to move towards NFV because of the agility and flexibility it promises. Extra operating cost reductions potentially can be gained by employing automation in multi-tenant orchestration cloud use cases.

The industry broadly views NFV as having applications in a diverse set of data center and transport network interconnect use cases. Carriers can use NFV to deliver services with increased depth and breadth, enabling them to broaden revenue opportunities. They also can reconfigure network infrastructure to satisfy elastic on-demand use cases with NFV service chains that combine network functions in different combinations per tenant, much like Lego building blocks. Meanwhile, data center and transport network operators can use NFV to solve shared cloud infrastructure multi-tenancy issues such as on-demand data center interconnects, tenant isolation/security risks, and noisy neighbor performance degradation.

Networking vendors and operators have many cloud orchestration and virtualization infrastructure options to choose from when deciding what and how to support and deploy NFV. These options include choices in cloud management stacks like Apache Cloud Stack or OpenStack, and hypervisors like KVM or Xen. Vendors and operators also must select compute platforms for high performance with choices on how to optimally assign compute resources to NFV instances. Networking performance in these diverse combinations varies widely.

The networking industry needs performance validation to help NFV stakeholders converge on carrier-viable combinations of hardware and NFV software using a mix of open-source and commercial solutions. Data center and transport service providers will not adopt NFV solutions without proof that it can be deployed at scale with the performance and reliability of dedicated networking hardware-based solutions.

Networking vendors must be able to demonstrate that dense, COTS computing hardware coupled with their NFV solutions within cloud orchestration and virtualization software stacks can exhibit performance in the same class as dedicated networking solutions. Performance tuning that yields optimal virtual infrastructure configurations and NFV topologies must be reported to help carriers build viable solution combinations.

2.2 Validation Goals

Brocade's Vyatta 5600 vRouter solution is targeted at NFV routing use cases with carriers, enterprise data centers, and cloud service providers. To demonstrate that high-end performance is achievable on common servers found in real-world deployments, as opposed to extremely high-end, cost-prohibitive server platforms, Brocade commissioned SDNCentral to drive a validation process that encompassed methodology definition, virtual/compute environment tuning, test execution and analysis, and result reporting. To perform the tests, SDNCentral partnered with one of its testing lab partners, Nitron Labs, LLC.

This initial validation round aims to demonstrate that commonly available COTS computing hardware coupled with the 5600 vRouter running on open-source cloud stacks are capable of achieving carrier-grade performance levels historically associated with dedicated hardware-based networking solutions.

The validation aimed to establish carrier-viable L3 forwarding performance up to 80Gbs and 60 million packets per second using a single dual-socket CPU compute node. This creates a baseline that demonstrates how a COTS computing platform can be used to cost effectively deliver NFV vRouting functions at scale. A normalized cost per gigabit analysis compared to an average of other industry solutions is included in Table 2 to characterize the potential for significant carrier CapEx reduction. To further demonstrate carrier readiness, L3 forwarding performance (data plane) was benchmarked across a range of IP traffic diversity including different frame sizes, emulated hosts, and IPv4/IPv6 ratios.

A secondary validation goal was to prove the Vyatta 5600 can deliver essential routing and security NFV functions with mid-range performance targets proportional to the overall vRouting instance size. The routing control-plane scalability was benchmarked by determining if route/peer scalability could meet practical carrier needs in the millions of routes range.

The stateful firewall function was validated using L4-7 application workloads to baseline TCP connection handling scalability and overall application quality of experience (QoE) throughput and response time metrics using different firewall rule set sizes.

The results of the validation showed that high-performance NFV routing can be achieved cost effectively using the Vyatta 5600 solution. The report presents the details of the compute environment configuration that yields these groundbreaking performance baselines in order for other parties to recreate the results.

2.3 Compute Environment Under Test

The compute environment under test (CEuT) consists of hardware and NFV components listed in Table 2.

CEuT Components	Model	Unit MSRP	Quantity	Cost
COTS Compute Hardware				
CPUs	Intel(R) Xeon(R) CPU E5-2670 v2 @ 2.50GHz	$1,500	2	$3,000
Memory	DDR3 1600 MHz Registered (buffered) 16GB	$160	4	$640
Hard Disk	1x 240GB SSD	$150	1	$150
10G Adaptors	Intel X520-DA2	$415	4	$1,660
Blade/Motherboard	Supermicro X9DRD-7LN4F(-JBOD)	$500	1	$500
Chassis/Power Supply		$550	1	$550
NFV vRouter				
Brocade Vyatta 5600 (v3.2)	BR-5655-PERP-VM	$6,995	2	$13990
Total				$20490

Table 2. Compute Environment under Test (CEuT) specifications with approximate pricing for compute hardware (current prices)

The test used a mid-tier, Intel E5 v2 Xeon platform with dual physical CPUs each having 10 physical cores, with each core rated at 2.5GHz clock speed. Intel has faster clock speed CPUs in the Xeon E5 v2 family, such as the E5-2687W v2, with 8 cores at 3.4GHz each, as well as other high-end models in the newer Xeon E5 v3 family. Additional performance most likely can be achieved by using these faster clock speed CPUs and will be considered for future phases of the Vyatta 5600 validation series.

Four Intel dual-port 10G network adaptors were installed into four available PCI-E 3.0x8 slots on the motherboard. A total of 64GB of memory was available on the compute node, but only 32MB was assigned to Vyatta 5600 instances.

2.3.1 Compute Setup to Harness Maximum Performance

This section and section 2.3.2 provide a brief overview of the tuning required to maximize system performance.

We used open source software for the virtual configuration. For the host operating system, we used Linux CentOS 7.0 with the KVM/QEMU hypervisor installed.

Two settings are required to garner maximum system performance in a virtualized setting. First, Intel Virtualization Technology for Directed I/O (VT-d) was enabled in the BIOS before the host OS boots. Second, a Linux grub configuration addition is required to ensure Input/Output Memory Management Unit (IOMMU) capability is enabled within this host OS.

The Vyatta 5600 requires hugepages to be setup in the host OS, which in turn requires another Linux grub configuration addition to provision a portion of the system memory for hugepages. The memory assigned to the two Vyatta 5600 instances must not exceed the total hugepages quota provisioned in the host OS.

2.3.2 Vyatta 5600 Virtual Setup

A block diagram of the physical CPU to PCI-E and memory interfaces of the compute platform is shown in Figure 1. For maximum performance, two Vyatta 5600 instances were used to handle traffic for each of the two physical CPUs.

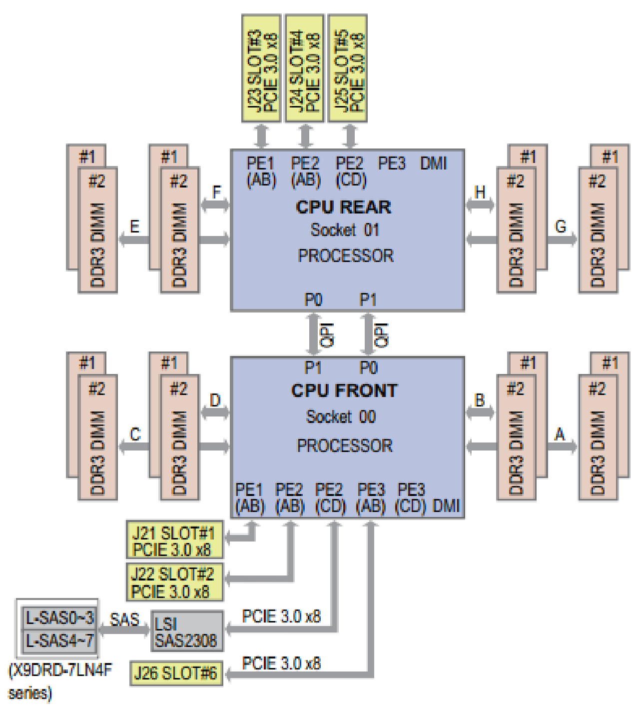

Figure 1. Compute Node CPU/Memory/PCI-E Interface Block Diagram

We took care to ensure the two dual-port NICs used for each virtual instance were evenly spread across the two physical CPUs. This meant one Vyatta 5600 instance would handle traffic for ports in PCI-E slots 1/2, and another instance would handle ports in PCI-E slots 3/4, as shown in Figure 1.

Four 10G PCI passthrough interfaces were added to each of the Vyatta 5600 instances. These allow the Vyatta 5600 instances to directly manage the NIC interfaces with minimum host OS interaction. In order for the KVM libvirtd daemon to load the PCI passthrough interfaces, the NICs had to be detached from the host OS kernel before the instances boot up. This was required with the CentOS Linux distribution and may not be required with Debian-based Linux distributions.

The PCI passthrough devices were added using the Linux virtual machine manager user interface, as shown in Figure 2.

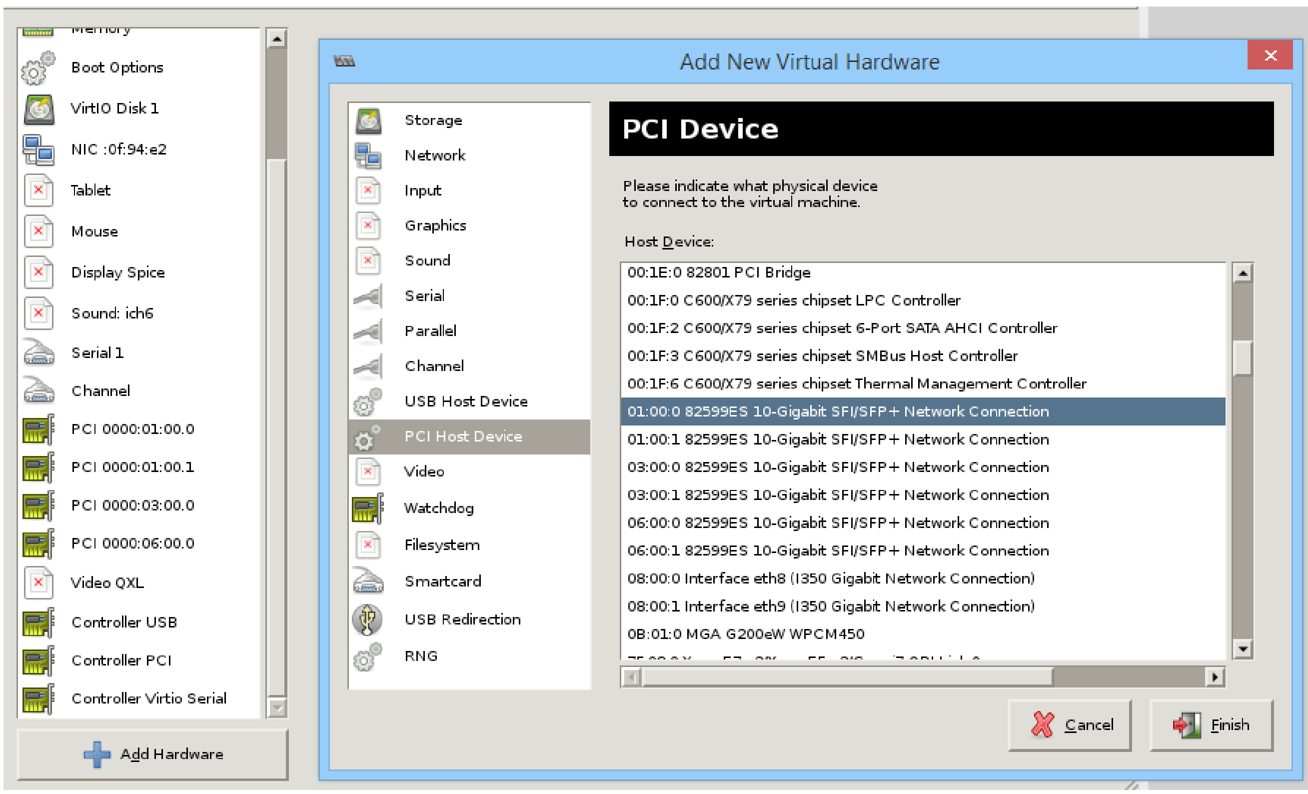

Figure 2. Linux virtual machine manager PCI passthrough configuration

Each Vyatta 5600 instance also was configured to make use of hugepages using the *libvirt virsh* command line tool.

Intel architecture supports hyper-threading, which provides two logical cores for each physical core. This provides a total of 40 logical cores (2 physical CPUs x 10 cores x 2 logical) with this particular compute platform available for assignment to virtual instances. The final key configuration to maximizing performance was to ensure each of two Vyatta 5600 instances used logical CPU cores that mapped to the physical CPU/PCI-E slot they were linked to without overlap.

Each Vyatta 5600 instance was assigned 16GB of memory and 16 logical cores.

2.4 Test Harness

For the test harness, Spirent Communications, a leading test and measurement vendor, agreed to collaborate with SDNCentral and Brocade. Tests were performed on Spirent's new Virtual Private Test Cloud (VPTC) environment. The recently unveiled Spirent VPTC service is designed for cloud technology vendors and network operators to validate NFV and cloud stack components without costly and time-consuming cloud infrastructure/test harness buildout. Spirent aims to support Continuous Cloud Validation (CCV) by empowering cloud stakeholders to automate and eliminate the uncertainty inherent in software-defined data centers due to their multi-layered complexity. The Spirent VPTC allows users to make use of a public cloud infrastructure to orchestrate real compute/virtual infrastructure and Spirent's test product breadth alongside NFV components under test.

We uploaded the Vyatta 5600 NFV software and remotely provisioned different configurations on the infrastructure, then assessed performance under each permutation.

We employed Spirent TestCenter L2/3 functionality to generate deterministic loads up to 80Gbps with all the IP traffic diversity we desired to assess L3 forwarding performance. The VPTC environment afforded us the use of a high-performance switch to create separate bridging domains to split traffic from two 40G interfaces into the eight separate 10G interfaces attached to the compute environment under test. We used Spirent TestCenter's BGP emulation capabilities to assess the routing control-plane scalability of the Vyatta 5600 solution.

We used Spirent's Avalanche product line to generate L4-7 application workloads and assess the performance of the Vyatta 5600 stateless/stateful firewall NFV security functions. Stateful connection handling scalability and QoE application metrics were collected using a mixture of common application workloads.

Test Harness Components	Model	Software Version
Hardware		
Spirent N4U Chassis	SPT-N4U	4.44.4246
SPIRENT DX2 8-PORT 40/10GBE QSFP+	DX2-40G-Q8	4.44.4246
SPIRENT MX2 2-PORT 40/10GBE QSFP+	MX2-40G-Q2	4.44.4246
Spirent C100	SPT-C100-MP-3	4.43.4136
Software		
Spirent TestCenter	Spirent TestCenter	4.44.4246
Spirent Avalanche	Spirent Avalanche	4.43.1126

Table 3. Spirent test hardware and software specifications used for the validation

3 TEST METHODOLOGY & RESULTS

Three major test categories comprise the methodology used to validate the Vyatta 5600 NFV performance including:

- L3 forwarding performance (RFC2544-based)
- Routing control-plane performance
- Firewall application performance

The overall methodology employed a progression through the test categories outlined above. L3 packet forwarding performance was established with the minimum functions enabled on the Vyatta 5600 instances. These forwarding rates were then used as maximums for benchmarking the routing control plane. The stateful firewall function was enabled for the final firewall/application performance test category, as it requires stateful TCP workloads.

The following sections summarize the category methodology and the measured results.

CEuT will be used to denote the dual-Vyatta 5600 instances and the compute environment outlined in section 2.3 that are essentially being tested as a system.

3.1 L3 Forwarding Performance

The goal of the L3 forwarding performance test category was to yield the highest possible aggregate IP forwarding rates across the eight 10G physical ports in the CEuT. This required all the CEuT tuning steps outlined in section 2.3.1 to arrive at these results. The frame-loss test from RFC2544 was used to determine forwarding rates and frame loss across the desired spectrum of traffic emulation and Vyatta 5600 functions.

Iteration was employed rerunning the RFC2544 tests across emulated IP traffic diversity conditions to determine if performance consistency could be maintained under various NFV deployment scenarios.

Test Topology

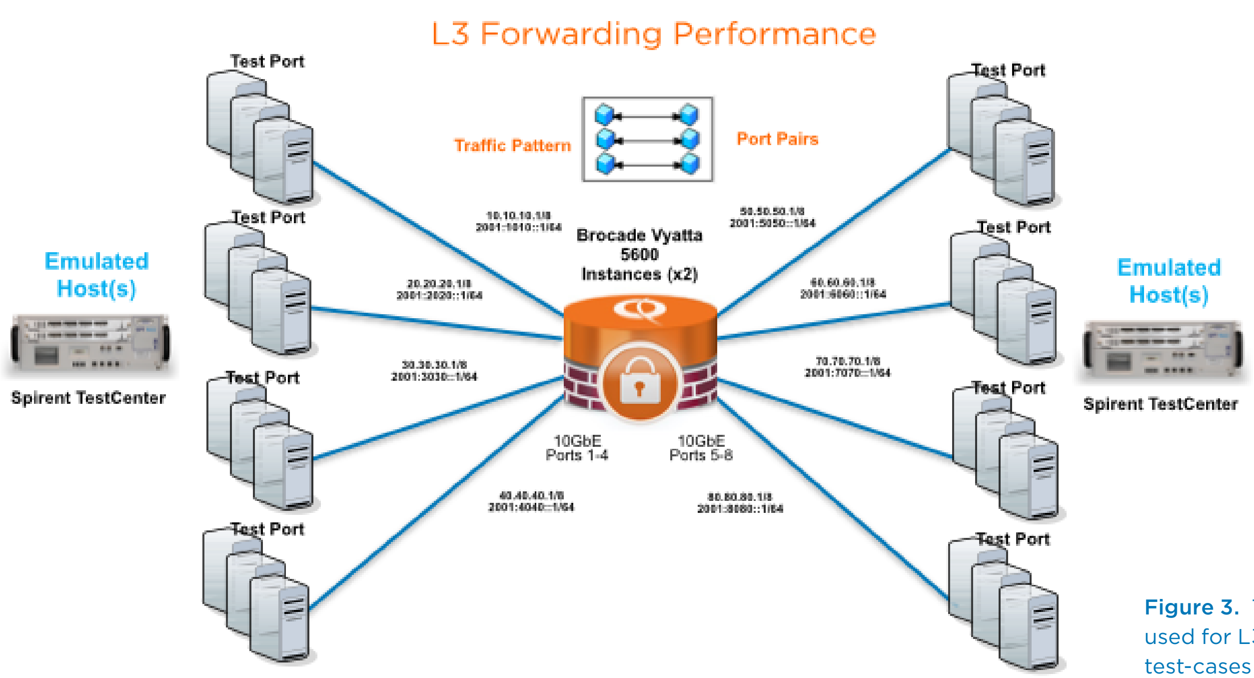

Figure 3. Topology used for L3 forwarding test-cases

The emulation used a varying number of emulated devices or hosts on each of eight network subnets configured on the dual Vyatta 5600 instances. The Vyatta 5600 was configured for dual-stack IPv4 and IPv6 operation. Emulated hosts were configured in different IPv4 and IPv6 ratios to evaluate whether the Vyatta 5600 is suitable for carrier IPv4/v6 co-existence deployments enabling carriers in the effort to become fully IPv6 ready. IMIX (distributions of commonly observed Internet frame sizes) test variants were included to validate that performance could be maintained with a more realistic set of mixed frames sizes.

The traffic pattern employed for all iterations was four port pairs per Vyatta 5600 instance. Other traffic patterns may yield different results but for the purposes of demonstrating maximum performance, a pair-wise pattern was selected. Table 3 includes some of the result highlights across all the iterations executed.

L2/3 Emulation				Aggregate Forwarding Rates (Mpps)	Aggregate Forwarding Rates (Gbps)	Frame Loss (%)	vDUT Functions
Emulated Hosts	IPv4/v6 Ratio	Frame Size (bytes)	Offered Load (%)				
80	100/0	64	60	60.934	40.947	14.693	L3 forwarding
80	100/0	64	70	66.533	44.71	20.161	L3 forwarding
80	100/0	64	80	69.689	46.831	26.827	L3 forwarding
80	100/0	64	90	69.711	46.846	34.936	L3 forwarding
80	100/0	64	100	69.871	46.954	41.308	L3 forwarding
80	100/0	IMIX1	100	31.531	77.97	2.537	L3 forwarding
80	100/0	IMIX2	90	20.58	71.999	0	L3 forwarding
80	100/0	IMIX2	100	22.371	78.265	2.168	L3 forwarding
80	100/0	IMIX3	90	19.861	72	0	L3 forwarding
80	100/0	IMIX3	100	21.564	78.171	2.286	L3 forwarding
80	100/0	128	100	60.223	71.304	10.87	L3 forwarding
80	100/0	256	100	36.232	79.999	0.001	L3 forwarding
80	100/0	512	100	18.797	79.999	0.001	L3 forwarding
80	100/0	1,024	100	9.579	80	0	L3 forwarding
80	100/0	1,280	100	7.692	80	0	L3 forwarding
80	100/0	1,518	100	6.502	80	0	L3 forwarding
80	0/100	IMIX3	60	8.444	47.964	0.075	L3 forwarding
80	0/100	IMIX3	70	9.85	55.95	0.09	L3 forwarding
80	0/100	IMIX3	80	11.257	63.945	0.086	L3 forwarding
80	0/100	IMIX3	90	12.622	71.694	0.425	L3 forwarding
80	0/100	IMIX3	100	13.664	77.616	2.98	L3 forwarding
160	50/50	IMIX3	60	8.447	47.983	0.036	L3 forwarding

160	50/50	IMIX3	70	9.854	55.972	0.049	L3 forwarding
160	50/50	IMIX3	80	11.262	63.971	0.045	L3 forwarding
160	50/50	IMIX3	90	12.608	71.615	0.535	L3 forwarding
160	50/50	IMIX3	100	13.649	77.532	3.085	L3 forwarding
2M flows	100/0	64	100	65.894	44.281	44.649	L3 forwarding
2M flows	100/0	128	100	53.579	63.438	20.703	L3 forwarding
2M flows	100/0	256	100	36.074	79.651	0.436	L3 forwarding
2M flows	100/0	512	100	18.797	79.998	0.002	L3 forwarding
2M flows	100/0	1,024	100	9.577	79.983	0.021	L3 forwarding
2M flows	100/0	1,280	100	7.69	79.977	0.029	L3 forwarding
2M flows	100/0	1,518	99.87	6.494	79.896	0	L3 forwarding

Table 4. L3 forwarding performance highlights across iterations

Close to full line rate throughput can be achieved with fixed frame sizes between 256-1518 byte frame sizes for both IPv4 and IPv6 traffic. The Vyatta 5600 can handle near line rate with IMIX frame sizes distributions, as shown in the figure below right.

AGGREGATE PORT L1 Rx RATE **AGGREGATE PORT L1 Rx RATE**

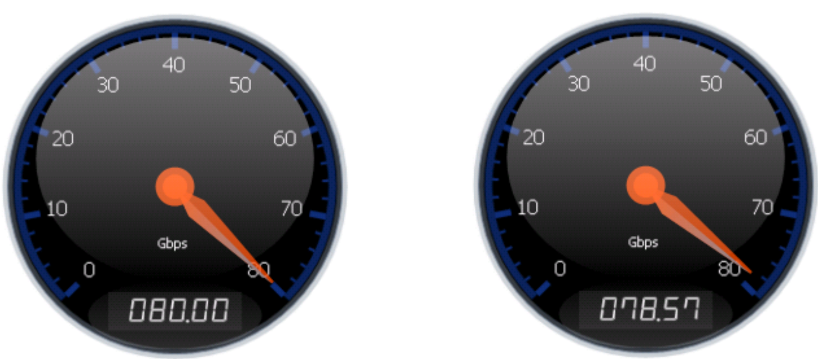

Figure 4. Left - 80Gbps for frame sizes between 256-1518 bytes (< 0.5% frame loss at 100% utilization);
Right - IMIX fowarding rates at 100% port utilization

Maximum aggregate forwarding rates of up to 70 million packets per second with 64 byte frame sizes were observed and shown below. We were consistently able to observe sustained aggregate forwarding rates of 65 million packets per second over the course of tests running for eight hours.

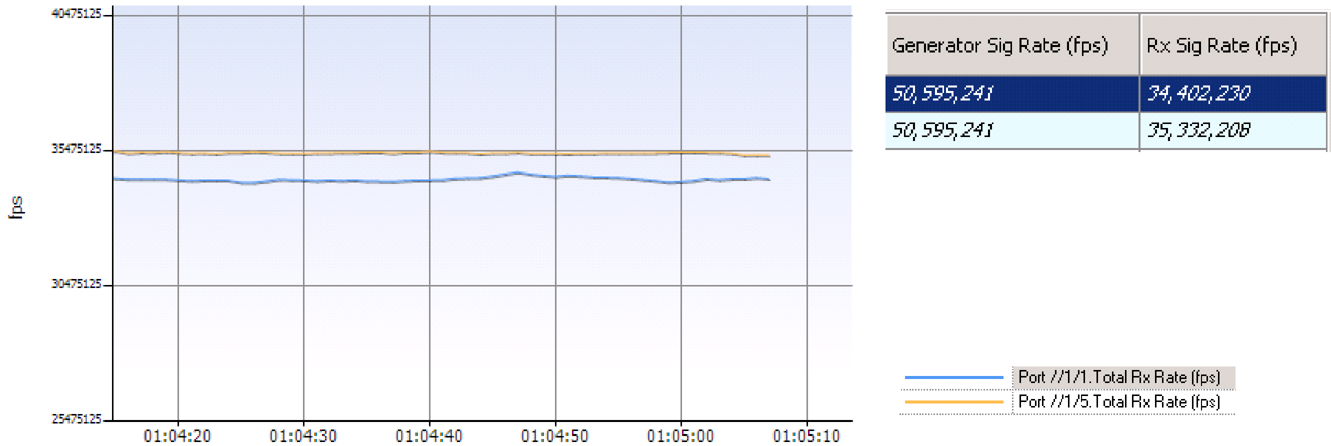

Figure 5. 70 million aggregate frames per second (fps) bi-directional L3 forwarding performance achieved using 64 byte frames at 85% port utilization

The results of the L3 forwarding performance test category demonstrate groundbreaking performance on a current, but modest COTS computing platform. Higher frequency CPUs will likely unlock even higher performance levels than measured in this report, as many of the logical cores were 100% utilized during heavy loads with small frame sizes.

Performance was maintained using IPv6 only and dual-stack IPv4/v6 scenarios between frame sizes of 256-1518 bytes. Test runs with IMIX frame size distributions (IMIX1, IMIX2, IMIX3 defined in section 5.2.1.1) yielded high aggregate forwarding rates at ~78Gbps and would be more representative of real-world conditions.

The forwarding information base (FIB) was stressed with up to 2 million (1 million per Vyatta 5600 instance) IP prefixes or overall flows. Testing was performed for IPv4 up to 2 million flows, and BGP emulation was employed for this phase of the validation with an emulated peer on each interface of the Vyatta 5600 instances. Route sizes of 100K and 250K routes were advertised from emulated routers with 1-to-1 mappings between routes on port pairs to create unique flows.

With 80Gbps forwarding rates under many scenarios and packet forwarding rates close to 70 million packets per second, the Vyatta 5600 with COTS computing hardware could provide cost-effective NFV vRouting for carriers.

3.2 Routing Control-Plane Performance

The goal of this test category is to determine the Vyatta 5600 NFV vRouting control-plane performance using a series of BGP route/peering scalability test scenarios. The methodology is by no means comprehensive around testing all aspects of routing performance. However, it does establish a baseline for NFV routing performance levels. Further validation of dynamic route flapping performance may be performed in future validation phases.

The methodology approach in this test category is to find the route scalability of the Vyatta 5600 from a single EBGP peer and then validate that route scalability can be maintained when routes are advertised from a wider set of emulated peers.

Test Topology

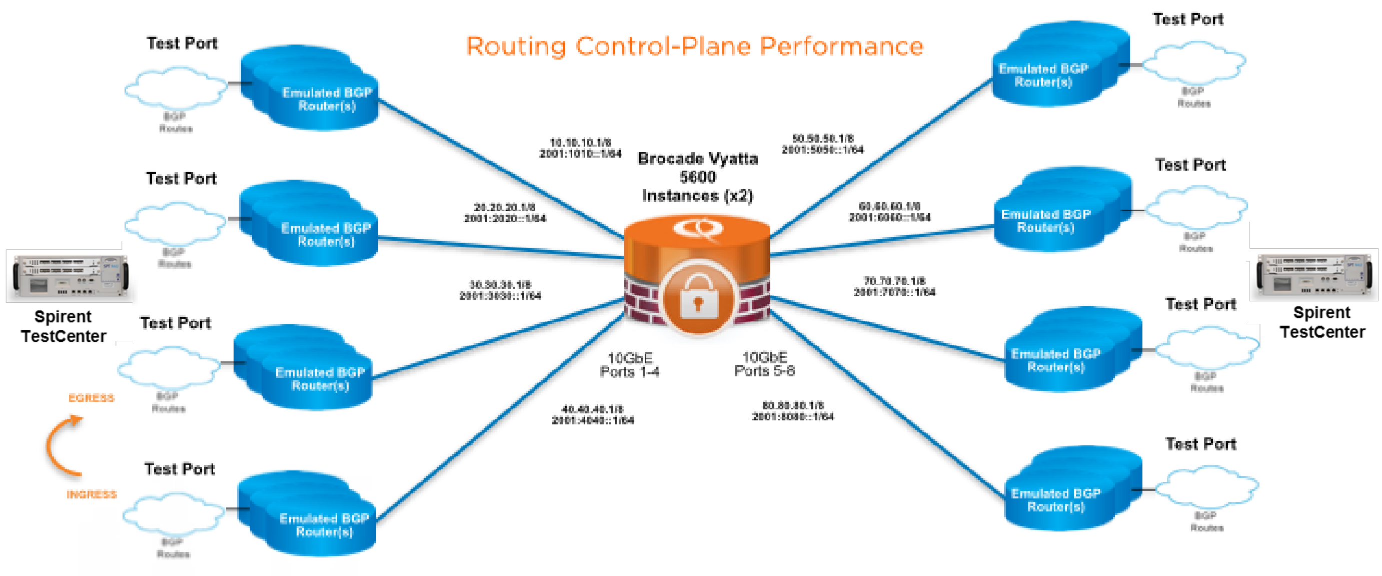

Figure 6. Topology used for routing control-plane performance test-cases

Distinct peers are configured on each of the test ports and establish peering sessions with the Vyatta 5600 instances. Different numbers of routes are advertised to the Vyatta 5600 from the emulated EBGP peers, and routes are advertised simultaneously on both of the Vyatta 5600 instances. Traffic is sent from ingress ports with destinations matching all of the routes advertised by emulated peers. Route learning by the Vyatta 5600 is validated by transmitting data-plane traffic below maximum forwarding rates found in the L3 forwarding test category for a given frame size and checking that traffic arrives at egress ports with no frame loss. The RFC2544 frame loss test was used for this purpose with one ingress port transmitting to all route prefixes advertised from an egress port. This ensures the learned routes are sustained in the Vyatta 5600's routing tables.

The table below includes some of the result highlights across the routing scalability test cases:

Routing & L2/3 Emulation				Routes Learned	Frame Loss %	vDUT Functions
Routes per Peer	EBGP Peers	IPv4/v6 Route Ratio	Frame Size (bytes)			
500K	2	100/0	256	1M	0	EBGP
4M	2	100/0	256	8M	0	EBGP
10K	8	100/0	256	80K	0	EBGP
100K	8	100/0	256	800K	0	EBGP
250K	8	100/0	256	2M	0	EBGP
500K	8	100/0	256	4M	0	EBGP
1M	8	100/0	256	8M	0	EBGP

* Routes learned were also verified using the Vyatta 5600 CLI in addition to data-plane verification

Table 5. Routing control-plane performance across iterations

Route scalability was verified up to 8 million routes (4 million per Vyatta 5600 instance) using a single EBGP peer per Vyatta 5600 instance, or using up to 8 peers per instance. Zero frame loss was observed when ingress ports sent traffic to egress ports with destination addresses that matched the entire set of advertised routes.

This test category showed that the Vyatta 5600 has the capacity to learn millions of routes and establish and maintain sessions with many BGP peers. The testing also demonstrates that the Vyatta 5600's routing information base (RIB) can hold a large number of routes. While mixed routing protocols scenarios were not tested in this validation, the Vyatta 5600 does support a breadth of interior and exterior routing protocols. Carriers can deploy the Vyatta 5600 flexibility as virtual routing NFV instances in cloud management platforms to offer agility to their customers without performance trade-offs.

3.3 Firewall Application Performance

The goal of the firewall application performance set of tests is to demonstrate the Vyatta 5600 vRouter's versatility in providing the NFV stateless/stateful firewall functions and its ability to deliver high-fidelity application QoE.

Standard firewall (RFC3511-based) benchmark methodology was used to evaluate the TCP connection scalability of the Vyatta 5600 with real application workloads. A stateful firewall must keep track of TCP state during connection establishment in order to block illegitimate traffic from flowing from secure public interfaces to private internal interfaces.

We used a connection capacity test case to measure the total number of open or concurrent connections the stateful firewall can maintain with application session transactions occurring. The rate at which the stateful firewall can process connections with application transactions is important to verify, as many carrier use cases have subscriber traffic that exhibits sudden spikes in connection rates, such as users logging in all at once in the morning or web search storms that occur during news-making events. A peak rate test case is used to measure the maximum rate of connections/transactions the stateful firewall can handle per second.

The final test case benchmarks the overall application performance of the stateful firewall from an end user perspective. It measures the aggregate network traffic rates for application traffic and average application response times for object transfers (web page views, file upload/downloads). The test category consists of three test cases including:

- Application connection capacity
- Application peak connection rate
- Mixed-application QoE

Test Topology

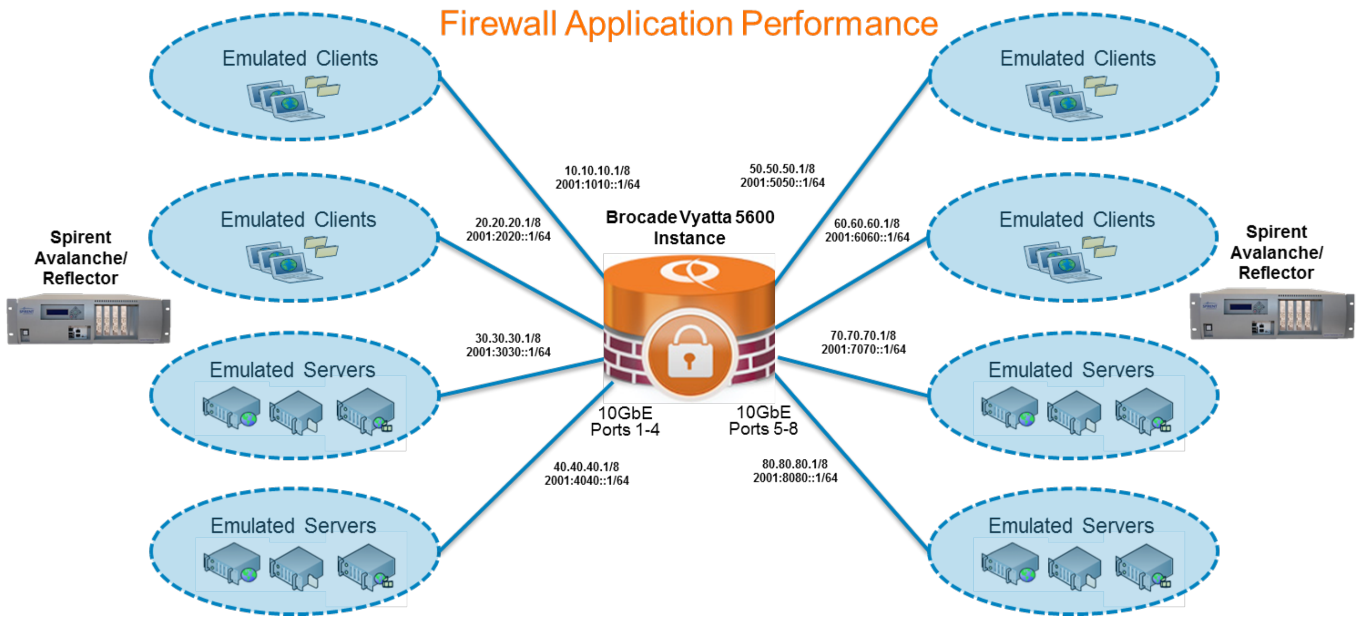

Figure 7. Topology used for firewall application performance test-cases

Emulated clients and servers are configured symmetrically on the eight test interfaces of the Vyatta 5600 instances. HTTP, FTP, and streaming (RTSP) application protocols are configured on clients and servers and used to perform upload/download transactions across the Vyatta 5600 instances.

Four of the Vyatta 5600 ports have firewall rule sets associated to them to represent interfaces on the public network. The Vyatta 5600 global stateful firewall NFV function was enabled for TCP and UDP. ACCEPT rules were configured to allow private-to-public connections to flow through the firewall, while DROP rules were set up to filter on IP addresses for some of the configured traffic. An implied global drop rule blocks all TCP connections originating from public interfaces. To mimic the conditions in the real world, each test employed background traffic sent continuously during measurement periods that the router had to block, while letting through valid traffic.

The three test cases were run with the stateless firewall enabled, then the stateful firewall enabled to establish performance baselines for each. The stateful firewall was provisioned with increasing rule set sizes to characterize the effect firewall rules have on performance.

The table below includes some of the result highlights for the application connection scalability test cases.

L4-7 Application Workload Emulation		Concurrent Connections	Peak Connection Rate (connections/ transactions per second)	vDUT Functions
Workload Protocols	**Load Profile**			
HTTP	Connections	2.5M	N/A	Stateless FW
HTTP	Connections	1M	N/A	Stateful FW
HTTP, FTP	Connections/sec	N/A	200K	Stateless FW
HTTP, FTP	Connections/sec	N/A	50K	Stateful FW

Table 6. Application TCP connection scalability test-cases

The stateless firewall has no issues sustaining up to 2.5 million concurrent TCP connections. The stateful firewall can maintain 1 million open TCP connections before new connections cannot be birthed reliably.

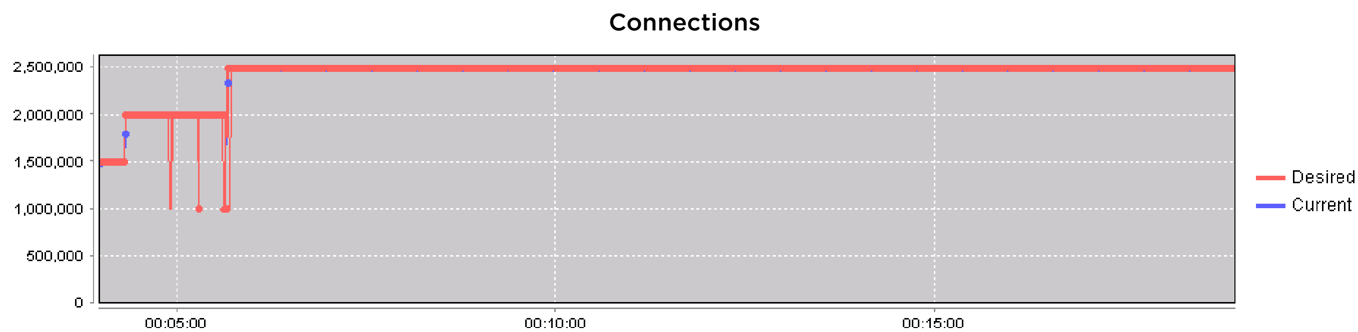

Figure 8. Connection Capacity - Stateless Firewall Function Enabled – 500 rule set size

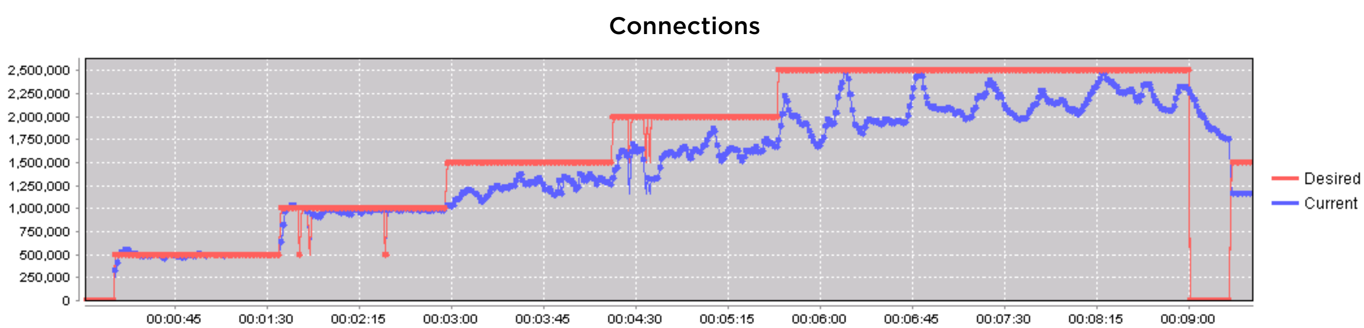

Figure 9. Connection Capacity - Stateful Firewall Function Enabled – 500 rule set size

The charts below demonstrate that the stateless firewall can sustain peak connection rates up to 200K connections/second, while the stateful firewall can maintain 50K connections/second before additional rate increments cannot be sustained reliably.

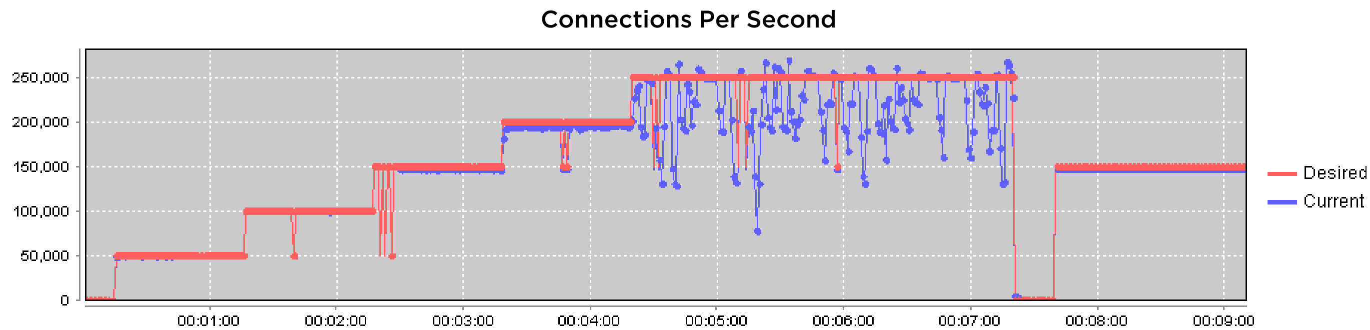

Figure 10. Peak Rate – Stateless Firewall Function Enabled – 500 rule set size

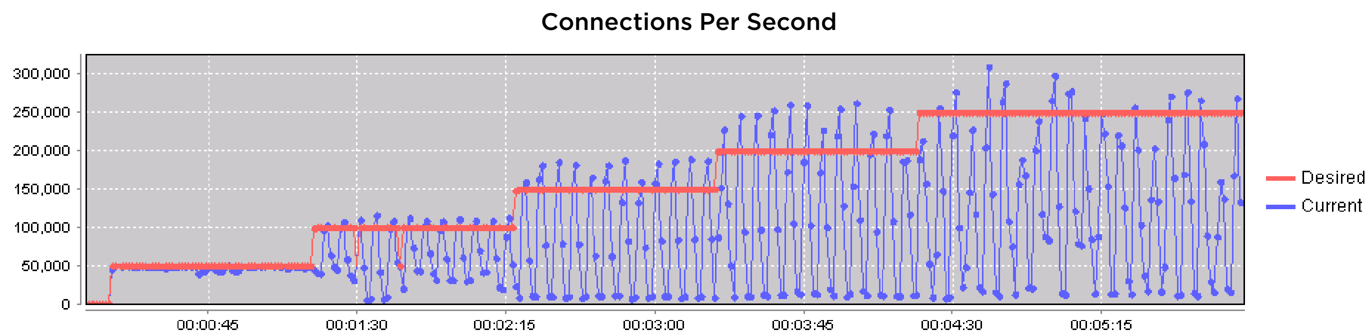

Figure 11. Peak Rate - Stateful Firewall Function Enabled – 500 rule set size

The Vyatta 5600 is able to deliver strong experience metrics between back-to-back baselines and with the stateful firewall enabled. The delta in network traffic rates between a back-to-back baseline and with the stateful firewall differs by only 6.9%. The average response times for application protocols increase marginally from the back-to-back baselines and are negligible, most likely not noticeable to end users.

L4-7 Application Workload Emulation				**Peak Aggregate Network Throughput**	**vDUT Functions**
Workload Protocols	**Upload Object Size (MB)**	**Download Object Sizes (MB)**	**Load Profile**		
HTTP, FTP, RTSP	10	7, 1, 1.225	SimUsers/sec	47.5 Gbps	Baseline (back-to-back)
HTTP, FTP, RTSP	10	7, 1, 1.225	SimUsers/sec	45 Gbps	Stateless FW
HTTP, FTP, RTSP	10	7, 1, 1.225	SimUsers/sec	44.2 Gbps	Stateful FW

Table 7. Peak aggregate, bi-directional network traffic rates for application QoE test-cases

	Average Response Time (ms)			
	Baseline	Stateless FW	Stateful FW	Max Delta
HTTP	12206	12276	12470	264
RTSP	41632	41646	41649	17
FTP	4520	5673	5902	1382

Table 8. Average per protocol response times for application QoE test-cases

The results of the firewall application performance test category show the Vyatta 5600 is able to perform the NFV virtual security function with high-performance stateful connection handling. With performance of up to 50K connections/second in stateful firewall operation, it should be able to handle heavy web transaction loads during peak usage intervals without putting strain on cloud infrastructure. It also can be used to maintain network perimeter security boundaries with large numbers of ACLs or to harden cloud infrastructure with stateful firewall functions for increased security against many stateless (network) but effective distributed denial of service (DDoS) attacks.

The Vyatta 5600 can be used in NFV service chains where it provides high-performance stateful connection handling in front of load balancing functions that reside in front of web, application, and database tiers – without being the bottleneck in the overall chain.

High-fidelity application QoE can be maintained with minimum TCP retransmits and high L4-7 application throughput when users are downloading or uploading large object sizes. The Vyatta 5600 has a marginal impact to application response times versus back-to-back baselines.

4 CONCLUSION

Server vendors now are shipping server platforms with next-generation Intel Haswell processors (E5-2600 v3 series). The cost and power consumption of the mid-range servers with Intel Ivy-Bridge processors (E5-2600 v2 series) put them in the sweet spot for cloud and NFV deployment. By choosing a typical cloud service provider 2-RU server with two E5-2670v2 processors, 64GB of RAM with 1600MHz for the validation, we were able to achieve remarkable results on a commonly available server platform.

The overall results of the validation demonstrate that in our test topology, up to 80Gbps and 65 million packets/second aggregate L3 forwarding rates can be sustained on a typical midrange x86 server (single dual-CPU compute node) using the Brocade Vyatta 5600 as a virtual router. We were pleasantly surprised by the outcome because these rates were achieved using a modest server with open source virtualization (KVM) and no special bare metal OS customizations. Results showed the Vyatta 5600 is NFV-ready and can be integrated into open-source cloud management stacks more easily than dedicated networking hardware solutions.

The measured performance levels of the Vyatta 5600 are normally seen only on dedicated ASIC/FPGA-based hardware or bare metal designs with highly optimized software tuned to the hardware. And while hardware-based systems might show stronger or more consistent performance across a larger set of scenarios, the results of this test demonstrate that a COTS platform with an NFV vRouter like the Vyatta 5600 is a worthy contender in replacing existing hardware routers across many use cases. And we expect vRouters on server platforms to continue to improve their performance as innovations in CPU packet processing continue and software architectures mature.

In addition, the test demonstrated the suitability of the Vyatta 5600 for meeting the needs of emerging VNF security functions in NFV deployments, with solid connection-handling performance and minimal effect on the application quality of experience.

The Vyatta 5600 solution has many benefits, including the ability to scale elastically for carrier cloud deployments, and the ease with which virtual instances and compute nodes can be added on demand. The solution also gives the enterprises and service providers flexible control over how to partition available compute resources for each tenant or subscriber.

4.1 Cost-Performance Analysis

To characterize the potential cost savings to carriers using the Vyatta 5600 running on COTS computing, we did an analysis using pricing for equivalent mid-range dedicated networking hardware routing products from other top industry vendors. We selected comparison classes of products suitable for data center interconnect and provider edge routing deployment scenarios that have 80Gbps forwarding capacities with similar IP service and routing functionality.

This analysis aims to show the rough orders of magnitude in potential cost savings to service providers. The comparison class includes products capable of packet forwarding at or beyond 60 Mpps with forwarding information base (FIB) scalability of 1 million entries and routing scalability of at least 4 million IP prefixes.

We averaged the pricing of products in the comparison class and arrived at an industry average of $75K. The actual selling prices will of course vary. We encourage readers to substitute their own average pricing estimates since comparing feature sets, performance (under customer-specific deployment scenarios), and pricing can be difficult. Likewise, performance under different scenarios and in different topologies can yield different results and readers should take care to evaluate the performance of vRouters in an environment that closely mirrors their specific use case.

Nevertheless, using the system cost of the compute environment under test used in this validation and tabulated in Table 2 (shown earlier in the document), which tallies to approximately $20K (includes compute costs plus the list price of two Vyatta 5600 licenses), we find the Vyatta 5600 vRouter could be up to 3.75x (~73%) less expensive than an industry average of $75K for comparable hardware-based solutions.

A final caveat: We did not validate against all traffic patterns, deployment scenarios, and validation dimensions. Some of the comparison products have proven deployment maturity, so this should be taken into account in a more detailed analysis. Even adjusting for these other factors, the price-performance value of the Vyatta 5600 running on COTS compute, coupled with its flexibility to deliver NFV functions, cannot be ignored by cost-conscious enterprises, cloud service providers, and carriers looking to reduce expenditures and increase revenue opportunities.

SDNCentral Labs, LLC
955 Benecia Avenue
Sunnyvale, CA 94085
USA
www.sdncentral.com

The Independent Community and #1 Resource for SDN and NFV