

## Lab Testing Summary Report

September 2014

Report 140920

Product Category:

**100GE  
Data Center  
Switch**

Vendor Tested:

**Cisco  
Systems, Inc.**

Product Tested:

**Nexus 7718**




### Key findings and conclusions:

- Testing found the Cisco Nexus 7718 can forward at line rate on all 192 of its 100GE ports – delivering over 38 Terabits/s of bidirectional traffic
- Testing confirmed the Cisco 7718 can distribute real-world IPv4 and IPv6 multicast traffic at wire speed, with each of 191 receiver ports handling 1,250 IGMPv2 groups
- The Cisco 7718 can process real-world MPLS traffic at line rate on all 192 of its 100GE ports, with no loss and low latency
- Testing confirmed that an active Supervisor module or fabric module can be replaced with no packet loss, with IPv4 and v6 traffic running at high capacity on all of its 100GE ports
- The Cisco 7718 executed an in-service software upgrade, with IPv4 and v6 traffic running at high capacity on all 192 of its 100GE ports, with no loss

Cisco engaged Miercom to independently evaluate the Cisco Nexus 7718 – among the most powerful data center switches on the market, which Cisco unveiled last year. Specifically, Miercom was to test-drive the 18-slot, L2 and L3 switch, applying realistic traffic loads and checking IPv4 and v6 throughput and latency, and assess the system's support for MPLS (Multiprotocol Label Switching), as well as its array of fail-safe, non-stop survivability features.

"When we tested the Cisco Nexus 7718 we were impressed by the incredible power and throughput of this system, coupled with consistent low latency and latency variation, as well as solid MPLS support and

**Figure 1: Nexus 7718 Throughput with No Loss  
Measured Throughput**



Source: Miercom, September 2014

**No packet loss.** Cisco Nexus 7718 exhibited maximum throughput of 19.157 Tbps while forwarding 9,216-byte frames from the Ixia test system. This graphic from Ixia IxNetwork software shows that the transmit (Tx) and receive (Rx) figures match.

high-availability features.” said Robert Smithers, CEO of Miercom. “The Cisco Nexus 7718 proves Cisco’s commitment for the Data Center market. This highly resilient switch, high performing, low latency, extremely robust product proved superior in testing and it was an honor to test it.”

## About the Cisco Nexus 7718 Switch

The Cisco Nexus 7718 is a Layer 2 and Layer 3 Ethernet switch designed for large data center deployments and organizations with core networks operating at 40 and 100 Gigabits per second.

The 18-slot Nexus 7718 represents the high end of the Cisco Nexus 7700 Series. There also are six- and ten-slot models in the line.

The Cisco Nexus 7718 offers the highest density of 100GE ports of any switch on the market today, up to 192 ports. Alternately, the switch can support up to 768 x 1GE or 10GE ports or 384 x 40GE ports via different interface modules.


For our evaluation the Cisco 7718 was fully loaded, with two Nexus 7700 Supervisor 2E modules and 16 Cisco Nexus 7700 F3-Series I/O modules. Each 7700 F3-100G module provided 12 x 100GE ports, collectively offering 192 x 100GE ports.

Based on a 'Switch on a Chip' (SoC) design, the F3 module series supports a feature set that combines

the best of the M and F series of I/O modules. From the F2 series modules there is Fibre Channel over Ethernet (FCoE) and Cisco FabricPath support. Supported features from the M series include Multiprotocol Label Switching (MPLS), Virtual Private LAN Services (VPLS), Overlay Transport Virtualization (OTV) and Location/ID Separation Protocol (LISP).



**The 18-slot Nexus 7718.** There are 16 payload slots for interface modules – eight above and eight below the two slots in the middle for redundant Supervisor modules. Up to 16 power supplies plug in at the bottom. Fabric modules plug in from the other side.



**Test Bed.** The device under test, in this case the 26-RU (45.5-inch-high) Cisco Nexus 7718, is in the bottom of the middle rack. Three of the five Ixia XG12 test-equipment chassis are in the rack on the left. Another is above the switch and the fifth is on the lower right. All performance testing for this review employed the Ixia test system, which connected 192 test ports directly to the Nexus 7718's 192 x 100GE ports.

The two Cisco Nexus 7700 Supervisor 2E modules operate redundantly. One is in active mode, handling control-plane and management functions, and providing centralized arbitration to the switch fabric for all line cards. The second Supervisor operates in hot-standby mode.

All of the Nexus 7000 switch models use a passive mid-plane architecture, providing physical connectors for interconnecting the fabric modules and I/O modules. All inter-module switching is performed via crossbar fabric ASICs on the individual fabric and I/O modules.

A scalable, fully distributed fabric architecture uses up to six fabric modules, each delivering up to 1.32 Tbps per slot of switching capacity. The mid-plane design supports flexibility and incremental upgrades of switching capacity.


All switches in the Nexus 7700 Series have true front-to-back airflow, from perforations in the front of the I/O modules and supervisor cards. This allows air to be drawn in through the front of the chassis, along the heat sinks and chips on the cards, and exhausted out the rear.

## Unicast Performance

In testing the Cisco Nexus 7718, traffic was all legitimate IP packets. Since the switch operates in store-and-forward mode, it examines the whole integrity of each packet and checks the forwarding table before appropriately routing each packet.

The voluminous test traffic required for this testing was generated by a battery of five Ixia XG12 test systems, which collectively delivered 100 Gbps of traffic to each of the 192 x 100GE ports of the Cisco Nexus 7718.


**Figure 2: Max Throughput with No Loss**



Source: Miercom, September 2014

The connections between the test system and the switch were full-duplex and test traffic flowed bi-directionally. Traffic was delivered to the switch, forwarded, and then returned back to the Ixia test gear, which meticulously checked for any loss and calculated packet latency – the time it took the packet to traverse the switch.

The Nexus 7718 could accept and forward all the IP traffic sent, using a pattern of mixed frame sizes simulating real-world traffic, called IMIX, at line speed with no loss.



Source: Ixia IxNetwork

**Unicast.** Throughput tests were run from port-to-port – that is, all traffic delivered on one port was forwarded to the same egress port across the switching fabric. Traffic in the opposite direction – since throughput testing was bi-directional -- followed the same return path through the switch.



The IMIX traffic streams that were used simulate Internet traffic that passes through a switch in a real-world network. Separate tests were run with Layer 2, IPv4 and IPv6 traffic.

The Layer 2 and IPv4 IMIX test traffic stream consisted of 64-, 570- and 1,518-byte packets. The distribution of each packet size was 7-4-1, respectively. Just slightly different, the IPv6 IMIX traffic consisted of 78-, 570- and 1,518-byte packets, with the same 7-4-1 distribution.

For the unicast throughput tests, port-pair combinations were assigned so that bi-directional test traffic was transmitted between line-card ports, across the fabric modules. This is done in accordance with the Internet specification for switch throughput measurement, RFC 2544.

## Latency

Latency – the time it takes a packet to traverse the switching system, in the absence of congestion and buffering – was also measured.

Once the maximum forwarding rate is ascertained, latency through the switch is then calculated by subtracting the transmit time stamp from the receive time stamp. Based on traffic tests that run for five minutes, three latency values are reported: minimum, maximum and average.

The Cisco Nexus 7718 operates in store-and-forward mode. As a result, the switch needs time to receive and validate the entire packet before forwarding it to the appropriate port to reach the destination address.

Latency values reported here do not include the time needed for the Cisco Nexus 7718 to receive and validate packets. Because the switch operates in store-and-forward mode, latency measurements use the LIFO (Last In, First Out) technique for calculation.

*Figure 3* shows minimum, average and maximum latency for unicast IMIX traffic through the Nexus 7718.

A latency metric just as important for switching systems is latency variance, also called jitter. This is the statistical difference in latency between packets of the same size during a test. There should not be much variance in latency, even at high traffic loads. The consequence could be packets being sent by the switch out of order, which can seriously affect the integrity of some data streams – especially real-time traffic like Voice-over-IP.

Our tests revealed that the Nexus 7718 induced remarkably little latency variance on passing data packets. Looking at the difference between average and maximum latency – a good gauge of variance – only a few microseconds separated the average and maximum latencies for the short packet size shown in *Figure 3*.

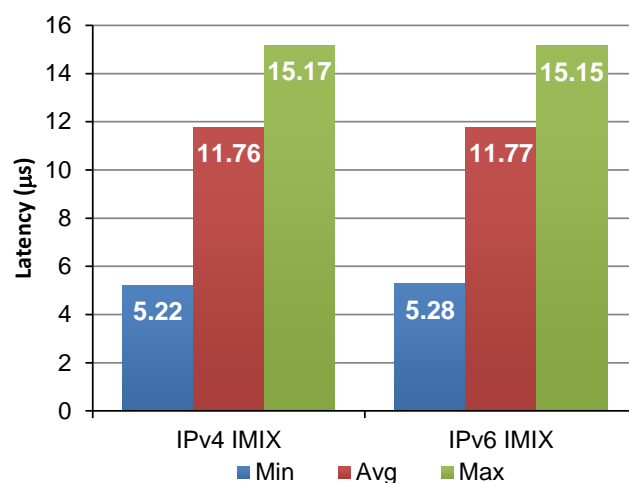
## Multicast Performance

IP multicast, where a router-switch replicates an IP packet stream to multiple receivers, is an efficient way of achieving one-to-many or many-to-many communications. Measuring multicast throughput and latency was done in accordance with Internet RFC 3918, which defines the processes for benchmarking IP multicast performance.

For our multicast testing the switch was configured to distribute the packet stream from one input port to 191 egress ports. The same IMIX distribution of packet sizes in the stream was used as in the unicast testing. As before, separate tests were run for IPv4 and then IPv6.

For IPv4 multicast, the 191 receiving ports all sent IGMPv2 joins and joined the same 1,250 IPv4 multicast groups. The total number of outgoing IPv4 interfaces (OIFs), then, was **238,750 (191 x 1,250)**. For IPv6 multicast testing, the 191 outgoing ports sent MLDv2 joins and joined 150 IPv6 multicast groups. The total number of outgoing IPv6 interfaces (OIFs), then, was **28,650 (191 x 150)**. These are noteworthy numbers and this high control plane

**Figure 3: Latency in Microseconds**



Source: Miercom, September 2014

scale is enabled by the powerful Supervisor 2E modules which have a two-quad-core Xeon processor with 32 GB of memory.

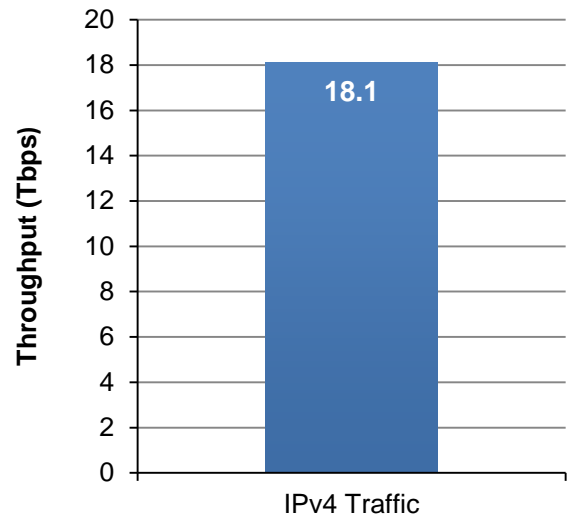
Impressively, we found that the Nexus 7718 could accept a line-rate, 100-Gbps, IMIX input stream – for IPv4, as well as IPv6 – and successfully replicate and distribute it to every other 191 ports, without losing a single packet.

Another aspect of multicast performance is latency. The replication and distribution process is demanding and most would expect considerably more latency on multicast packets than unicast. That turned out not to be the case with the 7718.

Given the mix of packet sizes in the IMIX distribution, per-packet latency values can be misleading. Still the results are impressive. Minimum IPv4 IMIX packet latency: 6.7 microseconds, and 10.7 microseconds average; nearly the same for IPv6: 6.7 microseconds minimum and 11.7 microseconds average.

Both IPv4 and IPv6 have protocol mechanisms for defining and joining multicasting groups, although they are subtly different. An IP multicast group address is used by sources and receivers to send and receive multicast messages. Sources use the group address as the IP destination address for multicast packets. Receivers use the group address to inform the network of interest in receiving packets sent to the group.

**Figure 4: Multicast Throughput**

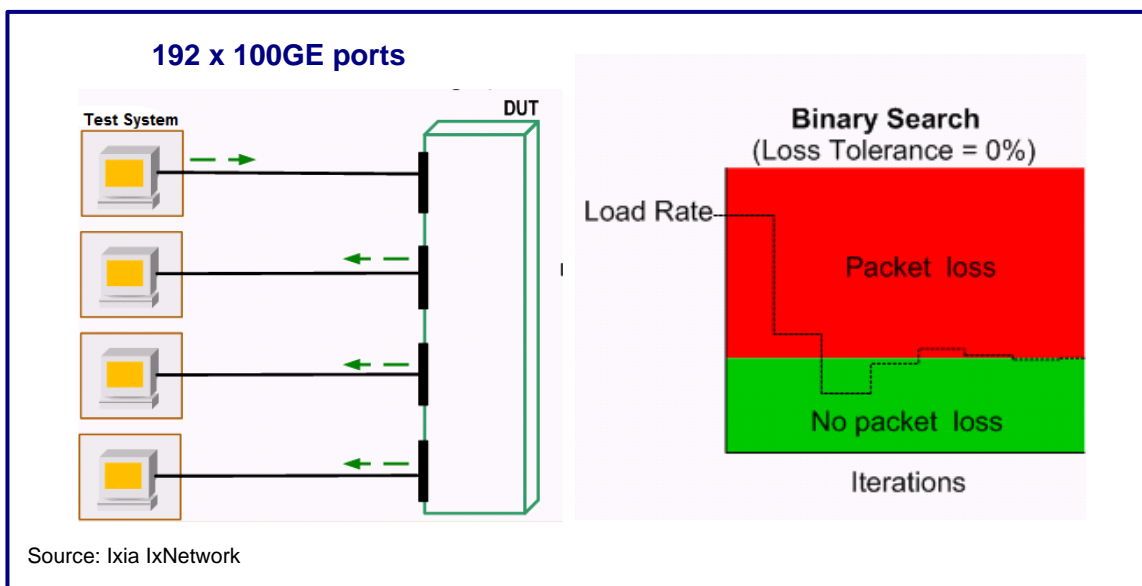


Source: Miercom, September 2014

## Multiprotocol Label Switching

MPLS is a high-performance packet-forwarding protocol that integrates the efficiency and traffic management of data-link-layer (Layer 2) switching, with the scalability and flexibility of network-layer (Layer 3) routing.

MPLS wraps almost any kind of protocol or service in labels, which are quickly appended and stripped off by MPLS routers as user data moves through a network. It enables organizations to quickly implement differentiated services, such as Layer 3



Source: Ixia IxNetwork

**Multicast.** Unlike with unicast, in multicast testing a single, uni-directional, high-volume traffic stream is delivered to one port, and is then replicated by the switch to every other port.

Virtual Private Network (VPN), within their existing network infrastructures.

For testing MPLS, a brief explanation of its operation is helpful. The Cisco Nexus 7718 was configured to perform as an LSR (Label Switch Router) using LDP (Label Distribution Protocol). LDP lets LSRs request, distribute and release label prefix bindings to peer routers in the network. LDP also lets LSRs discover potential peers and establish LDP sessions with peers to exchange label-binding information.

The primary job of an LSR is to swap the incoming label with a new label, and then forward the packet to the next hop that is determined by the new label.

OSPF and MPLS LDP were enabled on all 192 of the 100GE ports of the Nexus 7718 switch. The Ixia test system was then similarly configured, and LDP label bindings were exchanged between the Nexus 7718 and the Ixia system.

MPLS-labeled IMIX traffic was then sent from all 192 of the Ixia system's traffic-generator ports, at 100 percent line-rate, to the Cisco Nexus 7718. This 'label-swapping' test is formally described in RFC 2544 for MPLS port pairs.

The MPLS testing was conducted with the Cisco Nexus 7718 running a beta version of Cisco NX-OS operating system software 7.1(0) D1(1) [build 7.10) D1 (0.196)].

The results were impressive.

MPLS label swapping was performed by the Nexus 7718, with line-speed, 100-Gbps IMIX traffic

delivered on all on all 192 of the switch's 100GE ports, with no loss. The Ixia test system confirmed that every one of the 1.808 trillion frames sent to the Nexus 7718 were properly MPLS-processed and returned.

As shown in *Figure 5* below, this equated to an MPLS-processing throughput of 18.1546 terabits/second (Tbps) and 6.017 billion frames per second (fps).

What's more, the latency imposed on MPLS-processed data traversing the Cisco Nexus 7718 was remarkably low. While the IMIX traffic consists of several packet sizes, the traffic on the whole experienced minimum latency of just 5.3 microseconds, an average latency of 11.8 microseconds and a maximum latency of 15 microseconds – with full traffic load on all 192 of the switch's 100GE ports.


Cisco has reportedly implemented the Nexus 7718's MPLS processing in hardware, which helps explain the exceptional MPLS performance.

### Failover & High Survivability

Two different tests were conducted to confirm the Cisco Nexus 7718's ability to survive major module outages – and enable their in-service replacement – without missing a bit.

The first was a fabric module. There are six slots for fabric modules in the Cisco Nexus 7718. These are accessed from the rear side of the multislot switch chassis. Each fabric module gives 220 Gbps of capacity per slot, so when all six are present there is 1.32 Tbps of available bandwidth per slot.

**Figure 5: MPLS Performance – Throughput and Latency**



Source: Miercom, September 2014


If a fabric module fails, an unlikely but not inconceivable event, it would diminish overall switching capacity by about 17 percent. That leaves the system with about 83 percent capacity.

The Nexus 7718 was designed to survive such an event, and to permit such a failed module to be replaced with minimum impact on in-service activity. All fabric modules, line-card (I/O) modules and Supervisor Engine modules have ejector levers which aid in the graceful removal of a module without any packet loss. When the ejectors are engaged, a notification is sent to the system that the module is going to be removed. The module is disabled internally, software is reconfigured and traffic is redistributed to the other modules.

This automatic reconfigurability was tested by delivering a 75-percent load to the Nexus 7718 on all ports. A fabric module was ejected and removed. There was no data loss or disruption to traffic, since 87-percent capacity easily handled the 75-percent applied load. We proved this point by ejecting another fabric module, reducing system capacity to about two-thirds. Traffic loss did occur then, as designed, at about 6.3 percent. When one module was reinserted, the data loss stopped and traffic flow returned to normal.

A similar test was then conducted of a Supervisor module. The Cisco Nexus 7718, and the other 7700 Series models, are designed to be deployed with two Supervisor 2E modules. One is active, the other is in hot standby mode. Together, they form a fully redundant system.

The Supervisor 2E provides centralized arbitration to the switch fabric for all line cards, control-plane and management functions, controlling L2 and L3 services, redundancy capabilities, configuration, status monitoring, power and more.



With the Cisco Nexus 7718 handling a 75-percent load of IPv4 and IPv6 traffic, stateful switchover of the Supervisor 2E modules and the impact on the functioning of the switch were evaluated.

Manual removal of the active Supervisor module forced a stateful switchover. The Supervisor module that had been the hot-standby immediately assumed control of forwarding and management. The Cisco Nexus 7718 did not lose any traffic.

## **In-Service Software Upgrade (ISSU)**

ISSU is a key feature of the Cisco Nexus 7718, supported by all models of the Nexus 7700 Series that allows the system to remain in service during planned maintenance and service upgrades. It allows the user to upgrade the active system software while the system remains fully operational.

Executing ISSU was quick and easy. To launch it requires just entry of one CLI command:

**“install all system bootflash: versionname.bin  
kickstart bootflash: kickstartVersion.bin  
parallel.”**

ISSU then automatically executes, and in our case completed a software upgrade on the Cisco Nexus 7718 – migrating from version 6.2.8 to a newer version 6.2.8a – without loss or a single packet.

At the time, line-rate traffic loads, a mix of IPv4 and IPv6 traffic, were being applied to all 192 of the Nexus 7718's 100GE ports.

**No Frame Loss with ISSU.** The screenshot below shows successful completion of an In-Service Software Upgrade (ISSU). The highlighted area to the right shows that no frames were lost during the process.

The screenshot displays a remote desktop connection to a Nexus 7700 switch. The terminal window shows the output of the 'sh mod' command, listing 18 modules and their status. The status for all modules is 'ok'. Below the module list, there is a table showing the switch (Sw) and hardware version (Hw) for each module. The Sw is consistently 6.2 (8a) and the Hw is 1.1. At the bottom of the terminal, there is a table showing the MAC address and serial number for the first module.

The network management interface on the right shows a table of flow statistics. The table has columns for 'Flow Groups', 'Rx Ports', and 'Endpoint/Enc'. The 'Flow Groups' column contains the value '192'. The 'Rx Ports' column contains the value '192'. The 'Endpoint/Enc' column is empty. Below the table, there is a section for 'L2-L3 Test Summary Statistics' with a sub-section for 'Flow Statistics'. The 'Flow Statistics' table has columns for 'Tx Frame Rate' and 'Rx Frame Rate'. The 'Tx Frame Rate' is 6,393,611,001,500 and the 'Rx Frame Rate' is 6,393,611,102,000. The 'Loss %' column is 0.00. The 'Tx Frame Rate' and 'Rx Frame Rate' cells are highlighted with a yellow border.

## Miercom 100GE Switch Study

This report was sponsored by Cisco Systems, Inc. The data was obtained completely and independently as part of the Miercom 100GE Switch Study, an ongoing endeavor in which all vendors of data center switches with 100GE capability have an equal opportunity to participate and contribute to the test methodology.



## How We Did It

For switch performance testing the test bed was configured as shown in the diagram below. All performance testing of the Cisco Nexus 7718 data center switch employed the same Ixia test-system, running IxNetwork software, which controlled five XG12 chassis with a total of 48 Xcellon-Multis load modules. The test system had 192 bidirectional 100GE test ports, one for each corresponding port on the switch.

For throughput and latency measurement, bi-directional test traffic is applied on port pairs on the DUT so that test traffic is processed across the switch fabric. Also measured during these tests is the variation in delay, or jitter, of IP packets passing through the switch. Separate tests measured throughput and latency for IP multicast traffic.

### Device under Test (DUT)

Cisco Nexus 7718, 26RU, with:

- Cisco NX-OS software 6.2.8a
- 18 slots:
  - 2 Cisco Nexus 7700 Supervisor 2E Modules
  - 16 I/O modules: Cisco Nexus 7700 F3-Series 12x100GE
- Fiber interfaces for I/O modules: Cisco CPAK 100GBASE Modules



**Cisco Nexus 7718**

### Ixia Test System

Five XG12 multi-slot chassis with:

- 48 Xcellon-Multis XM100GE4CXP load modules
- IxNetwork software



**Ixia XG12 (5)**



**Ixia Xcellon-Multis XM100GE4CXP load module (48)**

192 x 100GE test ports



Bi-directional traffic



192 x 100GE switch ports

To help simulate real-world IP traffic, an Internet Mix (IMIX) traffic type was used. IMIX consists of a mix of different packet sizes, based on a statistical sampling of Internet traffic.

The maximum throughput achievable through the switch was validated using a jumbo frame size, 9,216 bytes. Testing also validated the maximum forwarding rate. The result is based 64-byte frames.

The tests in this report are intended to be reproducible by users who want to recreate them with the appropriate test and measurement equipment. Those interested in repeating these tests may contact Miercom at [reviews@miercom.com](mailto:reviews@miercom.com) for more details on setting up the test configurations we applied.

## Miercom Performance Verified

Miercom independently exercised and evaluated Cisco Systems' Nexus 7718 and was frankly stunned by the incredible power and throughput of this system, coupled with consistent low latency and latency variation, as well as solid MPLS support and high-availability features. The first switch we have independently tested with 192 x 100GE ports, the Cisco Nexus 7718 is awarded **Miercom Performance Verified** in our ongoing Data-Center-Class 100GE Switch Study.



**Cisco  
Nexus 7718**

---

**Cisco Systems, Inc.**  
170 West Tasman Drive  
San Jose, CA  
1-800-553-6387  
[www.cisco.com](http://www.cisco.com)

---

## About Miercom's Product Testing Services

Miercom has published hundreds of product-comparison analyses over the years, in leading network trade periodicals including *Network World*, *Business Communications Review*, *Tech Web - NoJitter*, *Communications News*, *xchange*, *Internet Telephony* and other publications. Miercom's reputation as the leading, independent product test center is unquestioned.

Miercom's private test services include competitive product analyses, as well as individual product evaluations. Miercom features comprehensive certification and test programs including: Certified Interoperable, Certified Reliable, Certified Secure and Certified Green. Products may also be evaluated under the Performance Verified program, the industry's most thorough and trusted assessment for product usability and performance.



 Before printing, please consider electronic distribution

Report 140919

**Miercom**

[reviews@miercom.com](mailto:reviews@miercom.com)  
[www.miercom.com](http://www.miercom.com)

Product names or services mentioned in this report are registered trademarks of their respective owners. Miercom makes every effort to ensure that information contained within our reports is accurate and complete, but is not liable for any errors, inaccuracies or omissions. Miercom is not liable for damages arising out of or related to the information contained within this report. Consult with professional services such as Miercom Consulting for specific customer needs analysis.