Cisco Integrated Services Routers Performance Overview

The Cisco® Integrated Services Routers Generation 2 (ISR G2) provides a robust platform for delivering WAN services, unified communications, security, and application services to branch offices. These platforms are designed to support existing WAN access circuits and offer the performance needed for the transition to Ethernet-based access services.

The Cisco ISR G2 series provides a highly secure and reliable platform family for scalable multiservice integration. The excellent service delivery on each ISR G2 platform regardless of size offers the ultimate user experience with the architectural scalability and investment protection needed to minimize overall deployment costs.

This document discusses the performance architecture of the Cisco ISR G2 family and provides specific performance information from a variety of service configurations and test use cases. The goal is to help you understand performance data points and how to use them.

Please keep in mind that performance should not be the only criteria when choosing the ISR G2 platform for a deployment. Other important requirements such as service readiness, modularity, high availability, etc. must be evaluated for the overall solution

Cisco ISRs are designed to deliver integrated services at high performance for the branch office. The platforms run Cisco IOS® Software on a central CPU using a shared memory pool, allowing the processor to dynamically allocate memory for required functions and services. The ISRs have three pieces of function-specific hardware:

Embedded encryption processor: The encryption processor provides hardware-based acceleration for IP Security (IPSec) (using Triple Digital Encryption Standard [3DES] or Advanced Encryption Standard [AES]) and Secure Sockets Layer (SSL) VPNs. For IPSec encryption, the acceleration chip performs the actual mathematical encryption, while relying on the router CPU to identify traffic for encryption, negotiate the security associations, and forward packets. Thus, the encryption chip offloads part of the overall encryption process but the CPU is still involved in the overall processing and forwarding of encrypted traffic.

Packet voice/fax DSP module 3s (PVDM3s): These chips provide dedicated resources for audio conferencing, transcoding, and public-switched-telephone-network (PSTN) connectivity. Again, the chips are specialized for these purposes, but still rely on the router CPU to forward packets to and from them. The multicore CPU on the Cisco ISR G2 platforms runs classic Cisco IOS Software. Since Cisco IOS Software is a single threaded operating system, only a single core is active. In most test cases, router performance is governed by a combination of available CPU cycles and how features are processed in the software.

Embedded Multi-Gigabit Switch fabric (MGF): The MGF is a multi-GE port layer 2 switch built into the internal architecture. Its GE-ports connects to the individual EHWIC, SM, ISM and PVDM3 slots with one GE link to each slot. MGF works autonomously from the host router and allows for switching traffic in Layer 2 with full Gigabit Ethernet line rate directly between modules without impacting routing performance.

Performance Positioning and Recommendations

Performance positioning is an attempt to account for common deployment scenarios and make a recommendation that will fit most requirements. The goal is to provide a recommendation that applies to 80 percent of customer use cases. It is not an all-inclusive metric, nor is it a performance limit of any kind. There will clearly be implementations where router performance can easily exceed the recommendations and others where specific configurations or services, extremely small packet sizes, or other factors can reduce performance below these thresholds.

Overall performance positioning is based on test results from a variety of these multi-service tests. By measuring performance across a variety of test cases, a median performance can be determined that fits many production networks. In simple configurations, customers may see significantly better performance. Customers are welcome to perform independent tests, and results may vary. However, this recommendation should be adequate for most customers.

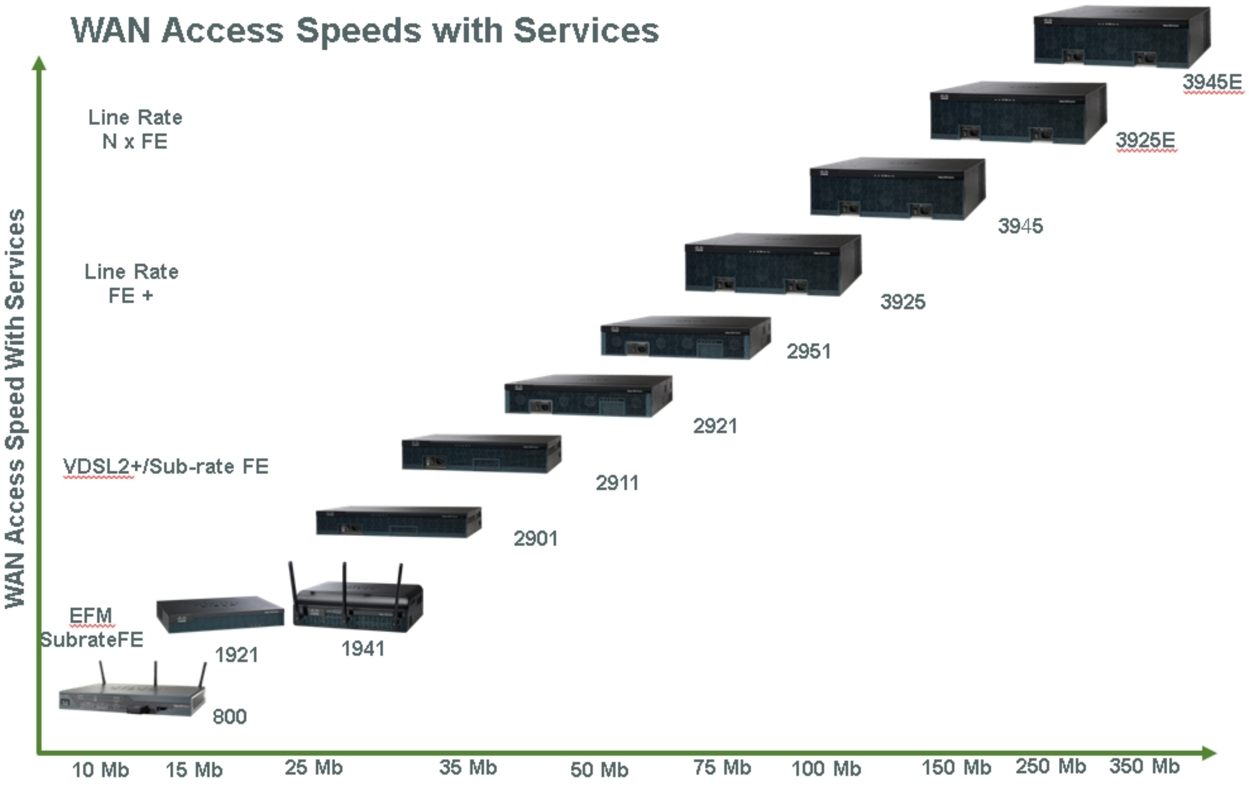

About Router Performance

Interpreting and comparing performance numbers can often become a complex task
Not only do we take different environment and test setups into account, but often also different units for presenting performance test results. This section aims at helping you understand and interpret the two most commonly units used for performance measurement, packets per second (pps) and bits per second (bps).

Packets Per Second (pps)

This is the number of packets a router is able to route for a given service and is generally referred to as “Routing capacity”

Bits Per Second (bps, Kbps, Mbps or Gbps)

This is the total number of actual bits a router is able to route for a given service and is generally referred to as “Speed” or “Bandwidth”

We need to bear in mind that the basic function of a router is to route packets at Layer 3. As such the fundamental router performance is not first and foremost about how many bits but rather about how many packets a router can route for a given time unit and service.

The number of packets a router is able to route per time unit depends on how much work (CPU instructions etc.) the router will have to put in for handling each packet. This will vary with the nature and complexity of the enabled services or service combinations.

The way a packet is switched through a router will vary depending on router family.
Decisions and handling is normally based on the IP header information of a packet and not on the payload. With most services on ISR G2, once a decision has been taken for a packet, the trailing payload of that same packet is more or less cut-through switched. This means that the payload size of a packet will not have any impact on the resulting number of CPU cycles or route decisions having to be dealt with by the router. The size of the payload will however impact number of total bits that gets switched through once the services have been applied. As illustrated in this table below

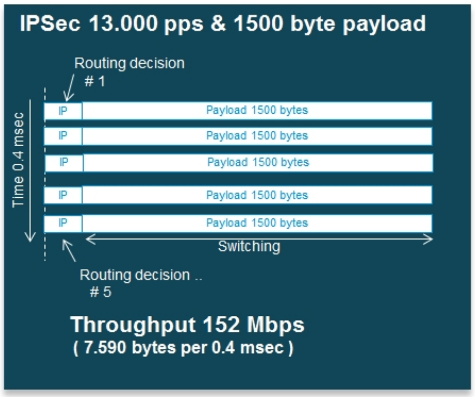

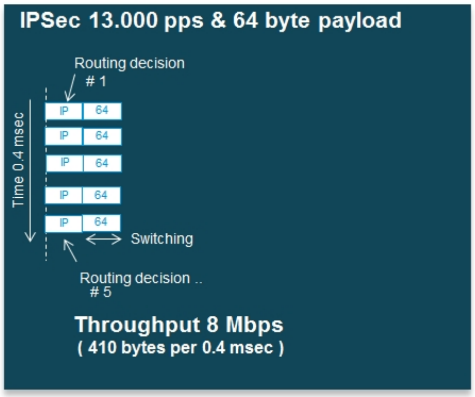

About Router Performance cont'd

This leaves us with two important facts
The number of packets per time unit will be more or less constant for a given service or service combination whereas the resulting number of total routed bits will vary with each routed packet's size

A good example of this can be had from an actual test result, in this case for Stateless (UDP) FW

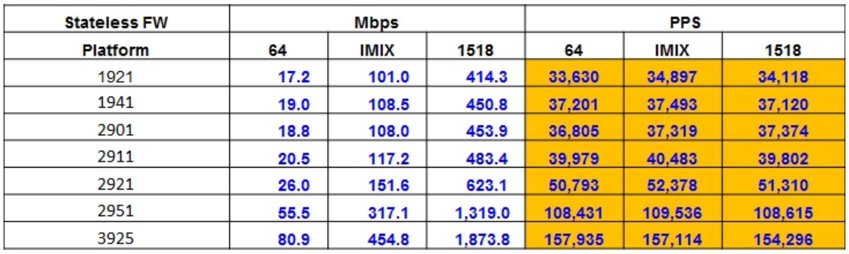

Stateless FW	Mbps			PPS		
Platform	64	IMIX	1518	64	IMIX	1518
1921	17.2	101.0	414.3	33,630	34,897	34,118
1941	19.0	108.5	450.8	37,201	37,493	37,120
2901	18.8	108.0	453.9	36,805	37,319	37,374
2911	20.5	117.2	483.4	39,979	40,483	39,802
2921	26.0	151.6	623.1	50,793	52,378	51,310
2951	55.5	317.1	1,319.0	108,431	109,536	108,615
3925	80.9	454.8	1,873.8	157,935	157,114	154,296

Let's take look at the numbers for the Cisco 1941 in the table above.
By measuring Mbps we see that the Cisco 1941 provides a throughput of 19 Mbps all the way up to 450 Mbps for the same FW service.
However if we look at the packets per second, we can see that this capacity stays constant at 37kpps for all three packet sizes. Thus regardless of the different results in Mbps, the Cisco 1941 has an actual routing capacity of 37 pps.

For ISR G2 as well as for many other router vendors, a performance number in Mbps is merely a product of resulting packet sizes switched through each time a packet is being routed. One drawback of this is that reports in Mbps are often tweaked to skew reported performance in one direction or the other, all depending on purpose of the report.

- Highest number in Mbps is achieved when all packets contains max # of allowed bytes (1518)
- Lowest number in Mbps is achieved when all packets contains min # of allowed bytes (64)

For this reason one might argue that bps is not an absolute measure of a router's routing capacity, whereas a number in pps is.

Performance numbers in bps do have their place though.
Measurement in bps is, in spite of the fact as per above, by far the most commonly known and used unit in the market for reporting and comparing routing performance. This partly because consumers are familiar with the term Mbps and partly because it compares nicely to physical data rate of Ethernet and WAN connections which are always measured in bps.

CPU Utilization

Most performance testing in a lab environment is performed between onboard Ethernet interfaces, although with the Cisco 3900 Series Integrated Services Routers more Ethernet interfaces had to be added to test platform capacity in some test cases.

Ethernet interfaces provide the least processor overhead, because the router must simply swap MAC headers. Serial interfaces, including T1/E1, dial, and others, require a new Layer 2 encapsulation and therefore require more CPU involvement. Serial interfaces, by definition, must also serialize the packet flow. Thus, when using a serial interface, the router passes less traffic at the same CPU usage than when using Ethernet interfaces.

Another focus in creating a production-network focused set of performance data is on router CPU. In order to provide comparable, architecture agnostic results, performance tests are carried out using a methodology called Non Drop Rate (NDR) as described below.
NDR tests generally push the router CPU to a very high percentage, sometimes up in the 98-99% range.

No production network is run at this type of CPU usage. Traffic on real-world networks is bursty, not smooth and consistent like a lab traffic profile. Lab routers do furthermore not have to converge routing protocols because of external events. In short, a router running at 99-percent CPU usage is not to be considered a working router in an operational branch deployment.

Most service providers set their CPU alarms to 60 or 65 percent while many enterprise customers are perfectly comfortable running production networks with a CPU load at around 70 or 75 percent.

In order to extrapolate a result that represents a load of around 75% CPU, subtract 23% of the RFC 2544 NDR result for a given test case. This will not provide an exact number but will represent a fair estimate of a platform’s throughput at 75% CPU.

Test Methodology

This section defines the test methodology that is utilized for performance testing on the selected platforms using traffic profiles as defined below.

Non Drop Rate and RFC-2544 Tests

NDR is a widely used method to measure throughput and is well known in the industry. It is based on RFC 2544 which is the "Benchmarking Methodology for Network Interconnect Devices". This methodology is available in the third party traffic generator tools such as Spirent and IXIA.

RFC 2544 employs a binary search for no drop throughput rate for a pre-defined packet size traffic stream, thereby presenting a fully architecture-agnostic test methodology through its independence of measured CPU load. This will provide comparable test results between single core and multi core architectures.

Traffic profiles

In this document, Cisco IMIX and Cisco IPSec-IMIX traffic profiles are used. The general convention of measuring performance with individual packet sizes such as 64, 256 and 1400 bytes provides control and stability but does not reflect real world traffic. Real network traffic or the Internet traffic does not consist of fixed packet sizes; it commonly consists of mixed packet sizes

Cisco IMIX

Cisco's standard IMIX uses 64, 594 and 1518 bytes packets in 7:4:1 ratio which relates to 7 streams of 64 bytes packet, 4 streams of 594 bytes packet and 1 stream of 1518 bytes packet. In terms of throughput bandwidth, this results 58.33% of 64 bytes, 33.33% of 594 bytes and 8.33% of 1518 bytes. This traffic definition is already included in third party Spirent and IXIA traffic tools and known as Cisco IMIX.

Cisco IPSec IMIX

For IPSec due to the additional overhead of its tunnel IP headers, 1518 packet size exceeds the maximum allowed MTU of most of the widely used links and results in fragmentation. In order to avoid fragmentation as far as possible, Cisco has developed a separate IMIX traffic profile for IPSec called Cisco IPSec IMIX. Cisco IPSec IMIX consists of 58.67% of 90 bytes, 2% of 92 bytes, 23.66% of 594 bytes and 15.67% of 1418 bytes. Cisco IPSec IMIX is also available in third party Spirent and IXIA traffic tools and known as Cisco IPSec IMIX.

Table 1. Number of tunnels, sessions and peers used in the below test cases

Platform	CEF Src/Dst IP address pairs	IPSec tunnels	GRE tunnels	FW sessions
891	75	15	15	50
1900	50	25	25	50
2901/2911	100	50	50	50
2921/2951	200	100	100	50
3925	300	200	200	50
3945	500	200	200	50
3925E	500	200	200	50
3945E	500	200	200	50

Important on how to read and understand test results

All test results in this document represent aggregate numbers of traffic going out as well as return traffic coming back in for each tested service. The actual split of outgoing and incoming traffic will vary between the test cases, all depending of the nature of the tested service.

Stateless Test Cases

With the exception of our Stateful HTTP test cases (tables 12 – 16) the Stateless test cases herein are using UDP based traffic profiles ranging from almost symmetric 50-50 splits to almost unidirectional splits of traffic going out and coming in.

Stateful Test Cases

Web based traffic patterns are becoming increasingly more common in today's Branches as more and more services are changed into centralized cloud based services. The Stateful HTTP tests provide performance data for these branch deployments with traffic patterns dominated by web traffic.

These tests are carried out using HTTP/TCP in which testers simulate 50 branch users, all of them constantly downloading 16 KB web pages during a period of 600 seconds. The resulting traffic profile becomes almost 95% unidirectional in the download direction since, according to standard HTTP web server operation, a web server is typically responding with maximized object sizes to relatively small queries from users.

Bidirectional numbers – Representing actual throughput?

Test numbers representing aggregate traffic are commonly referred to as "Bidirectional Traffic". Even though technically correct, the term bidirectional can be slightly misleading since it implies a 50-50 split, which in turn could imply that the true performance is actually only half of the posted number.
This is of course not correct.

Routers don't have separate CPUs for each direction of traffic hence the max throughput is based on what the CPU can process regardless of direction and traffic split.

For example; a number of 300 Mbps measured as bidirectional traffic for a certain service and packet sizes is indeed the router's total routing capacity. It shall not be erroneously understood as a 150 Mbps actual throughput because that 300 Mbps aggregate number could also represent 200 Mbps Down and 100 Mbps Up or perhaps 250 Mbps Down and 50 Mbps Up, etc.."

Referring to these kind of test results as "Aggregate" rather than "Bidirectional" could hence avoid unnecessary misperceptions.

Tested platforms

All modular ISR G2 platforms 1921 – 3945E have been tested and are included in this report.
Only two onboard Ethernet interfaces on each platform were used during tests.

- No modules were used
- Additional onboard Ethernet interfaces were not added when a platform performed at line rate

For some test cases we reached the max speed of an interface before any packet drops were recorded. These test cases are marked L/R, as in "Line Rate"

Cisco 891 is the only fixed configuration router included in these test to-date. Four more 800 fixed configuration router types are on the roadmap to get tested.

891 is a more recent platform than the modular ISR G2 platforms included in this report.
Since it has a newer architecture, some components are inevitably slightly more powerful than those of its modular counterparts. This is why in some test cases for single services this platform performs better than its intended positioning. It should however be noted that when multiple services are being combined like in a typical branch deployment, the 891 performs according to what it's positioned for.

Performance results are specific to the IOS releases and test cases listed in this document

IOS Releases used in this document

Stateless (IMIX) test cases

Cisco 891	15.0(1)M
Cisco CISCO1921/K9	15.0(1)M2
Cisco CISCO1941/K9	15.0(1)M
Cisco CISCO2901/K9	15.0(1)M
Cisco CISCO2911/K9	15.0(1)M
Cisco CISCO2921/K9	15.0(1)M
Cisco CISCO2951/K9	15.0(1)M
Cisco CISCO3925/K9	15.0(1)M
Cisco CISCO3925E/K9	15.1(1)T
Cisco CISCO3945/K9	15.0(1)M
Cisco CISCO3945E/K9	15.1(1)T

Stateful (HTTP 16K) test cases

All platforms	15.1(4)M

Single IOS Services

A recent analysis indicated that IOS release 12.4 contains more than 4000 named features. Most customers use a small subset of these, generally only 3-5 features. Most commonly used features are Access Control Lists (ACLs), Quality of Service (QOS), Generic Routing & Encapsulation (GRE), Network Address Translations (NAT) and two security features, encryption and firewall. The last two features will be addressed in the Security section of this document.

QOS is a technology that prioritizes latency-sensitive traffic over traffic that is not sensitive to latency. These tests include Hierarchical QOS (HQOS). HQOS uses a parent policy, usually a bandwidth shaper, and subsequent child policies to color and queue traffic within the parent policy. The tested configuration used 5 classes of traffic, of effectively 5 child policies.

When testing QOS, packets are marked for classification by the packet generator, and the router under test classifies the packets based on the markings. No congestion is generated to cause queuing, as, by definition, this would result in packet loss.

Typical / Basic HQOS performance by platform at NDR 2544

Table 2.

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	15	117	111	121	125	183	292	400	1494	482	1792
Kpps @ 64 bytes	80*	43	49	49	53	65	120	167	551	196	685

* 891 did 41 Mbps @ 64 byte packets.

Network Address Translation (NAT) translates IP addresses coming into and going out of the private network. It is used either to protect the private IP addresses of an organization from exposure to the Internet, or to extend the number of public addresses an organization has by using private address ranges. NAT can use dynamic or static mappings, or a combination of both. A one-to-many mapping of public to private IP addresses is called NAT overload, or Port Address Translation (PAT). This method was used for the performance tests.

Table 3. PAT performance by platform at NDR 2544

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	L/R	111	122	122	132	171	375	552	1400	665	1712
Kpps @ 64 bytes	76	37	41	41	46	59	130	191	499	232	611

IP tunneling with Generic Routing & Encapsulation (GRE) is the preferred method for tunneling IP packets when creating an overlaid network on top of an existing architecture. GRE is a light service with low CPU impact. More than two interfaces were used when measuring the E-series since both E-series routers hit line rate with two GE interfaces.

Table 4. GRE multi tunnel performance by platform at NDR 2544

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	L/R	330	386	380	423	528	966	1382	3489	1639	3596
Kpps @ 64 bytes	204	103	117	118	129	157	295	424	1117	503	1356

IOS Security Services and Performance

Security performance can be grouped into two categories – secure connectivity and threat defense. Secure connectivity includes IPSEC and SSL VPN technologies. Threat defense, from a performance perspective, would focus on firewall technology.

The United States government maintains very strict control on the export of strong cryptography, both from a technology and performance standpoint. As with many other products, the Cisco ISRs are subject to this regulation. In order to comply with this, both the temporary and permanent Security licenses have been limited in both performance and tunnel count. The limitation is applied to cumulative encrypted tunnel counts and concurrent throughput and will cap performance to a max of 85 Mbps. Encrypted tunnels are defined as IPSEC, SSL VPN or cRTP.

The HSEC license allows full scalability in both performance and connections.

Table 5. IPSEC throughput measured using Cisco IPSec with HSEC license on 2921 and up.

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Single Tunnel											
Mbps @ IMIX	75	51	58	58	64	82	150	212	633	244	800
Kpps @ 64 bytes	23	18	21	21	23	30	47	66	195	76	245
Multi tunnel See table 1 for # of tunnels tested on each platform											
Mbps @ IMIX	64	49	55	54	56	68	126	145	503	158	726
Kpps @ 64 bytes	20	17	20	18	21	25	38	44	156	49	214

A secondary data point for performance testing on secure connectivity technologies is maximum connections. This metric is not very applicable to the Cisco ISRs as they are primarily branch or access routers, deployed as Customer Premises Equipment (CPE) in managed service environments. This means the routers, in most deployments, will only be asked to support a few tunnels in a production environment. For IPSEC, a tunnel is represented on the router by configuration of a Virtual Tunnel Interface (VTI).

Table 6. Encrypted Tunnel Count by platform

Platform	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Encrypted Tunnels (SEC license)	150	150	225	225	225	225	225	225	225
Encrypted Tunnels (HSEC license)	N/A			900	1000	1500	3000	2000	3000

Firewall testing is much more complicated than any other test discussed in this document. Zone-Based Firewall is a Stateful application, maintaining and monitoring the state of all TCP connections through it. It has multiple Application Layer Gateways (ALGs) that allow it to inspect and monitor specific protocols and applications. Zone-based firewall also inspects traffic both within and between zones.

Thus, test methodology has a significant impact on performance. Testing different applications invokes specific ALGs, each of which may have a different impact on test results. Many test tools can generate packets with TCP headers, but never complete the handshake and establish state for monitoring. In some situations, the firewall may see this as a Denial of Service (DOS) attack, as it would rarely be encountered in a production network unless under attack. The use of pure UDP or other stateless traffic patterns can also produce varying results.

For the purposes of this document, firewall is configured with 2 zones, and all traffic is sent between zones. The traffic generated is stateless and uses the same UDP port number. Performance will be measured in maximum throughput and Maximum Concurrent Sessions. One element that influences the maximum sessions metric is the amount of installed memory in the platforms. These tests used default memory.

Table 7. Firewall performance by platform using Cisco at NDR 2544

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	L/R	101	108	108	117	151	317	454	1442	540	1765
Kpps @ 64 bytes	78	34	37	37	40	51	108	160	515	188	631

Heavy Service Combinations

Single service tests show the impact of IOS features on standard traffic flows. Most customers deploy multiple services. Testing with several IOS services configured shows the impact on performance of multiple algorithms running concurrently, and the impact they have on each other.

This test intends to provide performance indications from combining multiple very CPU intensive, or "Heavy ", features. We included features IPSec, Firewall, Complex QoS and Complex HQoS of which all are very CPU intensive by themselves. This test suite combines IPSec with each one of the three other intensive services.

Table 8. IPSec + Zone Based Firewall performance by platform at NDR 2544

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	37	35	40	41	44	55	88	120	378	135	485
Kpps @ 64 bytes	12	11	14	14	15	18	27	36	114	41	150

Table 9. IPSec + Complex QoS performance by platform at NDR 2544

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	28	27	35	35	35	44	71	98	341	115	445
Kpps @ 64 bytes	Not Applicable due to QoS settings in this test profile										

Table 10. IPSec + Complex HQoS performance by platform at NDR 2544

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps @ IMIX	14	20	22	22	24	30	43	58	204	65	264
Kpps @ 64 bytes	1,7	3,5	4,5	4,5	5,4	6,2	9,5	12	43	14	57

MPLS PE (Provider Edge) router

Tested platforms are only Cisco 2921 and up for which a PE deployment would make sense.

- Cisco 7200 was mixed in since this is a very common PE platform.
- Cisco 3845 was added for comparison.

Each PE carried 50 VRFs, with 2000 routes per VRF.

- 15% of those routes (300 / VRF) prefixes from Internal CPEs
- 85% (1700 / VRF) external prefixes from other PE's.

Test data is unidirectional traffic using 64 bytes fixed packet size.
With 362 byte (IMIX avg size) fixed size, all platforms from 3925 and up performed at line rate.

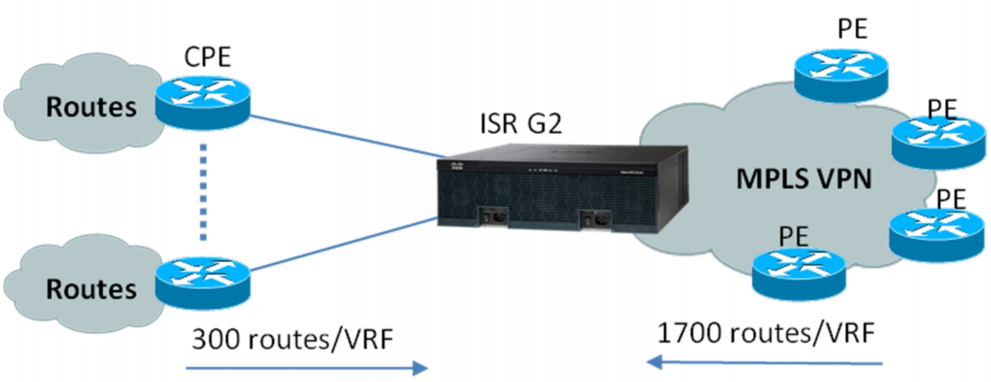

Table 11. MPLS VPN. ISR G2 as MPLS PE router at 64 bytes packet size – NDR 2544

Platform	2921	2951	3845	3925	3925E	3945	3945E	7200-NPEG1	7200-NPEG2
Mbps	159	236	197	350	712	428	712	360	616
Kpps	312	461	386	685	1391	801	1391	690	1205

Stateful HTTP tests (Cloud based traffic pattern)

HTTP tests provide performance data emulating a branch deployment with its traffic pattern dominated by web traffic. This pattern is becoming increasingly more common as more and more services are shifting into centralized cloud based services.

- Cloud based traffic consists at large of a web-based TCP based HTTP traffic pattern.
- Throughput in Mbps represents the average aggregated upload/download per second from 50 users constantly downloading 16 KB web pages over a 600 second test cycle.
- Due to the nature of web traffic, 95% of the measured traffic is unidirectional in the download direction.
- Packet per second was not measured in these tests.

Table 12. Cloud Intelligent Network traffic pattern - NAT

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps	100 L/R	287	315	315	335	435	927	1000 L/R	1000 L/R	1000 L/R	1000 L/R

Table 13. Cloud Intelligent Network traffic pattern - ZBFW

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps	100 L/R	240	260	260	280	357	461	627	1000 L/R	725	1000 L/R

Table 14. Cloud Intelligent Network traffic pattern - ZBFW + NAT

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps	97	135	151	153	165	212	238	313	1000 L/R	342	1000 L/R

Table 15. Cloud Intelligent Network traffic pattern - ZBFW + NAT + HQoS

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps	78	117	131	132	142	181	199	263	871	291	1000 L/R

Table 16. Cloud Intelligent Network traffic pattern - IPSec only

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Mbps	56	41	46	46	50	65	130	178	620	201	787

Max technical routing capacity (Drag Strip)

Providing performance information based on raw maximum transmission rate, a.k.a. “Drag Strip”, is a commonly used technique in the router market for making performance numbers look better than what they really would be in a production environment. Drag strip numbers are typically presented in either the maximum number of packets or the maximum number of bits that a router can forward with no additional services running.

- To get the highest number of packets we used the smallest possible packet size (64 bytes)
- To get the highest number of bits we used packets of maximized packet size (1500 bytes).

All platforms tested to the highest possible performance number by adding interfaces to the test traffic pattern until packet drops were observed. Cisco 3925E was tested using one less EHWIC module than 3925 thus yielding a slightly lower Mbps number.

These tests are only measuring technical platform internal performance and do not reflect how a router will perform in a production environment

Table 16.

Platform	891	1921	1941	2901	2911	2921	2951	3925	3925E	3945	3945E
Kpps (64-byte packets)	224	290	330	330	352	479	579	833	2179	982	2924
Mbps (1500-byte packets)	1400	2770	2932	3114	3371	3502	5136	6903	6703	8025	8675

Conclusion

Router performance can be measured using maximum transmission rates. This provides a measurement of the forwarding capability of the CPU, but no information about the effect of software algorithms, application awareness, or other services.

Another approach is to measure the performance of the router with multiple services enabled in a simulated production environment. This measurement produces performance data that network engineers can use for designing and upgrading customers’ networks in a real-world environment.

The Cisco ISR G2 routers provide an industry-leading ability to deploy integrated services into the branch office with world-class performance. Performance of the routers will vary depending on the services configured, IOS version used, packet mix, and available router CPU cycles.