

Cisco Integrated Services Routers—Performance Overview

What You Will Learn

The Cisco® Integrated Services Routers Generation 2 (ISR G2) provide 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.

This document discusses the performance architecture of the Cisco ISR G2 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.

The performance information in this document is divided into two sections. The first section provides details about some maximum performance values, and the second presents a set of data to be used for production network design.

Architecture for Integrated Services and Performance

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 two pieces of function-specific hardware:

- **An 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 process—the mathematically intensive part—but the CPU is still involved in the overall processing and forwarding of encrypted traffic.
- **Packet voice/fax DSP module 3s (PVD3s):** 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.

No Drop Rate and RFC-2544 Tests

Routers have traditionally been tested using RFC 2544 or similar types of performance tests. RFC 2544 requires tests to be run at a no drop rate (NDR). This testing is done by using a fixed packet size, usually 64-byte packets, and the results are usually published as a metric in kilopackets per second (kpps). The tests are designed to show the CPU power and processing power of the platform (Table 1).

Another popular technique for providing router performance information is also an NDR test, but it is performed with maximum packet size and presented as a throughput test. Results are delivered as megabits per second (Mbps). This test yields a maximum data-rate forwarding of specific features.

Table 1. Cisco ISR G2 RFC 2544-Based Performance (kpps and Mbps)

Platform	Cisco 860	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
kpps (64-byte packets)	25	50	100	290	330	330	352	479	579	833	1845	982	2924
Mbps (1500-byte packets)	197	198	1400	2770	2932	3114	3371	3502	5136	6903	6703	8025	8675

For NDR tests sometimes the platforms can process and forward packets faster than the aggregate bandwidth of the interfaces that the specific models can support. In this situation, all available interfaces are driven to line rate and CPU usage recorded.

What these tests do not provide is any indication of how the router will perform in a production environment. They assume that router CPUs scale linearly to the point where they drop packets. The tests provide no means for analyzing router services, software-based algorithms, or other features. There is no ability to account for real protocols, application layer gateways (ALGs), or other real-world traffic.

Also, production networks tend to have varied packet sizes. Voice traffic and TCP acknowledgements (ACKs) tend to be very small packets, generally 64 to 80 bytes. File transfers and some applications tend to use as large a packet size as they can negotiate. Thus, NDR tests with fixed packet sizes do not provide a very realistic look at router performance in a production environment.

Cisco IOS Software 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. From a performance perspective, threat defense focuses on firewall technology.

For IPSec, the focus is on throughput and scalability. IPSec throughput is measured using a single tunnel with 1400-byte packets, with no Secure Hash Algorithm (SHA) or Message Digest Algorithm 5 (MD5) authentication. The packet size must be reduced to account for the additional packet headers when using IPSec.

With regard to secure connectivity, the United States government maintains very strict control on the export of strong cryptography, from both technology and performance standpoints. As with many other products, the Cisco ISRs are subject to this regulation. In order to comply with this policy, both the temporary and permanent Security (SEC) licenses are limited in both performance and tunnel count. The limitation is applied to cumulative encrypted tunnel counts and concurrent throughput. Encrypted tunnels are defined as IPSec, SSL VPN, or Secure Real-Time Transport Protocol (SRTP). Currently that limitation is 170-Mbps throughput (85 Mbps in each direction) and 225 tunnels. This limit is enforced and cannot be exceeded with the SEC license. The High-Performance Security (HSEC) license allows full scalability in both performance and connections.

Table 2 gives performance information for IPSec and SSL VPN by platform.

Table 2. IPSec Maximum Performance by Platform

Platform	Cisco 860	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
IPSec Mbps (SEC license only, no HSEC needed)	46	102	125	149	170	170	170	—	—	—	—	—	—
IPSec Mbps (SEC + HSEC license)	—	—	—	—	—	—	—	207	282	770	1494	848	1503

A second data point for performance testing on secure connectivity technologies is maximum connections. This metric is not very applicable to the Cisco ISRs because they are primarily branch-office or access routers, deployed as customer premises equipment (CPE) in managed service environments, meaning that in most deployments the routers have to support only 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 3 gives information about encrypted tunnel count by platform.

Table 3. Encrypted Tunnel Count by Platform

Platform	Cisco 860*	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
Cumulative encrypted tunnels (SEC license)	5	20	50	150	150	150	225	225	225	225	225	225	225
SSL VPN tunnels	—	10	25	50	75	75	100	100	150	200	500	200	500
HSEC license IPSec VPN tunnels	—	—	—	—	—	—	—	900	1000	1500	3000	2000	3000

*The Cisco 860 models do not support SSL VPN

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

Thus, test methodology significantly affects performance. Testing different applications invokes specific ALGs, each of which may affect test results differently. 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 situation as a denial-of-service (DoS) attack, because it would rarely be encountered in a production network unless under attack. The use of pure User Datagram Protocol (UDP) or other stateless traffic patterns can also produce varying results.

For the purposes of this document, firewall is configured with two zones, and all traffic is sent between zones. The traffic generated is stateless and uses the same UDP port number. Performance is measured in maximum throughput and the number of 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 4 gives firewall performance information by platform.

Table 4. Firewall Performance by Platform

Platform	Cisco 860	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925 E	Cisco 3945	Cisco 3945 E
Maximum throughput (Mbps)	30	43	54	530	569	625	676	801	1350	2567	6112	3133	7473
Maximum number of concurrent sessions (1000s)	0.3	0.5	1.1	50	63	63	90	130	150	270	300	305	345

Again, the data presented in this section is for maximum performance and is not very valuable for use in a production network. Although a router may be able to forward more than 1 Gbps of encrypted traffic in a lab-based performance test, it should not be expected to perform at that level in a customer's network. Packet sizes will vary in a real network, and routers cannot be stressed to NDR.

Maximum tunnels are a specific point where performance derived in a lab situation varies from a production design. Although this number is easy to reproduce in a test environment, very little traffic will be forwarded over those tunnels during the test.

Firewall performance will vary depending on the nature of the traffic. Because ZBF monitors the state of traffic and monitors specific protocols and applications, actual application traffic will affect the throughput of the firewall.

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.

Testing for the positioning performance ranges was conducted using Internet mix (IMIX) traffic. IMIX is a packet mix that attempts to duplicate traffic bound for the Internet. Although it is not an industry standard, it is standardized within the test tool manufacturers. Every test tool manufacturer has its own version of IMIX, but the versions do not deviate significantly.

Packet Size

Two different IMIX traffic mixes were used for testing the Cisco ISR G2 routers. For test cases not involving encryption, the following traffic mix was used:

- 1518 bytes x 15 packets (15%)
- 594 bytes x 24 packets (24%)
- 64 bytes x 61 packets (61%)

The average packet size computes to 409 bytes.

$$[(1518 \times 15) + (61 \times 64) + (24 \times 594)/100] = 409$$

For test cases that used encryption, the maximum packet size needed to be reduced to avoid fragmentation. This reduction complies with best practices in VPN networks of setting the MTU to 1440 bytes on an interface to allow for IPSec headers.

- **1418 bytes x 15 packets (15%)**
- **594 bytes x 24 packets (24%)**
- **90 bytes x 61 packets (61%)**

The average packet size computes to 410 bytes.

$$[(1418 \times 15) + (61 \times 90) + (24 \times 594)/100] = 410$$

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 clearly had to be added to test platform capacity. 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, also must 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. As mentioned in the previous section, NDR tests generally push the router CPU to 99–100 percent, because this level is the limiting factor in raw performance. However, 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 test. Lab routers do not have to converge routing protocols because of real-world events. A router running at 99-percent CPU usage could never handle any of these events.

Most service providers set their CPU alarms to 60 or 65 percent. Many enterprise customers are comfortable running production networks with CPU around 70 or 75 percent. For performance positioning of the Cisco ISR G2 routers, the CPU threshold is set to 75-percent usage. This setting provides a valid metric for how the router will perform in a production environment, allowing for large traffic bursts and routing protocol convergence.

Single Cisco IOS Software Services

A recent analysis indicated that Cisco IOS Software Release 12.4 contains more than 4000 named features. Most customers use a small subset of these features, generally only 3 to 5 of them. Most commonly used features are access control lists (ACLs); quality of service (QoS); Network Address Translations (NAT); and two security features, encryption and firewall. The last two features are 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. In this testing we use 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 five classes of traffic, of effectively five 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, because, by definition, this congestion would result in packet loss.

Table 5 gives information about HQoS performance by platform.

Table 5. HQoS Performance by Platform; IMIX Traffic at 75-Percent CPU

Platform	Cisco 860*	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
Mbps	—	45	64	110	120	125	125	168	293	455	1368	572	1704

*HQoS is not supported on the Cisco 860 models

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 6).

Table 6. PAT Performance by Platform, IMIX Traffic at 75-Percent CPU

Platform	Cisco 860	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
Mbps	27	60	75	80	91	101	114	138	275	418	1067	496	1334

ACLs are used to classify and filter traffic. Most other Cisco IOS Software features use ACL technology to identify the traffic on which the service needs to be implemented. For this reason, ACLs are included by default in many of the test cases.

Recommended Security Positioning

Security services recommendations are based on similar considerations, including IMIX packet sizes; 75-percent CPU; and the effect of establishing security associations (SAs), various protocols used with ZBF, etc. (Table 7).

Table 7. Security Services Performance Recommendations (May Require SEC + HSEC Licenses)

Platform	Cisco 860	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
IPSec tunnels	5	20	50	150	150	150	225	400	500	750	1500	1000	2000
IPSec Mbps	6	20	30	35	48	53	61	72	103	154	477	179	670

Service Combinations and Overall Positioning

Single service tests show the effect of Cisco IOS Software features on standard traffic flows. Most customers deploy multiple services. Testing with several Cisco IOS Software services configured shows the effect on performance of multiple algorithms running concurrently—and the effect they have on each other. Table 8 shows the performance of the platforms when configured with QoS, ACLs and NAT.

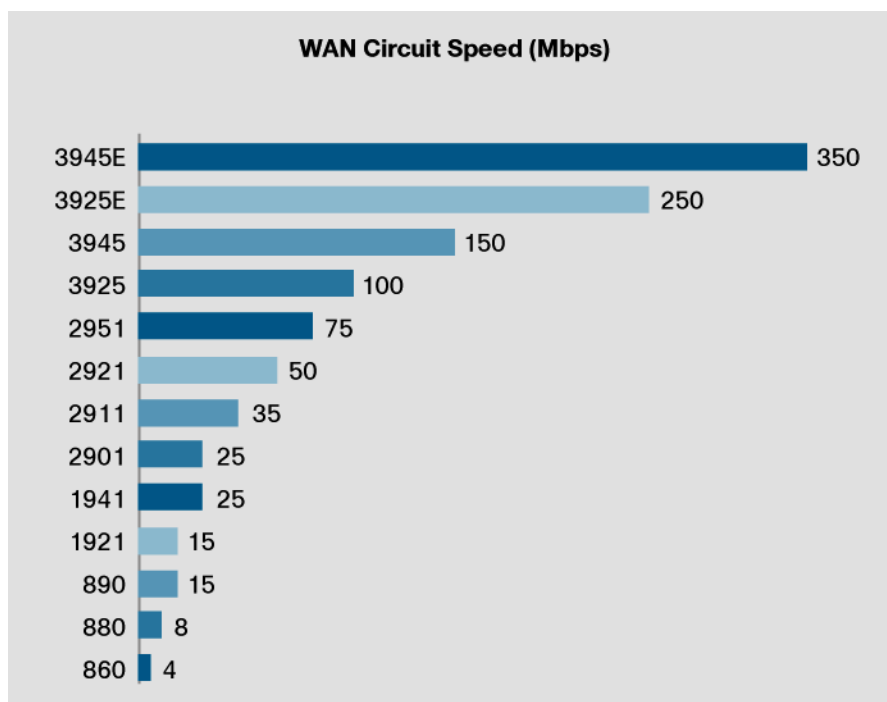
Table 8. NAT + QoS + ACL Performance by Platform, IMIX Traffic at 75-Percent CPU

Platform	Cisco 860	Cisco 880	Cisco 890	Cisco 1921	Cisco 1941	Cisco 2901	Cisco 2911	Cisco 2921	Cisco 2951	Cisco 3925	Cisco 3925E	Cisco 3945	Cisco 3945E
Mbps		30	45	68	76	77	81	105	114	198	534	223	668

Overall performance positioning is based on test results from a variety of these multiservice tests. By measuring performance across a variety of test cases, a median performance can be determined that fits many production networks. This metric is both a positioning metric as well as a recommendation. Additionally, the performance positioning is not any kind of limit. 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 enterprise and commercial customers (Figure 1).

Figure 1. Performance Positioning by Platform with Services



Conclusion

Router performance can be measured using maximum transmission rates, or NDR results. 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, packet mix, and available router CPU cycles.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)