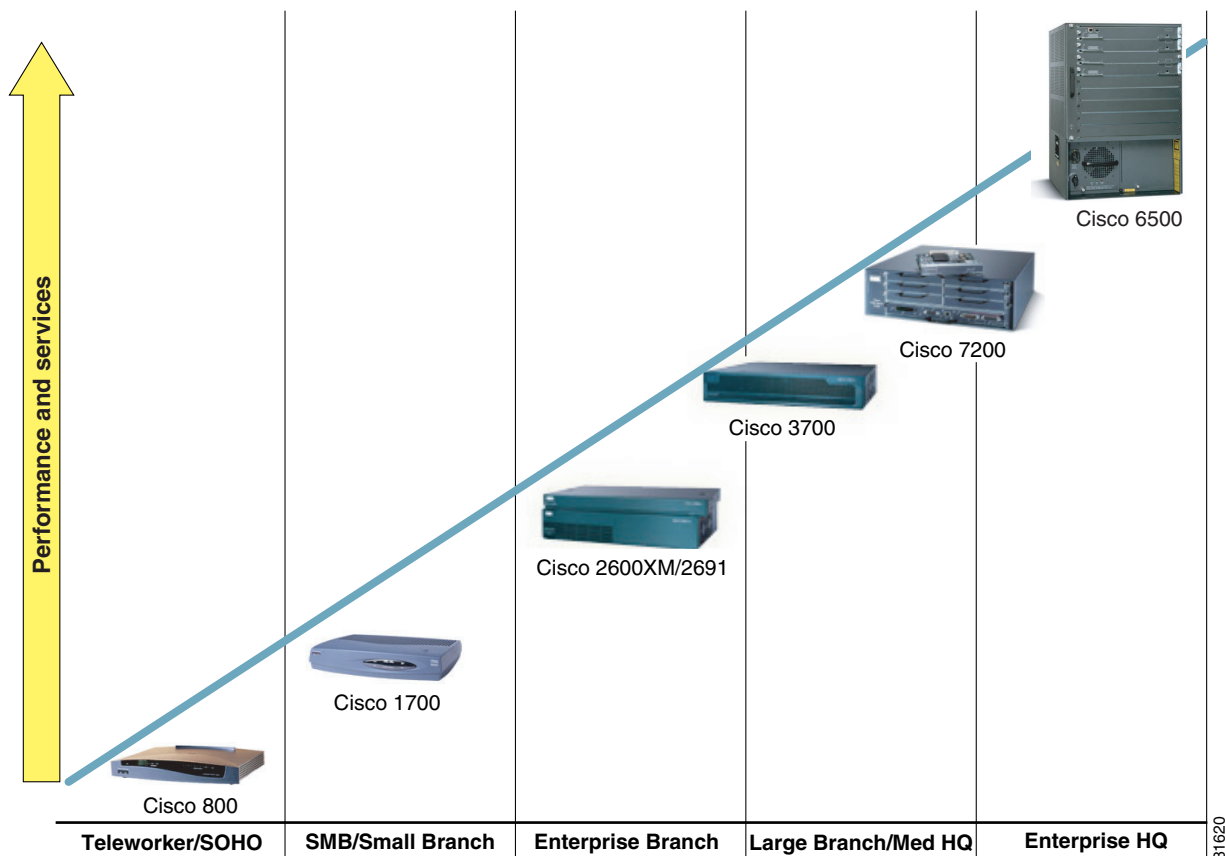


Product Selection

This chapter discusses the V³PN scalability and performance evaluation that was performed and gives product performance results, recommendations, and conclusions that can be used for design parameters when planning and implementing a V³PN deployment.

Cisco offers a whole line of VPN router products which range in application from small business up to large VPN tunnel aggregation points at a large enterprise central site. Figure 5-1 shows the range of products available, almost all of which were evaluated for V³PN performance and deployability.

Figure 5-1 Cisco IOS VPN Router Portfolio



The following topics are addressed within the individual sections of this chapter:

- [Scalability Test Methodology, page 5-2](#)
- [Traffic Profiles, page 5-3](#)
- [Head-end Product Selection, page 5-6](#)
- [Branch Office Product Selection, page 5-9](#)
- [Network Performance/Convergence, page 5-15](#)
- [Software Releases Evaluated, page 5-17](#)

Scalability Test Methodology

The Cisco Enterprise Solutions Engineering VPN performance and scalability lab uses test tools to generate data traffic simulating actual end-to-end applications running on Solaris and Red Hat Linux TCP/IP stacks. The traffic generated incorporates flow control inherit with a TCP implementation. These tools create a network environment that is fairly realistic in terms of how production networks perform.

NetIQ's Chariot test tool is used to generate network traffic. As NetIQ endpoints, SUN NETRA and Penguin Red Hat Linux servers are deployed. The Linux servers generate the simulated voice traffic, the SUN NETRA servers generate the data traffic. Solaris supports path MTU discovery by default.

More information on NetIQ Chariot can be found on the following NetIQ website:

<http://www.netiq.com/products/chr/default.asp>.

In addition to the Chariot test tool, portions of the test included implementing two CallManagers, a Survivable Remote Site Telephony (SRST) on a Cisco 2651, and three Cisco 7960 IP phones; one on a campus, one on the SRST Cisco 2651 and one in the core. Using actual phones and voice calls allows for subjective evaluation of the voice quality.

As shown in the diagram in [Appendix A, "Network Diagram Scalability Testbed and Configuration Files"](#), the scalability testbed included 240 branch offices aggregated to two head-end devices. The head-ends consisted of the Cisco 7200 VPN routers (refer to the ["Head-end Product Selection" section on page 5-6](#) for exact models tested). The branch offices consisted of Cisco VPN router products from the Cisco 800, Cisco 1700, Cisco 2600, Cisco 3600, and Cisco 3700 series (refer to the ["Branch Office Product Selection" section on page 5-9](#) for exact models tested).

Branch routers were evaluated at various link speeds ranging from 128 Kbps up to E1. Both head-end and branch router products were evaluated by raising traffic rates up and monitoring key performance parameters until the limitations of each platform were found. The sections that follow give specifics on the exact measurements for each platform. In general key parameters monitored included:

- CPU utilization
- Bi-directional throughput (in bits per second and packets per second)
- End-to-end peak and average latency (for voice traffic)
- End-to-end jitter (for voice traffic)
- Drop rates
- Network resiliency (in the case of head-ends)

All products were evaluated with hardware-accelerated encryption installed. All scalability testing for this Design Guide revision was obtained using IPSec Tunnel Mode, therefore the throughput results might differ in Transport Mode.

In addition to throughput performance testing, failover testing was also conducted. Please refer to the [“Network Performance/Convergence” section on page 5-15](#) for more information on the failover test scenario.

Traffic Profiles

Cisco Enterprise Solutions Engineering conducted solution testing of this design to validate scalability. Portions of the test plan were designed to simulate worst-case scenarios; intending to find the upper limit or breaking point of the devices under test. Other tests are intended to simulate traffic flows representative of actual production environments.

An initial baseline test was performed in which all traffic was RTP (G.729) streams, without QoS enabled. This was not intended to simulate a specific network, but rather to validate voice quality (latency, jitter and drops would be within acceptable limits) at the upper bounds of the CPU resources of each of the platforms under test.

Then V³PN performance and scalability tests were performed using a converged traffic profile (data and voice) that would be more representative of a real world implementation. In these tests, the traffic profile was as follows:

- UDP—Chariot DNS; script sends 100 bytes in both directions
- UDP—Chariot RTP (VoIPG729) script approximately 33 percent of link capacity (per call units)
- TCP—Chariot HTTPtext script; 100 bytes upstream and 1-to-3K bps downstream
- TCP—Chariot FTPGet and FTPPut script
- TCP—Chariot TN3270
- TCP—Chariot POP3 script; occurs every 1 minute

These different streams were then assigned to realistic QoS traffic classifications as shown in [Table 5-1](#).

Table 5-1 Test Traffic Streams to IP Precedence Mapping

Traffic Stream	IP Precedence
DNS, POP3, FTP	0
RTP (VoIP)	5
HTTP and TN3270	50 percent 2, 50 percent 0

In a production network, EMAIL sent with attachments has characteristics similar to FTP Put or Get, MTU sized packets being sent or received by the user. The POP3 script would represent the text message portion of EMAIL. Also, TN3270 supports screen sizes (lines by columns) of 24x80 (Mod2), 32x80 (Mod3), 43x80 (Mod4) and 27x132 (Mod5). Applications present varying amounts of text on a single screen (transaction), but it is common to see 1-to-2 Kbps downstream, as a 24x80 screen can contain 1920 characters.

Many published test results report performance for IMIX traffic, a traffic pattern developed from packet size samples of Merit’s (Internet) backbone—the “Internet mix” or IMIX. This packet size distribution is typically implemented as 1 packet at 1500 bytes, 4 packets at 512 bytes, and 7 packets at 64 bytes. While this traffic might be representative of the sampled data, it is important for the traffic generation tool to behave as a real TCP or UDP application would on the network. The Chariot test tool used in the

Cisco Enterprise Solutions Engineering lab runs on both Solaris and Linux platforms as end points, and thus uses a real TCP stack, which through adjusting the TCP window size, incorporates a flow control mechanism


Traffic generation tools that simply generate or replay TCP packets without regard to flow control might facilitate testing, but are not necessarily representative of real applications on a live network.

To monitor and represent the above application mix on a branch under test, a Netflow Protocol-Port-TOS aggregation was configured on a Cisco 2651 router with a WAN link speed of 512 Kbps. Flow exports were captured on a Penguin Linux server for 10 minutes. The raw data was captured and then summarized by application with a Perl script that created a CSV file, which was then graphed. The router under test was configured as follows.

```
ip flow-aggregation cache protocol-port-tos
cache timeout inactive 60
cache timeout active 1
export destination 10.254.0.100 7777
enabled
!
```

Using this captured data, the V³PN performance and scalability traffic profile was verified to be that shown in Figure 5-2. Shown are the different traffic streams, average packet sizes, and percentage of bandwidth consumed for that traffic stream. Note that overhead from IPSec and IP GRE are excluded due to the capture method.

Figure 5-2 V³PN Traffic Mix Sample for 512 Kbps Link Speed



The devices under test have a separate network interface which is used for collecting statistics and configuration. During the tests, the devices are polled via SNMP, are logging to a Syslog server, are configured for NTP, and have telnet sessions active. This management and data collection activity does influence the device's CPU and memory characteristics, but does not represent data traffic on the test network.

Additional Voice Quality Validation

To provide a subjective validation of the Chariot reported values an Agilent Technologies Telegra VQT 2.1 (Voice Quality Tester) was connected to the lab network. The handset cable from a Cisco 7960 IP Phone in a branch office location and the head-end location were connected to the Telegra VQT's handset/audio adapter. An audio file (.wav file) was played by the Telegra VQT through one handset adapter and the resulting voice stream was captured through the handset of the second Cisco 7960 phone.

A test was run on a Cisco 2611 branch router, with hierarchical CBWFQ shaping a T1 interface to a 512 Kbps profile:

```
policy-map 512kb
  class call-setup
    bandwidth percent 5
  class mission-critical
    bandwidth percent 22
  class voice
    priority 168
  class class-default
    fair-queue
    queue-limit 6
policy-map 512kb-shaper
  class class-default
    shape average 486400 4864 0
    service-policy 512kb
!
interface Serial0/0
  bandwidth 512
  ip address 192.168.124.6 255.255.255.252
  load-interval 30
  tx-ring-limit 1
  tx-queue-limit 1
  service-policy output 512kb-shaper
  crypto map static-map
!
```

The Chariot 512 Kbps branch traffic profile was run. This profile would normally include three G.729 RTP streams, but in this case, only two RTP streams were included with the third stream being provided by the Cisco 7960 phone. The test was run for 10 minutes, with the Cisco 7960 dialed using the keypad. Since VAD is by default disabled, the Cisco 7960 generates 50 pps for the entire test/call.

During the test, the blue “i” key (item number 1 in [Figure 5-3](#)) was pressed twice in succession to instruct the Cisco 7960 phone to display the call statistics. Average jitter on the two phones ranged from 8-to-20 msec, maximum jitter on one phone was 45 msec and the second 85 msec, RXDISC on both phones was zero and RXLOST was zero on one phone and five on the second.

Figure 5-3 Cisco 7960 Phone (Illustrating Key Pressed to Display Call Statistics Data)



The Telegra VQT can estimate latency including coder, de-jitter buffer and handset delay—*ear-to-mouth* delay is reported—Chariot does not include or report these values. In the test case the Telegra VQT reported one-way delay to be 122-to-130 msec. If this test configuration would be overlaid to an actual ISP which might include an additional 50-to-60 msec one way delay, it would be practical to assume one way delay less than 200 msec could be achieved.

The test was also run at a 1,024 Kbps data rates with the same Cisco 2611 router, five Chariot G.729 streams plus the Cisco 7960 call for a total of 6 voice streams. The CPU utilization during the test was in the 80 percent range with similar results to the 512 Kbps test.

Head-end Product Selection

This section provides recommendations based on the performance and scalability testing defined above to assist in selecting the appropriate head-end device to meet the VPN aggregation requirements offered by the branch sites.

Failover and Head-end Availability

The primary function of a router is to switch packets. Cisco IOS software has been enhanced to include value add features, for example IPSec, DLSw, TN3270 server. These features are either applications in themselves, like TN3270 server, or support and enhance end user applications—as does IPSec—by encrypting data. Functions such as path determination, which is the job of a routing protocol, are overhead, necessary to support packet switching, but overhead nevertheless.

In Cisco Enterprise Solutions Engineering lab performance testing of the design, a goal of the test was to identify at what rate the head-end routers could switch packets at the highest possible rate, yet still reserve sufficient CPU resources to maintain availability in the event of a network component failure. When redundant network devices fail or are taken out of service for maintenance or upgrades, the path determination process (the routing protocol) generates updates to indicate the network topology change. These updates must be processed by the remaining network components, and the act of processing the updates should not cause other outages in the network.

Consider results observed in one test. Two head-end routers were under test, one actively switching packets and second in reserve with no user data traffic. If the active head-end router was failed, the second router was able to pick up the offered load and maintain all EIGRP neighbor relationships. When the standby router was failed, the active router lost EIGRP neighbor adjacencies as there weren't sufficient reserve CPU cycles to process the topology change. Performance must be balanced with availability requirements.

Performance Under Converged V³PN Traffic Profile

Head-end platforms were configured to aggregate 120 branch offices to the device as primary tunnels and 120 branches as secondary tunnels. The converged (voice and data) V³PN traffic profile discussed in the “[Traffic Profiles](#)” section on page 5-3 was then applied and increased while performance of the platform was observed. The upper performance boundary was established by measuring five primary factors:

1. CPU utilization (less than 80 percent)
2. End-to-end RTP packet latency, peak and average (less than 50 msec)
3. End-to-end RTP packet jitter (less than 10 msec)
4. End-to-end RTP bytes lost (less than 0.5 percent)
5. Failover (no loss of peers with failover)

If any factor above was unacceptable, this was considered the breaking point for the product evaluated.

In addition, several other key parameters were also monitored during the evaluation, including:


- Percentage of packets dropped due to IPSec anti-replay and QoS interaction
- Percentage of packets dropped due to the QoS Service Policy
- Portion of packets being process switched
- Occurrences of crypto congestion

The results of this evaluation are summarized in [Figure 5-4](#).



Note

The scale in [Figure 5-4](#) was modified to accommodate the Cisco 6550 performance results.

Figure 5-4 Head-end Performance—Converged Traffic Case

Impact of QoS on VPN Head-end Performance

The design recommendation for V³PN head-end architecture is to deploy separate devices for WAN aggregation (ISP connection) and VPN tunnel aggregation. WAN aggregation devices would have QoS enabled. The VPN tunnel aggregation devices would typically have Fast Ethernet connectivity to the WAN aggregation devices (no interface congestion).

However, it is possible to have QoS enabled on the same head-end device which is providing VPN tunnel aggregation. This scenario was also evaluated during the scalability and performance testing.

Comparing the relative performance limits at approximately 80 percent CPU shows that having QoS enabled on the VPN head-end device in this evaluation resulted in an approximate 10 percent performance degradation for V³PN throughput. Slightly higher end-to-end latency was also observed, primarily due to the affects of congestion on the output interface since QoS is now operating on the device.

Head-End Scalability and Performance Observations

The following are a summary of the primary observations and findings from the V³PN scalability and performance evaluation of head-end products:

- For both traffic profiles, there is no significant difference in performance between the SA-ISA/ISM and SA-VAM hardware-accelerated encryption cards. The limitation on the Cisco 7200 (NPE-300/NPE-400) is the CPU utilization for forwarding packets.
- As expected, the 7200VXR NPE-G1 was the highest performing Cisco IOS VPN router platform of those evaluated, supporting up to 240 G.729 calls plus approximately 50 Mbps of data in a converged traffic configuration. It should be noted that adding a second VAM to the NPE-G1 platform had very little effect on overall throughput with the converged voice and data traffic profile.
- The Catalyst 6500 with VPN Service Module was also evaluated up to the current 1 Gbps traffic limit of the lab, handling 3,180 G.729 calls and 607 Mbps of simultaneous data traffic. This was in an IPSec-only configuration because GRE currently degrades overall performance of the Catalyst 6500.
- Discounting for the estimated WAN delay in the scalability lab of approximately 4-to-5 msec, end-to-end latency performance was very good: approximately 4-to-5 msec in the voice-only configuration and 15-to-20 msec in the converged traffic configuration.
- Although anticipated to be a major concern, the interaction between IPSec anti-replay and QoS resulted in less than 1 percent of data packets being dropped, primarily due to the positive affect of TCP flow-control. Tuning of the queue-limit parameter in the CBWFQ Service Policy can further reduce drops due to IPSec anti-replay to approximately 0.01 percent of packets.
- The number of VPN tunnels aggregated has significant impact on head-end platform performance. For example, a Cisco 3745 platform running an AIM-II encryption card can process up to 150 G.729 calls over a single tunnel with approximately the same CPU utilization required to process 60 G.729 calls over 60 separate tunnels.
- The impact of enabling QoS on the same VPN head-end device versus on a separate WAN aggregation device resulted in approximately 10 percent performance degradation.
- Monitoring of crypto congestion showed that only under the most extreme cases (for example running the Cisco 7200 above 90 percent CPU utilization) did crypto report congestion. Therefore it is believed with the traffic profile used for evaluation, the LLQ for Crypto feature would not engage, and would not provide additional benefit for head-end platforms.
- However, it is also believed that for organizations with a larger percentage of large packet sizes (such as video), or with applications that might cause *bursts* of large packet sizes, it might be possible to induce congestion on the crypto engine, and in these cases the LLQ for Crypto feature would provide additional insurance that rises in latency did not occur.
- The gating factor for router performance, whether packets are encrypted or clear text, is packets per second, not bytes per second. It requires a similar amount of CPU cycles to encrypt and switch a 64-byte packet as it does a 1400-byte packet. When VoIP is added to the traffic mix, there is a corresponding decrease in average packet size. Consequently, overall throughput in terms of megabits per second is lower with VoIP in the traffic mix because the average packet size decreases.

Branch Office Product Selection

This section provides recommendations based on the performance and scalability testing defined above to assist in selecting the appropriate branch device to meet the VPN requirements for the branch sites.

Due to the large number of branch devices deployed in enterprise networks, the time and expense of installation and upgrades must be balanced with the initial cost and ability to support future network capacity requirements. Many factors besides V³PN performance might need to be considered for branch office router requirements, including:

- What other services will the branch router be performing, such as Cisco IOS Firewall, WAN connection, voice gateway, SRST, or DLSw?
- What is the WAN connectivity, for example point-to-point, DSL, Cable, ISDN, etc.?
- What *future-proofing* factors need to be considered, such as growth potential, new applications on the horizon, etc.?

Product Applicability by Link Speed

This section provides a summary of the ability of different branch platforms to be able to support a given link speed. It can be used as a guideline for deploying the appropriate platform to match up with the network requirements for connection to the service provider.

Table 5-2, Table 5-3, and Table 5-4 summarize the link speeds tested, the maximum number of concurrent G.729 calls for a given link speed, and applicable Cisco VPN router platforms.

Table 5-2 Applicability of Current Products by Link Speed Less Than or Equal to E1

Link Speed (Kbps)	G.729 calls (at 33% of link bandwidth)	831	1751	1760	2611XM	2621XM	2651XM	3660 (AIM)	2691 (AIM)	3725 (AIM)	3745 (AIM)
128	1	x ¹	x	x	x	x	x	x	x	x	x
256	2	x	x	x	x	x	x	x	x	x	x
512	3	x	x	x	x	x	x	x	x	x	x
768	4		x	x	x	x	x	x	x	x	x
1024	6		x	x	x	x	x	x	x	x	x
1280	7		x	x	x	x	x	x	x	x	x
T1 (1536)	9			x		x	x	x	x	x	x
E1 (2048)	12							x	x	x	x

1. x = Supported

Table 5-3 Applicability of Current Products by Link Speed Greater than E1

Link Speed (Mbps)	G.729 calls (at 33% of link bandwidth)	3660 (AIM)	2691 (AIM)	3725 (AIM)	3745 (AIM)	3660 (AIM-II)	2691 (AIM-II)	3725 (AIM-II)	3745 (AIM-II)
3	18	x ¹	x	x	x	x	x	x	x
4	24			x	x	x	x	x	x

Table 5-3 Applicability of Current Products by Link Speed Greater than E1

Link Speed (Mbps)	G.729 calls (at 33% of link bandwidth)	3660 (AIM)	2691 (AIM)	3725 (AIM)	3745 (AIM)	3660 (AIM-II)	2691 (AIM-II)	3725 (AIM-II)	3745 (AIM-II)
5	30					x	x	x	x
10	60					x	x	x	x
15	90								x
20	120								x
25	150								x

1. x = Supported

Table 5-4 Applicability of Legacy Products by Link Speed

Link Speed (Kbps)	G.729 calls (at 33% of link bandwidth)	806	2611	2621	2651	3620	3640
128	1	x ¹	x	x	x	x	x
256	2		x	x	x	x	x
512	3		x	x	x	x	x
768	4			x	x		x
1024	6				x		x
1280	7						
T1 (1536)	9						
E1 (2048)	12						

1. x = Supported

Performance Under Converged V³PN Traffic Profile

Branch platforms were configured with a primary and secondary IPSec/GRE tunnel back to two separate head-ends. The converged (voice and data) V³PN traffic profile discussed in the [“Traffic Profiles” section on page 5-3](#) was then applied and increased while performance of the platform was observed. The upper performance boundary was established by measuring four primary factors:

- CPU utilization (less than 80 percent)
- End-to-end RTP packet latency, peak and average (less than 50 msec)
- End-to-end RTP packet jitter (less than 10 msec)
- End-to-end RTP bytes lost (less than 0.5 percent)

If any factor above was unacceptable, this was considered the “breaking” point for the product evaluated. In addition, several other key parameters were also monitored during the evaluation, including:

- Percentage of packets dropped due to IPSec anti-replay and QoS interaction
- Percentage of packets dropped due to the QoS Service Policy
- Portion of packets being process switched
- Occurrences of crypto congestion

The results of this evaluation are summarized in [Figure 5-5](#) through [Figure 5-8](#).

Figure 5-5 Cisco IOS High-End Branch Router Performance for V³PN




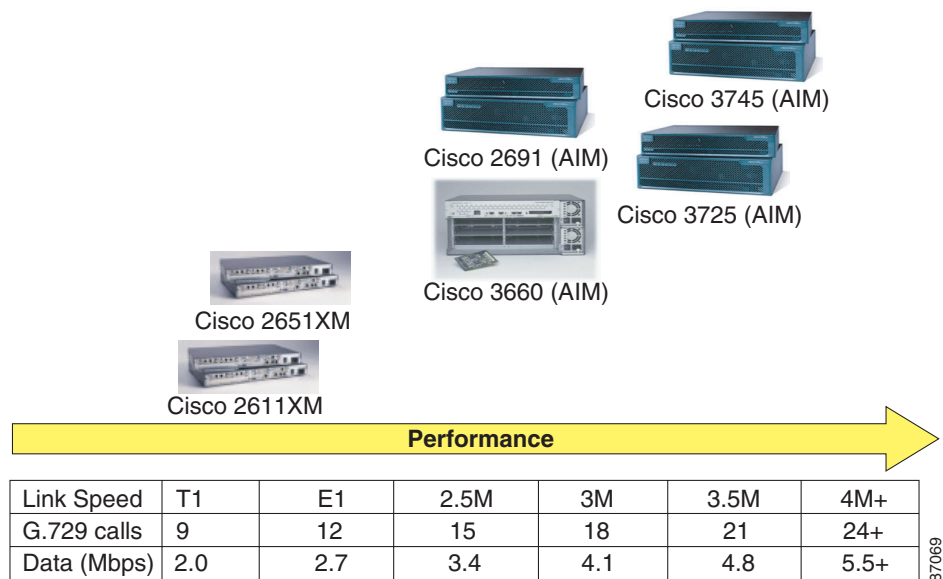
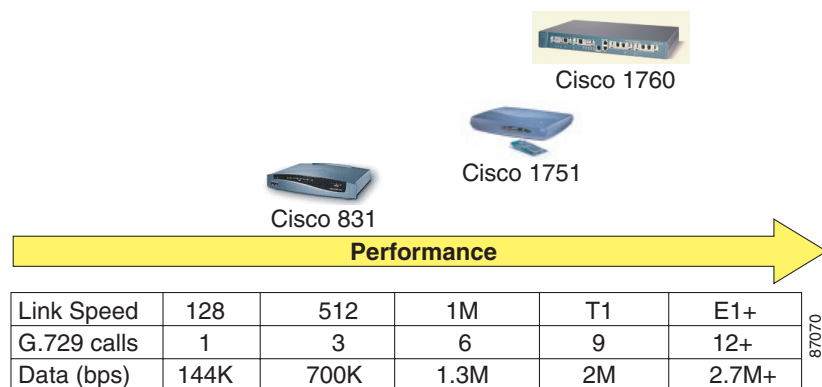
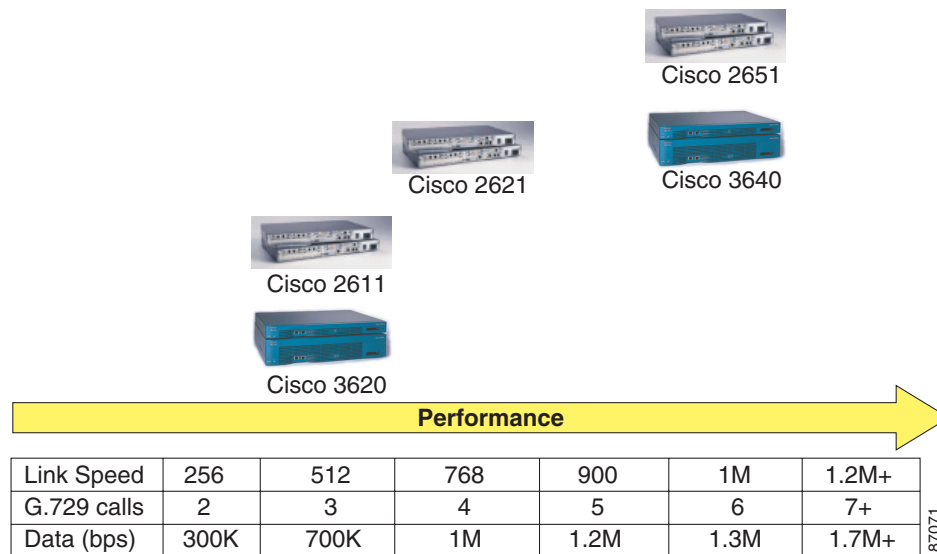
Figure 5-6 V³PN Performance—Current High/Mid-Range Branch Products**Figure 5-7 Cisco IOS VPN SMB/Small Branch Router—Performance for V³PN**

Figure 5-8 Cisco IOS VPN Legacy Branch Router—Performance for V³PN

Branch Scalability and Performance Observations

The following are a summary of the primary observations and findings from the V³PN scalability and performance evaluation of branch products:

- Measured end-to-end latency performance was of the solution was very good: approximately 8-to-9 msec in the voice-only configuration and 20-to-25 msec in the Converged traffic configuration (excluding the estimated WAN delay of 4-to-5 ms).
- Although anticipated to be a major concern, the interaction between IPSec anti-replay and QoS resulted in less than 1 percent of data packets being dropped, given the traffic profile discussed earlier, of predominantly TCP-based traffic. Tuning of the queue-limit parameter in the CBWFQ Service Policy can further reduce drops due to IPSec anti-replay to approximately 0.01 percent of packets.
- Monitoring of crypto congestion showed that in many cases crypto did not experience congestion. Therefore it is believed with the traffic profile used for evaluation, the LLQ for Crypto feature would not engage, and would not provide additional benefit for the majority of branch platforms (Cisco 1700, Cisco 2600 and Cisco 3600 VPN routers).
- The Cisco 3660, Cisco 2691 and Cisco 3700 VPN routers would, because of their much higher CPU capacity, induce congestion on the crypto engine (AIM). Therefore it is believed that the LLQ for Crypto feature would engage on these platforms providing additional insurance that rises in latency did not occur, and that if over-committed, lower priority traffic is dropped, preserving voice traffic. However, with the next generation AIM-II for these platforms, CPU will again be the bottleneck.
- It is also believed that for networks with a larger percentage of large packet sizes (such as video), or with applications that might cause *bursts* of large packet sizes, it might be possible to induce congestion on the crypto engine, and in these cases the LLQ for Crypto feature would provide additional insurance that rises in latency did not occur.
- The Cisco 3660, Cisco 2691 and Cisco 3700 VPN routers performed well. With current AIM hardware encryption accelerator cards, the Cisco 3660 and Cisco 2691 handled up to 3 Mbps link speeds, while the Cisco 3700s handled up to 4 Mbps link speeds.

- The Cisco 3660, Cisco 2691 and Cisco 3700 VPN routers with the latest AIM-II hardware encryption accelerator cards performed exceptionally well. The Cisco 3660, Cisco 2691, and Cisco 3725 handled up to 60 simultaneous calls on 10 Mbps link speeds, while the 3745 handled up to 150 simultaneous calls at a 25 Mbps link speed. Keep in mind that testing was performed on these platforms in a branch VPN router scenario, meaning two IPSec/GRE tunnels (primary and secondary) established to the head-end VPN routers. If deployed as a hub/head-end tunnel aggregation device, performance would be expected to decrease as the number of IPSec/GRE tunnels is increased.
- The Cisco 806 VPN router has some limitations if deployed as a V³PN device. Since the Cisco 806 does not support hardware-accelerated encryption, latency performance can quite easily be disrupted by data traffic or issuing CLI commands and cause voice latency spikes in a range of 500 msec to 1 second. Therefore, deploying this platform is not recommended. The Cisco 831 should be deployed instead.
- The Cisco 1751/1760 VPN routers performed quite well, handling up to a T1 link speed of converged traffic. However, the Cisco IOS software revision required (containing QoS Pre-Classify) is a special branch, 12.2(4)YB. Also, process switching was occurring on the majority of packets on the 1700.
- The Cisco 1721 VPN router was not evaluated; however, since it is similar in CPU and VPN hardware-acceleration, V³PN performance should be analogous to Cisco 1751 and Cisco 1760 performance.
- The legacy Cisco 2600/3600 VPN routers were found to have limitations when subscribed to higher link speeds (such as T1) due to the nature of their packet switching characteristics. Therefore the recommendation is to follow the guidance given in the [“Product Applicability by Link Speed” section on page 5-10](#).
- The Cisco 2600XM VPN routers were evaluated and do not appear to have the limitations identified on older Cisco 2600 VPN router platforms. Cisco 2600XM platforms handled up to T1 link speeds of converged V³PN traffic.
- The scalability evaluation was performed with both Frame Relay and HDLC as the L2 encapsulation. Performance was fairly comparable with HDLC providing slightly higher link utilization.

Network Performance/Convergence

Each organization might have different convergence time requirements. The design principles in this guide were used to perform a scalability test with 240 branch offices aggregated to two Cisco 7200 NPE-400 head-end devices.

Two aggregation configurations were evaluated:

- Active/Standby—One head-end device was loaded with all 240 primary IPSec/GRE tunnels (170 with active traffic streams), while the second is in a *standby* mode configured with all 240 secondary tunnels.
- Active/Active—Each of the two head-end devices is configured with 120 primary (85 with active traffic streams) and 120 secondary IPSec/GRE tunnels, fairly equal loading on both head-ends.

The test was performed by powering off one of the head-end devices to simulate a complete failure. In this test, the network fully converged after approximately 20-to-23 seconds. The starting and failover traffic/tunnel aggregation conditions are shown below in [Table 5-5](#).

Table 5-5 Head-end Failover Scenario

	Head-end 1	Head-end 2
Active-Standby Active Fails		
Starting Condition	52.0 Mbps total traffic 170 calls 240 primary tunnels 80 percent CPU	120 Kbps total traffic 0 calls 240 secondary tunnels 1 percent CPU
During Simulated Failure	Simulated Failure	61.2 Mbps total traffic 170 calls 240 primary tunnels 79 percent CPU
Active-Standby Standby Fails		
Starting Condition	51.7 Mbps total traffic 170 calls 240 primary tunnels 76 percent CPU	120 Kbps total traffic 0 calls 240 secondary tunnels 1 percent CPU
During Simulated Failure	60.3 Mbps total traffic 170 calls 240 primary tunnels 83 percent CPU	Simulated Failure
Active-Active Active Fails		
Starting Condition	25.9 Mbps total traffic 85 calls 120 primary tunnels 120 secondary tunnels 46 percent CPU	25.8 Mbps total traffic 85 calls 120 primary tunnels 120 secondary tunnels 46 percent CPU
During Simulated Failure	Simulated Failure	50.2 Mbps total traffic 170 calls 240 primary tunnels 80 percent CPU

**Note**

The traffic and voice call loads were intentionally raised near or slightly above recommended operating limits to characterize worst-case failover conditions. It is not recommended to operate platforms at the extremes shown in this table.

While all VoIP streams experienced packet loss during the failure, after routing convergence all simulated call streams continued to function within acceptable end-to-end latency, jitter and packet loss limits.

In all scenarios, the failed head-end device was then powered back on, resulting in the network re-converging in less than two seconds. The IPSec tunnels re-established a few at a time as their corresponding SAs were renegotiated. The last IPSec tunnels re-established connectivity after 3-to-4 minutes.

Software Releases Evaluated

The software releases shown in Table 5-6 were used in the V³PN scalability and performance evaluation.

Table 5-6 Software Releases Evaluated

Cisco Product Family	SW Release	Notes/Other Information
Cisco 7200VXR VPN Routers	Cisco IOS software 12.1(9)E	3DES IPsec support C7200-IK2S-M
Cisco 3700 Series VPN Routers	Cisco IOS software 12.2(11)T1	3DES IPsec support C3725-IK9O3S-M
Cisco 3600 Series VPN Routers	Cisco IOS software 12.2(11)T1	3DES IPsec support C3640-IK9O3S-M
Cisco 2600 Series VPN Routers	Cisco IOS software 12.2(11)T1	3DES IPsec support C2600-IK9O3S-M
Cisco 1700 Series VPN Routers	Cisco IOS software 12.2(4)YB	3DES IPsec support C1700-K9O3SY7-M
Cisco 800 Series VPN Routers	Cisco IOS software 12.2(8)YN	3DES IPsec support C831-K9O3SY6-M
Cisco 7500 WAN Aggregation Routers	Cisco IOS software 12.2(4)XV4	RSP-JSV-M
Cisco 6x00 Catalyst Switches	CatOS 5.5(3)	
Cisco 2948G Catalyst Switches	CatOS 4.5(9)	



Note

Several Cisco IOS software images exist, each configured with various levels of encryption technology. There are certain restrictions and laws governing the use and export of encryption technology.

Before selecting Cisco IOS software, perform the appropriate research on www.cisco.com and consult the appropriate support channels. It is important understand issues inherent to specific levels of Cisco IOS software code that might affect other features configured on the network's routers.

