

**Comparative Performance of Layer 2 and IPSec Encryption
on Ethernet Networks**

October 2006

Rochester Institute of Technology
Networking, Security and Systems Administration Department

Luther Troell
Bruce Hartpence
Seth Simons

Table of Contents

Table of Contents	2
List of Figures	3
List of Tables	3
Introduction.....	5
Introduction.....	5
Discussion of Previous Study	5
Technology Overview.....	5
IPsec	5
Ethernet Encryption	6
Equipment Overview	7
Cisco Systems Catalyst 6500	7
Cisco Systems IPsec VPN Services Module	7
SafeNet SafeEnterprise™ Ethernet Encryptor (SEE).....	7
Ixia Hardware and Software	8
Device Configuration and Topology	9
Scenario 1 – Infrastructure Baseline	9
Scenario 2 – SEE Unencrypted Baseline	10
Scenario 3 –SEE Encryption Activated	10
Test Methodology	11
Frame loss	11
Throughput.....	12
Latency.....	15
Discussion of Results	16
Frame loss	16
Throughput.....	16
Scenario 1 – Infrastructure Baseline	16
Scenario 2 – SEE Unencrypted Baseline	16
Scenario 3 – SEE Encryption Activated	16
Latency.....	16
Scenario 1 – Infrastructure Baseline	16
Scenario 2 – SEE Unencrypted Baseline	17
Scenario 3 – SEE Encryption Activated	17
Results of Previous IPsec Study	19
Frame Loss.....	20
Throughput.....	21
Latency.....	23
Comparison of Results	24
Frame Loss.....	24
Throughput.....	24
Latency.....	25
Conclusion	27
Recommendations for Future Studies.....	27

Acknowledgements.....	27
References.....	27
Appendices.....	29
Appendix A - Equipment List.....	29
Appendix B – Diagrams.....	30
Infrastructure Baseline Topology	30
Unencrypted Topology with SafeNet SafeEnterprise™ Ethernet Encryptors.....	30
Encrypted Topology with SafeNet SafeEnterprise™ Ethernet Encryptors	31
Encrypted Topology with Cisco IPsec Encryption.....	31
Appendix C – Infrastructure Configuration Information.....	32
Cisco 6500 Router - Timmy	32
Cisco 6500 Router - Jimmy	34
Safenet Ethernet Encryptor - SEE_Top_192.168.0.2	36
Safenet Ethernet Encryptor - SEE_Bottom_192.168.0.3.....	36

List of Figures

Figure 1 - IPsec Packet Header in Tunnel Mode	6
Figure 2 - Encrypted Ethernet Frame Format	6
Figure 3 - Cisco 6509 Catalyst Switching Chassis	7
Figure 4 - Cisco IPsec VPN Module.....	7
Figure 5 - SafeNet SafeEnterprise™ Ethernet Encryptor (SEE)	7
Figure 6 - Ixia 250 Test Platform.....	8
Figure 7 – Topology for Test Scenario 1	9
Figure 8 - Topology for Test Scenario 2.....	10
Figure 9 - Topology for Test Scenario 3.....	10
Figure 10 - Comparison of Theoretical Throughput.....	14
Figure 11 - Latency Added by SEE	18
Figure 12 - IPsec Test Topology from Previous Study.....	19
Figure 13 - IPsec Frame Loss by Percent Throughput and Frame Size.....	20
Figure 14 - IPsec Frame Loss	20
Figure 15 - IPsec Throughput Data.....	21
Figure 16 - Comparative Bandwidth Utilization of IPsec	22
Figure 17 - IPsec Latency	23
Figure 18 - Comparative Encrypted Throughput Data	24
Figure 19 - Comparative Latency of IPsec and Ethernet Encryption	26
Figure 20 - Infrastructure Baseline Topology.....	30
Figure 21 - Unencrypted Topology with SafeNet SafeEnterprise™ Ethernet Encryptors.....	30
Figure 22 - Encrypted Topology with SafeNet SafeEnterprise™ Ethernet Encryptors ...	31
Figure 23 - Encrypted Topology with Cisco IPsec Encryption	31

List of Tables

Table 1 - Theoretical Ethernet Performance.....	12
Table 2 - IPsec ESP Tunnel Mode Overhead	13

Table 3 - Theoretical IPsec Performance.....	13
Table 4 - Ethernet Latency Data.....	17
Table 5- IPsec Throughput Data.....	21
Table 6 - IPsec Latency Data.....	23
Table 7 - IPsec and Ethernet Encryption Latency Comaprison.....	26

Introduction

The emergence of Metropolitan Ethernet and the Ethernet Wide Area Network (WAN) services market, and the growth of global data security and privacy regulations has necessitated the need for strong, interoperable, high-speed encryption technology. As was demonstrated in a previous study, IPsec encryption can create significant network bottlenecks. This study examines the performance characteristics of native Layer-2 Ethernet encryption as a solution for high-speed network security in MAC-transparent Layer-2 networks.

The study, commissioned by SafeNet, was performed by examining a series of performance metrics for SafeNet's SafeEnterprise™ Ethernet Encryptor (SEE) and comparing them to the same performance characteristics for a leading IPsec solution over the same network topology. The metrics examined were latency, throughput, and available bandwidth

Discussion of Previous Study

This study builds upon the RIT study *Dedicated vs. Converged Encryption Appliance*, which compared the performance benefits of dedicated IPsec encryptors to IPsec routers in high-speed Ethernet networks. In the previous study, RIT examined the differences in performance characteristics between a dedicated IPsec encryption device and converged network equipment with integrated IPsec encryption capabilities. It was found that IPsec solutions, in general, can consume up to 40 percent of available bandwidth. As frame sizes decrease, the ratio of overhead to data traffic increases. With 64-byte frames for example, such as those used for real-time video and audio applications, every 64 bytes of data carries 57 bytes of overhead. Furthermore, it was found that the converged IPsec solution also added significant levels of latency to the unencrypted baseline, particularly at small frame sizes. This study uses the converged IPsec results from the previous study to make comparisons and draw conclusions between the performance characteristics of IPsec and Ethernet encryption.

Technology Overview

IPsec

IPsec is a standard for encryption of IP packets to provide confidentiality and integrity over networks (see RFC 4301, *Security Architecture for the Internet Protocol*). The IPsec standard defines two primary modes of operation: transport mode and tunnel mode.

In transport mode, IPsec has less data overhead but does not provide confidentiality for the Layer-3 IP header. This means that sensitive information about the addressing of the internal network can be maliciously acquired by monitoring the public network over which the traffic travels. This presents an unacceptable level of risk to many organizations.

Tunnel mode IPsec addresses this security concern by encrypting the entire IP packet and encapsulating it into another IP packet. This packet contains only the address of the encryption devices at either end-point and not of the actual hosts communicating on the internal network. While tunnel mode does address the security and privacy concerns of transport mode IPsec, it also adds a significant amount of data overhead. Processing of this extra IP header also has some performance issues in terms of latency, which will be illustrated later.




Figure 1 - IPsec Packet Header in Tunnel Mode

Ethernet Encryption

Ethernet encryption is a new technology designed by SafeNet, Inc. to address the issues encountered when using IPsec tunnel mode. This technology relies on a dedicated encryption device called the SafeNet SafeEnterprise™ Ethernet Encryptor (SEE). This device is placed on the network edge, similar to an IPsec solution, and encrypts the entire IP packet without adding the overhead of an additional IP header. The SEE operates on Ethernet based Layer-2 MAC transparent networks only. This means that as the Layer-2 Ethernet frame moves through the intermediary networks between the two primary sites, the MAC address of the original frame must not be altered. Because routers, operating at Layer-3, change the MAC address of the Ethernet frame, the encrypted frame can not pass through a router prior to being decrypted. This means that SEE technology is well suited for Metropolitan Ethernet or Ethernet WAN services, as well as remote backup, Storage Area Network, Data Center, and Business Continuity/Disaster Recovery applications. It is not applicable for Layer 3 applications, such as remote access over public networks.

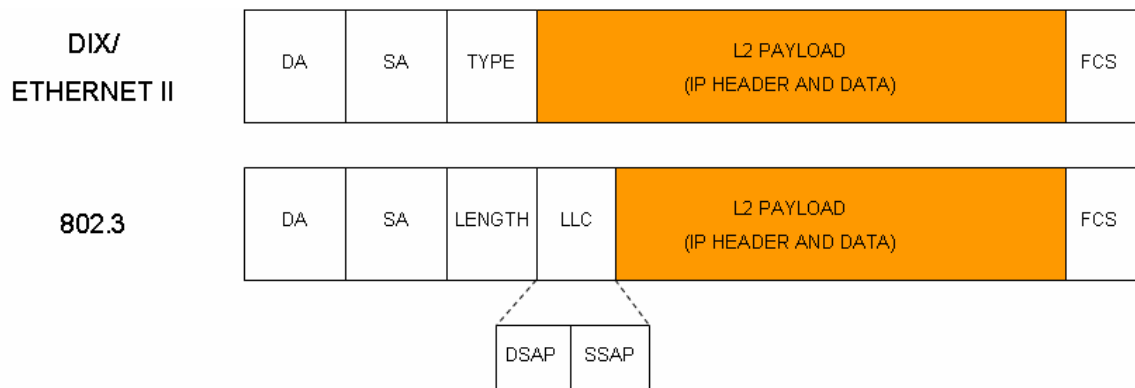


Figure 2 - Encrypted Ethernet Frame Format

Equipment Overview

This section provides a brief overview of the equipment used in this study, as well as the Cisco Systems IPsec blade in the last study, to solidify the reader's understanding of the test scenarios presented in the next section.

Cisco Systems Catalyst 6500

The Cisco Systems Catalyst 6500 series switching chassis is a widely used enterprise class high-performance switch. This platform is scalable and can accept expansion cards for multiple transport technologies and encryption acceleration capabilities. The model used in this study was the Cisco Catalyst 6509 Switching Chassis.



Figure 3 - Cisco 6509 Catalyst Switching Chassis

Cisco Systems IPsec VPN Services Module

The Cisco Systems IPsec VPN module provides IPsec encryption capabilities to the Cisco Catalyst 6500 series switching chassis.




Figure 4 - Cisco IPsec VPN Module

SafeNet SafeEnterprise™ Ethernet Encryptor (SEE)

The SafeNet SafeEnterprise™ Ethernet Encryptor (SEE) is a dedicated high-speed encryption appliance providing Layer-2 network traffic encryption on Ethernet networks.



Figure 5 - SafeNet SafeEnterprise™ Ethernet Encryptor (SEE)

Ixia Hardware and Software

Ixia is a leading provider of high-performance network test equipment. The Ixia 250 chassis is a dedicated hardware platform designed to run their software test suite.

IxExplorer is a graphical user interface to the Ixia hardware platform for performing throughput, latency, and other network performance tests.



Figure 6 - Ixia 250 Test Platform

Device Configuration and Topology

The test topologies were deliberately simple to limit the number of factors affecting the results. Each topology contained two (2) Cisco 6509 switching chassis, each with a single gigabit Ethernet line card, a supervisor module, and two power supplies, as well as the Ixia 250 testing and analysis hardware as shown below. All devices used were connected with UTP, CAT-5 copper cabling.

Each test scenario contained three IP networks routed over an Ethernet Data Link Layer (Layer-2) designed to simulate a basic Ethernet Wide Area Network (WAN). Networks A and C represent private networks which must communicate securely over network B which represents an untrusted WAN. The simplicity of this network topology allowed the configuration of the Ixia test suite to remain static, which simplified testing and helped prevent configuration inconsistencies that could have adversely affected the accuracy of the results.

Scenario 1 – Infrastructure Baseline

The first of the test topologies consisted of the two Cisco 6509 chassis and the Ixia 250 testing and analysis hardware. The purpose of this scenario was to generate a comparative baseline for the results of the encryption tests to be carried out later.




Figure 7 – Topology for Test Scenario 1

Scenario 2 – SEE Unencrypted Baseline

The second test topology added Ethernet encryptors configured in bypass mode. In bypass mode, the encryptors forward all traffic without encrypting or discarding anything. The purpose of this test was to ascertain the effect of adding additional hardware to the data path without encryption. It was anticipated that this data would be useful when analyzing the results later. It was hypothesized that the additional hardware would introduce a small amount of latency, as well some frame loss, leading to a possible decrease in throughput.

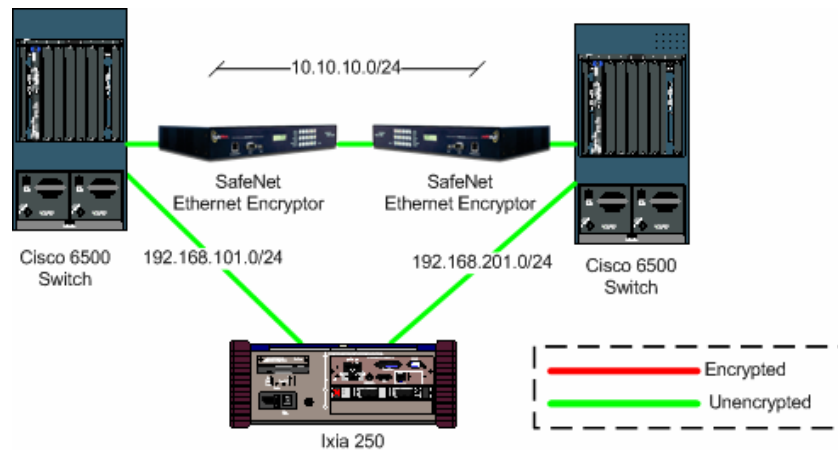


Figure 8 - Topology for Test Scenario 2

Scenario 3 –SEE Encryption Activated

The last test topology was identical to the Scenario 2 topology, except the Ethernet encryptors were configured to encrypt all data. This scenario was expected to add the greatest amount of latency and frame loss, resulting in diminished throughput.

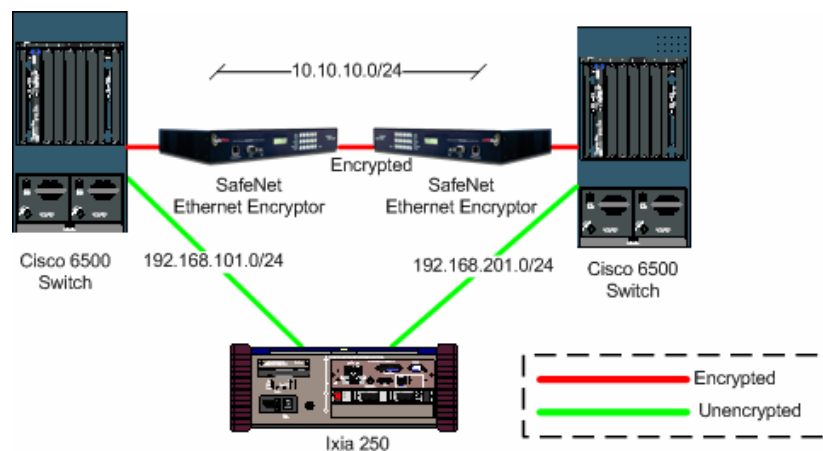


Figure 9 - Topology for Test Scenario 3

Test Methodology

The testing and data collection focused on measuring the performance characteristics of the baseline and SEE, Layer-2 encryption device. Tested metrics included latency, throughput, and frame loss. This data was the basis for a comparative performance analysis of SEE and IPsec encryption, using the results from our previous study. The comparative metrics examined were overhead, available bandwidth, and latency. Further analysis compared the actual measured results to the theoretical performance capabilities of each technology.

When examining throughput, it must be noted that the available bandwidth of a gigabit link is reduced by the overhead added by the Ethernet protocol. This is independent of the variables tested here and applies to each scenario. It is necessary to account for this overhead when comparing measured performance to the stated performance of the link. This issue is discussed in depth later.

All tests were performed at gigabit data rates over Category 5 UTP cables. Each test was bidirectional so that the full capacity of the link was tested up to two (2) gigabits. Each of the three test metrics, throughput, frame loss and latency, were examined at the following frame sizes: 64-byte, 128-byte, 256-byte, 512-byte, 1024-byte, 1280-byte and 1420-byte. These are the same frame sizes used in the previous RIT study and, as described in RFC 2544, *Benchmarking Methodology for Network Interconnect Devices*, provides a standard methodology for benchmarking networking devices. Note that frames over 1500 bytes, also known as “jumbo frames,” were not tested.

An equipment list is available in Appendix A and additional diagrams are available in Appendix B.

Frame loss

Simply put, frame loss is the difference between the number of frames transmitted by one interface and the number frames received by another. The purpose of the frame loss test was to determine the maximum speed at which the network devices could operate without frame loss. This would indicate the maximum throughput permissible at a particular frame size, as well as performance as throughput increases. For any frame size, frame loss was calculated using the following formula:

$$\text{frames/sec} = (([\text{frames transmitted (local)}]) - ([\text{frames received (remote)}]))/60$$

For any given frame size the methodology used was as follows:

:

1. Start the Ixia test engine at maximum line rate
2. Frame loss is encountered, reduced the data rate by 10% of line speed
3. Record the number of frames lost at each increment
4. Repeated ten times for each frame size
5. Graph the average

Throughput

A standard Ethernet frame has the capacity to carry a 1500-byte payload. These 1500-bytes do not include the CRC, MAC, Control Field, UDP, Link, or Ethertype data..

Ethernet includes a preamble of 8 bytes at the start of the frame and an inter-frame gap of 12 bytes which together reduce the maximum theoretical throughput of a given link. For a given frame size, the maximum theoretical throughput, and number of frames per second for a given link, can be calculated with the following formulas (note that this includes Ethernet, UDP, and IP headers):

$$\text{Maximum Theoretical frames/sec} = \text{link speed} / ((\text{frame size} + 12 \text{ bytes [inter-frame gap]} + 8 \text{ bytes [preamble]}) * 8[\text{convert bits to bytes}])$$

$$\text{Maximum Theoretical Throughput (Mbit/sec)} = \text{Maximum frames/sec} * \text{frame size} * 8[\text{convert bits to bytes}]$$

Applying these calculations to the frame sizes chosen for this study yields the following table.

Ethernet Frame Size	Max Ethernet (Mbps)	Max Ethernet Frames/Sec
64	762	1488095
128	865	844595
256	928	452899
512	962	234962
1024	981	119732
1280	985	96154
1420	986	86806

Table 1 - Theoretical Ethernet Performance

Adding IPsec encryption to a network link further reduces the maximum payload by adding a second IP header which encapsulates the IP header of the original packet as well as an additional authentication header and trailer. The previous study performed by RIT found that generally the cost is an additional 57 bytes of overhead. The table below shows how these additional bytes are allocated in the IPsec headers.

IPsec ESP Tunnel Mode Overhead (bytes):

IP header encapsulation	20
SPI	4
Sequence Number	4
Initialization Vector	8
Pad Size	1
Next Protocol	1
Padding	7
Authenticator	12
Total	57

Table 2 - IPsec ESP Tunnel Mode Overhead

When IPsec is used with the parameters described above, for a given frame size the actual maximum theoretical throughput and number of frames per second of a given link can be calculated with the following formulas (note that this includes Ethernet, UDP, IP, and IPsec headers).

$$\text{Maximum theoretical frames/sec} = \text{link speed} / ((\text{frame size} + 12 \text{ bytes [inter-frame gap]} + 8 \text{ bytes [preamble]} + 57 \text{ bytes [IPsec overhead]}) * 8[\text{convert bits to bytes}])$$

$$\text{Maximum Theoretical Throughput (Mbit/sec)} = \text{Maximum frames/sec} * \text{frame size} * 8[\text{convert bits to bytes}]$$

Applying these calculations to the frame sizes chosen for this study yields the following table.

Ethernet Frame Size	Max IPsec Tunnel (Mbps)	Max IPsec Frames/Sec
64	454	886525
128	624	609756
256	769	375375
512	869	212224
1024	930	113533
1280	943	92115
1420	949	83500

Table 3 - Theoretical IPsec Performance

As can be seen in Figure 10, the addition of 57 bytes has a significant impact at smaller frame sizes which are commonly found in real-time voice and data applications. This impact becomes less significant as the frame size increases. Note that these are theoretical numbers, not actual performance data.

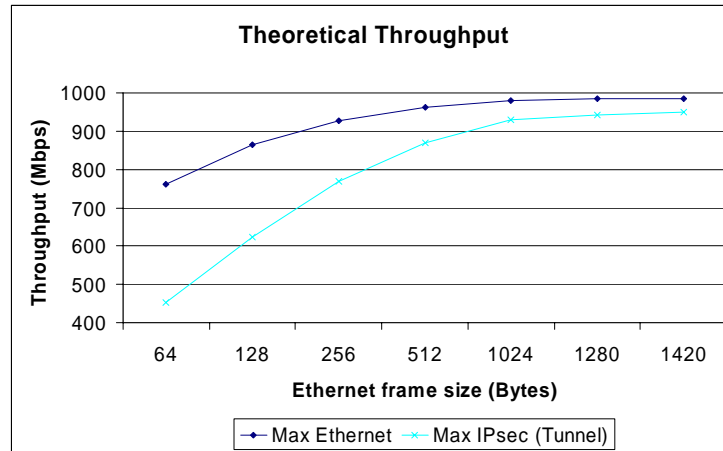


Figure 10 - Comparison of Theoretical Throughput

Maximum throughput is the maximum rate at which a given link can sustain data transfer with no data loss. Starting at the maximum theoretical line rate, this test was performed by sequentially reducing the output of the Ixia traffic engine until there was no frame loss for a given frame size. Again using the Ixia device, these results were then verified with one minute trials to ensure this was the correct data rate. The measured results were compared to the maximum theoretical throughput (MTT) calculations shown above. The relationship between MTT and the measured throughput without frame loss was graphed. Frame counts rather than throughput were measured because throughput in this case is calculated based on the number of frames transmitted and received to account for the overhead discussed above.

The specific methodology used was as follows:

1. Select a frame size
2. Begin a one minute test at maximum throughput with zero loss from the Ixia traffic engine
 - a. Collect bi-directional frame counts as the test progresses
 - b. Calculate lost frames (if any)
3. Repeat for a total of nine iterations
4. Compute an average for all of the test data
5. Calculate percentage of maximum theoretical throughput for that frame size for the 1 minute test using the formula below

Repeat for each frame size.

Below are the formulas used for calculating the percentage of MTT achieved. They can be used interchangeably.

$$\% \text{ of MTT} = [\textit{Throughput Measured}] / [\textit{Maximum Theoretical Throughput}]$$
$$\% \text{ of MTT} = [\textit{Frames Received}] / ([\textit{Maximum Frames/Sec}] \times [\textit{Test Duration in Seconds}])$$

Latency

RFC 2544 defines latency for a bit forwarding or cut-through device as “the time interval starting when the end of the first bit of the input frame reaches the input port and ending when the start of the first bit of the output frame is seen on the output port.” This study measures the cut through end-to-end system latency. This is the time it takes for the first bit of a frame to traverse the topology from source to destination.

It is expected that every networking device added to a system introduces additional latency. The extent to which latency is introduced is affected by factors such as the amount of frame data that must be processed, the memory and processor capabilities of the networking device, and the forwarding and queuing algorithms used, as well as other factors. The Ixia test device records minimum, maximum and average latency figures.

Latency was measured at the maximum throughput without frame loss for each frame size. Two minute trials were used and the average of ten trials was graphed for each frame size. This approach slightly modified the testing methodology used in the previous RIT study, which employed one minute trials rather than two. This modification was based on best practices as defined in RFC 2544.

Discussion of Results

Frame loss

The frame loss test revealed that for all three scenarios tested in this study there was less than 1/1000th of 1% frame loss for each given frame size. This amount of frame loss is statistically insignificant and resulted in an average frame loss of 0 for 100% of line speed. Once this was identified, the frame loss measured in the throughput test also served as the frame loss data. This was found to be true for all scenarios, all frame sizes, and in both directions.

Throughput

Once testing began, it was discovered that the remote interface was receiving more frames than were being transmitted and in some cases, was exceeding the calculated MTT for that frame size. This prompted further investigation which revealed that the Ixia test was actually fractions of a second longer than one minute and that the Cisco devices were generating spanning tree and Cisco discovery frames. Compensating for this restored the accuracy of the tests. During the testing, the throughput calculations were modified slightly based on the use of a BPDU filter which excluded the spanning tree frames from being counted.

Scenario 1 – Infrastructure Baseline

The Cisco 6509s were found to forward traffic on average, at 100% of MTT for all frame sizes. This was true for traffic flowing in both directions.

Scenario 2 – SEE Unencrypted Baseline

The addition of the SEEs in bypass mode (encryption turned off) had no impact on throughput. The encryptors forwarded traffic at 100% of MTT for all frame sizes in both directions.

Scenario 3 – SEE Encryption Activated

For this test, the topology remained the same as in Scenario 2, but encryption was activated on the encryptors causing them to encrypt or decrypt all frames sent and received. The results showed that the addition of encryption had no effect on throughput or frame loss, in either direction, regardless of frame size, and throughput remained at 100% of MTT.

Latency

Scenario 1 – Infrastructure Baseline

The average latency for traffic traversing the infrastructure baseline ranged from 1210 µsec (1.2 ms) for 64 byte frames up to 1308 µsec (1.3 ms) for 1420 byte frames. This was true for traffic flowing in both directions across the entire infrastructure.

Scenario 2 – SEE Unencrypted Baseline

It was hypothesized that adding an additional device would result in increased latency. This was the case, however the increase was small, ranging from 12-16 μ sec, irrespective of frame size.

Scenario 3 – SEE Encryption Activated

It was hypothesized that activating encryption on the SEEs would introduce additional latency beyond Scenario 2 but this was not the case. The SEEs added no additional latency as a result of the encryption/decryption process, regardless of frame size or direction.

Some minor amount of variation was found between the data sets collected on the two ports of the Ixia. This variation was consistent throughout the range of frame sizes and is attributed to low-level inconsistencies in the hardware throughout the testing topologies. Because this inconsistency was minimal, and because it was consistent throughout the range of tests, we considered it insignificant and the values of the two datasets were averaged.

It was found that on average the encryptors added less than 1% latency to the Infrastructure Baseline and that the encryptors actually performed slightly better when encrypting than bypassing traffic. Table 4 and Figure 13 show the additional system latency when the SEEs are added in encryption mode.

Frame Size	% Max Line Rate	Baseline Latency (μ sec)	SEE's in Bypass Mode (μ sec)	SEE with Encryption Latency (μ sec)	% Increase Over Baseline
64	100%	1216	1227	1227	0.93%
128	100%	1266	1278	1275	0.72%
256	100%	1272	1288	1285	0.95%
512	100%	1278	1293	1290	1.00%
1024	100%	1286	1302	1300	1.08%
1280	100%	1294	1306	1306	0.92%
1420	100%	1298	1312	1309	0.89%

Table 4 - Ethernet Latency Data

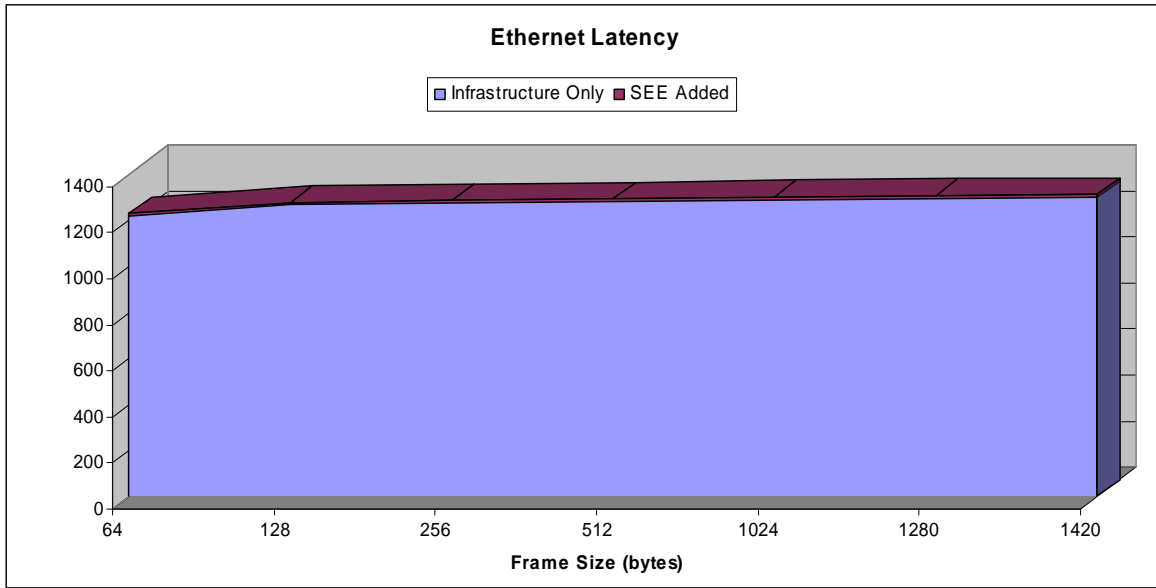


Figure 11 - Latency Added by SEE

Results of Previous IPsec Study

The previous study was conducted in a similar manner to this one. First, a battery of frame loss, throughput, and latency tests were run on a baseline that mirrors that of this study. That same physical topology was retested with the configuration of the Cisco IPsec Services Module to encrypt all traffic between the two endpoints. The diagram below represents the topology used and shows the Cisco IPsec acceleration blade in each 6509 chassis.

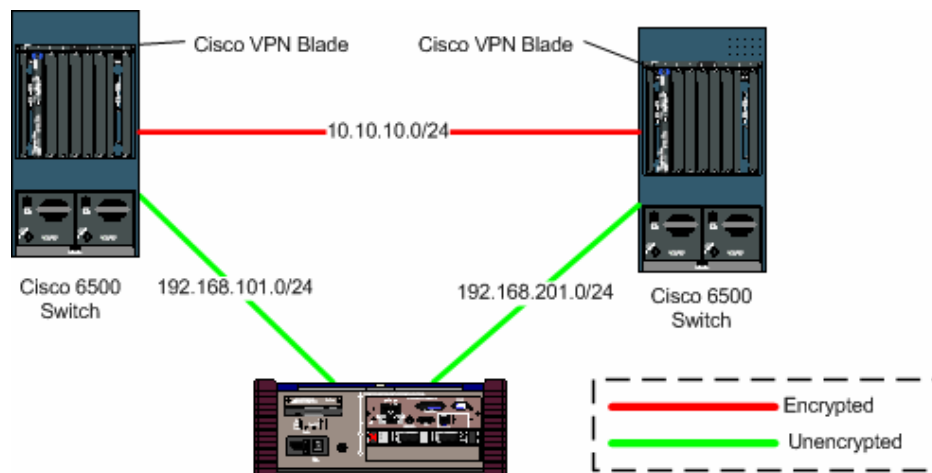


Figure 12 - IPsec Test Topology from Previous Study

Frame Loss

In the converged IPsec scenario, RIT found that for all frame sizes, as load increased, so did frame loss. Furthermore, it was found that as frame size decreased, frame loss increased in a logarithmic fashion. At 64-byte frame sizes, over 40 percent frame loss was encountered at only 30 percent of maximum theoretical throughput.

Frame Size	64	128	256	512	1024	1280	1420
10	0						
20	12.774						
30	41.876	0					
40	56.325	15.805					
50	65.056	32.729					
60	70.873	44.026	0				
70	75.029	52.21	13.008				
80	78.353	58.121	24.057				
90	80.857	62.671	32.437	0	0	0	0
100	82.505	66.309	39.031	9.505	5.078	4.118	4.06

Figure 13 - IPsec Frame Loss by Percent Throughput and Frame Size




Figure 14 - IPsec Frame Loss

Throughput

Table 5 and Figure 15 below show that the converged IPsec encryption solution was unable to achieve MTT at smaller frame sizes. At 512-byte, 1024-byte and 1280-byte frame sizes, approximately 100% MTT was achieved. At 256-byte frames, however, that performance was just 73 %, at 128-byte frames it dropped further to 47%, and at 64-byte frames it was just 27% of MTT.

Frame Size (bytes)	Measured Ipsec Infrastructure Baseline Throughput	Ipsec Max Theoretical Throughput (MTT)	Measured Ipsec Encryption Throughput	Percent of MTT for Measured IPsec Throughput
64	762	454	125	27%
128	865	624	291	47%
256	928	769	562	73%
512	962	869	871	100%
1024	981	930	931	100%
1280	985	943	944	100%
1420	986	949	940	99%

Table 5- IPsec Throughput Data

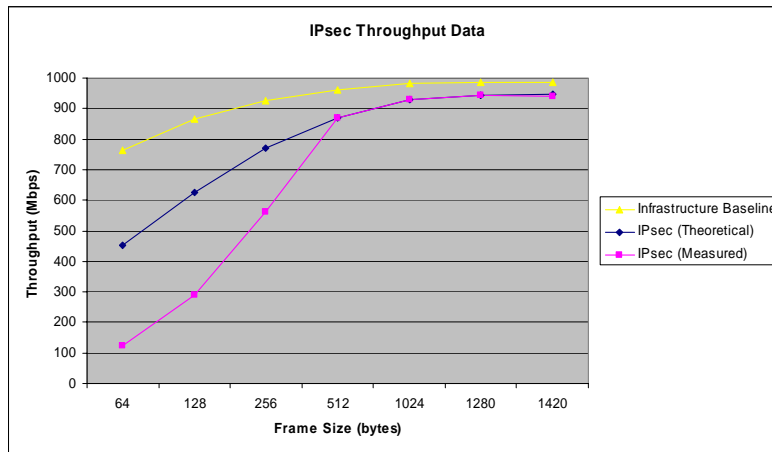


Figure 15 - IPsec Throughput Data

Figure 16 represents the throughput data another way, detailing the theoretical and measured bandwidth utilization of IPsec compared to the theoretical bandwidth available to Ethernet. Note that at small frame sizes, the IPsec performance is significantly less than the theoretical maximum. This we attributed to the limited processing power of the encryption device, which must process significantly more data at smaller frame sizes.

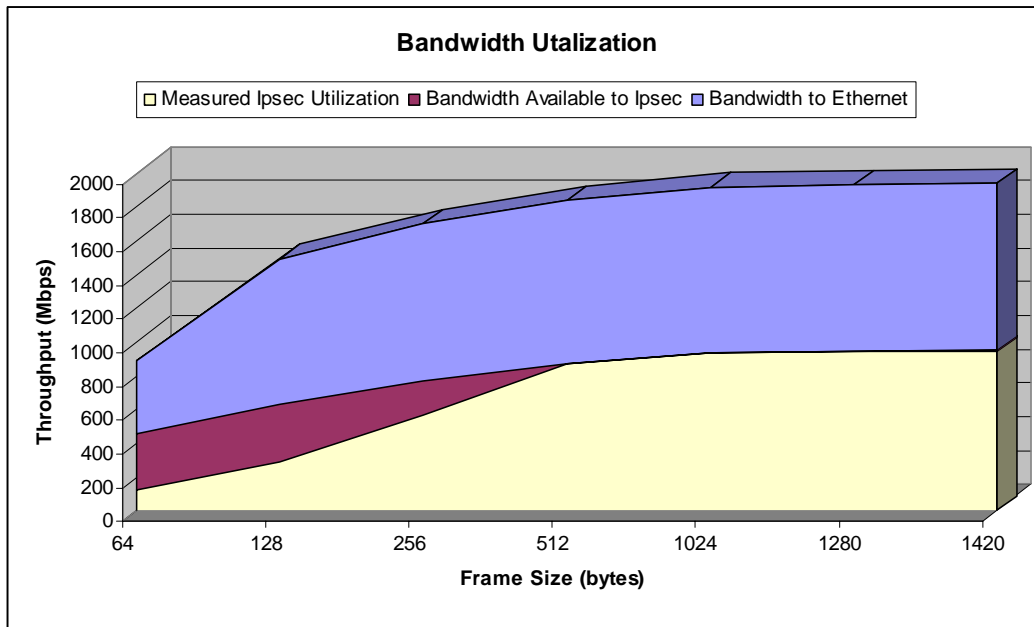


Figure 16 - Comparative Bandwidth Utilization of IPsec

Latency

In the aforementioned study latency was measured using the store-and-forward method, as defined in RFC 2544, rather than the cut-through method that was employed in this study. This makes direct comparison of the results difficult however, by comparing the results as a percentage increase over their respective baselines, a comparison can be drawn. On average the latency increased from 6 – 10 times depending on frame size.

Because frame loss was encountered at every frame size during the IPsec tests, the data rate used in each of the latency tests was reduced from the MTT. For example, Table 6 shows that the latency test for 64-byte frames is run at just 16% of MTT. Figure 17 shows the additional latency added to the unencrypted baseline for each frame size tested.

Frame Size	% Max Line Rate	Baseline Latency (ns)	IPsec Encryption Latency (ns)	% Increase Over Baseline
64	16%	12062	117594	975%
128	34%	13767	136633	992%
256	61%	17913	154689	864%
512	90%	24773	205958	831%
1024	95%	38594	264292	685%
1280	96%	45085	304925	676%
1420	95%	49078	340695	694%

Table 6 - IPsec Latency Data

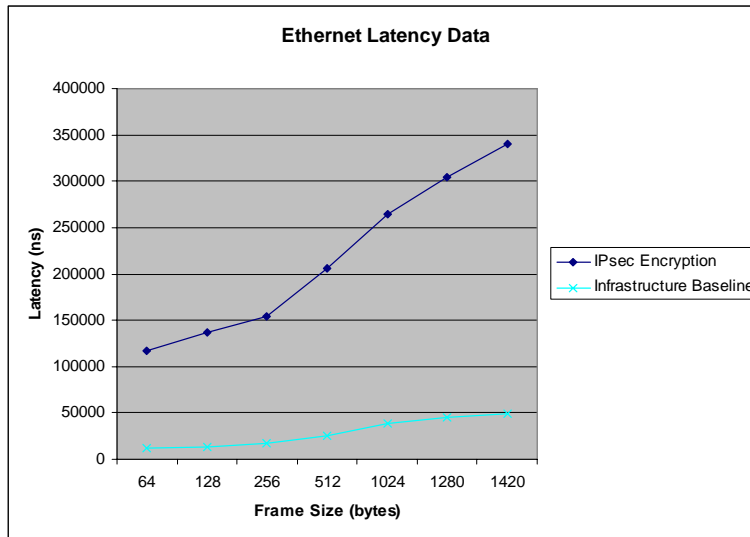


Figure 17 - IPsec Latency

Comparison of Results

Frame Loss

As previously stated, the Ethernet encryption solution experienced less than 1/1000th of 1% frame loss for each given frame size. Again, this amount of frame loss is statistically insignificant and resulted in an average frame loss of 0 for 100% of line speed. In comparison, the IPsec encryption study found that for all frame sizes loss was significant. Ultimately this significantly limits the available bandwidth that can be used with this solution as seen below.

Throughput

It was found the native Ethernet encryption solution was able to use 100% of the bandwidth theoretically available to Ethernet. Independent of hardware limitations, IPsec encryption further reduces the amount of bandwidth theoretically available to Ethernet. This effect is most dramatic at smaller frame sizes where IPsec overhead represents a higher percentage of the total frame size. The IPsec encryption study found that due to hardware limitations, the actual measured throughput was significantly less than the IPsec MTT. Figure 18 below shows the comparative performance of theoretical and measured results for the two technologies. Note that the differences in performance are most dramatic at smaller frame sizes which are common in mixed Internet traffic and real-time voice and data applications.

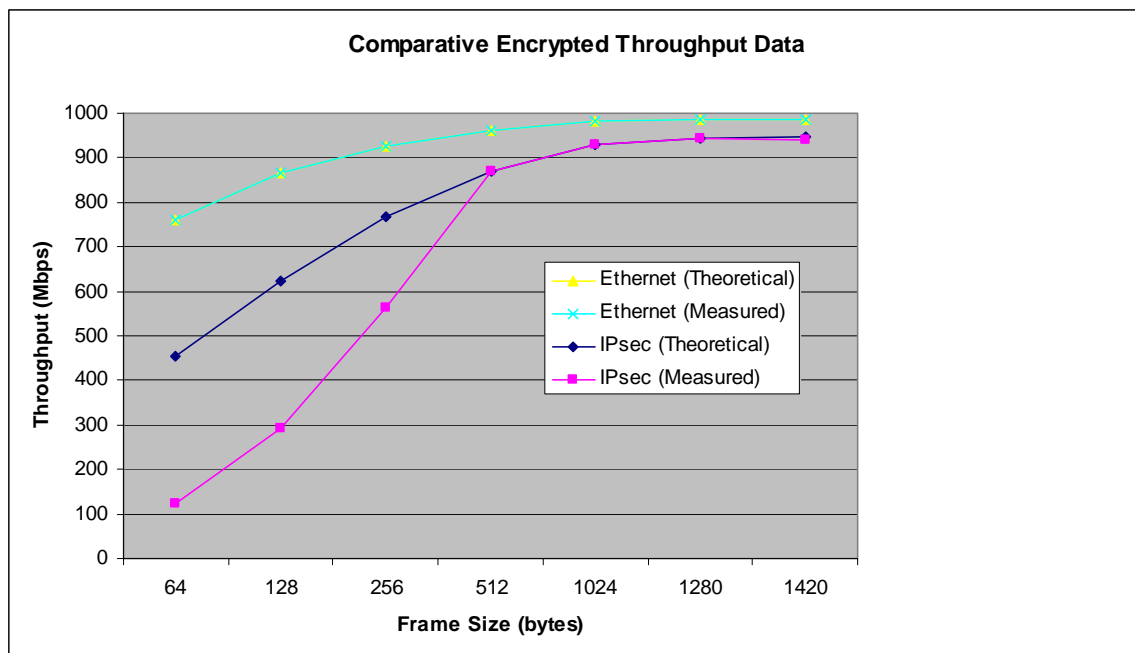


Figure 18 - Comparative Encrypted Throughput Data

Latency

The latency test used in both studies can not tolerate frame loss. This is because of the tagging method the Ixia uses to measure how long it takes a frame to reach its destination. If a frame is lost latency is calculated at 100%. Because in the original study frame loss was encountered at all frame sizes when encrypting with IPsec, it was necessary to reduce the bandwidth to a rate at which no frames were lost. In this study, because frame loss was negligible at all frame sizes at 100% MTT it was not necessary to reduce the throughput. The previous study compensated for this by recalibrating the baseline tests to the maximum data rate without frame loss achieved for the IPsec encryption test at a give frame size. That consideration allows this study to accurately analyze the comparative performance of the two technologies. The following is an explanation of this methodology from the previous study. Note that they used a different methodology for calculating latency as previously discussed:

The [Ixia] latency test measures the time in nanoseconds needed for a packet leaving one testing interface on the packet generator to enter the receiving testing interface. The packet generator measures this by tagging one frame every second with a time stamp. Because of the way frames are tagged, it is essential that the throughput test be ran first and each latency test is not run above the maximum throughput for that frame size. Frame loss during the latency test would interfere with the results. The latency test was run for one minute per trial, with ten trials for each of the following frame sizes in bytes: 64, 128, 256, 512, 1024, 1280, and 1420. The throughput selected for each test was the maximum throughput previously determined in testing for that combination of frame size and encryption method. In order to make the latency testing comparable between setups we calculated the throughput at which the encryption technology ran for a certain frame size, and then ran a latency test at the same percentage and frame size on the baseline setup with no encryption. This allows for a comparison of baseline latency to encrypted latency determining the amount of time that is added solely due to encryption.

As previously mentioned, comparison of the latency results was done as a percentage increase over the respective unencrypted baseline latency for each test. To compare the latency seen in each test, percentages over the unencrypted baseline were used rather than actual numbers. Table 7 and Figure 19 show the comparison between IPsec encryption and native Ethernet encryption to secure a gigabit Ethernet link. Note that this shows percent increase over the respective Infrastructure Baselines only, and does not represent throughput, which is significantly less at all frame sizes for the IPsec solution than the Ethernet solution, as discussed above. It was found that on average the IPsec solution added over 13 times the latency of the SEE solution as shown in Table 7.

Frame Size	% Increase Over Baseline Ipsec Encryption	% Increase Over Baseline Ethernet Encryption
64	975%	0.93%
128	992%	0.72%
256	864%	0.95%
512	831%	1.00%
1024	685%	1.08%
1280	676%	0.92%
1420	694%	0.89%
Average	817%	0.93%

Table 7 - IPsec and Ethernet Encryption Latency Comparison

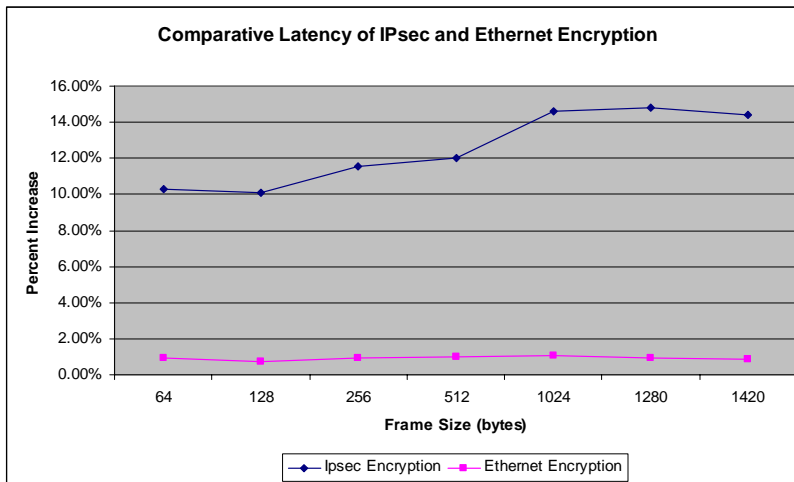


Figure 19 - Comparative Latency of IPsec and Ethernet Encryption

Conclusion

Overwhelming logic suggests that, from a theoretical performance perspective, Layer 2 Ethernet encryption should be superior to Layer 3 IPsec encryption. The current battery of tests, and data generated in the previous study, confirm the reality of the throughput and latency limitations induced by the IPsec wrapper overhead. Testing exposed the detrimental effect on network performance that is typically imposed by IPsec's innate processing requirements, as well as the processing limitations of the tested IPsec hardware.

In contrast, the SafeNet Ethernet Encryptor operates at line speed. Testing also revealed no significant frame loss with the SafeNet Ethernet encryption solution, whereas significant frame loss was encountered at comparatively low data rates with the IPsec solution. This resulted in a significant reduction of the data rate, as seen in throughput testing, and indicates that achieving line rate encryption, even at the reduced maximum theoretical throughput of IPsec, is impossible at regardless of frame size with the IPsec solution.

Finally, the measured latency of the Cisco IPsec encryption solution was found to be over 13 times that of the SafeNet Ethernet encryption solution. In environments where Ethernet encryption technology meets the needs of the organization, its performance is clearly superior to IPsec.

Recommendations for Future Studies

There are a number of factors that must be considered when choosing a WAN encryption technology. Performance is one such factor, however; a careful analysis must be performed to ensure the solution meets all of the needs of the organization. A future study might work to develop a quantitative framework for analyzing these non-performance related factors.

This study was performed using discrete frame sizes for each test run. A future study might analyze these performance characteristics with varying frame sizes as commonly found in mixed Internet traffic.

Acknowledgements

We are grateful to Ixia for the use of their testing equipment and for their consistent support as we developed and conducted our tests. We would also like to thank Lumarc for their responsiveness to our equipment needs.

References

[1] "Converged vs. Dedicated IPSec Encryption Testing in Gigabit Ethernet Networks", Troell, Luther, et al., 2005.

[2] "Proposal for Evaluating Ethernet Encryption in High Speed MAC Transparent Networks", Baker, Davin, 2006.

[3] "RFC 2544 - Benchmarking Methodology for Network Interconnect Devices"
The Internet Society, 1999.

Appendices

Appendix A - Equipment List

- (2) GbE SafeNet SafeEnterprise™ Ethernet Encryptor
- (2) Cisco 6509 Router Chassis with SUP module
- (4) Cisco 2500 Watt power supply for 6500 chassis
- (2) Ixia 250 Chassis
- (1) Ixia TXS 10/100/1000 Ethernet Load Module (LM1000TXS4-256)
- (1) IxExplorer Software

Appendix B – Diagrams

Infrastructure Baseline Topology

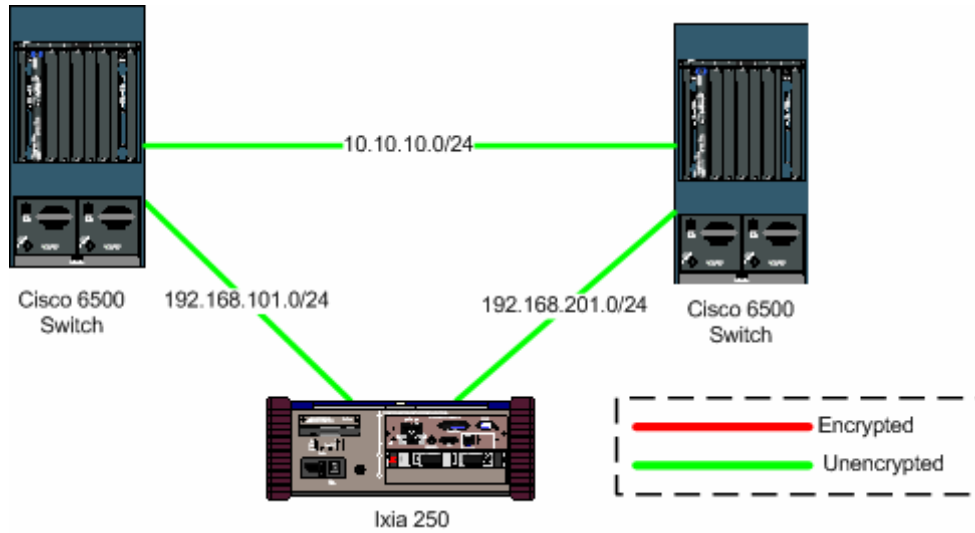


Figure 20 - Infrastructure Baseline Topology

Unencrypted Topology with SafeNet SafeEnterprise™ Ethernet Encrytors

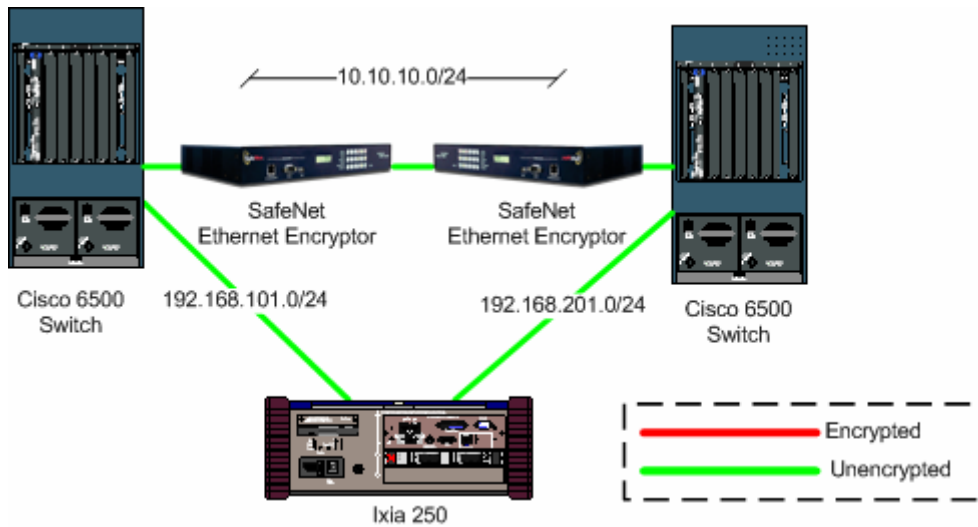


Figure 21 - Unencrypted Topology with SafeNet SafeEnterprise™ Ethernet Encrytors

Encrypted Topology with SafeNet SafeEnterprise™ Ethernet Encrytors

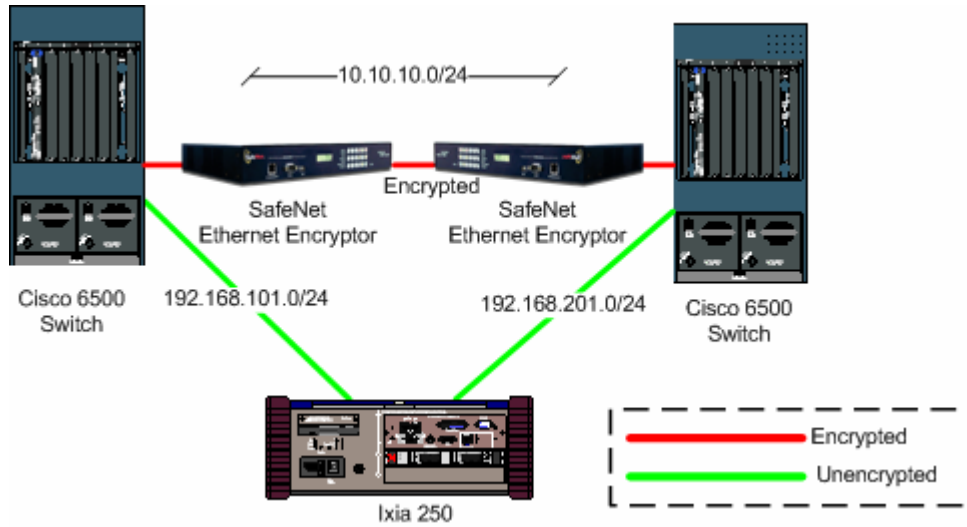


Figure 22 - Encrypted Topology with SafeNet Ethernet Encrytors

Encrypted Topology with Cisco IPsec Encryption

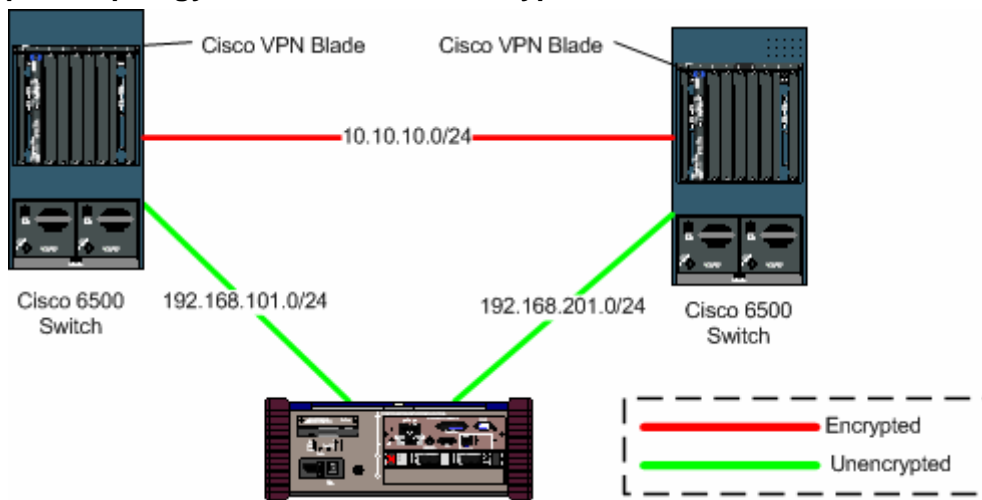


Figure 23 - Encrypted Topology with Cisco IPsec Encryption

Appendix C – Infrastructure Configuration Information**Cisco 6500 Router - Timmy**

```
!  
version 12.2  
service timestamps debug uptime  
service timestamps log uptime  
service password-encryption  
service counters max age 10  
!  
hostname timmy  
!  
boot system flash s72033-pk9sv-mz.122-18.SXD3.bin  
enable secret 5 $1$3kha$v7M6Xr2wK9XkbnoqOAMqc/  
!  
no aaa new-model  
clock timezone EST -5  
ip subnet-zero  
!  
ip domain-name sonetbonnet.com  
ip host jimmy 172.16.1.12  
ip host smcmgmt.sonetbonnet.com 172.16.1.23  
!  
mls ip multicast flow-stat-timer 9  
no mls flow ip  
no mls flow ipv6  
mls cef error action freeze  
!  
power redundancy-mode combined  
spanning-tree mode pvst  
no spanning-tree optimize bpdu transmission  
diagnostic cns publish cisco.cns.device.diag_results  
diagnostic cns subscribe cisco.cns.device.diag_commands  
!  
redundancy  
mode sso  
main-cpu  
auto-sync running-config  
!  
vlan internal allocation policy ascending  
vlan access-log ratelimit 2000  
!  
interface GigabitEthernet1/1  
no ip address  
switchport  
switchport access vlan 101  
switchport mode access  
!  
interface GigabitEthernet1/1
```

```
ip address 10.10.20.11 255.255.255.0
!  
interface GigabitEthernet5/1  
no ip address  
shutdown  
!  
interface GigabitEthernet5/2  
no ip address  
shutdown  
!  
interface Vlan1  
no ip address  
shutdown  
!  
interface Vlan101  
ip address 192.168.101.254 255.255.255.0  
!  
ip classless  
ip route 192.168.201.0 255.255.255.0 10.10.10.12  
no ip http server  
!  
control-plane  
!  
line con 0  
line vty 0 4  
password 7 111A160B120609030A242E30  
login  
!  
ntp clock-period 17181487  
ntp server 172.16.1.23  
end
```

Cisco 6500 Router - Jimmy

```
!  
version 12.2  
service timestamps debug uptime  
service timestamps log uptime  
service password-encryption  
service counters max age 10  
!  
hostname jimmy  
!  
boot system flash s72033-pk9sv-mz.122-18.SXD3.bin  
enable secret 5 $1$3kha$v7M6Xr2wK9XkbnqOAMqc/  
!  
no aaa new-model  
clock timezone EST -5  
ip subnet-zero  
!  
ip domain-name sonetbonnet.com  
ip host timmy 172.16.1.11  
ip host smcmgmt.sonetbonnet.com 172.16.1.23  
!  
mls ip multicast flow-stat-timer 9  
no mls flow ip  
no mls flow ipv6  
mls cef error action freeze  
!  
power redundancy-mode combined  
spanning-tree mode pvst  
no spanning-tree optimize bpdu transmission  
diagnostic cns publish cisco.cns.device.diag_results  
diagnostic cns subscribe cisco.cns.device.diag_commands  
!  
redundancy  
mode sso  
main-cpu  
auto-sync running-config  
!  
vlan internal allocation policy ascending  
vlan access-log ratelimit 2000  
!  
interface GigabitEthernet1/1  
no ip address  
switchport  
switchport access vlan 201  
switchport mode access  
!  
interface GigabitEthernet1/1  
ip address 10.10.20.12 255.255.255.0  
!
```

```
interface GigabitEthernet5/1
no ip address
shutdown
!
interface GigabitEthernet5/2
no ip address
shutdown
!
interface Vlan1
no ip address
shutdown
!
interface Vlan201
ip address 192.168.201.254 255.255.255.0
!
ip classless
ip route 192.168.101.0 255.255.255.0 10.10.10.11
no ip http server
!
control-plane
!
line con 0
line vty 0 4
password 7 111A160B120609030A242E30
login
!
ntp clock-period 17181487
ntp server 172.16.1.23
end
```

Safenet Ethernet Encryptor - SEE_Top_192.168.0.2

Software Version 3.2.0

IP Address 192.168.0.2
Net Mask 255.255.255.0
MAC Address 00:20:e2:30:03:87
Operating Mode: twistedPair

System Module Version: B2010A003
System Module Serial Number: 300387
1Gb Ethernet Crypto/Interface Card NETWORK Version: B2084A001
Serial Number: 10NWDW
1Gb Ethernet Crypto Engine Firmware Hardwarte version: B2084A001
Serial Number: 10NWDW

Current Certificate(RSA signed with 1024 bit internal CA)
Number 1
Version 1
Subject Name CN=SEE_192.168.0.2,OU=GCCIS,O=RIT,C=US
Issuer (CA) Name SafeNet CA
Valid After Sun Jul 23 19:08:15 2006
Expires Sat Jul 23 19:08:15 2016

Safenet Ethernet Encryptor - SEE_Bottom_192.168.0.3

Software Version 3.2.0

IP Address 192.168.0.3
Net Mask 255.255.255.0
Default Gateway 192.168.0.1
MAC Address 00:20:e2:30:03:88
Operating Mode: twistedPair

System Module Version: B2010A003
System Module Serial Number: 300388
1Gb Ethernet Crypto/Interface Card NETWORK Version: B2084A001
Serial Number: 10NWGK
1Gb Ethernet Crypto Engine Firmware Hardwarte version: B2084A001
Serial Number: 10NWGK

Current Certificate(RSA signed with 1024 bit internal CA)
Number 2
Version 1
Subject Name CN=SEE_192.168.0.3,OU=GCCIS,O=RIT,C=US
Issuer (CA) Name SafeNet CA
Valid After Sun Jul 23 19:06:21 2006
Expires Sat Jul 23 19:06:21 2016