Benchmarking Working Group   M. Kaeo
Internet-Draft   Double Shot Security
Expires: September 6, 2006   T. Van Herck
                                                           Cisco Systems
                                                           March 5, 2006

Methodology for Benchmarking IPsec Devices
draft-ietf-bmwg-ipsec-meth-01

Status of this Memo

Copyright Notice

Abstract

The purpose of this draft is to describe methodology specific to the benchmarking of IPsec IP forwarding devices. It builds upon the tenets set forth in [RFC2544], [RFC2432] and other IETF Benchmarking Methodology Working Group (BMWG) efforts. This document seeks to extend these efforts to the IPsec paradigm.

The BMWG produces two major classes of documents: Benchmarking

Terminology documents and Benchmarking Methodology documents. The Terminology documents present the benchmarks and other related terms. The Methodology documents define the procedures required to collect the benchmarks cited in the corresponding Terminology documents.

Table of Contents

1. Introduction

This document defines a specific set of tests that can be used to measure and report the performance characteristics of IPsec devices. It extends the methodology already defined for benchmarking network interconnecting devices in [RFC2544] to IPsec gateways and additionally introduces tests which can be used to measure end-host IPsec performance.

2. Document Scope

The primary focus of this document is to establish a performance testing methodology for IPsec devices that support manual keying and IKEv1. Both IPv4 and IPv6 addressing will be taken into consideration for all relevant test methodologies.

The testing will be constrained to:

- Devices acting as IPsec gateways whose tests will pertain to both IPsec tunnel and transport mode.

- Devices acting as IPsec end-hosts whose tests will pertain to both IPsec tunnel and transport mode.

Note that special considerations will be presented for IPsec end-host testing since the tests cannot be conducted without introducing additional variables that may cause variations in test results.

What is specifically out of scope is any testing that pertains to considerations involving NAT, L2TP [RFC2661], GRE [RFC2784], BGP/MPLS VPNs [RFC2547] and anything that does not specifically relate to the establishment and tearing down of IPsec tunnels.

3. Key Words to Reflect Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119. RFC 2119 defines the use of these key words to help make the intent of standards track documents as clear as possible. While this document uses these keywords, this document is not a standards track document.

4. Test Considerations

Before any of the IPsec data plane benchmarking tests are carried

out, a Baseline MUST be established. I.e. the particular test in question must first be measured for performance characteristics without enabling IPsec. Once both the Baseline clear text performance and the performance using an IPsec enabled datapath have been measured, the difference between the two can be discerned.

This document explicitly assumes that you MUST follow logical performance test methodology that includes the pre-configuration of routing protocols, ARP caches, IPv6 neighbor discovery and all other extraneous IPv4 and IPv6 parameters required to pass packets before the tester is ready to send IPsec protected packets. IPv6 nodes that implement Path MTU Discovery [RFC1981] MUST ensure that the PMTUD process has been completed before any of the tests have been run.

For every IPsec data plane benchmarking test, the SA database (SADB) MUST be created and populated with the appropriate SAs before any actual test traffic is sent, i.e. the DUT/SUT MUST have active tunnels. This may require a manual command to be executed on the DUT/SUT or the sending of appropriate learning frames to the DUT/SUT. This is to ensure that none of the control plane parameters (such as IPsec tunnel setup rates and IPsec tunnel rekey rates) are factored into these tests.

For control plane benchmarking tests (i.e. IPsec tunnel setup rate and IPsec tunnel rekey rates), the authentication mechanisms(s) used for the authenticated Diffie-Hellman exchange MUST be reported.

5. Test Topologies

The tests can be performed as a DUT or SUT. When the tests are performed as a DUT, the Tester itself must be an IPsec peer. This scenario is shown in Figure 1. When tested as a DUT where the Tester has to be an IPsec peer, the measurements have several disadvantages:

- The Tester can introduce interoperability issues and skew results.

- The measurements may not be accurate due to Tester inaccuracies.

On the other hand, the measurement of a DUT where the Tester is an IPsec peer has two distinct advantages:

- IPsec client scenarios can be benchmarked.

- IPsec device encryption/decryption abnormalities may be identified.

```
                              +------------+
                              |            |
                       +----[D]   Tester   [A]----+
                       |      |            |      |
                       |      +------------+      |
                       |                          |
                       |      +------------+      |
                       |      |            |      |
                       +----[C]    DUT     [B]----+
                              |            |
                              +------------+
```

Figure 1: Topology 1

The SUT scenario is depicted in Figure 2. Two identical DUTs are used in this test set up which more accurately simulate the use of IPsec gateways. IPsec SA (i.e. AH/ESP transport or tunnel mode) configurations can be tested using this set-up where the tester is only required to send and receive cleartext traffic.

```
                              +------------+
                              |            |
          +-----------------[F]   Tester   [A]-----------------+
          |                   |            |                   |
          |                   +------------+                   |
          |                                                    |
          |      +------------+            +------------+      |
          |      |            |            |            |      |
          +----[E]    DUTa    [D]--------[C]    DUTb    [B]----+
                 |            |            |            |
                 +------------+            +------------+
```

Figure 2: Topology 2

When an IPsec DUT needs to be tested in a chassis failover topology, a second DUT needs to be used as shown in figure 3. This is the high-availability equivalent of the topology as depicted in Figure 1. Note that in this topology the Tester MUST be an IPsec peer.

```
                              +------------+
                              |            |
                  +---------[F]   Tester   [A]---------+
                  |           |            |           |
                  |           +------------+           |
                  |                                    |
                  |           +------------+           |
                  |           |            |           |
                  |    +----[C]    DUTa    [B]----+    |
                  |    |      |            |      |    |
                  |    |      +------------+      |    |
                  +----+                          +----+
                       |      +------------+      |
                       |      |            |      |
                       +----[E]    DUTb    [D]----+
                              |            |
                              +------------+
```

Figure 3: Topology 3

When no IPsec enabled Tester is available and an IPsec failover scenario needs to be tested, the topology as shown in Figure 4 can be used. In this case, either the high availability pair of IPsec devices can be used as an Initiator or as a Responder. The remaining chassis will take the opposite role.

```
                              +------------+
                              |            |
       +--------------------[H]   Tester   [A]----------------+
       |                      |            |                  |
       |                      +------------+                  |
       |                                                      |
       |         +------------+                               |
       |         |            |                               |
       |   +---[E]    DUTa    [D]---+                         |
       |   |     |            |     |      +------------+     |
       |   |     +------------+     |      |            |     |
       +---+                        +----[C]    DUTc    [B]---+
           |     +------------+     |      |            |
           |     |            |     |      +------------+
           +---[G]    DUTb    [F]---+
                 |            |
                 +------------+
```

Figure 4: Topology 4

6. Test Parameters

For each individual test performed, all of the following parameters MUST be explicitly reported in any test results.

6.1. Frame Type

6.1.1. IP

Both IPv4 and IPv6 frames MUST be used. The basic IPv4 header is 20 bytes long (which may be increased by the use of an options field). The basic IPv6 header is a fixed 40 bytes and uses an extension field for additional headers. Only the basic headers plus the IPsec AH and/or ESP headers MUST be present.

It is recommended that IPv4 and IPv6 frames be tested separately to ascertain performance parameters for either IPv4 or IPv6 traffic. If both IPv4 and IPv6 traffic are to be tested, the device SHOULD be pre-configured for a dual-stack environment to handle both traffic types.

IP traffic with L4 protocol set to 'reserved' (255) MUST be used. This ensures maximum space for instrumentation data in the payload section, even with framesizes of minimum allowed length on the transport media.

6.1.2. UDP

It is also RECOMMENDED that the test is executed using UDP as the L4 protocol. When using UDP, instrumentation data SHOULD be present in the payload of the packet. It is OPTIONAL to have application payload.

6.1.3. TCP

It is OPTIONAL to perform the tests with TCP as the L4 protocol but in case this is considered, the TCP traffic is RECOMMENDED to be stateful. With a TCP as a L4 header it is possible that there will not be enough room to add all instrumentation data to identify the packets within the DUT/SUT.

6.2. Frame Sizes

Each test SHOULD be run with different frame sizes. The recommended cleartext layer 2 frame sizes for IPv4 tests over Ethernet media are 64, 128, 256, 512, 1024, 1280, and 1518 bytes, per RFC2544 section 9 [RFC2544]. The four CRC bytes are included in the frame size specified.

For GigabitEthernet, supporting jumboframes, the cleartext layer 2 framesizes used are 64, 128, 256, 512, 1024, 1280, 1518, 2048, 3072, 4096, 5120, 6144, 7168, 8192 and 9234 bytes

Since IPv6 requires that every link has an MTU of 1280 octets or greater, it is MANDATORY to execute tests with cleartext layer 2 frame sizes that include 1280 and 1518 bytes. It is RECOMMENDED that additional frame sizes are included in the IPv6 test execution, including the maximum supported datagram size for the linktype used.

6.3. Fragmentation and Reassembly

IPsec devices can and must fragment packets in specific scenarios. Depending on whether the fragmentation is performed in software or using specialized custom hardware, there may be a significant impact on performance.

In IPv4, unless the DF (don't fragment) bit is set by the packet source, the sender cannot guarantee that some intermediary device on the way will not fragment an IPsec packet. For transport mode IPsec, the peers must be able to fragment and reassemble IPsec packets. Reassembly of fragmented packets is especially important if an IPv4 port selector (or IPv6 transport protocol selector) is configured. For tunnel mode IPsec, it is not a requirement. Note that fragmentation is handled differently in IPv6 than in IPv4. In IPv6 networks, fragmentation is no longer done by intermediate routers in the networks, but by the source node that originates the packet. The path MTU discovery (PMTUD) mechanism is recommended for every IPv6 node to avoid fragmentation.

Packets generated by hosts that do not support PMTUD, and have not set the DF bit in the IP header, will undergo fragmentation before IPsec encapsulation. Packets generated by hosts that do support PMTUD will use it locally to match the statically configured MTU on the tunnel. If you manually set the MTU on the tunnel, you must set it low enough to allow packets to pass through the smallest link on the path. Otherwise, the packets that are too large to fit will be dropped.

Fragmentation can occur due to encryption overhead and is closely linked to the choice of transform used. Since each test SHOULD be run with a maximum cleartext frame size (as per the previous section) it will cause fragmentation to occur since the maximum frame size will be exceeded. All tests MUST be run with the DF bit not set. It is also recommended that all tests be run with the DF bit set.

Note that some implementations predetermine the encapsulated packet size from information available in transform sets, which are

configured as part of the IPsec security association (SA). If it is predetermined that the packet will exceed the MTU of the output interface, the packet is fragmented before encryption. This optimization may favorably impact performance and vendors SHOULD report whether any such optimization is configured.

6.4. Time To Live

The source frames should have a TTL value large enough to accommodate the DUT/SUT. A Minimum TTL of 64 is RECOMMENDED.

6.5. Trial Duration

The duration of the test portion of each trial SHOULD be at least 30 seconds. In the case of IPsec tunnel rekeying tests, the test duration must be at least two times the IPsec tunnel rekey time to ensure a reasonable worst case scenario test.

6.6. Security Context Parameters

All of the security context parameters listed in this section and used in any test MUST be reported.

6.6.1. IPsec Transform Sets

All tests should be done on different IPsec transform set combinations. An IPsec transform specifies a single IPsec security protocol (either AH or ESP) with its corresponding security algorithms and mode. A transform set is a combination of individual IPsec transforms designed to enact a specific security policy for protecting a particular traffic flow. At minumim, the transform set must include one AH algorithm and a mode or one ESP algorithm and a mode, as shown in Table 1:

Transform Set	AH Algorithm	ESP Algorithm	Mode
1	AH-SHA1	None	Tunnel
2	AH-SHA1	None	Transport
3	AH-SHA1	ESP-3DES	Tunnel
4	AH-SHA1	ESP-3DES	Transport
5	AH-SHA1	ESP-AES128	Tunnel
6	AH-SHA1	ESP-AES128	Transport
7	None	ESP-3DES	Tunnel
8	None	ESP-3DES-HMAC-SHA1	Tunnel
9	None	ESP-3DES	Transport
10	None	ESP-3DES-HMAC-SHA1	Transport
11	None	ESP-AES128	Tunnel
12	None	ESP-AES128-HMAC-SHA1	Tunnel
13	None	ESP-AES128	Transport
14	None	ESP-AES128-HMAC-SHA1	Transport

Table 1

Testing of all the transforms shown in Table 1 MUST be supported. Note that this table is derived from the updated IKEv1 requirements as described in [RFC4109]. Optionally, other AH and/or ESP transforms MAY be supported.

6.6.2. IPsec Topologies

All tests should be done at various IPsec topology configurations and the IPsec topology used MUST be reported. Since IPv6 requires the implementation of manual keys for IPsec, both manual keying and IKE configurations MUST be tested.

For manual keying tests, the IPsec SAs used should vary from 1 to 101, increasing in increments of 50. Although it is not expected that manual keying (i.e. manually configuring the IPsec SA) will be deployed in any operational setting with the exception of very small controlled environments (i.e. less than 10 nodes), it is prudent to test for potentially larger scale deployments.

For IKE specific tests, the following IPsec topologies MUST be tested:

- 1 IKE SA & 1 IPsec SA (i.e. 1 IPsec Tunnel)

- 1 IKE SA & {max} IPsec SA's

- {max} IKE SA's & {max} IPsec SA's

It is RECOMMENDED to also test with the following IPsec topologies in order to gain more datapoints:

- {max/2} IKE SA's & {(max/2) IKE SA's} IPsec SA's

- {max} IKE SA's & {(max) IKE SA's} IPsec SA's

6.6.3. IKE Keepalives

IKE keepalives track reachability of peers by sending hello packets between peers. During the typical life of an IKE Phase 1 SA, packets are only exchanged over this IKE Phase 1 SA when an IPsec IKE Quick Mode (QM) negotiation is required at the expiration of the IPSec Tunnel SAs. There is no standards-based mechanism for either type of SA to detect the loss of a peer, except when the QM negotiation fails. Most IPsec implementations use the Dead Peer Detection (i.e. Keepalive) mechanism to determine whether connectivity has been lost with a peer before the expiration of the IPsec Tunnel SA's.

All tests using IKEv1 MUST use the same IKE keepalive parameters.

6.6.4. IKE DH-group

There are 3 Diffie-Hellman groups which can be supported by IPsec standards compliant devices:

- DH-group 1: 768 bits

- DH-group 2: 1024 bits

- DH-group 14: 2048 bits

DH-group 2 MUST be tested, to support the new IKEv1 algorithm requirements listed in [RFC4109]. It is recommended that the same DH-group be used for both IKE Phase 1 and IKE phase 2. All test methodologies using IKE MUST report which DH-group was configured to be used for IKE Phase 1 and IKE Phase 2 negotiations.

6.6.5. IKE SA / IPsec SA Lifetime

An IKE SA or IPsec SA is retained by each peer until the Tunnel lifetime expires. IKE SA's and IPsec SA's have individual lifetime parameters. In many real-world environments, the IPsec SA's will be configured with shorter lifetimes than that of the IKE SA's. This will force a rekey to happen more often for IPsec SA's.

When the initiator begins an IKE negotiation between itself and a remote peer (the responder), an IKE policy can be selected only if the lifetime of the responder's policy is shorter than or equal to the lifetime of the initiator's policy. If the lifetimes are not the same, the shorter lifetime will be used.

To avoid any incompatibilities in data plane benchmark testing, all devices MUST have the same IKE SA and IPsec SA lifetime configured and they must be configured to a time which exceeds the test duration timeframe or the total number of bytes to be transmitted during the test.

Note that the IPsec SA lifetime MUST be equal to or less than the IKE SA lifetime. Both the IKE SA lifetime and the IPsec SA lifetime used MUST be reported. This parameter SHOULD be variable when testing IKE rekeying performance.

6.6.6. IPsec Selectors

All tests MUST be performed using standard IPsec selectors.

7. Capacity

7.1. IKE SA Capacity

Objective:

Measure the maximum number of IKE SA's that can be sustained on an IPsec Device.

Procedure:

The IPsec Device under test initially MUST NOT have any Active IPsec Tunnels. The Initiator (either a tester or an IPsec peer) will start the negotiation of an IPsec Tunnel (a single Phase 1 SA and a pair Phase 2 SA's).

After it is detected that the tunnel is established, a limited number (50 packets RECOMMENDED) SHALL be sent through the tunnel. If all packet are received by the Responder (i.e. the DUT), a new IPsec Tunnel may be attempted.

This proces will be repeated until no more IPsec Tunnels can be established.

At the end of the test, a traffic pattern is sent to the initiator that will be distributed over all Active IPsec Tunnels, where each tunnel will need to propagate a fixed number of packets at a minimum rate of 5 pps. When all packets sent by the Iniator are being received by the Responder, the test has succesfully determined the IKE SA Capacity. If however this final check fails, the test needs to be re-executed with a lower number of Active IPsec Tunnels. There MAY be a need to enforce a lower number of Active IPsec Tunnels i.e. an upper limit of Active IPsec Tunnel SHOULD be defined in the test.

Reporting Format:

The reporting format SHOULD be the same as listed in 7.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

7.2. IPsec SA Capacity

Objective:

Measure the maximum number of IPsec SA's that can be sustained on an IPsec Device.

Procedure:

The IPsec Device under test initially MUST NOT have any Active IPsec Tunnels. The Initiator (either a tester or an IPsec peer) will start the negotiation of an IPsec Tunnel (a single Phase 1 SA and a pair Phase 2 SA's).

After it is detected that the tunnel is established, a limited number (50 packets RECOMMENDED) SHALL be sent through the tunnel. If all packet are received by the Responder (i.e. the DUT), a new pair of IPsec SA's may be attempted. This will be achieved by offering a specific traffic pattern to the Initiator that matches a given selector and therfore triggering the negotiation of a new pair of IPsec SA's.

This proces will be repeated until no more IPsec SA' can be established.

At the end of the test, a traffic pattern is sent to the initiator that will be distributed over all IPsec SA's, where each SA will need to propagate a fixed number of packets at a minimum rate of 5 pps. When all packets sent by the Iniator are being received by the Responder, the test has succesfully determined the IPsec SA

Capacity. If however this final check fails, the test needs to be re-executed with a lower number of IPsec SA's. There MAY be a need to enforce a lower number IPsec SA's i.e. an upper limit of IPsec SA's SHOULD be defined in the test.

Reporting Format:

The reporting format SHOULD be the same as listed in 7.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

8. Throughput

This section contains the description of the tests that are related to the characterization of the packet forwarding of a DUT/SUT in an IPsec environment. Some metrics extend the concept of throughput presented in RFC 1242. The notion of Forwarding Rate is cited in RFC2285.

A separate test SHOULD be performed for Throughput tests using IPv4/UDP, IPv6/UDP, IPv4/TCP and IPv6/TCP traffic.

8.1. Throughput baseline

Objective:

Measure the intrinsic cleartext throughput of a device without the use of IPsec. The throughput baseline methodology and reporting format is derived from [RFC2544].

Procedure:

Send a specific number of frames that matches the IPsec SA selector(s) to be tested at a specific rate through the DUT and then count the frames that are transmitted by the DUT. If the count of offered frames is equal to the count of received frames, the rate of the offered stream is increased and the test is rerun. If fewer frames are received than were transmitted, the rate of the offered stream is reduced and the test is rerun.

The throughput is the fastest rate at which the count of test frames transmitted by the DUT is equal to the number of test frames sent to it by the test equipment.

Reporting Format:

The results of the throughput test SHOULD be reported in the form of a graph. If it is, the x coordinate SHOULD be the frame size, the y coordinate SHOULD be the frame rate. There SHOULD be at least two lines on the graph. There SHOULD be one line showing the theoretical frame rate for the media at the various frame sizes. The second line SHOULD be the plot of the test results. Additional lines MAY be used on the graph to report the results for each type of data stream tested. Text accompanying the graph SHOULD indicate the protocol, data stream format, and type of media used in the tests.

We assume that if a single value is desired for advertising purposes the vendor will select the rate for the minimum frame size for the media. If this is done then the figure MUST be expressed in packets per second. The rate MAY also be expressed in bits (or bytes) per second if the vendor so desires. The statement of performance MUST include:

- Measured maximum frame rate

- Size of the frame used

- Theoretical limit of the media for that frame size

- Type of protocol used in the test

Even if a single value is used as part of the advertising copy, the full table of results SHOULD be included in the product data sheet.

8.2. IPsec Throughput

Objective:

Measure the intrinsic throughput of a device utilizing IPsec.

Procedure:

Send a specific number of cleartext frames that match the IPsec SA selector(s) at a specific rate through the DUT/SUT. DUTa will encrypt the traffic and forward to DUTb which will in turn decrypt the traffic and forward to the testing device. The testing device counts the frames that are transmitted by the DUTb. If the count of offered frames is equal to the count of received frames, the rate of the offered stream is increased and the test is rerun. If fewer frames are received than were transmitted, the rate of the

offered stream is reduced and the test is rerun.

The IPsec Throughput is the fastest rate at which the count of test frames transmitted by the DUT/SUT is equal to the number of test frames sent to it by the test equipment.

For tests using multiple IPsec SA's, the test traffic associated with the individual traffic selectors defined for each IPsec SA MUST be sent in a round robin type fashion to keep the test balanced so as not to overload any single IPsec SA.

Reporting format:

The reporting format SHOULD be the same as listed in 7.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

8.3. IPsec Encryption Throughput

Objective:

Measure the intrinsic DUT vendor specific IPsec Encryption Throughput.

Procedure:

Send a specific number of cleartext frames that match the IPsec SA selector(s) at a specific rate to the DUT. The DUT will receive the cleartext frames, perform IPsec operations and then send the IPsec protected frame to the tester. Upon receipt of the encrypted packet, the testing device will timestamp the packet(s) and record the result. If the count of offered frames is equal to the count of received frames, the rate of the offered stream is increased and the test is rerun. If fewer frames are received than were transmitted, the rate of the offered stream is reduced and the test is rerun.

The IPsec Encryption Throughput is the fastest rate at which the count of test frames transmitted by the DUT is equal to the number of test frames sent to it by the test equipment.

For tests using multiple IPsec SA's, the test traffic associated with the individual traffic selectors defined for each IPsec SA MUST be sent in a round robin type fashion to keep the test balanced so as not to overload any single IPsec SA.

Reporting format:

The reporting format SHOULD be the same as listed in 7.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

8.4. IPsec Decryption Throughput

Objective:

Measure the intrinsic DUT vendor specific IPsec Decryption Throughput.

Procedure:

Send a specific number of IPsec protected frames that match the IPsec SA selector(s) at a specific rate to the DUT. The DUT will receive the IPsec protected frames, perform IPsec operations and then send the cleartext frame to the tester. Upon receipt of the cleartext packet, the testing device will timestamp the packet(s) and record the result. If the count of offered frames is equal to the count of received frames, the rate of the offered stream is increased and the test is rerun. If fewer frames are received than were transmitted, the rate of the offered stream is reduced and the test is rerun.

The IPsec Decryption Throughput is the fastest rate at which the count of test frames transmitted by the DUT is equal to the number of test frames sent to it by the test equipment.

For tests using multiple IPsec SAs, the test traffic associated with the individual traffic selectors defined for each IPsec SA MUST be sent in a round robin type fashion to keep the test balanced so as not to overload any single IPsec SA.

Reporting format:

The reporting format SHOULD be the same as listed in 7.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

9. Latency

This section presents methodologies relating to the characterization of the forwarding latency of a DUT/SUT. It extends the concept of

latency characterization presented in [RFC2544] to an IPsec environment.

A separate tests SHOULD be performed for latency tests using IPv4/UDP, IPv6/UDP, IPv4/TCP and IPv6/TCP traffic.

In order to lessen the effect of packet buffering in the DUT/SUT, the latency tests MUST be run at the measured IPsec throughput level of the DUT/SUT; IPsec latency at other offered loads is optional.

Lastly, [RFC1242] and [RFC2544] draw distinction between two classes of devices: "store and forward" and "bit-forwarding". Each class impacts how latency is collected and subsequently presented. See the related RFCs for more information. In practice, much of the test equipment will collect the latency measurement for one class or the other, and, if needed, mathematically derive the reported value by the addition or subtraction of values accounting for medium propagation delay of the packet, bit times to the timestamp trigger within the packet, etc. Test equipment vendors SHOULD provide documentation regarding the composition and calculation latency values being reported. The user of this data SHOULD understand the nature of the latency values being reported, especially when comparing results collected from multiple test vendors. (E.g., If test vendor A presents a "store and forward" latency result and test vendor B presents a "bit-forwarding" latency result, the user may erroneously conclude the DUT has two differing sets of latency values.).

9.1. Latency Baseline

Objective:

Measure the intrinsic latency (min/avg/max) introduced by a device without the use of IPsec.

Procedure:

First determine the throughput for the DUT/SUT at each of the listed frame sizes. Send a stream of frames at a particular frame size through the DUT at the determined throughput rate using frames that match the IPsec SA selector(s) to be tested. The stream SHOULD be at least 120 seconds in duration. An identifying tag SHOULD be included in one frame after 60 seconds with the type of tag being implementation dependent. The time at which this frame is fully transmitted is recorded (timestamp A). The receiver logic in the test equipment MUST recognize the tag information in the frame stream and record the time at which the tagged frame was received (timestamp B).

The latency is timestamp B minus timestamp A as per the relevant definition from RFC 1242, namely latency as defined for store and forward devices or latency as defined for bit forwarding devices.

The test MUST be repeated at least 20 times with the reported value being the average of the recorded values.

Reporting Format

The report MUST state which definition of latency (from [RFC1242]) was used for this test. The latency results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the latency test was run for that frame size, for the media types tested, and for the resultant latency values for each type of data stream tested.

9.2. IPsec Latency

Objective:

Measure the intrinsic IPsec Latency (min/avg/max) introduced by a device when using IPsec.

Procedure:

First determine the throughput for the DUT/SUT at each of the listed frame sizes. Send a stream of cleartext frames at a particular frame size through the DUT/SUT at the determined throughput rate using frames that match the IPsec SA selector(s) to be tested. DUTa will encrypt the traffic and forward to DUTb which will in turn decrypt the traffic and forward to the testing device.

The stream SHOULD be at least 120 seconds in duration. An identifying tag SHOULD be included in one frame after 60 seconds with the type of tag being implementation dependent. The time at which this frame is fully transmitted is recorded (timestamp A). The receiver logic in the test equipment MUST recognize the tag information in the frame stream and record the time at which the tagged frame was received (timestamp B).

The IPsec Latency is timestamp B minus timestamp A as per the relevant definition from [RFC1242], namely latency as defined for store and forward devices or latency as defined for bit forwarding devices.

The test MUST be repeated at least 20 times with the reported value being the average of the recorded values.

Reporting format:

The reporting format SHOULD be the same as listed in 8.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

9.3. IPsec Encryption Latency

Objective:

Measure the DUT vendor specific IPsec Encryption Latency for IPsec protected traffic.

Procedure:

Send a stream of cleartext frames at a particular frame size through the DUT/SUT at the determined throughput rate using frames that match the IPsec SA selector(s) to be tested.

The stream SHOULD be at least 120 seconds in duration. An identifying tag SHOULD be included in one frame after 60 seconds with the type of tag being implementation dependent. The time at which this frame is fully transmitted is recorded (timestamp A). The DUT will receive the cleartext frames, perform IPsec operations and then send the IPsec protected frames to the tester. Upon receipt of the encrypted frames, the receiver logic in the test equipment MUST recognize the tag information in the frame stream and record the time at which the tagged frame was received (timestamp B).

The IPsec Encryption Latency is timestamp B minus timestamp A as per the relevant definition from [RFC1242], namely latency as defined for store and forward devices or latency as defined for bit forwarding devices.

The test MUST be repeated at least 20 times with the reported value being the average of the recorded values.

Reporting format:

The reporting format SHOULD be the same as listed in 8.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

9.4. IPsec Decryption Latency

Objective:

Measure the DUT Vendor Specific IPsec Decryption Latency for IPsec protected traffic.

Procedure:

Send a stream of IPsec protected frames at a particular frame size through the DUT/SUT at the determined throughput rate using frames that match the IPsec SA selector(s) to be tested.

The stream SHOULD be at least 120 seconds in duration. An identifying tag SHOULD be included in one frame after 60 seconds with the type of tag being implementation dependent. The time at which this frame is fully transmitted is recorded (timestamp A). The DUT will receive the IPsec protected frames, perform IPsec operations and then send the cleartext frames to the tester. Upon receipt of the decrypted frames, the receiver logic in the test equipment MUST recognize the tag information in the frame stream and record the time at which the tagged frame was received (timestamp B).

The IPsec Decryption Latency is timestamp B minus timestamp A as per the relevant definition from [RFC1242], namely latency as defined for store and forward devices or latency as defined for bit forwarding devices.

The test MUST be repeated at least 20 times with the reported value being the average of the recorded values.

Reporting format:

The reporting format SHOULD be the same as listed in 8.1 with the additional requirement that the Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

10. Time To First Packet

Objective:

Measure the time it takes to transmit a packet when no SAs have been established.

Procedure:

Determine the IPsec throughput for the DUT/SUT at each of the listed frame sizes. Start with a DUT/SUT with Configured Tunnels. Send a stream of cleartext frames at a particular frame size through the DUT/SUT at the determined throughput rate using frames that match the IPsec SA selector(s) to be tested.

The time at which the first frame is fully transmitted from the testing device is recorded as timestamp A. The time at which the testing device receives its first frame from the DUT/SUT is recorded as timestamp B. The Time To First Packet is the difference between Timestamp B and Timestamp A.

Note that it is possible that packets can be lost during IPsec Tunnel establishment and that timestamp A & B are not required to be associated with a unique packet.

Reporting format:

The Time To First Packet results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the TTFP test was run for that frame size, for the media types tested, and for the resultant TTFP values for each type of data stream tested. The Security Context parameters defined in 5.6 and utilized for this test MUST be included in any statement of performance.

11. Frame Loss Rate

This section presents methodologies relating to the characterization of frame loss rate, as defined in [RFC1242], in an IPsec environment.

11.1. Frame Loss Baseline

Objective:

To determine the frame loss rate, as defined in [RFC1242], of a DUT/SUT throughout the entire range of input data rates and frame sizes without the use of IPsec.

Procedure:

Send a specific number of frames at a specific rate through the DUT/SUT to be tested using frames that match the IPsec SA selector(s) to be tested and count the frames that are transmitted by the DUT/SUT. The frame loss rate at each point is calculated

using the following equation:

```
( ( input_count - output_count ) * 100 ) / input_count
```

The first trial SHOULD be run for the frame rate that corresponds to 100% of the maximum rate for the frame size on the input media. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used and then for 80% of this rate. This sequence SHOULD be continued (at reducing 10% intervals) until there are two successive trials in which no frames are lost. The maximum granularity of the trials MUST be 10% of the maximum rate, a finer granularity is encouraged.

Reporting Format:

The results of the frame loss rate test SHOULD be plotted as a graph. If this is done then the X axis MUST be the input frame rate as a percent of the theoretical rate for the media at the specific frame size. The Y axis MUST be the percent loss at the particular input rate. The left end of the X axis and the bottom of the Y axis MUST be 0 percent; the right end of the X axis and the top of the Y axis MUST be 100 percent. Multiple lines on the graph MAY used to report the frame loss rate for different frame sizes, protocols, and types of data streams.

11.2. IPsec Frame Loss

Objective:

To measure the frame loss rate of a device when using IPsec to protect the data flow.

Procedure:

Ensure that the DUT/SUT is in active tunnel mode. Send a specific number of cleartext frames that match the IPsec SA selector(s) to be tested at a specific rate through the DUT/SUT. DUTa will encrypt the traffic and forward to DUTb which will in turn decrypt the traffic and forward to the testing device. The testing device counts the frames that are transmitted by the DUTb. The frame loss rate at each point is calculated using the following equation:

```
( ( input_count - output_count ) * 100 ) / input_count
```

The first trial SHOULD be run for the frame rate that corresponds to 100% of the maximum rate for the frame size on the input media. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used and then for 80% of this rate. This sequence SHOULD be continued (at reducing 10% intervals) until there are two successive trials in which no frames are lost. The maximum granularity of the trials MUST be 10% of the maximum rate, a finer granularity is encouraged.

Reporting Format:

The reporting format SHOULD be the same as listed in 10.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

11.3. IPsec Encryption Frame Loss

Objective:

To measure the effect of IPsec encryption on the frame loss rate of a device.

Procedure:

Send a specific number of cleartext frames that match the IPsec SA selector(s) at a specific rate to the DUT. The DUT will receive the cleartext frames, perform IPsec operations and then send the IPsec protected frame to the tester. The testing device counts the encrypted frames that are transmitted by the DUT. The frame loss rate at each point is calculated using the following equation:

((input_count - output_count) * 100) / input_count

The first trial SHOULD be run for the frame rate that corresponds to 100% of the maximum rate for the frame size on the input media. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used and then for 80% of this rate. This sequence SHOULD be continued (at reducing 10% intervals) until there are two successive trials in which no frames are lost. The maximum granularity of the trials MUST be 10% of the maximum rate, a finer granularity is encouraged.

Reporting Format:

The reporting format SHOULD be the same as listed in 10.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

11.4. IPsec Decryption Frame Loss

Objective:

To measure the effects of IPsec encryption on the frame loss rate of a device.

Procedure:

Send a specific number of IPsec protected frames that match the IPsec SA selector(s) at a specific rate to the DUT. The DUT will receive the IPsec protected frames, perform IPsec operations and then send the cleartext frames to the tester. The testing device counts the cleartext frames that are transmitted by the DUT. The frame loss rate at each point is calculated using the following equation:

((input_count - output_count) * 100) / input_count

The first trial SHOULD be run for the frame rate that corresponds to 100% of the maximum rate for the frame size on the input media. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used and then for 80% of this rate. This sequence SHOULD be continued (at reducing 10% intervals) until there are two successive trials in which no frames are lost. The maximum granularity of the trials MUST be 10% of the maximum rate, a finer granularity is encouraged.

Reporting format:

The reporting format SHOULD be the same as listed in 10.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

11.5. IKE Phase 2 Rekey Frame Loss

Objective:

To measure the frame loss due to an IKE Phase 2 (i.e. IPsec SA) Rekey event.

Procedure:

The procedure is the same as in 10.2 with the exception that the IPsec SA lifetime MUST be configured to be one-third of the trial test duration or one-third of the total number of bytes to be transmitted during the trial duration.

Reporting format:

The reporting format SHOULD be the same as listed in 10.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

12. Back-to-back Frames

This section presents methodologies relating to the characterization of back-to-back frame processing, as defined in [RFC1242], in an IPsec environment.

12.1. Back-to-back Frames Baseline

Objective:

To characterize the ability of a DUT to process back-to-back frames as defined in [RFC1242], without the use of IPsec.

Procedure:

Send a burst of frames that matches the IPsec SA selector(s) to be tested with minimum inter-frame gaps to the DUT and count the number of frames forwarded by the DUT. If the count of transmitted frames is equal to the number of frames forwarded the length of the burst is increased and the test is rerun. If the number of forwarded frames is less than the number transmitted, the length of the burst is reduced and the test is rerun.

The back-to-back value is the number of frames in the longest burst that the DUT will handle without the loss of any frames. The trial length MUST be at least 2 seconds and SHOULD be repeated at least 50 times with the average of the recorded values being reported.

Reporting format:

The back-to-back results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size and for the resultant average frame count for each type of data stream tested. The standard deviation for each measurement MAY also be reported.

12.2. IPsec Back-to-back Frames

Objective:

To measure the back-to-back frame processing rate of a device when using IPsec to protect the data flow.

Procedure:

Send a burst of cleartext frames that matches the IPsec SA selector(s) to be tested with minimum inter-frame gaps to the DUT/SUT. DUTa will encrypt the traffic and forward to DUTb which will in turn decrypt the traffic and forward to the testing device. The testing device counts the frames that are transmitted by the DUTb. If the count of transmitted frames is equal to the number of frames forwarded the length of the burst is increased and the test is rerun. If the number of forwarded frames is less than the number transmitted, the length of the burst is reduced and the test is rerun.

The back-to-back value is the number of frames in the longest burst that the DUT/SUT will handle without the loss of any frames. The trial length MUST be at least 2 seconds and SHOULD be repeated at least 50 times with the average of the recorded values being reported.

Reporting Format:

The reporting format SHOULD be the same as listed in 11.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

12.3. IPsec Encryption Back-to-back Frames

Objective:

To measure the effect of IPsec encryption on the back-to-back frame processing rate of a device.

Procedure:

Send a burst of cleartext frames that matches the IPsec SA selector(s) to be tested with minimum inter-frame gaps to the DUT. The DUT will receive the cleartext frames, perform IPsec operations and then send the IPsec protected frame to the tester. The testing device counts the encrypted frames that are transmitted by the DUT. If the count of transmitted encrypted frames is equal to the number of frames forwarded the length of the burst is increased and the test is rerun. If the number of forwarded frames is less than the number transmitted, the length of the burst is reduced and the test is rerun.

The back-to-back value is the number of frames in the longest burst that the DUT will handle without the loss of any frames. The trial length MUST be at least 2 seconds and SHOULD be repeated at least 50 times with the average of the recorded values being reported.

Reporting format:

The reporting format SHOULD be the same as listed in 11.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

12.4. IPsec Decryption Back-to-back Frames

Objective:

To measure the effect of IPsec decryption on the back-to-back frame processing rate of a device.

Procedure:

Send a burst of cleartext frames that matches the IPsec SA selector(s) to be tested with minimum inter-frame gaps to the DUT. The DUT will receive the IPsec protected frames, perform IPsec operations and then send the cleartext frame to the tester. The testing device counts the frames that are transmitted by the DUT. If the count of transmitted frames is equal to the number of frames forwarded the length of the burst is increased and the test is rerun. If the number of forwarded frames is less than the number transmitted, the length of the burst is reduced and the test is rerun.

The back-to-back value is the number of frames in the longest burst that the DUT will handle without the loss of any frames. The trial length MUST be at least 2 seconds and SHOULD be repeated at least 50 times with the average of the recorded values being reported.

Reporting format:

The reporting format SHOULD be the same as listed in 11.1 with the additional requirement that the Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

13. IPsec Tunnel Setup Behavior

13.1. IPsec Tunnel Setup Rate

Objective:

Determine the rate at which IPsec Tunnels can be established.

Procedure:

Configure the DUT/SUT with n IKE Phase 1 and corresponding IKE Phase 2 policies. Ensure that no SA's are established and that the DUT/SUT is in configured tunnel mode for all n policies. Send a stream of cleartext frames at a particular frame size through the DUT/SUT at the determined throughput rate using frames with selectors matching the first IKE Phase 1 policy. As soon as the testing device receives its first frame from the DUT/SUT, it knows that the IPsec Tunnel is established and starts sending the next stream of cleartext frames using the same frame size and throughput rate but this time using selectors matching the second IKE Phase 1 policy. This process is repeated until all configured IPsec Tunnels have been established.

The IPsec Tunnel Setup Rate is determined by the following formula:

Tunnel Setup Rate = n / [Duration of Test - (n *
frame_transmit_time)]

The IKE SA lifetime and the IPsec SA lifetime MUST be configured to exceed the duration of the test time. It is RECOMMENDED that n=100 IPsec Tunnels are tested at a minimum to get a large enough sample size to depict some real-world behavior.

Reporting Format:

The Tunnel Setup Rate results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the test was run for that frame size, for the media types tested, and for the resultant Tunnel Setup Rate values for each type of data stream tested. The Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

13.2. IKE Phase 1 Setup Rate

Objective:

Determine the rate of IKE SA's that can be established.

Procedure:

Configure the DUT with n IKE Phase 1 and corresponding IKE Phase 2 policies. Ensure that no SAs are established and that the DUT is in configured tunnel mode for all n policies. Send a stream of cleartext frames at a particular frame size through the DUT at the determined throughput rate using frames with selectors matching the first IKE Phase 1 policy. As soon as the Phase 1 SA is established, the testing device starts sending the next stream of cleartext frames using the same frame size and throughput rate but this time using selectors matching the second IKE Phase 1 policy. This process is repeated until all configured IKE SAs have been established.

The IKE SA Setup Rate is determined by the following formula:

IKE SA Setup Rate = n / [Duration of Test - (n *
frame_transmit_time)]

The IKE SA lifetime and the IPsec SA lifetime MUST be configured to exceed the duration of the test time. It is RECOMMENDED that n=100 IKE SAs are tested at a minumum to get a large enough sample size to depict some real-world behavior.

Reporting Format:

The IKE Phase 1 Setup Rate results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the test was run for that frame size, for the media types tested, and for the resultant IKE Phase 1 Setup Rate values for each type of

data stream tested. The Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

13.3. IKE Phase 2 Setup Rate

Objective:

Determine the rate of IPsec SA's that can be established.

Procedure:

Configure the DUT with a single IKE Phase 1 policy and n corresponding IKE Phase 2 policies. Ensure that no SAs are established and that the DUT is in configured tunnel mode for all policies. Send a stream of cleartext frames at a particular frame size through the DUT at the determined throughput rate using frames with selectors matching the first IPsec SA policy.

The time at which the IKE SA is established is recorded as timestamp A. As soon as the Phase 1 SA is established, the IPsec SA negotiation will be initiated. Once the first IPsec SA has been established, start sending the next stream of cleartext frames using the same frame size and throughput rate but this time using selectors matching the second IKE Phase 2 policy. This process is repeated until all configured IPsec SA's have been established.

The IPsec SA Setup Rate is determined by the following formula:

IPsec SA Setup Rate = n / [Duration of Test - {A +((n-1) *
frame_transmit_time)}]

The IKE SA lifetime and the IPsec SA lifetime MUST be configured to exceed the duration of the test time. It is RECOMMENDED that n=100 IPsec SAs are tested at a minumum to get a large enough sample size to depict some real-world behavior.

Reporting Format:

The IKE Phase 2 Setup Rate results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the test was run for that frame size, for the media types tested, and for the resultant IKE Phase 2 Setup Rate values for each type of data stream tested. The Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

14. IPsec Rekey Behavior

The IPsec Rekey Behavior test all need to be executed by an IPsec aware test device since the test needs to be closely linked with the IKE FSM and cannot be done by offering specific traffic pattern at either the Initiator or the Responder.

14.1. IKE Phase 1 Rekey Rate

Objective:

Determine the maximum rate at which an IPsec Device can rekey IKE SA's.

Procedure:

The IPsec Device under test should initially be set up with the determined IKE SA Capacity number of Active IPsec Tunnels.

The IPsec aware tester should then perform a binary search where it initiates an IKE Phase 1 SA rekey for all Active IPsec Tunnels. The tester MUST timestamp for each IKE SA when it initiated the rekey and MUST timestamp once more once the FSM declares the rekey is completed. Once the itteration is complete the tester now has a table of rekey times for each IKE SA. The reciproce of the average of this table is the IKE Phase 1 Rekey Rate.

This is obviously granted that all IKE SA were able to rekey succesfully. If this is not the case, the IPsec Tunnels are all re-established and the binary search goes to the next value of IKE SA's it will rekey. The process will repeat itself until a rate is determined at which a all SA's in that timeframe rekey correctly.

Reporting Format:

The IKE Phase 1 Rekey Rate results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the test was run for that frame size, for the media types tested, and for the resultant IKE Phase 1 Rekey Rate values for each type of data stream tested. The Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

14.2. IKE Phase 2 Rekey Rate

Objective:

Determine the maximum rate at which an IPsec Device can rekey IPsec SA's.

Procedure:

The IPsec Device under test should initially be set up with the determined IKE SA Capacity number of Active IPsec Tunnels.

The IPsec aware tester should then perform a binary search where it initiates an IKE Phase 2 SA rekey for all IPsec SA's. The tester MUST timestamp for each IPsec SA when it initiated the rekey and MUST timestamp once more once the FSM declares the rekey is completed. Once the itteration is complete the tester now has a table of rekey times for each IPsec SA. The reciproce of the average of this table is the IKE Phase 2 Rekey Rate.

This is obviously granted that all IPsec SA were able to rekey succesfully. If this is not the case, the IPsec Tunnels are all re-established and the binary search goes to the next value of IPsec SA's it will rekey. The process will repeat itself until a rate is determined at which a all SA's in that timeframe rekey correctly.

Reporting Format:

The IKE Phase 2 Rekey Rate results SHOULD be reported in the format of a table with a row for each of the tested frame sizes. There SHOULD be columns for the frame size, the rate at which the test was run for that frame size, for the media types tested, and for the resultant IKE Phase 2 Rekey Rate values for each type of data stream tested. The Security Context parameters defined in 6.7 and utilized for this test MUST be included in any statement of performance.

15. IPsec Tunnel Failover Time

This section presents methodologies relating to the characterization of the failover behavior of a DUT/SUT in a IPsec environment.

In order to lessen the effect of packet buffering in the DUT/SUT, the Tunnel Failover Time tests MUST be run at the measured IPsec throughput level of the DUT. Tunnel Failover Time tests at other offered constant loads are OPTIONAL.

Tunnel Failovers can be achieved in various ways like :

- Failover between two or more software instances of an IPsec stack.

- Failover between two IPsec devices.

- Failover between two or more crypto engines.

- Failover between hardware and software crypto.

In all of the above cases there shall be at least one active IPsec device and a standby device. In some cases the standby device is not present and two or more IPsec devices are backing eachother up in case of a catastrophic device or stack failure. The standby (or potential other active) IPsec Devices can back up the active IPsec Device in either a stateless or statefull method. In the former case, Phase 1 SA's as well as Phase 2 SA's will need to be re-established in order to guarantuee packet forwarding. In the latter case, the SPD and SADB of the active IPsec Device is synchronized to the standby IPsec Device to ensure immediate packet path recovery.

Objective:

Determine the time required to fail over all Active Tunnels from an active IPsec Device to its standby device.

Procedure:

Before a failover can be triggered, the IPsec Device has to be in a state where the active stack/engine/node has a the maximum supported number of Active Tunnnels. The Tunnels will be transporting bidirectional traffic at the Tunnel Throughput rate for the smallest framesize that the stack/engine/node is capable of forwarding (In most cases, this will be 64 Bytes). The traffic should traverse in a round robin fashion through all Active Tunnels.

It is RECOMMENDED that the test is repeated for various number of Active Tunnels as well as for different framesizes and framerates.

When traffic is flowing through all Active Tunnels in steady state, a failover shall be triggered.

Both receiver sides of the testers will now look at sequence counters in the instrumented packets that are being forwarded through the Tunnels. Each Tunnel MUST have it's own counter to keep track of packetloss on a per SA basis.

If the tester observes no sequence number drops on any of the Tunnels in both directions then the Failover Time MUST be listed as 'null', indicating that the failover was immediate and without any packetloss.

In all other cases where the tester observes a gap in the sequence numbers of the instrumented payload of the packets, the tester will monitor all SA's and look for any Tunnels that are still not receiving packets after the Failover. These will be marked as 'pending' Tunnels. Active Tunnels that are forwarding packets again without any packetloss shall be marked as 'recovered' Tunnels. In background the tester will keep monitoring all SA's to make sure that no packets are dropped. If this is the case then the Tunnel in question will be placed back in 'pending' state.

Note that reordered packets can naturally occur after en/decryption. This is not a valid reason to place a Tunnel back in 'pending' state. A sliding window of 128 packets per SA SHALL be allowed before packetloss is declared on the SA.

The tester will wait until all Tunnel are marked as 'recovered'. Then it will find the SA with the largest gap in sequence number. Given the fact that the framesize is fixed and the time of that framesize can easily be calculated for the initiator links, a simple multiplication of the framesize time * largest packetloss gap will yield the Tunnel Failover Time.

If the tester never reaches a state where all Tunnels are marked as 'recovered', the the Failover Time MUST be listed as 'infinite'.

Reporting Format:

The results shall be represented in a tabular format, where the first column will list the number of Active Tunnels, the second column the Framesize, the third column the Framerate and the fourth column the Tunnel Failover Time in milliseconds.

16. Acknowledgements

The authors would like to acknowledge the following individual for their help and participation of the compilation and editing of this document: Michele Bustos, Ixia. ; Paul Hoffman, VPNC

17. References

17.1. Normative References

[RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.

[RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2285] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.

[RFC2393] Shacham, A., Monsour, R., Pereira, R., and M. Thomas, "IP Payload Compression Protocol (IPComp)", RFC 2393, December 1998.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC2402] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.

[RFC2403] Madson, C. and R. Glenn, "The Use of HMAC-MD5-96 within ESP and AH", RFC 2403, November 1998.

[RFC2404] Madson, C. and R. Glenn, "The Use of HMAC-SHA-1-96 within ESP and AH", RFC 2404, November 1998.

[RFC2405] Madson, C. and N. Doraswamy, "The ESP DES-CBC Cipher Algorithm With Explicit IV", RFC 2405, November 1998.

[RFC2406] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

[RFC2407] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.

[RFC2408] Maughan, D., Schneider, M., and M. Schertler, "Internet Security Association and Key Management Protocol (ISAKMP)", RFC 2408, November 1998.

[RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

[RFC2410] Glenn, R. and S. Kent, "The NULL Encryption Algorithm and Its Use With IPsec", RFC 2410, November 1998.

[RFC2411] Thayer, R., Doraswamy, N., and R. Glenn, "IP Security Document Roadmap", RFC 2411, November 1998.

[RFC2412] Orman, H., "The OAKLEY Key Determination Protocol", RFC 2412, November 1998.

[RFC2432] Dubray, K., "Terminology for IP Multicast Benchmarking", RFC 2432, October 1998.

[RFC2451] Pereira, R. and R. Adams, "The ESP CBC-Mode Cipher Algorithms", RFC 2451, November 1998.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

[RFC2547] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March 1999.

[RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.

[RFC4109] Hoffman, P., "Algorithms for Internet Key Exchange version 1 (IKEv1)", RFC 4109, May 2005.

[I-D.ietf-ipsec-ikev2] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", draft-ietf-ipsec-ikev2-17 (work in progress), October 2004.

[I-D.ietf-ipsec-properties] Krywaniuk, A., "Security Properties of the IPsec Protocol Suite", draft-ietf-ipsec-properties-02 (work in progress), July 2002.

17.2. Informative References

[FIPS.186-1.1998] National Institute of Standards and Technology, "Digital Signature Standard", FIPS PUB 186-1, December 1998, <http://csrc.nist.gov/fips/fips1861.pdf>.

Authors' Addresses

Merike Kaeo
Double Shot Security
520 Washington Blvd #363
Marina Del Rey, CA 90292
US

Phone: +1 (310)866-0165
Email: kaeo@merike.com

Tim Van Herck
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
US

Email: herckt@cisco.com

Intellectual Property Statement

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.