Analyze | Assure | Accelerate™

Spirent Communications Test Methodologies

IPv6 Edition

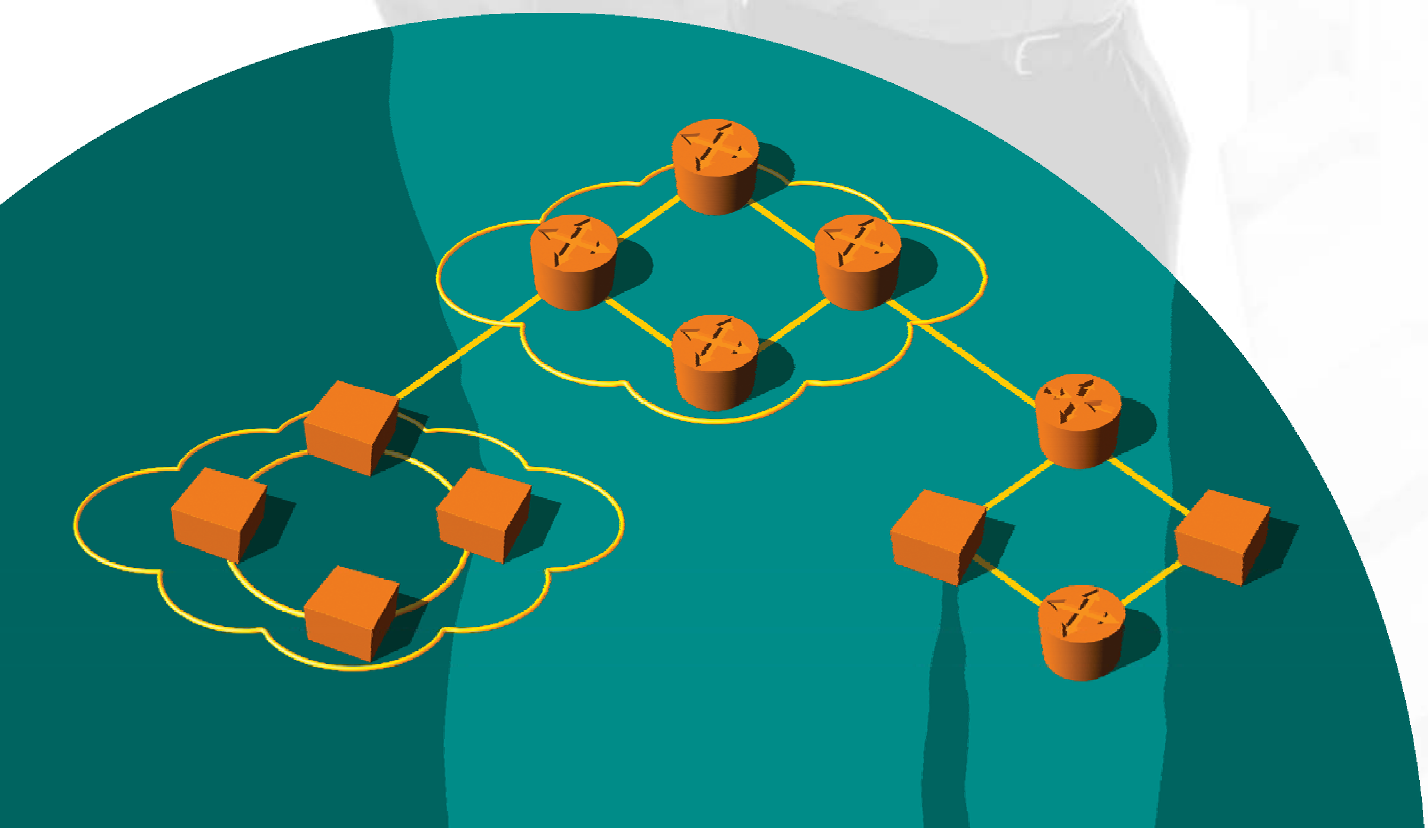

Introduction

Welcome to the IPv6 Edition of the Spirent Communications Test Methodologies series of journals. In this edition, you will find valuable IPv6 test methodologies to use in the lab or with deployed networks.

Spirent Communications is a worldwide provider of integrated performance analysis and service assurance systems for next generation network technologies. Spirent is also the leader in comprehensive test methodologies.

Testing is the only way to understand the diverse conditions that networks and their components will accept and continue to operate efficiently. By testing early and often, your technical staff can monitor the pulse of equipment and determine why systems behave or fail under load.

To supplement your knowledge of IPv6 and learn more about Spirent's IPv6 solutions, visit our IPv6 testing site at ‹www.spirentcom.com/ipv6›.

If you want to review additional test methodologies, visit ‹http://scdn.spirentcom.com›. After registering, you can download free scripts, share information about developing the scripts and obtain help from a technical community comprising industry members from all over the world.

You will find a list of IPv6 test methodologies below. Best wishes with your testing efforts.

Spirent Communications

www.spirentcom.com

IPv6 Test Methodologies

- IP Throughput and Latency for IPv4 and IPv6 (TM-0100) 1
- Packet Loss and Latency for IPv4 and IPv6 (TM-0101) 3
- Back to Back (Burst Size) for IPv4 and IPv6 (TM-0102) 5
- IPv4/IPv6 Dual Stack Performance (TM-0120) 7
- IPv6/IPv4: 6to4 Tunneling Performance Test (TM-0122) 9
- IPv6 Packet Loss and Latency for NAT-PT (TM-0123) 11
- IPv6 Prefix Length Performance (TM-0121) 13
- Router Performance Impact of IPv6 FIB Size (TM-0124) 15
- Glossary 17

IP Throughput and Latency for IPv4 and IPv6

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
- RFC 2544: Benchmarking Methodology for Network Interconnect Devices

Objective

This data plane test is designed to provide throughput and latency information for a single router (device under test or DUT) or for a system of interconnected routers (system under test or SUT).

The throughput of a device or system is the maximum packet-forwarding rate for which the device or system will not drop any of the offered packets. Any packet loss can induce significant delays in the execution of higher layer applications; thus, knowing the maximum data rate a device or system can support without any packet loss is of crucial importance when judging the performance of a router or system of interconnected routers.

This test also determines the latency of a device or system (the time it takes a packet to travel through the device or the system), calculated at the maximum forwarding rate for which no packet loss is experienced (throughput rate).

The test methodology is applicable to IPv4 and IPv6 configurations.

Overview

To determine the throughput and latency of a device under test (DUT) or system under test (SUT), a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and will offer traffic to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that traffic offered by the data transmitters will be forwarded toward the data receivers.

From the transmitter test ports, a predetermined number of packets is offered to the DUT/SUT. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports is compared to the number of packets offered. If packet loss occurs, then the offered load is decreased and the test is repeated. If packet loss is not observed, the test is repeated with an increased number of offered packets. By implementing a binary search pattern, the maximum rate for which no packet loss occurs is recorded. This rate corresponds to the DUT/SUT’s throughput, or the first measurement of the test.

To calculate latency, a predefined test stream is delivered to the DUT/SUT from the transmitter test ports at the calculated throughput rate. The transmitting timestamp, corresponding to when the test packet is sent by a transmitter test port, is subtracted from the receiving timestamp that corresponds to when the packet arrives at a receiver test port. The difference between the receiving timestamp and the transmitting timestamp gives the latency for a test packet. The test is run for a determined period of time, which should be long enough to collect sufficient data for an accurate representation of the system’s latency.

Setup

IP Throughput and Latency Test for IPv4 and IPv6

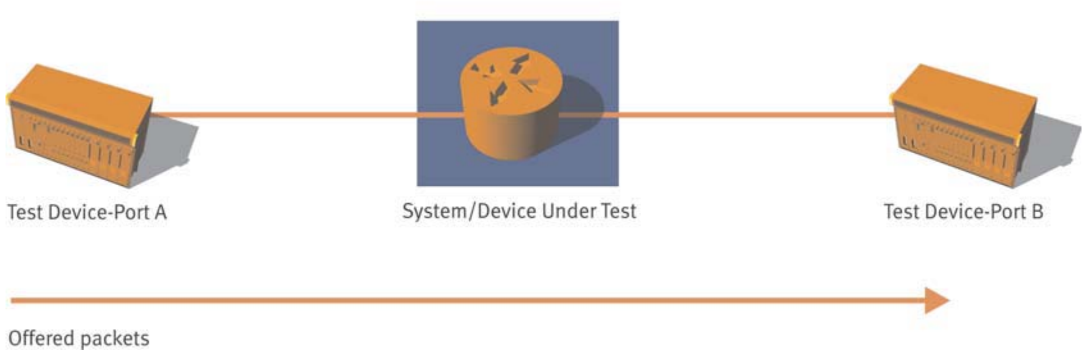

The configuration of the test is presented above. In the baseline test, two test ports are used: Test Port A is the data transmitter, while Test Port B is the data receiver. For more complex tests, multiple data transmitters and multiple data receivers can be used.

Test Steps

1. Configure the DUT/SUT so the traffic offered from transmitter port(s) will be forwarded to the receiver port(s). This can be accomplished via a routing protocol or statically configured routes.
2. Configure the test parameters:
   a. Initial packet rate.
   b. Packet size.
   c. Resolution rate for the binary search.
   d. Latency test run time.
3. Send packets from transmitter port(s):
   Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs, the throughput rate is equal to the offered data rate. Stop the test.
5. If packet loss occurs, decrease the packet rate and repeat the test.
6. Continue the binary search algorithm until the maximum packet rate with no packet loss is calculated (throughput rate).
7. Send a packet stream at the calculated throughput rate and measure the latency for the specified duration of time.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Resolution rate for the binary search.
- Latency test run time.

Test Outcome

Throughput, latency.

Packet Loss and Latency for IPv4 and IPv6

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
- RFC 2544: Benchmarking Terminology for Network Interconnect Devices

Objective

This data plane test determines the packet loss rate and latency for a single router (device under test or DUT) or for a system of interconnected routers (system under test or SUT) for various input data rates and packet sizes.

The packet loss rate of a device or system is the percentage of Layer 3 frames that were offered at the input of the device or system but were not forwarded by the device or system due to hardware and software limitations. Calculating the packet loss rate of a system under different load conditions (input data rate and packet size) serves to evaluate how the system will perform under similar conditions in real-life operation.

This test also determines the latency of a device or system (the time it takes a packet to travel through the device or the system) calculated for the various input data rates, for which packet loss may or may not be experienced.

The test methodology is applicable to IPv4 and IPv6 configurations.

Overview

To determine the packet loss rate and latency of a DUT or SUT, a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and will offer traffic to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that the traffic offered by the data transmitters will be forwarded toward the data receivers.

From the data transmitters, a predetermined number of packets is offered to the DUT/SUT for a given amount of time. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports (RxPacketCount) is compared to the number of packets offered from the transmitter ports (TxPacketCount).

The packet loss rate is calculated using the following formula:

$$[(TxPacketCount - RxPacketCount) \times 100] \div TxPacketCount$$

The packet loss is calculated for input data rates starting with 100 percent of the maximum rate that can be offered from the transmitter test ports. The input data rate is decremented and the test repeated, until there are two successive trials where there is no packet loss. The test may be stopped when the input data rate reaches a user-selected threshold beyond which no measurements are required. The amount by which the input data rate is decremented should be at most 10 percent of the maximum input data rate and can be as low as 1 percent.

System latency is measured at every trial; thus, for every value of the input data rate.

To calculate latency, a predefined test stream is delivered to the DUT/SUT from the transmitter test ports at the calculated throughput rate. The transmitting timestamp, corresponding to when a test packet is sent by a transmitter test port, is subtracted from the receiving timestamp that corresponds to when the packet arrives at a receiver test port. The difference between the receiving timestamp and the transmitting timestamp give the latency for a test packet. The test is run for a determined period of time, which should be long enough to collect sufficient data for an accurate representation of the system's latency.

Setup

Packet Loss and Latency Test for IPv4 and IPv6

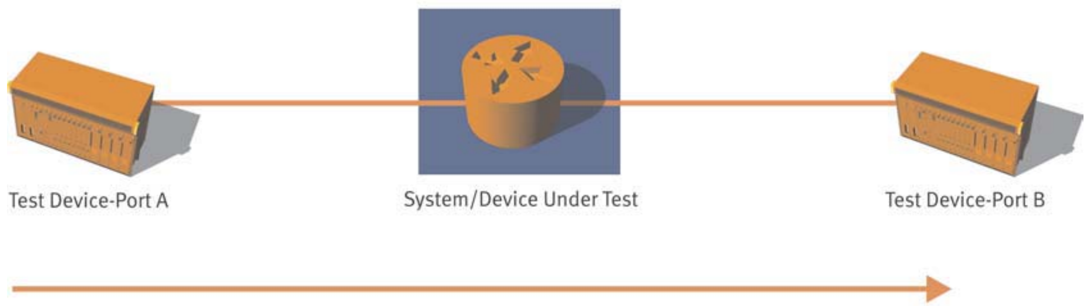

The configuration of the test is presented in the setup diagram. In the baseline test, two test ports are used: Test Port A is the data transmitter, while Test Port B is the data receiver. For more complex tests, multiple data transmitters and multiple data receivers can be used.

Test Steps

1. Configure the DUT/SUT so the traffic offered from the transmitter port(s) is forwarded to the receiver port(s). This can be accomplished via a routing protocol or statically configured routes.
2. Configure the test parameters:
   a. Initial packet rate.
   b. Packet size.
   c. Resolution by which the input data rate is decremented.
   d. Run time for every trial.
3. Send packets from transmitter port(s). Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs in two successive trials, stop the test.
5. If packet loss occurs, calculate the packet loss rate and the latency.
6. Decrement the input data rate and repeat test.
7. The test stops when no packet loss is recorded in two successive trials, or when the input data rate reaches a user-selected threshold.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Test time.
- Packet rate.
- Rate increment/decrement resolution.

Test Outcome

Packet loss, latency.

Back-to-Back (Burst Size) Test for IPv4 and IPv6

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
- RFC 2544: Benchmarking Terminology for Network Interconnect Devices

Objective

This data plane test determines the maximum number of packets a router (device under test or DUT), or system of interconnected routers (system under test or SUT), can forward “back-to-back” without packet loss. The number of packets in the longest burst that does not cause packet loss is the back-to-back value.

In a back-to-back test, packets are delivered at full line rate with no pause between successive packets, except the required “legal” separation for a given technology and physical medium.

The test methodology is applicable to IPv4 and IPv6 configurations.

Overview

To determine the back-to-back value of a device under test (DUT) or system under test (SUT), a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and offer traffic bursts to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that the traffic offered by the data transmitters will be forwarded toward the data receivers.

From the transmitter test ports, a burst of back-to-back packets is offered to the DUT/SUT. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports is compared to the number of packets offered. If packet loss occurs, then the burst size is decreased (the number of back-to-back packets is decreased) and the test is repeated. If no packet loss is observed in the next iteration, then the burst size is increased and the test is repeated. By implementing a binary search, the back-to-back value is determined.

Setup

Back-to-Back Test for IPv4 and IPv6 Capable Routers

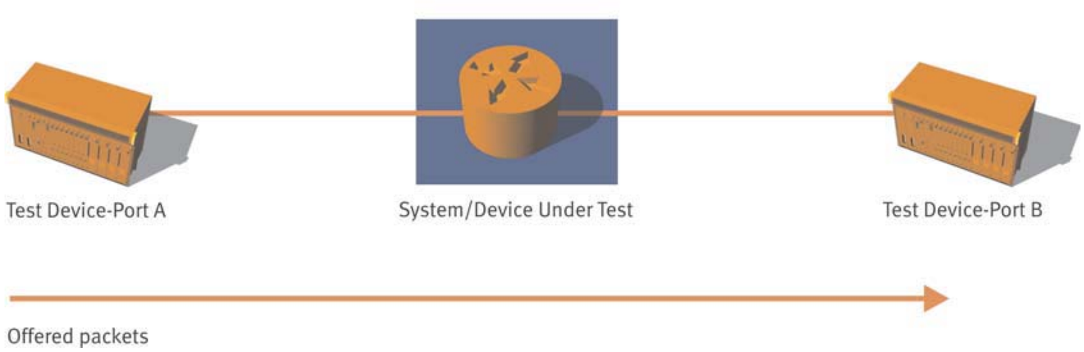

In the baseline test, two test ports are used: Test Port A is the data transmitter, while Test Port B is the data receiver. For more complex tests, multiple data transmitters and multiple data receivers can be used.

Test Steps

1. Configure the DUT/SUT so traffic offered from transmitter port(s) is forwarded to the receiver port(s). This can be accomplished via a routing protocol or statically configured routes.
2. Configure the test parameters:
   a. Packet size.
   b. Burst size (number of packets).
   c. Packet rate.
   d. Burst increment/decrement resolution (for the binary search).
3. Send packets from transmitter port(s) for a determined period of time. Start at the full line rate. Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs, the DUT/SUT can handle back-to-back packets coming at full line rate; thus, there is no burst size limitation. Stop the test.
5. If packet loss occurs, decrease the burst size and repeat the test.
6. Continue the binary search algorithm until the maximum burst size with no packet loss is determined (the back-to-back value).

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Burst size.
- Packet rate.
- Burst increment/decrement resolution.

Test Outcome

Back-to-back value.

IPv4/IPv6 Dual Stack Performance

- RFC 2893: Transition Mechanisms for IPv6 Hosts and Routers
- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
- RFC 2544: Benchmarking Methodology for Network Interconnect Devices

Objective

This data plane test determines the ability of the device under test (DUT) to forward both IPv4 and IPv6 packets. Devices operating in this mode are commonly referred to as running a dual stack.

RFC 2893 also covers "6over4" tunneling. This topic is dealt with in a separate performance test in this IPv6 journal.

The key measurement areas are packet loss and latency. The packet loss rate of a device or system is the percentage of Layer 3 frames that were offered at the ingress of the device or system but were not forwarded by the device or system due to hardware and software limitations. Calculating the packet loss rate of a system under different load conditions (input data rate and packet size) serves to evaluate how the system will perform under similar conditions in real-life operation.

Packet latency is also determined for a device or system (the time it takes a packet to travel through the device or system). Packet latency is calculated for the various input data rates and packet sizes, for which packet loss may or may not be experienced.

Overview

Currently, many routers and switches are capable of running a dual stack. Early adopters of IPv6 should, however, be aware of possible performance limitations when operating in this mode. Not all devices currently exploit their hardware to assist in switching or routing. Many of these functions could be provided by software only. For some time to come, it would be prudent to characterize the performance of devices operating in this mode before deployment.

To determine the packet loss rate and latency of a DUT or system under test (SUT), a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and will offer traffic to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that the traffic offered by the data transmitters will be forwarded to the data receivers.

From the data transmitters, a predetermined number of packets is offered to the DUT/SUT for a given amount of time. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports (RxPacketCount) is compared to the number of packets offered from the transmitter ports (TxPacketCount).

The packet loss rate is calculated using the following formula:

$$[(\text{TxPacketCount} - \text{RxPacketCount}) \times 100] \div \text{TxPacketCount}$$

The packet loss is calculated for input data rates starting with 100 percent of the maximum rate that can be offered from the transmitter test ports. The input data rate is decremented and the test repeated until there are two successive trials where there is no packet loss. The test may be stopped when the input data rate reaches a user-selected threshold beyond which no measurements are required. The amount by which the input data rate is decremented should be at most 10 percent of the maximum input data rate and can be as low as 1 percent.

System latency is measured at every trial; thus, for every value of the input data rate.

To calculate latency, a predefined test stream is delivered to the DUT/SUT from the transmitter test ports at the calculated throughput rate. The transmitting time stamp, corresponding to when a test packet is sent by a transmitter test port, is subtracted from the receiving time stamp that corresponds to when the packet arrives at a receiver test port. The difference between the receiving time stamp and the transmitting time stamp give the latency for a test packet. The test is run for a determined period of time, which should be long enough to collect sufficient data for an accurate representation of the system's latency.

Setup (see diagram)

Packet Loss Rate and Latency Test
for Dual Stack Capable Routers

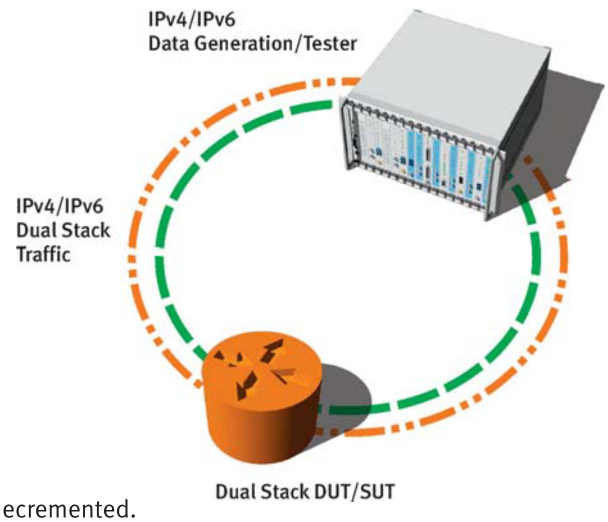

Test Steps

1. Configure the DUT/SUT for dual stack:
   Typically, this is achieved by simply enabling IPv6 on the port(s) to be tested.
2. Configure the test parameters:
   a. Initial packet rate.
   b. Packet size.
   c. Resolution by which the input data rate is decremented.
   d. Run time for every trial.
   e. Set the test equipment to share the load between IPv6 and IPv4.
3. Send packets from transmitter port(s):
   Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
   a. If no packet loss occurs in two successive trials, stop the test.
   b. If packet loss occurs, calculate the packet loss rate and the latency.
   c. Decrement the input data rate and repeat the test.
   d. Stop the test when no packet loss is recorded in two successive trials, or when the input data rate reaches a user-selected threshold.

Note: Depending on the implementation of IPv6 on the DUT, it may not be possible to share the load equally between IPv6 and IPv4. In this situation, it may be best to set the traffic disproportionately. For example, run the test with 10 percent IPv6 traffic and 90 percent IPv4. Other ratios should be used to fully characterize the DUT/SUT.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Burst size.
- Packet rate.
- Burst incremental/decrement resolution.
- Ratio of IPv6 traffic to IPv4 traffic.

Test Outcome

Packet loss, latency.

IPv6/IPv4: 6to4 Tunneling Performance

- RFC 3056: Connection of IPv6 Domains via IPv4 Clouds
- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
- RFC 2544: Benchmarking Methodology for Network Interconnect Devices

Objective

This data plane test determines the ability of the device under test (DUT) or system under test (SUT) to encapsulate IPv6 packets on the ingress port with an IPv4 header for transport across an IPv4 network. Prior to arriving at the destination IPv6 network, this IPv4 header will need to be removed by a router to expose the original IPv6 packet.

The key measurement areas are packet loss and latency. The packet loss rate of a device or system is the percentage of Layer 3 frames offered at the ingress of the device or system but not forwarded due to hardware and software limitations. Calculating the packet loss rate of a system under different load conditions (input data rate and packet size) serves to evaluate how the system will perform under similar conditions in real-life operation.

Packet latency is also determined for a device or system (the time it takes a packet to travel through the device or system). Packet latency is calculated for the various input data rates and packet sizes, for which packet loss may or may not be experienced.

Overview

There are a number of tunneling mechanisms available for IPv6. This test focuses on “6to4.” This mechanism has a well known prefix of 2002 followed by the IPv4 address of the tunnel destination. A common syntax for this address is 2002:V4ADDR::/48. The 2002 prefix has been assigned by the Internet Assigned Numbers Authority (IANA). This is a well known address that routers will recognize, indicating this packet will need to be encapsulated with IPv4.

To determine the packet loss rate and latency of a DUT or SUT, a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and will offer traffic to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that traffic offered by the data transmitters will be forwarded toward the data receivers.

From the data transmitters, a predetermined number of packets is offered to the DUT/SUT for a given amount of time. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports (RxPacketCount) is compared to the number of packets offered from the transmitter ports (TxPacketCount).

The packet loss rate is calculated using the following formula:

$$[(\text{TxPacketCount} - \text{RxPacketCount}) \times 100] \div \text{TxPacketCount}$$

The packet loss is calculated for input data rates starting with 100 percent of the maximum rate that can be offered from the transmitter test ports. The input data rate is decremented and the test repeated, until there are two successive trials where there is no packet loss. The test may be stopped when the input data rate reaches a user-selected threshold beyond which no measurements are required. The amount by which the input data rate is decremented should be at most 10 percent of the maximum input data rate and can be as low as 1 percent.

System latency is measured at every trial; thus, for every value of the input data rate.

To calculate latency, a predefined test stream is delivered to the DUT/SUT from the transmitter test ports at the calculated throughput rate. The transmitting time stamp, corresponding to when a test packet is sent by a transmitter test port, is subtracted from the receiving time stamp that corresponds to when the packet is accepted by a receiver test port. The difference between the receiving time stamp and the transmitting time stamp give the latency for a test packet. The test is run for a determined period of time, which should be long enough to collect sufficient data for an accurate representation of the system's latency.

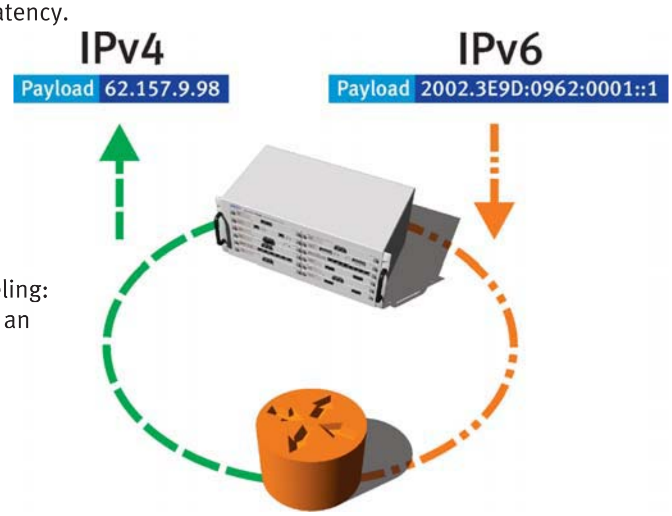

Setup (see diagram)

Packet Loss Rate and Latency Test for 6to4 Capable Routers

Test Steps

1. Configure the DUT/SUT for 6to4 tunneling: Traffic offered at the ingress port with an IPv6 packet header needs to be encapsulated with an IPv4 packet to tunnel through to the receiver port(s).
2. Configure the test parameters:
   a. Initial packet rate.
   b. Packet size.
   c. Resolution by which the input data rate is decremented.
   d. Run time for every trial.
3. Send packets from transmitter port(s): Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs in two successive trials, stop the test.
5. If packet loss occurs, calculate the packet loss rate and the latency.
6. Decrement the input data rate and repeat test.
7. Stop the test when no packet loss is recorded in two successive trials, or when the input data rate reaches a user-selected threshold.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Burst size.
- Packet rate.
- Burst incremental/decrement resolution.

Test Outcome

Packet loss, latency.

IPv6 Packet Loss and Latency for NAT-PT

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
- RFC 2544: Benchmarking Methodology for Network Interconnect Devices
- RFC 2766: Network Address Translation - Protocol Translation (NAT-PT)

Objective

This data plane test measures the number of packets a device under test (DUT) will lose (packet loss) and the increase in delay (latency) when Network Address Translation - Protocol Translation (NAT-PT) is used. The measurements used in this test are based on formulas in TM-0101 (Packet Loss and Latency for IPv4 and IPv6) found in this IPv6 test methodology journal and the previous issue, “Next Generation Network Test Methodologies: Router Performance Testing.”

Overview

NAT-PT is an enabler technology in which IPv4 and IPv6 devices can communicate. Based on IPv4 NAT, NAT-PT takes packets from either an IPv4 or IPv6 network and creates the appropriate address to communicate on the other network. As with NAT, this process involves time to make the conversion. However, unlike NAT, NAT-PT always requires the additional processing required from the larger address (IPv4 = 32 bits, IPv6 = 128 bits) and recalculation of header and checksum information.

To determine the impact of NAT-PT on the DUT, a minimum of two test ports will be required with one port being able to transmit IPv6 and the other port being able to receive IPv4. Both ports will need the ability to track packets using a “signature tag” since the IP address will be different on each side.

One or more ports will act as data transmitter(s) and will offer traffic to the DUT. The other port(s) will act as data receivers and will receive traffic from the DUT. It is important that the DUT is configured to allow traffic to be forwarded between the test ports.

From the data transmitter test port(s), IPv6 packets are offered to the DUT. The DUT will replace the IPv6 addresses with IPv4 addresses and then forward them to the receiver port(s). The number of IPv4 packets accepted at the receiver port(s) is compared to the number of IPv6 packets offered. Any packet loss and the amount of latency are recorded as the test is repeated while adjusting variables such as packet rate and packet size.

Setup

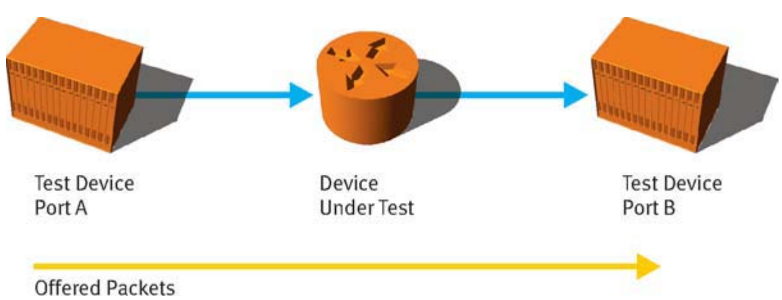

NAT-PT Test of IPv4 and IPv6 Traffic

The configuration of the test is presented in the setup diagram.

In the baseline test, two test ports are used: Test Port A is the data transmitter and is offering IPv6 traffic while Test Port B is the data receiver and accepting IPv4 traffic. For more complex tests, multiple data transmitters and multiple data receivers can be used.

Test Steps

1. Configure the DUT so the traffic offered from the transmitter port(s) is forwarded to the receiver port(s).
2. Configure the test parameters:
   a. Initial packet rate.
   b. Packet size.
   c. Resolution by which the input data rate is decremented.
   d. Run time for every trial.
   e. Number of flows/streams.
3. Send packets from transmitter port(s). Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs in two successive trials, stop the test.
5. If packet loss occurs, calculate the packet loss and the latency.
6. Decrement the input data rate and repeat test.
7. Stop the test when no packet loss is recorded in two successive trials, or when the input data rate reaches a user-selected threshold.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Test time.
- Packet rate.
- Rate increment/decrement resolution.
- Number of flows/streams.

Test Outcome

Packet loss, latency.

IPv6 Prefix Length Performance

- RFC 2460: Internet Protocol Version 6 (IPv6) Specification
- RFC 2374: IPv6 Aggregatable Global Unicast Address Format
- RFC 2373: IP Version 6 Addressing Architecture

Objective

This data- and control-plane test measures the IPv6 data forwarding performance of a device under test (DUT) that has an IPv6 route table of varying prefix lengths.

IPv6 routes with varying prefix lengths will be sent into the DUT using an IPv6 routing protocol from a test device port. A second test device port will be used to send data packets through the DUT. These data packets will utilize each of the different prefix lengths. Performance measurements per prefix length are then taken.

Overview

Routing protocols have been modified to carry the longer 128-bit IPv6 address. Routers now also need to look further into IP address to determine which network each packet is destined for. With IPv6, network prefix lengths can range from /4 to /64. Data forwarding tests, using a routing table with varying IPv6 network prefix lengths, are required.

Router manufacturers can use BGP-4+, OSPFv3, IS-ISv6 and RIPng to advertise the IPv6 routes. Users should be aware of possible performance limitations, especially on the longer prefix lengths. Heavier processing burdens are placed on the router to determine longest prefix match, and more memory is required to store the longer addresses. Not all devices have implemented hardware assist for IPv6 packet forwarding. Many of these functions could be provided by software only. For some time to come, it would be prudent to characterize the performance of devices before deployment.

Measurements of the IPv6 forwarding rate, latency and packet loss per prefix length will be taken to see if they are as consistent in conformance, functionality and performance as an IPv4 addressing scheme.

Setup

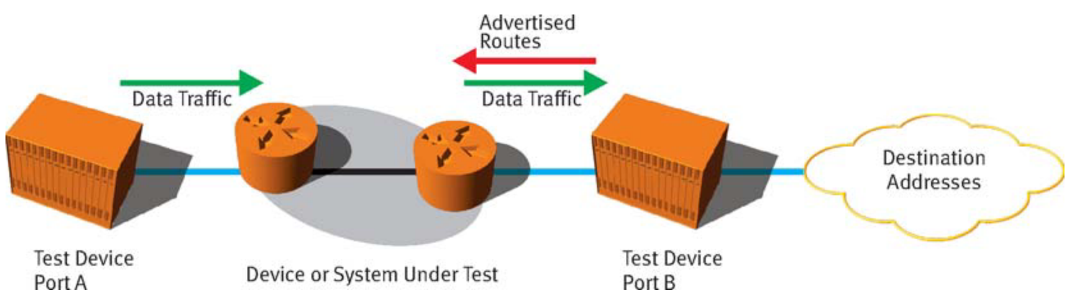

Test Steps

1. Determine which IPv6 routing protocol you will be using for this test: BGP-4+, OSPFv3, IS-ISv6 or RIPng.
2. Configure the DUT for the protocol chosen in Step 1, and establish a peering session from Test Port B to the DUT.
3. Advertise a fixed number of routes per prefix length from Test Port B to the DUT. The routes (destination prefixes) should have 61 unique prefix lengths from /4 to /64.
4. Configure the traffic parameters on Test Port A. Make sure all prefix lengths will be utilized by the traffic.
   a. Initial packet rate.
   b. Packet size.
   c. Run time for every trial.
5. Send packets from Test Port A to the destination prefixes that were previously advertised by Test Port B. Start with the maximum packet rate (wire rate) supported by the DUT test port.
6. If no packet loss occurs in two successive trials, stop the test. The test is complete.
7. If packet loss occurs, observe the following for each prefix length:
   a. Packet loss percentage.
   b. Forwarding rate.
   c. Number of transmit and receive packets.
   d. Average latency.
8. Decrement the data packet rate from Test Port A and repeat the test.
9. Stop the test when no packet loss is recorded in two successive trials, or when the input data rate reaches a user-selected threshold.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Number of IPv6 routes per prefix length.
- Number of prefix lengths used.
- Test data packet length.
- Test data packet rate (incrementing and/or decrementing).
- Test run time.

Test Outcome

For each prefix length, the statistic of total packets transmitted and received at each port, forwarding rate, average latency and packet loss.

Router Performance Impact of IPv6 FIB Size

RFC 2460: Internet Protocol Version 6 (IPv6) Specification

Objective

This data- and control-plane test measures the effect of data packet performance when the device under test (DUT) has different amounts of IPv6 routes in its Forwarding Information Base (FIB) table.

A fixed number of IPv6 routes will be sent into the DUT using an IPv6 routing protocol from a test device port. A second test device port will be used to send data packets through the DUT that utilize routes in the FIB table. Performance measurements of throughput, latency and packet loss of the data packets are taken. The test is then repeated using more IPv6 routes in the router FIB table.

Overview

Today's routers need more memory and processing power to both store the IPv6 network prefixes and parse the IPv6 data packets to determine longest prefix match, and then forward them.

Makers and testers of IPv6 routers that use protocols BGP-4+, OSPFv3, IS-ISv6 and RIPng should be aware of possible performance limitations, as many of these IPv6-related functions could be provided by software only. Tests need to ensure IPv6 routing protocols and IPv6 data forwarding engines are consistent in conformance, functionality and performance as their IPv4 counterparts.

Setup

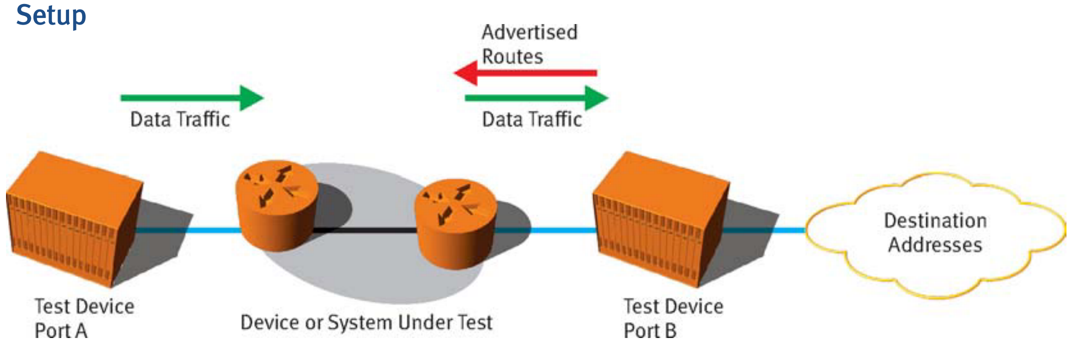

Test Steps

1. Determine which IPv6 routing protocol you will use for this test. Options are BGP-4+, OSPFv3, IS-ISv6 or RIPng.
2. Configure the DUT for the protocol chosen in Step 1 and establish a peering session from Test Port B to the DUT.
3. Advertise a fixed number of IPv6 routes from Test Port B to the DUT. The routes (destination prefixes) should use the same prefix length but should not be contiguous (sequential).

   Note: Disable route aggregation on the DUT. This will prevent the router from aggregating routes that would create a smaller than intended FIB size.

4. Configure the traffic parameters on Test Port A. Make sure all prefixes used in Step 3 will be utilized by the traffic.
   a. Initial packet rate.
   b. Packet size.
   c. Run time for every trial.
5. Send packets from Test Port A to the routes previously advertised by Test Port B. Start with the maximum packet rate (wire rate) supported by the DUT test port.
6. If no packet loss occurs in two successive trials, stop the test. The test is complete.
7. If packet loss occurs, observe the following:
   a. Packet loss percentage.
   b. Forwarding rate.
   c. Number of transmit and receive packets.
   d. Average latency.
8. Decrement the data packet rate from Test Port A and repeat the test.
9. Stop the test when no packet loss is recorded in two successive trials; or, when the input data rate reaches a user-selected threshold.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Number of IPv6 routes.
- Different IPv6 prefix lengths (/4 to /64).
- Test data packet length.
- Test data packet rate (incrementing and/or decrementing).
- Test run time.

Test Outcome

Total packets transmitted and received, forwarding rate, average latency and packet loss.

Glossary

IPv6/IPv4 Definitions

6Bone – The 6Bone is a test network for IPv6, based on a worldwide backbone connecting sites in more than 42 countries that run IPv6 implementations from some 30+ vendors and institutes. The first 16 bits of a 6Bone address are set to 3FFE.

6over4 – 6over4 provides a solution for environments where a number of IPv6 hosts are distributed in an IPv4 domain, and none of them have a direct IPv6 connectivity. The hosts perform the tunneling. By providing a router with a native IPv6 connection, which also understands 6over4, the 6over4 hosts can also connect to native IPv6 hosts. IPv6 packets are encapsulated in IPv4 packets sent over the IPv4 network. IPv6 packets are transmitted in IPv4 packets with an IPv4 protocol type of 41.

6to4 – The 6to4 transition mechanism provides a way to connect IPv6 end-site networks by automatically tunneling over the intervening IPv4 Internet. A special IPv6 routing prefix (2002::/16) is used to indicate that the remaining 32 bits of the external routing prefix contain the IPv4 end-point address of a boundary IPv6 router for that site that will respond to IPv6 in IPv4 encapsulation. Thus, another 6to4 end-site can automatically discover this tunnel end point by using a DNS lookup for any host on that network.

Dual-Stack – A term used to describe a network node running both IPv4 and IPv6 protocol stacks (or possibly others) at the same time. Such a machine can act as a protocol converter between the two networks.

IPv6 – Next generation Internet protocol. The current version is IPv4.

IPv6/IPv4 (Dual Stack) Note – A host or router that implements both IPv4 and IPv6 as well as other transition mechanisms such as tunneling.

IPv4-Compatible IPv6 Address – An address given to an IPv6 node that can be used in both IPv6 and IPv4 packets. An IPv4-compatible IPv6 address holds an IPv4 address in the low-order 32-bits. The high-order 96 bits bear the prefix 0:0:0:0:0:ffff. The entire 128-bit address can be used when sending IPv6 packets. The low-order 32-bits can be used when sending IPv4 packets.

IPv4-Mapped IPv6 Address – The address of an IPv4-only node represented as an IPv6 address. The IPv4 address is stored in the low-order 32-bits of an IPv4-mapped IPv6 address. The high-order 96 bits bear the prefix 0:0:0:0:0:0. The address of any IPv4-only node may be mapped into the IPv6 address space by adding the prefix 0:0:0:0:0:0 to its IPv4 address. IPv4-mapped IPv6 addresses always identify IPv4-only nodes; they never identify IPv6/IPv4 or IPv6-only nodes.

IPv6-Only Address– An IPv6 address that does not necessarily hold an IPv4-address embedded in the low-order 32-bits. IPv6-only addresses bear prefixes other than 0:0:0:0:0:0 and 0:0:0:0:0:ffff. IPv6-only addresses always identify IPv6/IPv4 or IPv6-only nodes; they never identify IPv4-only nodes.

Jitter – The variation in time between packets when they arrive at the destination.

Latency – The time it takes a data packet to travel through a network or the DUT.

Longest Prefix Match – Determines the best next-hop route for a packet depending solely on the destination address in the packet header.

NAT – Network Address Translation provides allows a single IP address to be used by multiple PCs. This is accomplished by assigning unique IP addresses to the PCs on the private network and then replacing the source IP address before going to the public network. NAT is used for IPv4 networks.

NAT-PT – Network Address Translation-Protocol Translation uses the same principle as NAT but adds the ability to bind protocol between IPv6 and IPv4.

Packet Loss – Packet loss means the dropped data packets by the DUT. Packet loss measurements can be obtained along with the throughput measurement.

Throughput – Data received per second. It also can be described as the maximum data packet-forwarding rate for which the device will not drop any data packets. For voice, the time a signal (voice sound) requires to travel through a device, including processing time.

Spirent Communications
Test Methodologies Journals

The Spirent Communications Test Methodologies series of journals has been developed to make network and equipment testing easier. Spirent provides a Web site for scripts and other resources at ‹http://scdn.spirentcom.com/welcome.asp›. Registration is required to access scripts and other tools to help you expedite testing.

Spirent offers test methodology documents and support data on the following topics:

Router Performance Edition — Router Performance Testing
Volume II — IPv6, Wireless, Edge Router, Metro Optical, VoIP, SS7, QoS
Layer 4-7 Edition — Network and Application Testing
IPv6 Edition — IPv6 and IPv4 Test Methodologies
Optical Edition — Near End Performance Testing
Edge Router Edition — Edge Router Testing
MPLS Edition — MPLS Network Testing
IPSec Edition — IP Security Testing

Analyze | Assure | Accelerate™

Spirent Communications is a worldwide provider of integrated performance analysis and service assurance systems for next-generation network technologies. Our solutions accelerate the profitable development and deployment of network equipment and services by emulating real-world conditions in the lab and assuring end-to-end performance of large-scale networks.

Spirent performance analysis solutions include instruments and systems that measure and analyze the performance of network equipment, particularly the devices that route voice and data messages to their destination. Our service assurance solutions include remote test, fault and performance management systems that let network service providers quickly identify network faults and monitor real-time performance.

Spirent's integrated performance analysis and service assurance solutions enable our customers to more rapidly develop and certify new devices, lowering the cost of widespread deployment and operation of new network services.

Spirent Communications is a wholly owned subsidiary of Spirent plc, an international network technology company.

Spirent Communications
15200 Omega Drive
Rockville, MD USA
20850-3240
Tel: +1 301.417.1224
info@spirentcom.com

Sales and Information

Americas
Tel: +1 800.927.2660
productinfo@spirentcom.com

Europe, Middle East, Africa
Tel: +33 1 6137.2250
salesemea@spirentcom.com

Asia Pacific
Tel: +852.2166.8382
spirentasia@spirentcom.com

Spirent Communications Developers Network
http://scdn.spirentcom.com

www.spirentcom.com

P/N 340-5022-001 Rev B, 12/03