RFC 2544 Testing with Aurora Tango

RFC 2544 describes a set of tests for measuring the performance of Ethernet protocol equipment.
This Application Note describes how Aurora Tango can be used to perform the RFC 2544 tests to measure network performance.

Application Note AuroraTango 2544

Testing the World's Digital Networks

TrendCommunications

INTRODUCTION

RFC 2544 (Request For Comments Document) is a document that is designed to provide benchmarking tests for network devices. The aim of the authors of this document was to reduce the 'smoke and mirrors' and 'specmanship' that are sometimes employed by equipment vendors. The tests they describe aim to provide a measure of how a device would perform in the real world.

RFC 2544 describes OOS (Out Of Service) tests. If a customer is already having problems with a network, flooding it with test traffic will only add to the problem; in this case it would be more appropriate to monitor the pattern of real network traffic.

The RFC defines a test as being made up of multiple trials. Each trial provides a piece of data, for example the loss rate at a particular input frame rate. There may therefore be many trials in a test - with the device set up in different ways.

TEST CONDITIONS

Tester Used

A tester with both transmitting and receiving ports is recommended. The tester must include sequence numbers in the frames it transmits, so that it can check that all frames are received back.

Traffic Used in the Test

Traffic Pattern - the traffic on a real network is not constant, it occurs in bursts. The RFC suggests that the tests should be carried out using constant traffic and with Test Conditions traffic - repeated bursts of frames, the frames within the bursts separated by the minimum inter-frame gap.

Protocol Addresses - the simplest way to perform these tests is to use a single stream of data. Networks in the real world do not have just one stream of data. The RFC suggests that after the tests have been run in this way they should be re-run using a random destination address. For routers the RFC suggests that the addresses used should be random, and evenly distributed over a range of 256 networks. For bridges it recommends that the range should be uniformly distributed over the full MAC range.

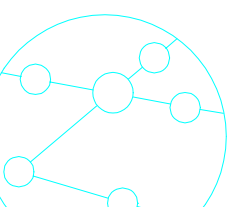

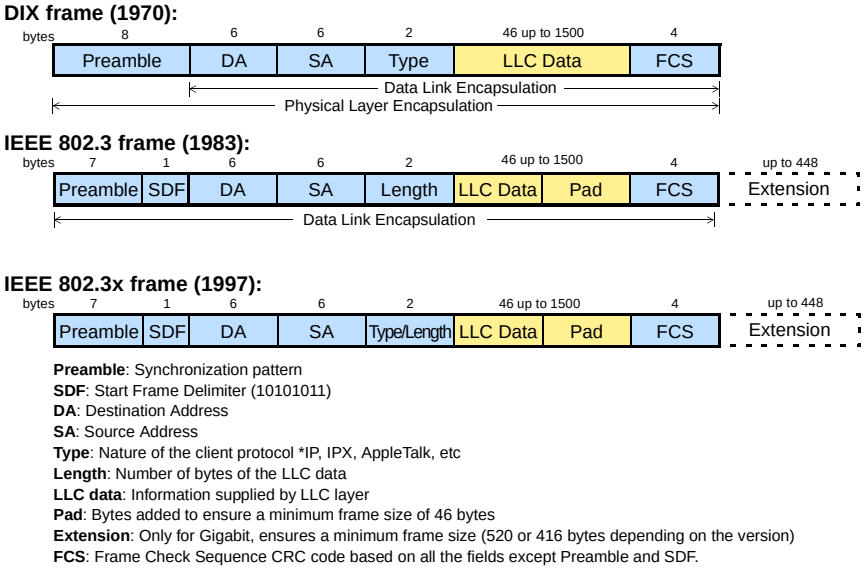

Preamble: Synchronization pattern
SDF: Start Frame Delimiter (10101011)
DA: Destination Address
SA: Source Address
Type: Nature of the client protocol *IP, IPX, AppleTalk, etc
Length: Number of bytes of the LLC data
LLC data: Information supplied by LLC layer
Pad: Bytes added to ensure a minimum frame size of 46 bytes
Extension: Only for Gigabit, ensures a minimum frame size (520 or 416 bytes depending on the version)
FCS: Frame Check Sequence CRC code based on all the fields except Preamble and SDF.

Figure 1 802.3 MAC frame formats. Major changes in IEEE 802.3 frame format are the SDF field which is basically a name change, the Length field enables internal management of the padding field, so that higher protocols no longer need to provide their own padding mechanism. In 1997 Full Duplex support was implemented and it was necessary to use the DIX Type field to differentiate the MAC control protocol.

Maximum Frame Rate - when testing on a LAN the maximum frame rate for the medium and frame size being used should be used for the test. When testing on a WAN a rate greater than the maximum theoretical rate for the medium and frame size should be used.

Frame Sizes - the RFC recommends that the tests are carried out at a range of frame sizes - 64, 128, 256, 512, 1024, 1280, 1518 bit. This covers the range of frame sizes that are typically transmitted.

Frame formats - the format of the frames of TCP/IP over Ethernet are specified in appendix C of the RFC.

Test duration

These tests are designed to measure how a device will perform under continuous operation. The test time must be a compromise between this and the time available to complete a test suite. The RFC recommends that the duration of each trial should be at least 60 seconds per test.

RFC 2544 was designed for laboratory testing of equipment, the tests as described may take several days to complete. This duration is unlikely to be possible or necessary when testing a network in the field. The time taken for the test can be reduced by selecting the tests to be run and reducing the number of repetitions.

Test setup

The aim of this set of tests is to evaluate the performance of equipment under real-world conditions. The RFC states that all the protocols supported by the device must be enabled during a test. The equipment must be set up using the instructions supplied to the user. The only changes allowed between tests are those necessary to perform the different tests. It is not acceptable, for example, to change the size of the frame-handling buffer between tests of frame handling rates.

Reports

The RFC recommends that, in addition to the test results the following should be included in test reports:

- DUT setup - which functions disabled, which used
- DUT software version
- Frame formats
- Filter setups

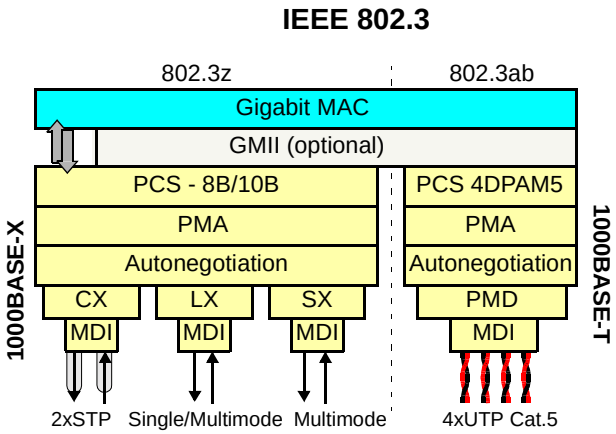

Figure 2 Gigabit Ethernet defines several transmission media: 802.3z (1000BASE-X) based on the existing Fibre Channel technolog and 802.3ab (1000BASET) which uses UTP.

RUNNING THE TRIAL

The RFC defines a test as being made up of multiple trials. Each trial gives a piece of data, for example the loss rate at a particular input frame rate. The procedure below describes the steps for a single trial:

1. If the device you are testing is a router, sending the routing update to the input port, then wait two seconds.
2. Send the learning frames to the output port.

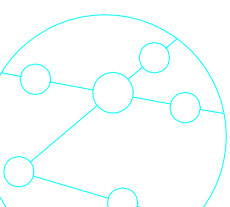

3. Run the trial.
4. Wait for two seconds to receive all the data sent.
5. Wait at least five seconds before starting the next trial.

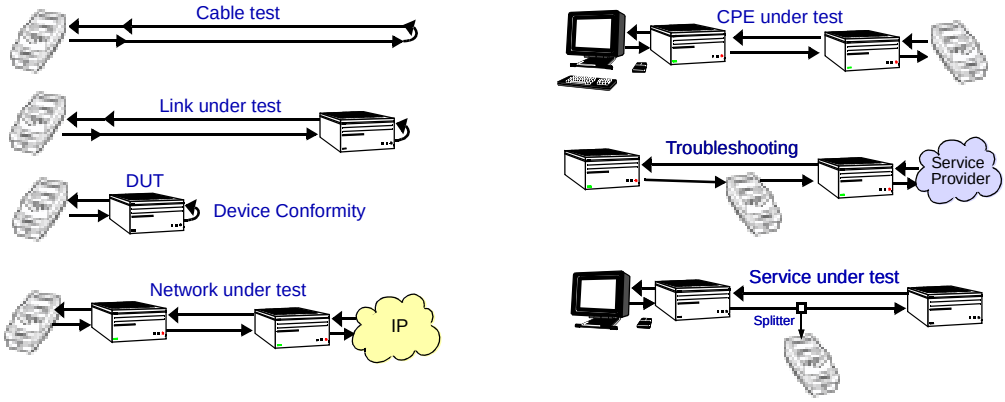

Figure 3 Gigabit Ethernet testing.

THE TESTS

RFC 2544 discusses and defines a number of tests that may be used to describe the performance characteristics of a network interconnecting device. In addition to defining the tests this document also describes specific formats for reporting the results of the tests.

Throughput

Objective: To determine the DUT throughput as defined in RFC 1242.

Procedure:

1. Send a specific number of frames at a specific rate through the DUT and then count the frames that are transmitted by the DUT.
2. If the count of transmitted frames is equal to the count of received frames, increase the throughput and rerun the test.
3. Rerun the test until fewer frames are transmitted by the DUT than were received by the DUT.
   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.

Latency

Objective: To determine the latency as defined in RFC 1242.

Procedure:

1. Measure the throughput for DUT at each of the listed frame sizes.
2. Send a stream of frames at a particular frame size through the DUT at the determined throughput rate to a specific destination.
   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.

Frame loss rate

Objective: To determine the frame loss rate, as defined in RFC 1242, of a DUT throughout the entire range of input data rates and frame sizes.

Procedure:

1. Send a specific number of frames at a specific rate through the DUT to be tested and count the frames that are transmitted by the DUT.
   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.
   The frame loss rate at each point is calculated using the following equation:

   *((input_count - output_count) * 100) / input_count*
2. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used and then for 80% of this rate.
3. 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.

Back-to-back frames

Objective: To characterize the ability of a DUT to process back-to-back frames as defined in RFC 1242.

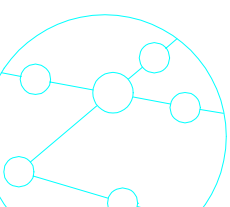

Procedure:

1. Send a burst of frames with minimum inter-frame gaps to the DUT and count the number of frames forwarded by the DUT.
2. If the count of transmitted frames is equal to the number of frames forwarded increase the length of the burst and rerun the test.

   OR

   If the number of forwarded frames is less than the number transmitted, reduce the length of the burst and rerun the test. 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.

System recovery

Objective: To characterize the speed at which a DUT recovers from an overload condition.

Procedure:

1. First measure the throughput for a DUT at each of the listed frame sizes.
2. Send a stream of frames at a rate 110% of the recorded throughput rate or the maximum rate for the media, whichever is lower, for at least 60 seconds.
3. At Timestamp A reduce the frame rate to 50% of the above rate and record the time of the last frame lost (Timestamp B). The system recovery time is calculated by subtracting Timestamp B from Timestamp A.
   The test must be repeated a number of times and the average of the recorded values being reported.

Reset

Objective: To characterize the speed at which a DUT recovers from a device or software reset.

Procedure:

1. First measure the throughput for the DUT for the minimum frame size on the media used in the testing.
2. Send a continuous stream of frames at the determined throughput rate for the minimum sized frames.

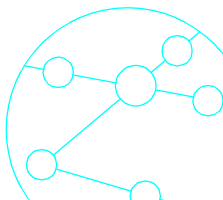

3. Cause a reset in the DUT.
4. Monitor the output until frames begin to be forwarded and record the time that the last frame (Timestamp A) of the initial stream and the first frame of the new stream (Timestamp B) are received. A power interruption reset test is performed as above except that the power to the DUT should be interrupted for 10 seconds in place of causing a reset. This test should only be run using frames addressed to networks directly connected to the DUT so that there is no requirement to delay until a routing update is received. The reset value is calculated by subtracting Timestamp A from Timestamp B. Hardware and software resets, as well as a power interruption should be tested.

PERFORMING A RFC 2544 TEST USING AURORA TANGO

Performing a RFC 2544 test with Aurora Tango is very easy. Aurora Tango enables you to choose all the settings for the different parts of the test. A full RFC 2544 test, performed at all the frame sizes can take a long time (up to 2.5 days). Aurora Tango enables you to select the tests that you want to perform, the frame sizes that are used for the tests and the tests settings. This enables you to reduce the time the test takes while performing the correct tests for the situation.

The Aurora Tango Gigabit Ethernet software is very simple to use. To perform a test, from the main screen, you just choose the test and then the green button.

The software enables you to set up all the conditions for the test:.

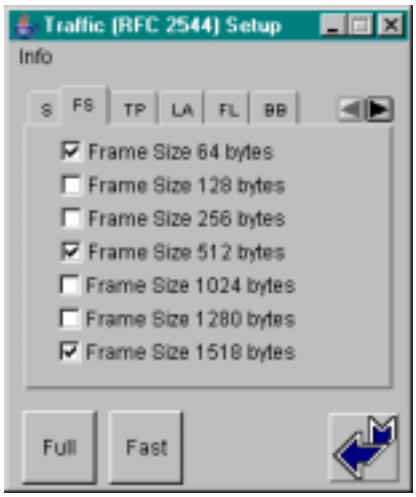

- The **Full** and **Fast** buttons, enable you to choose the type of test you want to run. A **Full** RFC 2544 test may take up to 2.5 days to complete. The **Fast** test enables you to perform a test in less time.
- The **S** (summary) tab displays the settings you have chosen and an estimate of the time the test will take to run. As you make changes to the test setup on the other tabs the estimated time is updated.
- The **FS** (frame size) tab enables you to choose the frame sizes used for all the tests on the following tabs. Obviously the more frame sizes you choose the longer the tests will take to complete.
- The **TP** (throughput) tab enables you to choose the **Duration** of the test and the **Resolution** of the test. A higher number in the **Resolution** field will make the test quicker, and a lower number will improve the accuracy of the results but make the test slower.You cannot run a latency or System recovery test without running a throughput test.
- The **LA** (latency) tab enables you to choose the number of times the test is run (**Repetitions**), the **Duration** of each repetition of the test, and the **Interval** between the tagged frames transmitted by Aurora Tango. When you choose to run a latency test the throughput test is also switched on.
- The **FL** (frame loss) tab enables you to choose the **Duration** of each frame loss test, and the **Granularity** of the test. The first test is performed at 100% of the maximum rate for each of the frame sizes you have selected on the **FS** tab. The rate is then reduced by the **Granularity** you have set, for example 10%. This is repeated until no frames are lost in two successive tests.
- The **BB** (back to back frames) tab enables you to choose the size of the **Initial Burst** of frames, the number of **Repetitions** of the test, and the

Resolution of the test. When the test is run, a minimum initial burst of 2 seconds of frames is sent.

- The **SR** (system recovery) tab enables you to choose the **Duration** of each system recovery test, and the number of **Repetitions** of the test. The traffic rate used for the test is based on the rate measured in the throughput test. When you choose to run a system recovery test the throughput test is also switched on.

  At the end of the test the results are clearly displayed. They can easily be printed or saved on the Pocket PC or PC for later review.

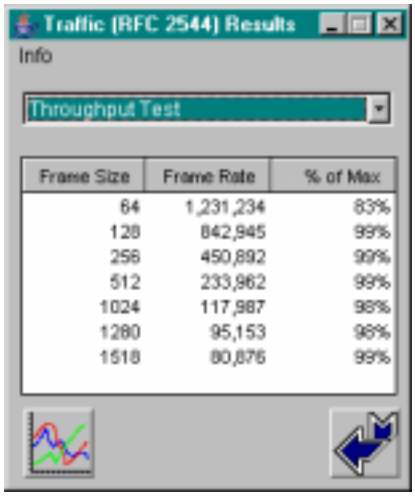

TrendCommunications

Trend Communications Ltd
Knaves Beech Estate
Loudwater, High Wycombe
Bucks HP10 9QZ UK
www.trendcomms.com
infoline@trendcomms.com

Trend Communications SL
Pujades, 60
08005 Barcelona
(Spain)

International: .... +44 1628 524977
España:..... 93 300 3313
Deutschland: 089 32 30 09 11
US: 256 461 0790

UK:..... 01628 524977
France: 01 69 35 54 70
India:..... 22 8597 463/4