

Ethernet RFC-2544 explained

The RFC 2544 describes benchmarking tests for network devices. Vendors can use these tests to measure and outline the performance characteristics of their equipment. As these tests follow standard procedures, they also make it easier for customers to make sense of the glitzy marketing-speak employed by most vendors.

The tests described in the document aim to evaluate how a device would act in a real situation. The RFC describes six out-of-service tests, which means that real traffic must be stopped and the tester will generate specific frames to evaluate throughput, latency, frame loss rate, burst tolerance, overload conditions recovery and reset recovery. The document also describes specific formats for reporting the results of these tests.

1. TEST CONDITIONS

All Ethernet tests should be run consistently without changing the configuration of the device, and without running a specific protocol or feature. The DUT should include the normal routing update intervals and keep frequency alive.



Figure 1 ALBEDO Ether.Giga is an Ethernet & IP tester equipped with all the features of legacy testers such as BER and RFC2544, plus the new ones like Y.1564, Y.1731, and FCS error insertion in pass mode.

Which Tester to Use?

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

The RFC 2544 can be used to test layer-2 and layer-3 devices. For layer-3 testing, IP packets need to be configured, including parameters such as mask and subnetworks that can be understood by routers. MAC frames must always be programmed, including parameters such as frame size, bit rate, or traffic shape.

Traffic Used in the Test

- **Traffic pattern** – The traffic on a real network is not constant, but occurs in bursts. The RFC suggests that the tests should be carried out using constant traffic and with test conditions traffic, i.e., 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 the range should be uniformly distributed over the full MAC range.
- **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 bits. 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.

RFC 2544 was designed for laboratory testing of equipment, which is why 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 by reducing the number of repetitions.

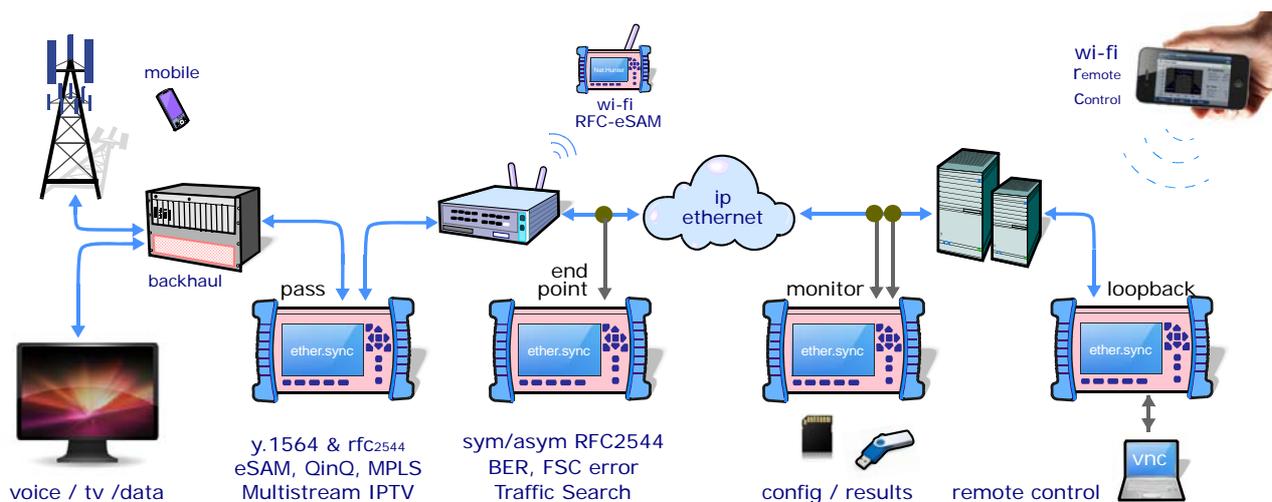


Figure 2 *Ether.Giga is test solution that seamless adapts operating needs without sacrificing portability, speed or cost. Perform the RFC 2544 test option, testing throughput, frame loss, latency, jitter and burst is straightforward. Ether.Giga can execute it both in symmetric and asymmetric way and with the far-end device in loopback mode or peer-to-peer mode.*

Test Setup

The aim of this set of tests is to evaluate the performance of equipment in real-world situations. The RFC states that all the protocols supported by the device must be enabled when testing, and the equipment must be set up according to the instructions supplied to the user. The only changes allowed between tests are those needed 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.

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

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

2. 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 following procedure describes the steps for a single trial:

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

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

3. THE TESTS

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

Throughput

The aim of a *throughput test* is to determine the maximum number of frames per second that the device can process and forward without dropping or losing any. The procedure is as follows:

1. Send a certain number of frames at a specific rate through the DUT and count the frames transmitted by the DUT.
2. If the count of transmitted frames is equal to the count of received frames, increase the throughput and re-run the test.
3. Re-run the test until fewer frames are transmitted than 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.

If a single value (minimum frame size) is desired, it must be expressed in frames per second, or alternatively in bits or bytes per second.

The statement of performance must include the following information:

- the measured maximum frame rate
- the size of the frame used
- the theoretical limit of the media for that frame size
- the type of protocol used in the test.

Latency

This test determines the *latency* inherent in the DUT. The initial data rate is based on the results of the throughput test. Typically, a packet is time stamped and inserted in the middle of a burst, and the time the stamped packet takes to travel through the DUT is measured.

You must first measure the throughput for the DUT at each of the defined frame sizes, and send a stream of frames through the DUT at the determined throughput rate to a specific destination. The duration of the stream should be at least 120 seconds. After 60 seconds, an identifying tag should be included in one frame.

The time at which this frame is completely transmitted is recorded, and this will be timestamp A. The receiver of the test equipment must recognize the tag information in the frame stream and record the reception time of the tagged frame. This will be timestamp B.

The latency is the difference between timestamp B and timestamp A, according to the definition found in RFC 1242.

Figure 3 consists of two screenshots of test results. The left screenshot, titled 'Throughput test', shows a 'FAIL' status. The right screenshot, titled 'Latency test', shows a 'PASS' status.

Size	Theor.max (fr/s)	Max.rate (fr/s)	Max.rate (%)
128	844,594	844,594	100.00
256	452,898	452,898	100.00
512	234,962	234,962	100.00
1024	119,731	119,731	100.00
1280	96,153	96,153	100.00
1518	81,274	81,274	100.00
2000	61,881	0	0.00

Size	Throughput (%)	Delay (µs)
64	3.28	50.42
128	5.78	50.89
256	10.00	51.85
512	10.00	54.23
1024	10.00	58.04
1280	10.00	60.31
1518	10.00	62.26

Figure 3 RFC 2544 test results: (a) Throughput results table, (b) Latency results table.

Frame Loss Ratio

The aim of this test is to determine the *frame loss ratio* throughout the entire range of input data rates and frame sizes. The procedure is the following:

1. Send a certain number of frames at a specific rate through the DUT, counting the frames transmitted.
The first trial should be run for the frame rate that is 100% of the maximum rate for the frame size on the input media.
The frame loss rate at each point is calculated as follows:
$$((input_count - output_count) * 100) / input_count$$
2. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used, and then for 80% of this rate.
3. Continue this sequence (at reducing 10% intervals) until there are two consecutive trials where no frames are lost. The maximum granularity of the trials must be 10% of the maximum rate, although you may want to define a finer granularity.

In this test the test frame is addressed to the same destination as the rest of the data stream, and each test frame addressed to a new destination network.

The figure shows three screenshots of test results from a network testing tool. Each screenshot has a status bar at the top indicating 'Stopped' and a time. Below the status bar is a table of results.

(a) Frame loss test (FAIL)

Throughput (%)	Frame loss (%)
100.00	32.797
90.00	25.330
80.00	15.996
70.00	3.996
60.00	0.000
50.00	0.000

(b) Back-to-back test (FAIL)

Size	Burst length (fr)
64	206
128	214
256	222
512	253
1024	344
1280	417
1518	500

(c) System recovery test (PASS)

Size	Throughput (%)	Test rate (%)	Recovery time (μs)
64	3.28	3.60	30.99
128	5.78	6.35	14.53
256	10.00	11.00	0.00
512	10.00	11.00	0.00
1024	10.00	11.00	0.00
1280	10.00	11.00	0.00
1518	10.00	11.00	0.00

Figure 4 RFC 2544 test results: (a) Frame loss result table, (b) Back-to-back frames result table, (c) System recovery time results table

The test report must state which definition of latency (from RFC 1242) was used for this test. The latency results should be reported as a table, with a row for each tested frame size. 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 latency values for each type of data stream tested.

Back-to-Back Frames

A back-to-back frame test determines the *node buffer capacity* by sending bursts of traffic at the highest theoretical rate, and then measuring the longest burst where no packets are dropped. This is done to check the speed at which a DUT recovers from an overload condition, and the procedure is as follows:

1. Send a burst of frames with minimum inter-frame gaps to the DUT, and count the number of frames forwarded.
2. If the count of transmitted frames is equal to the number of frames forwarded, increase the length of the burst and re-run the test. If the number of forwarded frames is less than the number transmitted, reduce the length of the burst and re-run the test.

The back-to-back value is the number of frames in the longest burst that the DUT can handle without losing any frames. The trial length must be at least 2 seconds, and it should be repeated at least 50 times with the average of the recorded values being reported.

The back-to-back results should be reported as 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.

System Recovery

This test determines the node speed at which a DUT recovers from an overload condition. The procedure is as follows:

1. Measure the throughput for a DUT at each of the listed frame sizes.
2. Send a stream of frames at a rate that is 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 is reported.

The system recovery results should be reported as a table, with a row for each of the tested frame sizes. There should be columns for the frame size, the frame rate used as the throughput rate for each type of data stream tested, and for the measured recovery time for each type of data stream tested.

Reset

This test is intended to characterize the speed at which a DUT recovers from a device or software reset. The procedure is as follows:

1. 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. Reset 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 described above, except that instead of resetting, the power to the DUT should be interrupted for 10 seconds.

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.

4. TRAFFIC GENERATION

A major factor when determining hardware and software needs for the Giga-bit Ethernet DUT test approach is the *internal data handling capability* of the DUT. When testing the product, it is essential to pass data through the product. Some products can be put into a mode where an unframed PRBS stream can be passed without any issues. Other products must receive and transmit actual IP data, which may be looped back internally or externally.

Still other DUTs that require actual IP data also contain smart routing algorithms built into the on-board firmware. This adds additional constraints, since if the IP data is looped back to the same port, the router knows this is not correct, and the data will be destroyed. In some cases a DUT may even have the built-in ability to internally generate its own data stream and to verify that data when it returns to the Rx port.

Because of these varying internal operating characteristics, it is imperative to know the internal hardware and software design of the DUT. In the first case above, a simple 1-giga PRBS BER test instrument can be used. The second case calls for a tester that supports IP traffic. In the third case, not only will IP traffic support be needed, but the test set-up will also need either a golden module or chassis that can accept the test traffic and route it legally without losing any data. This requirement makes a stand-alone fixture approach impractical. In the final case where no external traffic test equipment is required, traffic testing can be accomplished using only loop-back cabling. To support traffic tests, there must be some means of loop-back to provide a complete path through the product under test.

These loop-backs are most effective when they are achieved externally to the product under test, either via an external jumper or other modules capable of handling the test traffic. This approach ensures that the module's backplane connections are tested. To provide a more granular level of fault isolation, some products are designed with internal loop-back capabilities. If a test fails, these localized loop-backs can be employed to isolate the fault location. If this capability is desired, the test engineer will need to contact the product design group during the design phase, as this capability must be designed into the product.

There is not much difference between testing done by using traffic streams of mixed protocols and testing with individual protocols. That is, if protocol A testing and protocol B testing give two different performance results, mixed protocol testing appears to give a result which is the average of the two [4].

The maximum frame rate to be used for LAN-WAN-LAN configurations is a judgment that can be based on known characteristics of the overall system, including compression effects, fragmentation, and gross link speeds. Practice suggests that the rate should be at least 110% of the slowest link speed.



ALBEDO Telecom

ALBEDO Telecom designs, manufactures, and delivers solutions that enable Telecom organizations of all sizes to test, measure, troubleshoot, monitor, and migrate mission critical networks and multiplay applications.

On local segments and across distributed networks, ALBEDO enable Organizations, Installers, Operators, Service Providers and Suppliers to quickly check the health of Network Architectures, Service Agreements (SLA), IP Quality (QoS), or fix any issue.

Your Business Partner

Results. ALBEDO Telecom helps the industry to make the most of the investment on infrastructure.

Expertise. ALBEDO Telecom engineers and consultants provide industry leading knowledge in hand-held TAPs and WAN emulators, IPTV, VoIP, Carrier-Ethernet, Synchronization, Jitter, Wander, SyncE, PTP, E1, and Datacom to address customers unique needs.

Integration. ALBEDO Telecom integrates disparate telecom technologies and applications, facilitating new business efficiencies.

Agility. ALBEDO Telecom increases the ability of customers to respond quickly to new market opportunities and requirements.

Coverage. ALBEDO Telecom offers solutions that facilitates the migration and the roll-out to new architectures.

ALBEDO

Telecom

the Path to Excellence

Joan d'Austria, 112 - Barcelona - 08018 - SP

Chalfont St Peter - Bucks - SL9 9TR - UK

www.albedotelecom.com

aims
eius

- + UNDERSTAND causes of telecom interoperability issues
- + EXPERIENCE the best quality in unified networking
- + ASSESS different hardware, firmware, and software solutions
- + LEARN from experts by means of professional services and consultancy