How-To Compare Performance of Data Plane Devices

Matthias Holdorf

Advisors: Daniel Raumer and Florian Wohlfart
Innovative Internet Technologies and Mobile Communications WS2015/2016
Chair for Network Architectures and Services
Department of Informatics, Technische Universität München
Email: matthias.holdorf@tum.de

ABSTRACT

Over the past years network technologies have grown and evolved fast. The requirements on network devices have become more complex and diverse. In order to measure and compare the performance of data plane devices, benchmarks are used. This paper describes a taxonomy of benchmark methodologies as well as the requirements and test setups to assess them. It is surveyed which of these benchmarks are used by test equipment vendors. Based on this survey, the methodologies are discussed in usefulness and feasibility of their measurements.

Keywords

Performance Testing, Benchmarks, Data Plane Devices, RFC-2544, RFC-2889, ITU-T Y.1564

1. INTRODUCTION

The convergence of several traffic types and the increasing data rates lead to higher performance requirements on networks. With new technologies emerging, such as Cloud Computing and Software-Defined Networking (SDN), today's networking architecture is transforming. [2, 3] The performance of data plane devices influences the network size, stability and reliability. The definition and tasks of a data plane device are described in chapter 2. In order to compare these devices, benchmarks are used. The term benchmark refers to a test, used to evaluate the performance of a network device under predefined starting conditions, relative to the performance of another device under the same conditions. The goal of a benchmark is to enable a meaningful comparison between two devices. [1] There are different benchmarking recommendations for data plane devices, as discussed in chapter 3.

Benchmarking was developed to measure the performance of a Device under Test (DUT). But there are additional aspects that need to be considered when defining a benchmark. One of them is comparability. In order for a benchmark to deliver a fair comparison between two DUTs, independently executed test runs should yield comparable results. This is of course, given the circumstances that both setup environments are identical. The comparability is also affected by the repeatability of a benchmark on an identical DUT in different moments in time. In order to achieve a fair comparison, the test setup and the experimentation methodologies need to be well defined. [2] The definition of the requirements, test setups and the execution of benchmark tests are depicted in chapter 4.

In chapter 5 the benchmarking methodology as described in the RFC-2544 is outlined as well as the reporting format. The usage of benchmark methodologies by test equipment vendors are investigated in chapter 6 and 7. Based on this survey, the usefulness and feasibility of the described methodologies are discussed in chapter 8.

2. NETWORK TRAFFIC PLANES

Every network device can be partitioned into three basic elements with distinct activities: (1) the data plane, (2) the control plane and (3) the management plane. Each logically separated plane classifies a different type of traffic in the network. Every plane has its own distinctive characteristic and security requirements. These three planes are the components of the layered architecture that networks have evolved to today. [4, 5, 6, 9] In traditional networking, the three planes are implemented in the firmware of routers and switches. [7, 8] In the following chapters the three network traffic planes are described in more detail. The task of each device is explained, their distinction in responsibilities as well as their interdependencies.

2.1 Management Plane

The management plane handles the administrative interface into the overall system. It is also associated with monitoring, configuration and maintenance of a system. The management plane is often considered as a subset of the control plane. [9, 10]

2.2 Control Plane

The control plane is responsible for routing decisions. It is the *Signalling* of the network. Therefore, it comprises the protocols by which routers learn the forwarding topologies and the state of the network. Implementing these complex protocols in the data plane would lead to poor forwarding performance. Thus, it maintains the information necessary for the data plane to operate. [9, 10, 20, 21] This information is collected by routing protocols like Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP) or Border Gateway Protocol (BGP). [10]

The control plane informs the data plane about the collected information. This updates the Routing Information Base (RIB) or a separate Forwarding Information Base (FIB) of the data plane. End users rarely interact with the control plane. One exception is the ICMP ping, where a control plane protocol can be directly employed. [9]

2.3 Data Plane

The data plane is also known as forwarding plane. It is defined as the part of the network that carries the traffic. It enables data transfer to and from clients, handling multiple conversations through multiple protocols. Data Plane traffic should not have destination IP addresses that refer to any networking device. It should rather be sourced from and destined to a certain devices e.g. a server or client. The main task of a router in the case of the data plane is to merely forward a packet. [9]

Under normal circumstances transit packets constitute a large amount of traffic that enters the data plane. This is the reason why routers use specialized forwarding hardware, such as Application-Specific Integrated Circuits (ASIC), to accomplish this forwarding as fast as possible. [9, 10] However, there are exceptions that need to be taken into account. Not every transit packet belongs to the data plane and not only transit traffic is forwarded by the data plane. In the case of such an exception, additional router processing resources are consumed to forward a packet. The data plane should be focused on forwarding packets but is yet commonly burdened by other activities: NAT session creation and NAT table maintenance, NetFlow Accounting, Access Control List (ACL) logging and error signalling (ICMP). [5, 9]

In order to define data plane devices, the definition that every networking device consists of the three layered plane architecture (including the data plane) is used. Thus, in order to define a data plane device, we first need to define a network device. [4, 5, 6, 9]

After the examination of additional sources [11, 12, 13, 14, 15, 16], the following devices can be classified as data plane devices: Network Interface Card, Repeater, Hub, Bridge, Switch, Router, Firewall and Gateway. Since the survey of all devices is beyond the scope of this paper, only the performance benchmarking of routers and switches are examined.

3. BENCHMARK STANDARDS

The Benchmarking Methodology Working Group (BMWG) is a working group of the Internet Engineering Task Force (IETF). It proposes recommendations concerning the performance of networking devices and services. The BMWG defines benchmark methodologies for the management, control and data plane. [17, 18, 19]

The Request for Comments (RFC) 2544: *Benchmarking methodology for network interconnect devices* [23], which is proposed by the BMWG, is widely accepted in the industry. The proposal became an international standard for testing the performance of networking devices. [49, 50] It provides the benchmarking tests for the performance indices defined in the RFC-1242 [22], as well as the test setup conditions to apply and report format to document the tests. [1] Since the RFC-2544 does not address some of the specificities of IPv6, a new recommendation RFC-5180: *IPv6 Benchmarking Methodology for Network Interconnect Devices* [45] was proposed by the BMWG.

The BMWG further proposed the RFC-2889: *Benchmarking Methodology for LAN Switching Devices* [24], which extends the general methodology for benchmarking tests de-

fined in the RFC-2544 to specific properties of LAN switching devices. This proposal primarily focuses on devices which can be assigned to the OSI-Layer 2.

There are additional recommendations from other organizations which can be applied to performance test data plane devices. [41] One of them is the ITU-T Y.1564: *Ethernet Service Activation Test Methodology* [42], which is also called *EtherSAM*. [43] This recommendation was developed to address drawbacks of the RFC-2544 in terms of testing today's Ethernet services and validating service-level agreements (SLA). [43] The MEF 14, proposed by the Metro Ethernet Forum (MEF), defines test procedures that may be specified as part of a Service Level Specification (SLS) for an Ethernet service. [44]

Moreover, network testing service providers define their own benchmark recommendations. The switch testing test plan from Ixia [48] builds on the RFC-2544 and RFC-2889, but extends them by additional benchmarking tests.

4. PREREQUISITE FOR TESTING

This chapter describes the prerequisites that need to be defined prior to the application of performance testing. In particular, these are the requirements, the test setup and the test execution. In order to have a fair comparison between two or more DUTs, it is essential that these topics are well defined.

4.1 Requirements

The DUT must be configured by the provided instructions. Particularly, it is anticipated that all of the supported protocols are configured and enabled. It is expected that all performance benchmarking tests run without altering the configuration of the DUT in any way other than specified in the requirements for the specific test. This should prevent manipulation to enhance the test results. For example, it is not allowed to disable all transport protocols but one; while testing that specific transport protocol. Further, the DUT should include the usual recommended routing update intervals and keep alive frequency. This procedure should ensure transparency and therefore a fair comparison between DUTs. [23]

In order to facilitate this transparency, well specified frame formats and sizes should be used while performing the tests. In addition, it is of interest to know the performance of a DUT under a number of different conditions. The performance tests should be applied under as many conditions as the test equipment can simulate. Therefore, the test suite should first be run without any modification. After that, the test should be repeated under each available condition separately. If the number of conditions or interesting condition combinations is feasible, the tests may also be performed while successively adding conditions. [23]

While not all possible manipulation can be covered, and the single DUTs vary in complexity and setting options, the exact configuration of the DUT and software, including which functions are disabled, have to be included as part of the report of a performance benchmark test. [23]

4.2 Test Setup

The RFC-2544 document moreover explains how the defined benchmark tests may be set up. Ideally, the series of tests is performed with a tester with both transmitting and receiving ports. Consequently a connection is made from the sending ports of the tester to the receiving ports of the DUT and also another connection from the sending ports of the DUT to the receiving ports of the tester (see Figure 1). This way the tester can determine how many packets received by the DUT were transmitted. In addition, the tester can verify that the correct packets were received. [23]
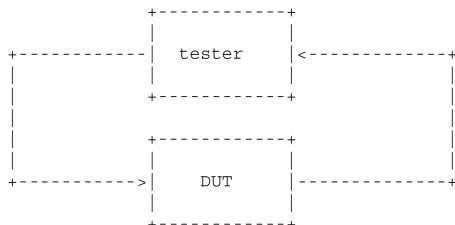
```
                +------------+
                |            |
+-----------+   |   tester   |<--------------+
|           |   |            |               |
|           |   +------------+               |
|           |                                |
|           |   +------------+               |
|           |   |            |               |
+---------->|   |    DUT     |---------------+
                |            |
                +------------+
```

Figure 1: Test Setup [23]

4.3 Test Execution

The execution of a performance test as described in [23] consists of multiple trials. Each of these trials returns a specific piece of information, for example the throughput at a particular input frame rate. There are five phases which each trial undertakes according to the RFC-2544:

1. In case the DUT is a router, send the routing update and pause two seconds to ensure the update is done.

2. Send the learning frames to the output port and wait for two seconds. The learning frames to be used should be uniformly specified.

3. Run the performance test trial.

4. Wait for two seconds for any remainder frames to be received.

5. Wait for at least five seconds for the DUT to stabilize.

The objective for benchmarking tests is to determine the results which can be continuously expected by the DUT. Therefore, the duration of a trial must be a compromise between the accuracy and the duration of a trial. While more trials yield to better statistical evaluable results, it might not be feasible to run a lot of trials with a long duration, especially when different conditions should be taken into account. However, the duration of a trial should be at least 60 seconds.

5. TAXONOMY OF PERFORMANCE TEST

This chapter describes the performance benchmark tests and the reporting formats which are defined in the broadly accepted standard document RFC-2544. The performance indexes for these tests are derived from the RFC-1242. [23]

5.1 Throughput

The objective is to determine the packet forwarding capability as defined in the RFC-1242. This refers to the maximum amount of packets per second that a DUT can forward without losing any packet. [1] In order to determine the throughput of a DUT, the following procedure is applied: Send a specific number of frames at a particular rate through the DUT. Afterwards, the number of transmitted frames by the DUT is counted. If fewer frames are transmitted than sent to the DUT, the test is rerun with a reduced number of frames sent to the DUT. The throughput is defined as "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." [23]

The achieved results of this benchmark test should be reported as a graph. The X coordinate of the axis describes the frame size and the Y coordinate represents the frame rate. Further, there should be one line which shows the theoretical frame rate for the media at the specific frame size. A second line will represent the actual test findings. Additionally, the protocol, data stream format and type of media used in the test should be described. This will improve transparency even further. [23]

5.2 Latency

The purpose of this test is to determine the latency as defined in the RFC-1242. The latency test measures the time a frame needs to travel from the DUT through the network to the destination device. At first the throughput of the DUT needs to be measured, to ensure the frames are transmitted without being discarded. The second step is for the packet generator to send traffic for 120 seconds. Every 60 seconds an identifying tag is included in one frame. The time at which this frame is fully transmitted is recorded (Timestamp A). The receiver logic in the test equipment recognizes the tagged frame and records the time at which it was received (Timestamp B). The latency value is the result of the subtraction of Timestamp B and Timestamp A. [22, 23, 32] This test can be configured to measure the round-trip time. [32, 43]

The report of the result must specify which definition of latency according to the RFC-1242 was used. The latency should be reported as a table where the rows contain the frame size. The columns of the table should represent the rate at which the latency test was run, the media type and the resultant latency value. For each frame size, the measurement must be conducted 20 times. The reporting value is the average of these measurements. [23, 32]

5.3 Packet Loss Rate

The target is to determine the frame loss rate as defined in the RFC-1242. The test equipment sends a specific number of frames at maximum line rate and then measures whether the network dropped any frames. If this is the case, the values are recorded and the test is restarted at a slower rate. The granularity of reducing the frame rate must be at least 10%, while a finer granularity is encouraged. This test is repeated until there are two successive trials in which no frames are lost. [1, 22, 23, 32]

The achieved results should be reported as a graph. The X axis constitutes the input frame rate as a percentage of the theoretical rate of the media. The Y axis depicts the percent loss at the given input rate. [23]

5.4 Back-to-Back Frame

The object of this test is to characterize the ability of a DUT to process back-to-back frames as defined in the RFC-1242. It assesses the buffering capability of a DUT. The test determines the maximum number of frames received before a frame is lost. An increasing number of devices can produce bursts of back-to-back frames. Since the MTU of Ethernet networks is relatively small, many fragments have to be transmitted. The loss of even one fragment can cause an endless loop as the sender continuously attempts to send the data again.

For the execution of the test, a burst of back-to-back frames with minimum inter-frame gap is sent to the DUT. Should a frame be dropped, the burst length is decreased. If the frames are received without any errors, the burst length will be increased. Each trial should be at least two seconds long. The measurement should be repeated 50 times. The back-to-back value is the average of the recorded values being reported for each frame size. [22, 23, 32]

The back-to-back frame benchmark should be reported as a table. In that case the rows represent the tested frame size. The columns show the average frame count for each type of data stream tested. Additionally, the standard deviation for each measurement can be reported. [23]

5.5 System Recovery

The purpose of this test is to characterize the speed at which a DUT recovers from an overload condition. The test procedure starts by measuring the throughput of a DUT. Afterwards a stream of frames at a rate of the minimum of 110% of the assessed throughput or maximum rate for the media is sent to the DUT for at least 60 seconds. At a tagged frame (Timestamp A) the frame rate is reduced to 50% of the initial rate. After that change in the frame rate, the time of the last frame lost is recorded (Timestamp B). The system recovery time is obtained by subtracting Timestamp B from Timestamp A. [23]

The benchmark test should be reported in the format of a table. Whereby the rows specify the tested frame sizes. The columns of the table constitute the frame rate used as well as the measured recovery time for each type of data stream tested. [23]

5.6 Reset

The target is to determine the speed at which a DUT recovers from a device or software reset. First, the throughput benchmark needs to be performed for the minimum frame size on the media used in the trial. After that, a continuous stream of data is sent to the DUT at the recorded throughput rate. Then a reset is caused in the DUT. The time of the last frame of the initial stream being received (Timestamp A) and the first frame of the new stream being received (Timestamp B) needs to be recorded. The reset value is determined by subtracting Timestamp A from Timestamp B. The report format is a simple set of statements, one for each reset type. [23]

6. HARDWARE TEST EQUIPMENT

This chapter examines the usage of performance methodologies in hardware test equipment and describes the adjustments and notes from vendors. By that, we can make conclusions about their usefulness and applicability.

6.1 Agilent

Agilent Technologies [25] uses the RFC-2544 methodology in two of their products which are called Agilent FrameScope™ and Agilent FrameScope™ pro. [26] The company emphasizes that the RFC-2544 is not a standard, but did became increasingly popular and well-accepted to determine the performance of network devices and therefore is used for benchmarking performance tests. As described by Agilent Technologies: In order to meet the requirements of the RFC-2544, a considerable amount of configuration needs to be done. The execution as well as the set up for testing is very time consuming. The Agilent Agilent FrameScope™ therefore incorporates several efficiency improvements to the testing as defined in the RFC-2544. Out of the six performance tests defined in the RFC, only throughput, latency, frame loss rate and back-to-back frames are supported. Hence, the system recovery and reset tests are omitted. [26]

The test parameter setup of the FrameScope™ allows for greater flexibility as defined in the RFC-2544. The tester has a choice of whether the testing will be done upstream only or downstream only or both. Moreover, the performance tests can be extended to frame sizes outside the range specified by the RFC-2544. [26]

The test suite of the product allows for configuring and saving all RFC-2544 testing parameters. This enables for testing under different conditions by modifications of the given parameters in a reasonable amount of time. The RFC-2544 specifies that one trial should last at least 60 seconds. [23] This makes testing time consuming. For this reason the test equipment allows for automated testing. Another crucial requirement in order to allow for transparency and fair comparison of DUTs is a meaningful reporting format. Therefore Agilent Technologies implements a web based reporting tool which satisfies the reporting requirements as defined in the RFC-2544. This document can also serve for Service Level Agreement (SLA) verification between a service provider and a customer. [26]

6.2 Albedo

Albedo [27] implements the benchmarking tests as defined in the RFC-2544 and ITU-T Y.1564 in their test equipment Ether.Giga. [28] The test device supports up to 10 Gigabit Ethernet. In the document *Ethernet RFC-2544 explained* [29], the company describes their motive and how they apply the performance tests.

A tester consists of both transmitting and receiving ports. The tester includes sequence numbers in the frames it transmits, in order to check that all frames transmitted are also received back. The test equipment can be used to test OSI-Layer 2 and OSI-Layer 3 data plane devices. However, one criterion that is not matched by the test equipment is the test duration as defined in the RFC-2544. Since the RFC was designed for laboratory testing [23, 42], the trials may take several days to complete. This duration is not feasible when applied in practice. The time can be reduced by the selection of certain tests to be run as well as reducing their duration. This violates the requirement that every

possible condition should be tested, which can be supported by a DUT. Further, it violates the requirement that a trial should run at least 60 seconds. [23, 29]

6.3 Exfo
Exfo [30] uses the RFC-2544 and ITU-T Y.1564 as performance benchmark methodologies in their test equipment Power Blazer, which supports 100 Gigabit Ethernet. [31] The motivation behind this decision is described in [32] as following: The customer's SLA dictate certain performance criteria which must be met. However, Ethernet performance criteria are difficult to prove and cannot be accomplished accurately by bit-error-rate (BER) anymore.

The portable RFC-2544 test equipment provided by Exfo enables performance testers to immediately test and demonstrate that the data plane device meets the customer's SLA. The results captured this way may serve for further comparisons of different devices. [32] The test equipment makes an addition to the defined benchmark tests in the RFC-2544, by the measurement of packet jitter. This is crucial, because exorbitant jitter can cause failures in real-time applications. This can cause dropout effects in VoIP applications. For video applications this can cause images to falsify. [32] As criticized in chapter 6.1, the duration of a trial as defined in the RFC-2544 leads to a problem. Every test should be performed by each defined frame sizes as defined in [23]. Further, each test trial has 20 iterations. This will yield to a length of almost five hours, which is not feasible. Therefore customization is needed. This can be accomplished by testing only two out of the seven defined frame sizes or conducting only two out of the six defined performance tests. It depends on the type of data plane device and the area in which it will be applied. [32]

6.4 Spirent
The Spirent TestCenter 2.0 as described in [34] is a comprehensive test suite which provides OSI-Layer 2-7 testing for up to 10 Gigabit Ethernet. The test suite implements both RFC-2544 and RFC-2889. Spirent [33] describes these RFCs as industry-standard. [34] It extends the benchmark methodology defined in the RFC-2544 similar to [26] by jumbo Ethernet frames in order to test streaming and conformance testing for certain protocols. The test suite contains six major fields of testing: Ethernet Switch Testing, Enterprise/Metro Router Testing, Carrier Ethernet Testing, Broadband Access Testing, Layer 4-7 Testing as well as IPTV and Video Quality Testing. [34]

Another test equipment product developed by Spirent is the Router Performance Tester AX/4000 (RPT). [35] The RPT allows for customizing of IP test packets and traffic generation through the DUT at full line rate and reports in real-time. The device supports testing data plane devices according to the RFC-2544. However, as in [26] and [32] the two tests of system recovery and reset as defined in the RFC-2544 are not supported. [35]

6.5 Viavi
An Ethernet testing solution provided by Viavi [36], described in the product note [38], is building on the RFC-2544 as well. The test equipment further supports the Y.1564 standard mentioned in chapter 3 and can saturate rates up to 10 Gigabit. The two standards can be tested asymmetrically with a unidirectional upstream and downstream traffic. Further test applications supported are mobile Ethernet backhaul up to LTE, cloud connectivity and fault isolation. [38]

6.6 Ixia
The IxAutomate test suite [47] from Ixia [46] implements both RFC-2544 and RFC-2889 recommendations. While the RFC-2544 was designed as a general methodology for networking devices of all types, the RFC-2889 was written specifically to benchmark the data plane performance of OSI-Layer 2 LAN switching devices. Additionally, the new recommendation RFC-2544-IPv6 [45] is integrated into the test equipment. While all benchmarking tests of the RFC-2889 are implemented, the system recovery and reset test as defined in the RFC-2544 are omitted. [47]

6.7 Xena Networks
The company Xena Networks [39] offers with XenaBay and XenaCompact two OSI-layer 2 and 3 test chassis which support up to 100 Gigabit Ethernet. These chassis can be configured with different test software provided by Xena. Specifically, software is provided for the following test methodologies: RFC-2544, RFC-2889, RFC-3918 and ITU-T Y.1564. Further, a software packet is provided for scripting test automation. [40]

7. SOFTWARE TEST EQUIPMENT
Besides testing data plane devices with hardware equipment there is the possibility to measure their performance with software. Approaches exist in PacketExpert [53], LAN Tornado RFC 2544 [54], iPerf [55], Ostinato [56] or MoonGen [57]. PacketExpert, LAN Tornado RFC 2544 and MoonGen come with direct support for the RFC-2544.

Both hardware and software testing procedures have to face the same challenges: performance, flexibility, and precision in timestamping and rate control. The advantages of software traffic generators are their flexibility and their low costs. The MoonGen [59] software e.g., uses Lua Scripts which allow for modification of each single packet and runs on Intel commodity NICs. Hardware test equipment on the other hand have advantages in performance and precision. [57]

It is important for test equipment to saturate high rates, while still being precise to assure repeatability of the same experiment in different moments in time and a fair comparison between different DUTs. If this can not be assured, the outcome of the tests have no value. This is the major downside of software packet generators, which often lack performance capabilities and precision. [58]

Botta *et al.* [58] discuss the inherent problems of software packet generator in more detail. While the MoonGen [59] project presents new approaches which overcome some of these disadvantages.

8. USEFULNESS AND FEASIBILITY
Based on the evaluation in chapter 6 and 7, we can conclude that test equipment vendors are using the RFC-2544, RFC-2889 and ITU-T Y.1564 as standards to build their

test equipment. Further, academic research is interested in developing network testing equipment that supports the RFC-2544. [1, 51, 52, 57]

However, there were also criticism and disadvantages mentioned from these parties. Since the RFC-2544 was written over a decade ago, in 1999, it is no longer sufficient in terms of fully validating today's Ethernet services. The RFC-2544 does not satisfy all requirements anymore, such as packet jitter, QoS measurement or multiple concurrent service levels. Also, the method of sequential testing takes several hours to complete, as mentioned by manufacturers. [25, 27] This test method is both time consuming and costly. Furthermore, the system recovery and reset test as defined in the RFC-2544 were rarely implemented by any test equipment vendor.

In the throughput test, the RFC-2544 makes no distinction between committed and excess traffic. This is not sufficient for testing SLA. The frame delay is determined based on the assessment of a single frame during a trial. This approach does not take into account any fluctuation or peak that may occur during the testing. Furthermore, the RFC-2544 does not measure the inter-frame delay variation (IFDV). [43, 63] Most of the tests defined in the RFC-2544 are performed with one endpoint generating traffic and another endpoint placed in loopback. While this is the simplest and fastest way to perform a test trial, there are disadvantages to this test setup. When a test fails, there is no information on where packets are being dropped or where delay is being introduced. [37] A solution to that is the testing in an asymmetric mode, as adapted by manufacturers. [25, 36]

Additionally, vendors of test equipment like Exfo, Ixia and MRV have concerns about the feasibility of the test methodologies described in the RFC-2544. [63, 64, 65] This is due to the reason of time consumption of the test trials as well as the lack of validating certain features. Ixia [64] states that the RFC-2544 and RFC-2889 are good for testing best case scenarios in a laboratory environment, but they do not provide an insight into the device performance under a real-world data center traffic load. Additionally, they do not assess performance of mixed frame sizes. Both [64] and [65] therefore suggest the usage of the ITU-T Y.1564. [42] This is reflected by the usage of methodologies by test equipment vendors as surveyed in chapter 6. The evaluation (see Table 1) shows, that all examined test equipment, developed after the approval of the ITU-T Y.1564 in March 2011, implement the suggest methodology.

Test Equipment	RFC-2544	Y.1564	GbE	Year
Agilent FrameScope	yes	no	1	2006
Albedo Ether.Giga	yes	yes	10	2012
Exfo Power Blazer	yes	yes	100	2015
Spirent TestCenter	yes	no	10	2006
Viavi QT-600-10	yes	yes	10	2015
Ixia IxNetwork	yes	no	1	2007
Xena XenaBay	yes	yes	100	2014

Table 1: Usage of benchmarking methodologies

The BMWG is continuously working on the mentioned disadvantages. The proposal of the RFC-4814 [62] considers the overlooked factors in network device benchmarking and addresses the issues in the RFC-2544 and RFC-2889. In the RFC-6815 [61], an applicability Statement for the RFC-2544, the BMWG states that actual methods may vary from the methodology described in the RFC-2544. Another proposal [60] considers the issues related to conducting tests similar to the RFC-2544 in a production network. The IETF has addressed the scenario of production network performance testing by commissioning a new working group by the name of IP Performance Metrics (IPPM). [60]

9. CONCLUSION

The RFC-2544 defines a methodology for benchmarking the data plane performance of networking devices. It measures a networking device's throughput, latency and frame loss rate for specified frame sizes and further defines a system recovery and reset test. On the basis of the RFC-2544, the RFC-2889 defines benchmarking tests for an OSI-layer 2 LAN switching device. While the ITU-T Y.1564 methodology addresses issues of today's Ethernet Services and SLA validations. The definition of the requirements, the test setup as well as the reporting format ensures transparency and therefore a fair comparison of different data plane devices. These methodologies are widely accepted for benchmarking data plane devices and for this reason established themselves as a standards.

Since the RFC-2544 was written over a decade ago in 1999, the methodology does not assess all Ethernet services that are available today. Further, the sequential approach of testing is very time consuming. For this reason new standards were developed. The ITU-T Y.1564 and the MEF 14 methodologies attempt to overcome the disadvantages of the RFC-2544. Both standards are used by test equipment vendors. The BMWG itself is aware of the disadvantages of the RFC-2544 and continuous to work at the topic of benchmarking data plane devices. The RFC-6815 and RFC-4814 reference the RFC-2544 and are addressing the concerns of industry and academic research.

Even if some sources testify that the RFC-2544 is deprecated in terms of testing today's Ethernet services, it still remains a foundation for the development of performance benchmarking methodologies and will therefore continue to contribute on how the performance of data plane devices is measured.

10. REFERENCES

[1] L. Niu, G. Feng and M. Duan: *Implementation of Instrument for Testing Performance of Network Based on RFC2544 Test*, In International Journal of Hybrid Information Technology Vol. 8, No. 2, pages 323-332, 2015.

[2] S. Bouckaert, J. V. V. Gerwen, I. Moerman, S. C. Phillips, and J. Wilander: *Benchmarking computers and computer networks*, EU FIRE White Paper, 2010.

[3] Meru Networks: *Demystifying Software-Defined Networking for Enterprise Networks*, 2013.

[4] M. Brown and R., Burns: *Cisco CCNA data center DCICT 640-916 official certification guide*, 2013.

[5] I. Pepelnjak: *Management, Control and Data Planes in Network Devices and Systems*, http://blog.ipspace.net/2013/08/management-

control-and-data-planes-in.html, 2013.

[6] B. Salisbury: *The Control Plane, Data Plane and Forwarding Plane in Networks*, `http://networkstatic.net/the-control-plane-data-plane-and-forwarding-plane-in-networks`, 2012.

[7] Open Networking Foundation: *Software-Defined Networking: The Norm for Networks*, 2012.

[8] D. Kreutz, F. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig: *Software-Defined Networking: A Comprehensive Survey*, In Proceedings of the IEEE Vol. 103, No. 1, 2015.

[9] G. Schudel and D. J. Smith: *Router Security Strategies: Securing IP Network Traffic Planes*, Cisco Press, pages 24-30, 2008.

[10] N. Shamsee,Klebanov D., H. Fayed, A. Afrose and O. Karakok: *CCNA Data Center DCICT 640-916 Official Cert Guide*, 2015.

[11] Savvius: *Network Interconnect Devices: Repeaters, Bridges, Switches, Routers*, `http://www.wildpackets.com/resources/compendium/interconnect_devices/overview`, 2015.

[12] SCO OpenServer: *Networking Guide*, 2003.

[13] C. Wells: *Network Interconnection Devices, TechnologyUK*, `http://www.technologyuk.net/telecommunications/networks/interconnection_devices.shtml`, 2001.

[14] A. Groessler: *Internetworking: Devices*, `http://pi4.informatik.uni-mannheim.de/pi4.data/content/courses/1996-ss/rn96/CN-Title/form/intdevie.htm`, 1995.

[15] ComputerNetwokringNotes: *Networking Devices: Hub, Switch, Router, Modem, Bridges, and Brouters Gateways*, `http://computernetworkingnotes.com/comptia-n-plus-study-guide/network-devices-hub-switch-router.html`, 2013

[16] CISCO: *Layer 2 and Layer 3 Switch Evolution*, The Internet Protocol Journal Vol. 1, No. 2, September 1998.

[17] H. Alvestrand: *RFC-3935: A Mission Statement for the IETF*, *https://www.ietf.org/rfc/rfc3935.txt*, 2004.

[18] E. Huizer: *RFC-1603: IETF Working Group Guidelines and Procedures*, `https://tools.ietf.org/rfc/rfc1603.txt`, 1994.

[19] BMWG: Charter for Working Group, `http://datatracker.ietf.org/wg/bmwg/charter/`, 2015.

[20] L. Yang, R. Dantu, T. Anderson and R. Gopal: *RFC-3746: Forwarding and Control Element Separation (ForCES) Framework*, IETF Network Working Group, 2004.

[21] M. Fratto: *What is the difference between Control plane and Data plane in the context of networking equipment like routers/switches?*, `https://www.quora.com/What-is-the-difference-between-Control-plane-and-Data-plane-in-the-context-of-networking-equipment-like-routers-switches`, 2012.

[22] S. Bradner: RFC-1242: *Benchmarking Terminology for Network Interconnection Devices*, IETF Benchmarking Methodology Group, July 1991.

[23] S. Bradner and J. McQuaid: *RFC-2544: Benchmarking methodology for network interconnect devices*, IETF Benchmarking Methodology Group, March 1999.

[24] R. Mandeville and J. Perser: *RFC-2889: Benchmarking Methodology for LAN Switching Devices*, IETF Benchmarking Methodology Group, August 2000.

[25] Agilent Technologies: `http://www.agilent.de/home`, November 2015.

[26] H. Pandya: *RFC-2544 Network Performance Testing with the Agilent FrameScope^{TM}*, Agilent Technologies, January 2006.

[27] Albedo: `http://www.albedo.com/`, November 2015.

[28] Albedo: *Ether.Giga*, `http://www.albedotelecom.com/pages/fieldtools/src/ethergiga.php`, February 2016.

[29] P. Fan: *Ethernet RFC-2544 explained*, Albedo, April 2003.

[30] Exfo: `http://www.exfo.com/`, November 2015.

[31] Exfo: Power Blazer `http://www.exfo.com/products/field-network-testing/bu2-transport-datacom/ethernet-testing/ftb-88100nge-power-blazer`, February 2016.

[32] B. Giguère: *RFC 2544: How it helps qualify a Carrier Ethernet Network*, Exfo, August 2004.

[33] Spirent: `http://www.spirent.com/`, November 2015.

[34] Spirent: *Get There Faster with Spirent TestCenter^{TM}*, May 2007.

[35] Spirent: *Router Performance Tester: AX/4000*, September 2005.

[36] Viavi: `http://www.viavisolutions.com/en-us`, November 2015.

[37] Viavi: *Viavi Solutions^{TM} Ethernet Testing*, November 2015.

[38] Viavi: *QT-600-10 10 Gigabit Ethernet Test Head*, June 2015.

[39] Xena Networks: `http://www.xenanetworks.com/`, November 2015.

[40] Xena Networks: *Xena Scripting*,`http://www.xenanetworks.com/test-software/xena-management-sw/xenascripting/`, February 2016.

[41] Ixia: *Ultra Low Latency (ULL) Testing*, June 2011.

[42] ITU-T: *Ethernet service activation test methodology*, March 2011.

[43] T. Diallo and M. Dorais: *EtherSAM: The new Standard in Ethernet Service Testing*, January 2014.

[44] Metro: *Abstract Test Suite for Traffic Management Phase 1*, November 2005.

[45] C. Popoviciu, A. Hamza, G. V. de Velde and D. Dugatkin: *RFC5180: IPv6 Benchmarking Methodology for Network Interconnect Devices*, IETF Benchmarking Methodology Group, May 2008.

[46] Ixia: `http://www.ixiacom.com/`, November 2015.

[47] Ixia: *IxAutomate RFC Benchmarking Test Suites*,

July 2007.

[48] Ixia: *Switch Testing: IxAutomate*, Test Plan, April 2006.

[49] Eantc: *Huawei Technologies Service Activation Using RFC 2544 Tests*, Mai 2008.

[50] Tolly: *Araknis Networks AN-300 Series Gigabit PoE Switch: Performance and Feature Comparison Versus Cisco Systems and Pakedge Device & Software*, Test Report, March 2014.

[51] C. Both, C. Battisti and F. Kuentzer: *FPGA implementation and performance evaluation of an RFC 2544 compliant Ethernet test set*, In Int. J. High Perform Systems Architecture Vol. 2, No. 2, pages 107-115, 2009.

[52] Y. Wang, Y. Liu, X. Tao and Q. He: *An FPGA-based high-speed network performance measurement for RFC 2544*, In EURASIP Journal on Wireless Communications and Networking, No. 1, pages 1-10, 2015.

[53] GL Communications Inc.: *PacketExpert*, `http://www.gl.com/packetexpert-rfc-2544-ber-loopback-testing.html`, February 2016.

[54] SoftDevTeam: *LAN Tornado RFC 2544*, `http://lan-tornado-rfc-2544.soft112.com/`, August 2012.

[55] iPerf: `https://iperf.fr/iperf-doc.php`, February 2016.

[56] Ostinato: `http://ostinato.org/`, February 2016.

[57] P. Emmerich, F. Wohlfart, D. Raumer and G. Carle: MoonGen: *A Scriptable High-Speed Packet Generator, ArXiv e-prints*, October 2014.

[58] A. Botta, A. Dainotti and A. Pescapé: *Do You Trust Your Software-Based Traffic Generator?*, IEEE Communications Magazine, IEEE 48 No. 9, pages 158-165, 2010.

[59] MoonGen:`https://github.com/emmericp/MoonGen/tree/master/rfc2544/benchmarks`, December 2015.

[60] R. Bonica and S. Bryant: *RFC2544: Testing in Production Networks, IETF Benchmarking Methodology Group*, October 2012.

[61] S. Bradner, K. Dubray, J. McQuaid and A. Morton: *RFC6815: Applicability Statement for RFC 2544: Use on Production Networks Considered Harmful*, IETF Benchmarking Methodology Group, November 2012.

[62] D. Newman and T. Player: RFC-4814: *Hash and Stuffing: Overlooked Factors in Network Device Benchmarking*, IETF Benchmarking Methodology Group, March 2007.

[63] Exfo: *Are You Still Testing to RFC 2544?*, `http://www.exfo.com/corporate/blog/2013/still-testing-rfc-2544-really`, December 2015.

[64] Ixia: *Is RFC2544 Enough to Benchmark the Data Center Switching Fabric Performance?*, `http://www.ixiacom.com/about-us/news-events/corporate-blog/rfc2544-enough-benchmark-data-center-switching-fabric`, December 2015.

[65] MRV: *Why RCF2544 is not sufficient anymore*, `http://www.mrv.com/blog/why-rcf2544-not-sufficient-anymore`, December 2015.