

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: September 8, 2011

M. Hamilton
BreakingPoint Systems
S. Banks
Cisco Systems
March 7, 2011

Benchmarking Terminology for Content-Aware Network Devices
draft-hamilton-bmwg-ca-bench-term-00

Abstract

The purpose of this document is to define and outline the terminology necessary to appropriately follow and implement "Benchmarking Methodology for Content-Aware Network Devices". Relevant terms will be defined and discussed throughout this document in order to ensure the comprehension of the previously mentioned methodology.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 8, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as

described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	4
2.	Scope	4
3.	Definitions	4
3.1.	Application Flow	5
3.2.	Application Throughput	5
3.3.	Average Time to TCP Session Establishment	6
3.4.	Content-Aware Device	6
3.5.	Deep Packet Inspection	7
3.6.	Network 5-Tuple	7
3.7.	Session Establishment Rate	8
3.8.	Session Establishment Time	8
3.9.	Simultaneous TCP Sessions	9
3.10.	Time To SYN	9
4.	IANA Considerations	10
5.	Security Considerations	10
6.	References	10
6.1.	Normative References	10
6.2.	Informative References	11
	Authors' Addresses	11

1. Introduction

Content-aware and deep packet inspection (DPI) device penetration has grown significantly over the last decade. No longer are devices simply using Ethernet headers and IP headers to make forwarding decisions. Devices that could historically be classified as 'stateless' or raw forwarding devices are now seeing more DPI functionality. Devices such as core and edge routers are now being developed with DPI functionality to make more intelligent routing and forwarding decisions.

The Benchmarking Working Group (BMWG) has historically produced Internet Drafts and Requests for Comment that are focused specifically on creating output metrics that are derived from a very specific and well-defined set of input parameters that are completely and unequivocally reproducible from testbed to testbed. The end goal of such methodologies is to, in the words of the BMWG charter "reduce specmanship" from network equipment manufacturers (NEM's). Existing BMWG work has certainly met this stated goal.

Today, device sophistication has expanded beyond existing methodologies, allowing vendors to reengage in specmanship. In order to achieve the stated BMWG goals, the methodologies designed to hold vendors accountable must evolve with the enhanced device functionality.

The BMWG has historically avoided the use of the term "realistic" throughout all of its drafts and RFCs. While this document will not explicitly use this term, the end goal of the terminology and methodology is to generate performance metrics that will be as close as possible to equivalent metrics in a production environment. It should be further noted that any metrics acquired from a production network MUST be captured according to the policies and procedures of the IPPM or PMOL working groups.

An explicit non-goal of this document is to replace existing methodology/terminology pairs such as RFC 2544 [1]/RFC 1242 [2] or RFC 3511 [3]/RFC 2647 [4]. The explicit goal of this document is to create a methodology and terminology pair that is more suited for modern devices while complementing the data acquired using existing BMWG methodologies. Existing BMWG work generally revolves around completely repeatable input stimulus, expecting fully repeatable output. This document departs from this mantra due to the nature of modern traffic and is more focused on output repeatability than on static input stimulus.

Some of the terms used throughout this draft have previously been defined in "Benchmarking Terminology for Firewall Performance" RFC

2647 [4]. This document SHOULD be consulted prior to using this document.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [5].

2. Scope

Content-aware devices take many forms, shapes and architectures. These devices are advanced network interconnect devices that inspect deep into the application payload of network data packets to do classification. They may be as simple as a firewall that uses application data inspection for rule set enforcement, or they may have advanced functionality such as performing protocol decoding and validation, anti-virus, anti-spam and even application exploit filtering.

This document is strictly focused on examining performance and robustness across a focused set of metrics that may be used to more accurately predict device performance when deployed in modern networks. These metrics will be implementation independent.

It should also be noted that the purpose of this document is not to perform functional testing of the potential features in the Device/System Under Test (DUT/SUT)[4] nor specify the configurations that should be tested. Various definitions of proper operation and configuration may be appropriate within different contexts. While the definition of these parameters are outside the scope of this document, the specific configuration of both the DUT and tester SHOULD be published with the test results for repeatability and comparison purposes.

While a list of devices that fall under this category will quickly become obsolete, an initial list of devices that would be well served by utilizing this type of methodology should prove useful. Devices such as firewalls, intrusion detection and prevention devices, application delivery controllers, deep packet inspection devices, and unified threat management systems generally fall into the content-aware category.

3. Definitions

3.1. Application Flow

Definition:

An application flow is the virtual connection between two network hosts that is used to exchange user data above the transport layer.

Discussion:

Content-aware devices may potentially proxy session-layer connections, acting as a virtual server to the client and a virtual client to the server. In this mode, the SUT/DUT may modify members of the network 5-tuple or act on their behalf, thus each end host is actually disconnected at the session layer. Application flows are virtual connections that are between the two hosts, irrespective of the nature of the session layer semantics.

Unit of Measurement:

N/A

Issues:

N/A

See Also:

5-Tuple

3.2. Application Throughput

Definition:

The rate at which data associated with an application flow is transmitted through the SUT/DUT.

Discussion:

Throughput metrics may be calculated at various layers in the network protocol stack. Each layer does contain associated overhead necessary to maintain that layer. Application throughput is the number of bits transmitted through a SUT/DUT, not including the overhead associated with lower layer protocols. Measurement should be taken at the receiver side to minimize the impact of session layer retransmissions.

Unit of Measurement:

N/A

Issues:

Some applications may not rely on session layer reliability mechanisms. This definition does not cover the case where an application may utilize its own specific reliability/retransmission algorithm.

See Also:
N/A

3.3. Average Time to TCP Session Establishment

Definition:

The average time that a SUT/DUT requires to complete the TCP session establishment process.

Discussion:

The average time to TCP session establishment is calculated by taking the sum of all "TCP Session Establishment Time" values acquired in the specified time frame and divide by the total number of sessions established within that timeframe. The timeframe in which the average is taken will depend on the methodology itself and what is trying to be measured.

Unit of Measurement:

Seconds.

Issues:

Depending on how the DUT/SUT handles TCP session establishment, the client and server may have different values for the same TCP session. A client-side session may be established prior to the server-side session being established.

See Also:
See Also.

3.4. Content-Aware Device

Definition:

A networking device which performs deep packet inspection.

Discussion:

For a more detailed discussion, please see "deep packet inspection".

Unit of Measurement:

Not Applicable.

Issues:

Not Applicable.

See Also:
Deep Packet Inspection

3.5. Deep Packet Inspection

Definition:

The process by which a network device inspects layer 7 payload as well as protocol headers when making processing decisions.

Discussion:

Deep packet inspection (DPI) has grown from a feature reserved for Intrusion Prevention Devices into functionality that is shared across many next generation networking devices. Devices traditionally classified as firewalls are now looking at layer 7 payloads to make decisions, whether it is classification, rate-shaping, or actually deeming whether a flow is allowed. Many deep-packet inspection devices utilize proxy behavior as a functional choice for performing inspection.

Unit of Measurement:

Not Applicable.

Issues:

Not Applicable.

See Also:

Content-Aware Device

3.6. Network 5-Tuple

Definition:

The set of 5 metrics which distinguish two session layer connections from each other.

Discussion:

When discussing data transfer between hosts, a Network 5-tuple is typically used to differentiate between multiple session layer connections. Source and destination IP addresses, source and destination session-layer ports, and the session layer protocol make up the network 5-tuple. The session layer protocol is typically TCP or UDP, but may be SCTP or another session layer protocol.

Unit of Measurement:

N/A

Issues:

N/A

3.7. Session Establishment Rate

Definition:

The rate at which TCP sessions may be established through a given DUT/SUT.

Discussion:

The session establishment rate is a measurement of how many TCP sessions the DUT/SUT is able to establish in a given unit of time. If within a 1 second time interval the tester is able to establish 10,000 sessions, that rate will be measured at 10,000 sessions per second. The session must be established in accordance with the policy set forth in "Session Establishment Time".

Unit of Measurement:

TCP session(s) per second

Issues:

Issues.

See Also:

See Also.

3.8. Session Establishment Time

Definition:

Session establishment time is the difference in time between the first TCP SYN packet sent from the client and when TCP ACK packet's arrival at the server interface.

Discussion:

This metric is calculated between the time the first bit of the TCP SYN packet is sent from the client and the time the last bit of the TCP ACK packet arrives on the server interface.

Unit of Measurement:

Seconds.

Issues:

Depending on how the DUT/SUT handles TCP session establishment, the client and server may have different values for the same logical TCP session. A client-side session may be established prior to the server-side session being established.

See Also:

3.9. Simultaneous TCP Sessions

Definition:

The number of TCP sessions which are in the 'Established State' as defined by RFC 793 [6].

Discussion:

This measurement counts the number of TCP sessions which are in the 'Established State'. Sessions which are in this state must be able to maintain data transfer between client and server, bi-directionally.

Unit of Measurement:

Sessions.

Issues:

Depending on the nature of the SUT/DUT, the number of simultaneous sessions may instantaneously be different when counted from the client and server sides of the SUT/DUT.

See Also:

See Also.

3.10. Time To SYN

Definition:

The Time to SYN is a one-way metric, which is the difference between the time that the first TCP SYN packet is sent by the client and the time at which the server receives the TCP SYN packet from the client.

Discussion:

This metric is more important with content-aware devices due to the potential proxying issues. Content-aware devices may proxy a TCP session on behalf of the server. Many times, the client will receive the SYN/ACK from the DUT/SUT and complete the TCP handshake before the SYN has been forwarded to the server. This measurement is actually a proxy measure for client-side session establishment time through the DUT/SUT, if the session is in fact proxied.

Unit of Measurement:

Seconds.

See Also:

4. IANA Considerations

This memo includes no request to IANA.

All drafts are required to have an IANA considerations section (see the update of RFC 2434 [9] for a guide). If the draft does not require IANA to do anything, the section contains an explicit statement that this is the case (as above). If there are no requirements for IANA, the section will be removed during conversion into an RFC by the RFC Editor.

5. Security Considerations

Benchmarking activities as described in this memo are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the other constraints RFC 2544 [1].

The benchmarking network topology will be an independent test setup and MUST NOT be connected to devices that may forward the test traffic into a production network, or misroute traffic to the test management network

6. References

6.1. Normative References

- [1] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [2] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [3] Hickman, B., Newman, D., Tadjudin, S., and T. Martin, "Benchmarking Methodology for Firewall Performance", RFC 3511, April 2003.
- [4] Newman, D., "Benchmarking Terminology for Firewall Performance", RFC 2647, August 1999.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [6] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

- [7] Popoviciu, C., Hamza, A., Van de Velde, G., and D. Dugatkin, "IPv6 Benchmarking Methodology for Network Interconnect Devices", RFC 5180, May 2008.
- [8] Brownlee, N., Mills, C., and G. Ruth, "Traffic Flow Measurement: Architecture", RFC 2722, October 1999.

6.2. Informative References

- [9] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

Authors' Addresses

Mike Hamilton
BreakingPoint Systems
Austin, TX 78717
US

Phone: +1 512 636 2303
Email: mhamilton@breakingpoint.com

Sarah Banks
Cisco Systems
San Jose, CA 95134
US

Email: sabanks@cisco.com

