

Internet Engineering Task Force (IETF)
Request for Comments: 6802
Category: Informational
ISSN: 2070-1721

S. Baillargeon
C. Flinta
A. Johnsson
Ericsson
November 2012

Ericsson Two-Way Active Measurement Protocol (TWAMP) Value-Added Octets

Abstract

This memo describes an extension to the Two-Way Active Measurement Protocol (TWAMP). Specifically, it extends the TWAMP-Test protocol, which identifies and manages packet trains, in order to measure capacity metrics like the available path capacity, tight section capacity, and UDP delivery rate in the forward and reverse path directions.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6802>.

Copyright Notice

Copyright (c) 2012 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	2
1.1. Requirements Language	3
2. Purpose and Scope	3
3. Capacity Measurement Principles	4
4. TWAMP-Control Extensions	5
4.1. Additional Considerations	6
5. Extended TWAMP-Test	6
5.1. Sender Behavior	6
5.1.1. Packet Timings	7
5.1.2. Session-Sender Packet Format	7
5.2. Reflector Behavior	12
5.2.1. Session-Reflector Packet Format	13
5.3. Additional Considerations	13
6. Experiments	14
7. Security Considerations	14
8. Acknowledgements	14
9. References	15
9.1. Normative References	15
9.2. Informative References	15

1. Introduction

The notion of embedding a number of meaningful fields in the padding octets has been established as a viable methodology for carrying additional information within the TWAMP-Test protocol running between a Session-Sender and a Session-Reflector [RFC5357] [RFC6038].

This memo describes an optional extension to the Two-Way Active Measurement Protocol (TWAMP) [RFC5357]. It is called the Ericsson TWAMP Value-Added Octets feature. This memo defines version 1.

This feature enables the controller host to measure capacity metrics like the IP-type-P available path capacity (APC) [RFC5136], IP-layer tight section capacity (TSC) [Y1540], and UDP delivery rate on both forward and reverse paths using a single TWAMP test session. The actual method to calculate the APC, TSC, or the UDP delivery rate from packet-level performance data is not discussed in this memo.

The Valued-Added Octets feature consists of new behaviors for the Session-Sender and Session-Reflector and a set of value-added octets of information that are placed at the beginning of the Packet Padding [RFC5357] or immediately after the Server Octets in the Packet Padding (to be reflected) [RFC6038] by the Session-Sender and are reflected or returned by the Session-Reflector. The length of the value-added octets in version 1 is 10 octets. The Valued-Added

Octets feature does not change the basic roles and functions of the TWAMP hosts, which are still responsible to include timestamp(s) and sequence number(s) in the test packets.

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 [RFC2119].

2. Purpose and Scope

The purpose of this memo is to describe the Ericsson TWAMP Valued-Added Octets feature (version 1) for TWAMP [RFC5357].

The scope of the memo is limited to specifications of the following enhancements:

- o The definition of a structure for embedding a sequence of value-added fields at the beginning of the Packet Padding [RFC5357] or Packet Padding (to be reflected) [RFC6038] in the TWAMP-Test packets
- o The definition of new Session-Sender and Session-Reflector behaviors

The motivation for this feature is to enable the measurement of capacity metrics on both the forward and reverse paths using a single TWAMP test session. Multiple TWAMP test sessions between a controller and a responder with different Diffserv Code Points (DSCPs) may also be used to evaluate the QoS impacts on the capacity metrics.

This memo captures the prototype presented and demonstrated at IETF 80. It may be used as a reference for future work or may be used during benchmark analysis to compare the accuracy or performance of the available path capacity estimates under various condition or use cases.

This memo does not extend the standard modes of operation through assignment of a new value in the Modes field (see Section 3.1 of [RFC4656] for the format of the Server Greeting message). This memo does not define a vendor-specific or experimental mode since the Modes field as currently defined does not explicitly reserve a value or range of values for this purpose.

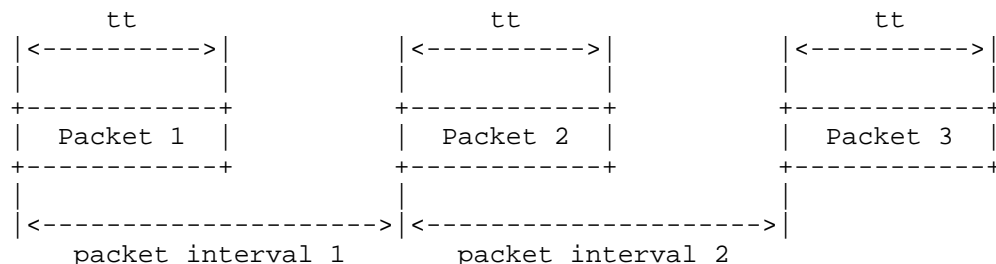
This memo assumes the TWAMP controller is capable to send test packets with value-added padding octets and the TWAMP responder is configured to interpret the value-added padding octets embedded in each TWAMP-Test packet. Bootstrapping such behavior at the TWAMP responder is implementation specific. By default, the feature MUST be disabled on the TWAMP host. The Value-Added Octets feature MUST be deployed in an environment where both controller and responder are managed by the same administrative entity and such entity has established an agreement to operate the Value-Added Octets feature between the pair of hosts or between specific UDP endpoints between the pair of hosts. See Sections 4 and 5.3 for additional considerations.

The Value-Added Octets Version 1 feature is intended to work in conjunction with any TWAMP modes. When the Server and Control-Client are configured or have agreed to use the Value-Added Octets Version 1 feature, then the Control-Client, the Server, the Session-Sender, and the Session-Reflector must all conform to the requirements of that feature, as identified below.

3. Capacity Measurement Principles

Most capacity estimation methods for APC [RRBNC] [PDM] [ENHJMME] [SBW] and for UDP delivery rate need to send and receive packets in groups, called "packet trains" or simply "trains". Each train is sent at a specific transmission rate in a given direction. These trains must be identified within each bidirectional test session stream.

The first measurement principle is to send multiple trains within a test session stream from one IP node to another IP node in order to estimate the APC, TSC, or UDP delivery rate in the forward direction. Each train consists of a group of test packets that are separated from each other by a packet interval, as shown in the figure below. The packet interval is measured using either the first bit or the last bit of two consecutive packets.



The test packet size and interval between consecutive packets for each train sent by the Session-Sender and reflected by the Session-Reflector MUST be calculated and determined by the controller or an application or entity communicating with the controller. The packet size and interval MAY vary within a train and/or between trains. Determination of the packet size and interval is implementation specific.

The transmission time tt to send one packet (i.e., determined by the interface speed and the IP packet size) is also shown in the figure above. Observe that the packet interval MUST be larger than or equal to tt .

At the Session-Reflector, each received test packet within a forward train is time stamped. This provides a second set of packet interval values. Methods for measuring the APC, TSC, and UDP delivery rate use the packet intervals obtained from both endpoints in the estimation process. The method of measuring the UDP delivery rate may also require the rate of packet loss. The estimation process itself, as well as any requirements on software or hardware, is implementation specific.

The second measurement principle is referred to as "self-induced congestion". According to this principle, in order to measure APC, TSC, and UDP delivery rates, some trains MUST cause momentary congestion on the network path. In essence, this means that some trains MUST be sent at a higher rate than what is available on the network path.

In order to fulfill the above measurement principles and to measure the APC, TSC, and UDP delivery rates in the reverse direction, the test packets at the Session-Reflector MUST be regrouped into trains and then transmitted back to the Session-Sender with a provided packet interval.

4. TWAMP-Control Extensions

TWAMP connection establishment follows the procedure defined in Section 3.1 of [RFC4656] and Section 3.1 of [RFC5357]. The TWAMP-Control protocol [RFC5357] uses the Modes field to identify and select specific communication capabilities. According to the standard specifications, the Value-Added Octets feature requires one new bit position (and value) to identify the ability of the Server/Session-Reflector to read and act upon the new fields in the value-added octets. Such bit position (and value) is not defined in this memo. Bootstrapping the TWAMP Value-Added Octets Version 1 feature is implementation specific.

Both the Reflect Octets mode and Symmetrical Size mode MAY be selected to ensure the reflection of the value-added padding octets by the Session-Reflector and symmetrical size TWAMP-Test packets in the forward and reverse directions of transmission.

4.1. Additional Considerations

In the TWAMP control architecture, the TWAMP reflector (server) signals the modes it wishes to operate and the TWAMP controller (control-client) selects the mode or modes supported by the responder. This feature is designed to retain backward compatibility with the original TWAMP-Control and TWAMP-Test protocols. As an alternative, the user may opt for TWAMP Light architecture, which does not require the TWAMP-Control protocol.

The methods to determine if the Value-Added Octets feature is supported on a TWAMP reflector is implementation specific. When the Value-Added Octets feature is not supported on a TWAMP reflector, the TWAMP controller MUST NOT select the Value-Added Octets feature and MUST NOT include any value-added octets in the test packets. If the TWAMP controller inadvertently sends value-added octets in the test packets to a TWAMP responder that does not support such feature, the TWAMP responder shall treat the value-added octets as regular padding octets and return the test packets as quickly as possible to the Session-Sender as defined in [RFC5357].

5. Extended TWAMP-Test

The forward and reverse APC, TSC, and UDP delivery rate measurement characteristics depend on the size and packet intervals of the test packets. This memo allows variable packet sizes and packet intervals between trains and even between packets in the same train. The functionality is described below.

The TWAMP-Test protocol carrying the value-added padding octets is identical to TWAMP [RFC5357] except for the definition of the first 10 octets in the Packet Padding that the Session-Sender expects to be reflected. The new octets define fields for Value-Added Octets Version, Flags, Last Sequence Number in Train, and Desired Reverse Packet Interval. Each of these fields are described in detail below.

The Session-Sender and Session-Reflector behaviors are also modified.

5.1. Sender Behavior

This section describes the extensions to the behavior of the TWAMP Session-Sender.

5.1.1. Packet Timings

The Send Schedule is not utilized in TWAMP, and this is unchanged in this memo.

5.1.2. Session-Sender Packet Format

The Session-Sender packet format follows the same procedure and guidelines as defined in TWAMP [RFC5357] and TWAMP Reflect Octets and Symmetrical Size Features [RFC6038].

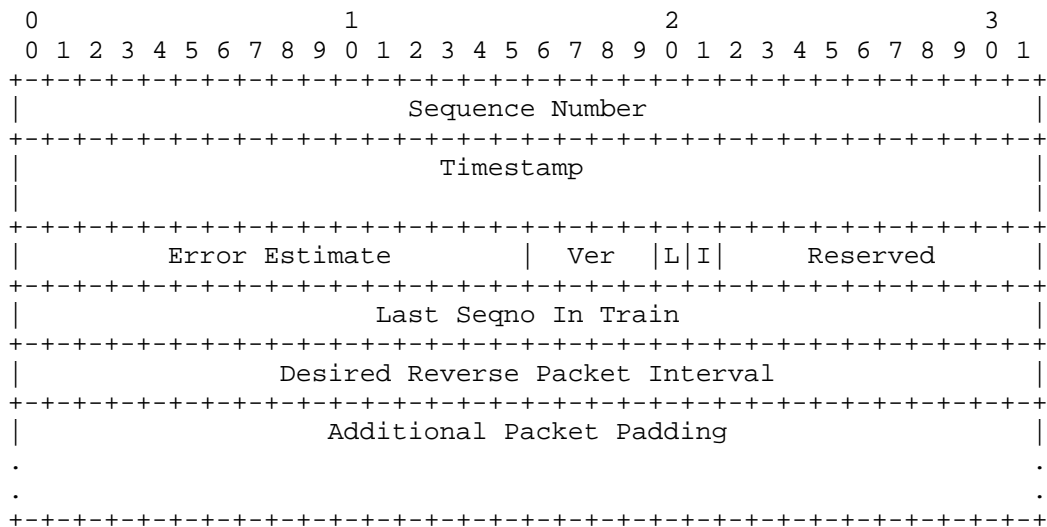
This feature allows the Session-Sender to set the first few octets in the TWAMP-Test Packet Padding with information to communicate value-added padding version number, flag bits, sequence number of the last packet in a train, and desired reverse packet interval (or per-packet waiting time) for the reverse path direction of transmission.

The Valued-Added Octets feature must be placed immediately after the TWAMP header or immediately after any new field that could be added to the TWAMP header or added to the beginning of the padding octets in the future. Therefore, the placement of the first bit from the valued-added octets depends on the mode(s) being selected.

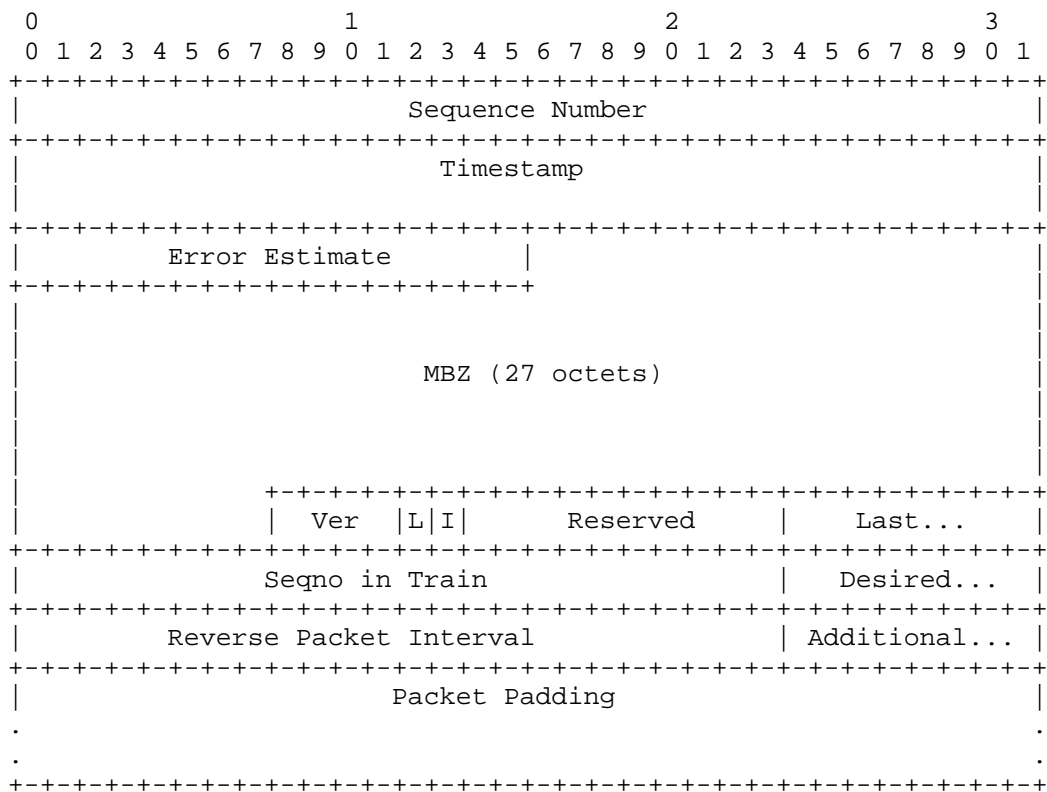
A version number and a sequence of flag bits are defined at the very beginning of the value-added padding octets. The version number identifies the version of the value-added padding octets and meaning of the flag bits and corresponding fields. Each flag bit indicates if a specific field is used in the valued-added padding octets. The version number and flag bits provide an effective method for extracting information at Session-Reflector and Session-Sender. This document defines version 1 with two flag bits: L and I.

The format of the test packet depends on the TWAMP modes. The Value-Added Octets Version 1 feature is intended to work with any TWAMP modes.

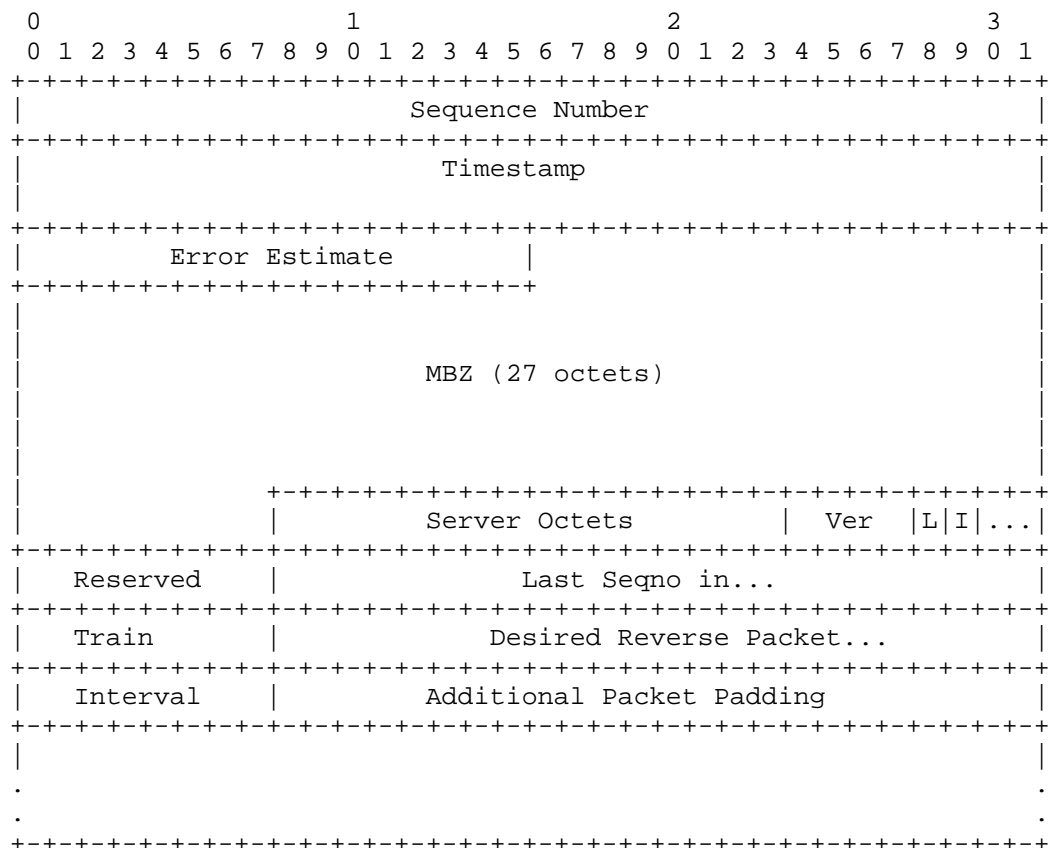
The Session-Sender SHALL use the following TWAMP-Test packet format when the Value-Added Octets Version 1 feature is selected in conjunction with the Unauthenticated mode:



The Session-Sender SHALL use the following TWAMP-Test packet format when the Value-Added Octets Version 1 feature is selected in conjunction with the Unauthenticated mode, Symmetrical Size mode, and Reflect Octets mode:



The Session-Sender SHALL use the following TWAMP-Test packet format when the Value-Added Octets Version 1 feature is selected in conjunction with the Unauthenticated mode, Symmetrical Size mode, and Reflect Octets mode with a non-zero value in the Server Octets field:



In the mode using Reflect Octets illustrated above, the value-added padding octets are embedded in the Packet Padding (to be reflected).

The Version (Ver) field MUST be encoded in the first 4 bits. It identifies the version number of the value-added padding octets and meaning of the flag bits and the corresponding fields. This memo defines version 1 with two flag bits: L and I. When the Value-Added Octets Version 1 feature is selected, the Session-Sender MUST set the Ver field to 1.

The 2 bits after the Version field are used for flags: L and I.

The Last Seqno in Train bit (L) is the first flag. When the Value-Added Octets Version 1 feature is selected, the Session-Sender MAY set the Last Seqno in Train bit L to 1.

The Desired Reverse Packet Interval bit (I) is the second flag. When the Value-Added Octets Version 1 feature is selected, the Session-Sender MAY set the Desired Reverse Packet Interval bit I to 1.

The Reserved field is reserved for future use. All 10 bits of the Reserved field MUST be transmitted as zero by the Session-Sender.

If the Last Seqno in Train bit is set to 1, then the Last Seqno in Train field MUST contain an unsigned 32-bit integer generated by the Session-Sender. It MUST indicate the expected sequence number of the last packet in the train. It SHOULD be used by the Session-Sender and Session-Reflector to identify the train to which a test packet belongs. The packets belonging to a train are determined by observing the test packet Sequence Number in relation to the Last Seqno in Train. The Last Seqno in Train MUST be higher or equal to Sequence Number of the packet. It must also be higher than the Last Seqno in Train for the previous train. If the L bit is set to 0, the Session-Sender shall set all the bits in the Last Seqno in Train field to zero.

If the Desired Reverse Packet Interval bit is set to 1, then the Desired Reverse Packet Interval field MUST contain an unsigned 32 bit integer generated by the Session-Sender. It MUST indicate the desired packet interval (or the waiting time) that the Session-Reflector SHOULD use when transmitting the reflected test packets towards the Session-Sender. The value 0 means the Session-Reflector SHOULD return the test packet to the Session-Sender as quickly as possible. The format of this field MUST be a fractional part of a second as defined in the One-Way Active Measurement Protocol (OWAMP) [RFC4656]. If the I bit is set to 0, the Session-Sender shall set all the bits in the Desired Reverse Packet Interval field to zero.

The values of the above fields are usually provided by a measurement method, tool, or algorithm. This measurement algorithm is outside the scope of this specification.

5.2. Reflector Behavior

The TWAMP Session-Reflector follows the procedures and guidelines in Section 4.2 of [RFC5357], with some changes and additional functions.

When the Value-Added Octets Version 1 feature is selected, the behavior of the Session-Reflector SHALL be as follows:

- o The Session-Reflector MUST read the Version field. If Ver = 1, the Session-Reflector MUST read the L and I flag bits.
- o If L=1 and I=1, the Session-Reflector MUST read and extract the information from the Last Seqno in Train field and the Desired Reverse Packet Interval field in the value-added padding octets.
 - The Last Seqno in Train MUST be compared to Sequence Number in the same packet in order to determine when a complete train has been collected. The Session-Reflector SHOULD buffer the packets belonging to the current train (or store the packet-level performance data). After the last packet of the train has been received, the Session-Reflector SHOULD transmit the packets belonging to a reverse train with a waiting time (packet interval) for each packet indicated in the Desired Reverse Packet Interval field. If the Desired Reverse Packet Interval field is set to zero, then the Session-Reflector SHOULD transmit the packet as quickly as possible. The last packet within a train has Sender Sequence Number = Last Seqno in Train.
 - The Last Seqno in Train of a packet MUST also be compared to the Last Seqno in Train of the previous packet in order to determine if a new train needs to be collected. In case of packet loss, the Session-Reflector MUST transmit the incomplete train when it receives a packet with a Last Seqno in Train belonging to another train (e.g., next train) of the test session or after a timeout. The timeout MAY be the REFWAIT timer specified in section 4.2 of [RFC5357].
 - Packets arriving out-of-order within a train MUST be buffered at the Session-Reflector if the train is not yet transmitted to the Session-Sender. If the train is already transmitted, the test packet SHOULD be returned to the Session-Sender as quickly as possible. The Session-Reflector MUST NOT reorder the test packets if they happen to arrive out-of-sequence.
 - Duplicate packets within a train MUST be buffered at the Session-Reflector if the train is not yet transmitted to the Session-Sender. If the train is already transmitted, the

duplicate test packet SHOULD be returned to the Session-Sender as quickly as possible. The Session-Reflector MUST NOT discard duplicate test packets.

For any other combinations of the Version field and the L and I flags, the Session-Reflector SHOULD return the test packet to the Session-Sender as quickly as possible.

The Session-Reflector MUST implement the changes described above when the Value-Added Octets Version 1 feature is selected.

5.2.1 Session-Reflector Packet Format

The Session-Reflector packet format follows the same procedure and guidelines as defined in TWAMP [RFC5357] and TWAMP Reflect Octets and Symmetrical Size Features [RFC6038], with the following changes:

- o The Session-Reflector MUST reuse (reflect) the value-added padding octets (10 octets) provided in the Sender's Packet Padding.
- o The Session-Reflector MAY reuse the rest of the padding octets in the Sender's Packet Padding.

The truncation process [RFC5357] is recommended when the Symmetrical mode is not used. The Session-Reflector MUST truncate exactly 27 octets of padding in Unauthenticated mode and exactly 56 octets in Authenticated and Encrypted modes.

5.3. Additional Considerations

The Session-Reflector supporting the Value-Added Octets feature should revert back to the standard Session-Reflector behavior if it cannot interpret the value-added padding octets in a given test packet. Section 5.2 also describes such behavior. For instance, the test packet is returned as quickly as possible to the Session-Sender when the Last Seqno in the Train is not what is expected.

Capacity measurements introduce an additional consideration when the test sessions operate in TWAMP Light. When the Session-Reflector does not have knowledge of the session state, the measurement system may be restricted to estimating or calculating the capacity metrics in the forward path direction of transmission only. Capacity measurements in the reverse path direction is best handled with a Session-Reflector supporting knowledge of the session state and being capable of identifying the test packets belonging to a specific test session. A method for creating a session state from the initial test packet may be implemented on the TWAMP Light Session-Reflector. This is outside the scope of this specification.

6. Experiments

This memo describes the protocol used in the current working prototype implementation of the Value-Added Octets feature in the Ericsson lab. The prototype has been tested in real network environments. The conclusion from these tests is that the Value-Added Octets feature is able to enable estimation of capacity metrics such as available path capacity in both the forward and reverse directions of the network path.

During the experiments with the protocol described in this memo, we have identified a need for the controller and responder to use the same maximum train length. The reflector must be able to buffer the whole train received from the controller. In order to reduce the risk for buffer overrun, the maximum train length should be negotiated. This can be resolved through configuration, introduction of a new field in the value-added octets, or a new maximum train length field in the Request-TW-Session message.

The Sender Discriminator (SD) field, which was proposed in an early draft of this document, was removed because of complications with different Session-Reflector implementations. A Session-Reflector may not be able to easily identify the SD field or associate it with a specific Session-Sender, which may skew the test results.

The flags defined in the value-added octets now indicate the usage of fields and not the presence of fields. This modification was needed to simplify the responder implementation in the working prototype.

7. Security Considerations

The value-added padding octets permit DoS attacks on the responder host communicating with core TWAMP [RFC5357]. For instance, a DoS condition could arise when the Last Seqno in Train is too large to handle, potentially causing undesirable processing delay or discard of the TWAMP-Test packets. The responder host MUST provide a mechanism to protect or limit the use of its local memory, buffer space, or maximum transmission time for a train.

The security considerations that apply to any active measurement of live networks are relevant here as well. See [RFC4656] and [RFC5357].

8. Acknowledgements

The authors thank Svante Ekelin for providing direction and comments on this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.
- [RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, February 2008.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.
- [RFC6038] Morton, A. and L. Ciavattone, "Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features", RFC 6038, October 2010.

9.2. Informative References

- [ENHJMMB] Ekelin, S., Nilsson, M., Hartikainen, E., Johnsson, A., Mangs, J., Melander, B., and M. Bjorkman, "Real-Time Measurement of End-to-End Available Bandwidth Using Kalman Filtering", Proceedings to the IEEE/IFIP Network Operations and Management Symposium, 2006.
- [PDM] Dovrolis, C., Ramanathan, P., and D. Moore, "Packet-Dispersion Techniques and a Capacity-Estimation Methodology", IEEE/ACM Transactions on Networking, December 2004.
- [RRBNC] Ribeiro, V., Riedi, R., Baraniuk, R., Navratil, J., and L. Cottrel, "pathChirp: Efficient Available Bandwidth Estimation for Network Paths", Passive and Active Monitoring Workshop, 2003.
- [SBW] Sommers, J., Barford, P., and W. Willinger, "Laboratory-based Calibration of Available Bandwidth Estimation Tools", Microprocessors and Microsystems, 2007.
- [Y1540] International Telecommunications Union, "Internet protocol data communication service - IP packet transfer and availability performance parameters", ITU-T Recommendation Y.1540, 2011.

Authors' Addresses

Steve Baillargeon
Ericsson
3500 Carling Avenue
Ottawa, Ontario K2H 8E9
Canada

EMail: steve.baillargeon@ericsson.com

Christofer Flinta
Ericsson
Farogatan 6
Stockholm, 164 80
Sweden

EMail: christofer.flinta@ericsson.com

Andreas Johnsson
Ericsson
Farogatan 6
Stockholm, 164 80
Sweden

EMail: andreas.a.johnsson@ericsson.com

