Network Working Group
INTERNET-DRAFT
Intended Status: Standards Track
Expires: March 11, 2012

S. Baillargeon
C. Flinta
A. Johnsson
S. Ekelin
Ericsson

September 8, 2011

TWAMP Value-Added Octets
draft-baillargeon-ippm-twamp-value-added-octets-02.txt

Abstract

This memo describes the optional extensions to the standard TWAMP test protocol for identifying test sessions and packet trains, and for measuring capacity metrics like the available path capacity, tight section capacity and UDP throughput in the forward and reverse path directions.

Status of this Memo

Copyright and License Notice

Table of Contents

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 feature for the Two-Way Active Measurement Protocol [RFC5357]. It is called the Value-Added Octets feature.

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 datagram delivery rate on both forward and reverse paths using a single test session. With this feature, it is also possible to setup multiple test sessions and improve the demultiplexing of test packets to the correct test sessions running on the controller and responder hosts when methods solely based on IP and UDP header information is not desirable or insufficient.

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 field [RFC5357] or at the beginning of the Packet Padding (to be reflected) field [RFC6038] by the Session-Sender, and are reflected or returned by the Session-Reflector. The length of the value-added octets varies in size between 6, 10 and 14 octets depending on the setting of the flag bits specified at the beginning of the value-added octets.

This memo is an update to the TWAMP core protocol specified in [RFC5357]. Measurement systems are not required to implement the feature described in this memo to claim compliance with [RFC5357].

IP-Type-P APC metric is defined in Defining Network Capacity [RFC5136]. IP-layer TSC metric is defined in IP Packet Transfer and Availability Performance Parameters [Y1540]. The actual method to calculate the available path capacity, the tight section capacity or the UDP datagram delivery rate from packet-level data performance data is not discussed in this memo.

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 define the OPTIONAL Valued-Added Octets feature for TWAMP [RFC5357].

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

- The extension of the modes of operation through assignment of a new value in the Mode field to communicate feature capability and use,

- The definition of a structure for embedding a sequence of value-added fields at the beginning of the Packet Padding field [RFC5037] or Packet Padding (to be reflected) field [RFC6038] in the TWAMP test packets and,

- The definition of new Session-Sender and Session-Reflector behaviors

The motivation for this feature is to enable the measurements of capacity metrics on both the forward and reverse paths, and to improve the demultiplexing of test packets to the correct test session at both endpoints.

This memo extends the modes of operation through assignment one new value in the Modes field (see Section 3.1 of [RFC4656] for the format of the Server Greeting message), while retaining backward compatibility with the core TWAMP [RFC5357] implementations. The new value correspond to the Valued-Added Octets Version 1 feature defined in this memo.

When the Server and Control-Client have agreed to use the Value-Added Octets Version 1 mode during control connection setup, then the Control-Client, the Server, the Session-Sender, and the Session-Reflector MUST all conform to the requirements of that mode, as identified below.

The OPTIONAL packet padding octets are designed to retain backward compatibility with the original TWAMP test protocol [RFC5357].

3 Capacity Measurement Principles

Most capacity estimation methods for available path capacity [RRBNC][PDM][ENHJMMB][SBW] and for UDP datagram 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 bi-directional 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 available path capacity, tight section capacity or UDP datagram delivery rate in the forward direction. Each train consists of a group of test packets which are separated from each other by a packet interval, as shown in the picture below.

```
      tt                      tt                      tt
|<---------->|          |<---------->|          |<---------->|
|            |          |            |          |            |
+------------+          +------------+          +------------+
|  Packet 1  |          |  Packet 2  |          |  Packet 3  |
+------------+          +------------+          +------------+
|                       |                       |
|<--------------------->|<--------------------->|
    packet interval 1       packet interval 2
```

The packet 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 interval MAY be constant within a train. Determination of the packet interval within a train as well as for consecutive trains for a given test session 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 picture. 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 available path capacity, tight section capacity and UDP datagram delivery rate use the packet intervals obtained from both end points in the estimation process. The method to measuring the UDP datagram delivery rate may also require the packet loss at the receiving end. 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 the available path capacity, tight section capacity and UDP datagram delivery rate, 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. The congestion is only transient, for the duration of the train which is typically short.

In order to fulfill the above measurement principles and to measure the available path capacity, tight section capacity and UDP datagram delivery rate in the reverse direction, the reflected test packets MUST be re-grouped into trains at the Session-Reflector.

4 Test packet Demultiplexing Principles

The controller (or the Session-Sender) requires a method for demultiplexing the received test packets to the correct test session especially when it manages multiple active test sessions. The responder also requires a method for demultiplexing the received test packets from multiple active test sessions originating from the same controller or from different controllers.

The purpose of this section is to provide some basic principles for identifying the test packets and to clarify the optional usage of the Sender Discriminator (SD) field described in this memo. It is important to note the actual method for identifying a test packet and the process for mapping it to the correct test session are implementation-specific. They may differ between various controllers and responders.

In general, the methods are based on fields available in the various headers of the TWAMP test packet (e.g. Ethernet, IP, UDP and TWAMP headers). Note the SID [RFC4656] is generally not used for identification purpose since it does not normally appear in the TWAMP test packets. As an example, a measurement system (controller or responder) may use the source IP address of the incoming test packet in order to associate it to the correct test session. This method is valid but has a number of limitations. It is simple and effective when each measurement system only requires a single test session for each peer but fails when multiple test sessions (with different characteristics) are running between the same pair of controller and responder.

Another approach is to use a combination of the source IP address, destination IP address, source UDP port and destination UDP port.

This method is also valid but to work effectively, it requires that the controller allocates multiple UDP ports (one for each test session for instance) and/or the responder listens on multiple ports. Ideally, a measurement system should limit the number of UDP ports for sending and receiving test packets. This approach may be improved by using a combination of the IP addresses, UDP ports and DSCP codepoint. This method also has its limitations. For instance, it cannot identify test packets from different test sessions running between the same pair of controller and responder if they are using the same UDP endpoints and the same DSCP codepoint.

This memo introduces a new field, the Sender Discriminator (SD) field intended to simplify the identification of the test packets at the controller and responder. It is especially useful when multiple test sessions with different DSCP codepoints and/or test packet sizes are expected to be running between the same pair of UDP endpoints. As described in 6.1.2, the SD is a number generated by the Session-Sender that uniquely identifies a test session on its system. With this field, the controller can explicitly identify the test packets belonging to a test session. When provided, the responder MAY use the SD field in combination of the source IP address for instance to identify the test packets belonging to a test session.

5 TWAMP Control Extensions

TWAMP-Control protocol [RFC5357] uses the Modes field to identify and select specific communication capabilities, and this field is a recognized extension mechanism.

TWAMP connection establishment follows the procedure defined in Section 3.1 of [RFC4656] and Section 3.1 of [RFC5357]. The new feature require 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. See the IANA section for details on the assigned value and bit position.

The Server sets the new bit position in the Modes field of the Server Greeting message to indicate its capability to operate in this new mode.

Both the Reflect Octets mode and Symmetrical Size mode SHOULD 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.

The forward and reverse APC, TSC and UDP datagram delivery rate measurement characteristics depend on the size of the test packets.

All test packets (forward and reverse test packets) belonging to a specific test session responsible to measure the available path capacity, tight section capacity and/or UDP datagram delivery rate MUST have the same IP packet size.

6 Extended TWAMP Test

The TWAMP-test protocol carrying the value-added padding octets is identical to TWAMP [RFC5357] except for the definition of first 6, 10 or 14 octets in the Padding Octet field that the Session-Sender expects to be reflected.

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

6.1 Sender Behavior

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

When the Value-Added Octets Version 1 mode is selected, the Session-Sender MAY set the Sender Discriminator Present bit to 1. If it is set to 1, the Session-Sender MUST generate and transmit a unique nonzero discriminator value in the Sender Discriminator field.

When the Value-Added Octets Version 1 mode is selected, the Session-Sender MAY set the Last Seqno in Train Present bit to 1. If it is set to 1, the Session-Sender MUST generate and transmit a valid sequence number in the Last Seqno in Train field. The Session-Sender MUST also group the test packets in trains and send the trains towards the Session-Reflector at the desired forward packet intervals.

When the Value-Added Octets Version 1 mode is selected, the Session-Sender MAY set the the Desired Reverse Packet Interval Present bit to 1. If it is set to 1, the Session-Sender MUST generate and transmit a valid inter-packet time interval in the Desired Reverse Packet Interval field.

The desired forward and reverse rate interval parameters are usually provided by a measurement method, tool or algorithm. This measurement algorithm is outside the scope of this specification.

6.1.1 Packet Timings

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

6.1.2 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 field with information to communicate value-added padding version number, flag bits, sender discriminator, sequence number of the last packet in a train and desired inter-packet time interval (or per-packet waiting time) for the reverse path direction of transmission.

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. Each flag bit indicates if a specific field of a specific size is present in the valued-added padding octets. The flag bits are designed to accommodate different combinations of fields while reducing padding overhead when certain fields are not needed. The version number and flag bits also provide an effective method for parsing information at Session-Reflector and Session-Sender. This document defines version 1 with three flag bits: S, L and D.

The format of the test packet depends on the TWAMP mode and flag bits being used. The Value-Added Octets Version 1 mode is intended to work with any TWAMP modes.

When the Value-Added Octets Version 1 is selected with S=1, L=0 and D=0, the Session-Sender SHALL use the following TWAMP test packet format in unauthenticated mode:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Sequence Number                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                          Timestamp                            |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |         Error Estimate        |  Ver  |S|L|D|   Reserved      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                     Sender Discriminator                      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                   Additional Packet Padding                   |
    .                                                               .
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

When the Value-Added Octets Version 1 is selected with S=0, L=1 and D=0, the Session-Sender SHALL use the following TWAMP test packet format in unauthenticated mode:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Sequence Number                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                          Timestamp                            |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |         Error Estimate        |  Ver  |S|L|D|   Reserved      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                      Last Seqno in Train                      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                   Additional Packet Padding                   |
    .                                                               .
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

When the Value-Added Octets Version 1 is selected with S=1, L=1 and D=0, the Session-Sender SHALL use the following TWAMP test packet format in unauthenticated mode:

```
        0                   1                   2                   3
        0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                        Sequence Number                        |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                          Timestamp                            |
       |                                                               |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |         Error Estimate        |  Ver  |S|L|D|   Reserved      |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                      Sender Discriminator                     |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                       Last Seqno in Train                     |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                   Additional Packet Padding                   |
       .                                                               .
       .                                                               .
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

When the Value-Added Octets Version 1 is selected with S=0, L=1 and D=1, the Session-Sender SHALL use the following TWAMP test packet format in unauthenticated mode:

```
        0                   1                   2                   3
        0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                        Sequence Number                        |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                          Timestamp                            |
       |                                                               |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |         Error Estimate        |  Ver  |S|L|D|   Reserved      |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                      Last Seqno In Train                      |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                Desired Reverse Packet Interval                |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                   Additional Packet Padding                   |
       .                                                               .
       .                                                               .
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

When the Value-Added Octets Version 1 is selected with S=1, L=1 and D=1, the Session-Sender SHALL use the following TWAMP test packet format in unauthenticated mode:

```
       0                   1                   2                   3
       0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                        Sequence Number                        |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                          Timestamp                            |
      |                                                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |         Error Estimate        |  Ver  |S|L|D|   Reserved      |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                      Sender Discriminator                     |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                      Last Seqno In Train                      |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                Desired Reverse Packet Interval                |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                   Additional Packet Padding                   |
      .                                                               .
      .                                                               .
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

When the Value-Added Octets Version 1 is selected with S=1, L=1 and

D=1, the Session-Sender SHALL use the following TWAMP test packet format in conjunction with the unauthenticated mode, Symmetrical Size mode and Reflect Octets mode:

```
       0                   1                   2                   3
       0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                        Sequence Number                        |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                          Timestamp                            |
      |                                                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |         Error Estimate        |                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
      |                                                               |
      |                                                               |
      |                         MBZ (27 octets)                       |
      |                                                               |
      |                                                               |
      |                                                               |
      |               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |               |  Ver  |S|L|D|    Reserved     |    Sender...  |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                Discriminator                  |     Last...   |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |               Seqno in Train                  |   Desired...  |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |         Reverse Packet Interval               | Additional... |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                       Packet Padding                          |
      .                                                               .
      .                                                               .
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

When the Value-Added Octets Version 1 is selected with S=1, L=1 and

D=1, the Session-Sender SHALL use the following TWAMP test packet format in conjunction with the unauthenticated mode, Symmetrical Size mode and Reflect Octets mode with a non-zero value in the Server octets field:

```
       0                   1                   2                   3
       0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                        Sequence Number                        |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                          Timestamp                            |
      |                                                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |         Error Estimate        |                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
      |                                                               |
      |                                                               |
      |                         MBZ (27 octets)                       |
      |                                                               |
      |                                                               |
      |                                                               |
      |               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |               |         Server octets         |  Ver  |S|L|D| |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |    Reserved   |             Sender Discriminator...           |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |               |              Last Seqno...                    |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |   In Train    |            Desired  Reverse Packet...         |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |   Interval    |              Additional Packet...             |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                              Padding                          |
      .                                                               .
      .                                                               .
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

In the combined mode including Reflect Octets, the value-added padding octets are embedded in the Packet Padding (to be reflected) field.

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. This document defines version 1.

The Sender Discriminator Present bit (S) MUST be the first flag. If the Sender Discriminator Present bit is set to 1, then a Sender Discriminator field MUST be present and MUST contain valid information.

The Last Seqno in Train Present bit (L) MUST be the second flag. If the Last Seqno in Train Present bit is set to 1, then the Last Seqno in Train field MUST be present and MUST contain valid information.

The Desired Reverse Packet Interval Present bit (D) MUST be the third flag. If the Desired Reverse Packet Interval Present bit is set to 1, then Desired Reverse Packet Interval Present field MUST be present and MUST contain valid information.

The Reserved field is reserved for future use. All 9 bits of the Reserved field MUST be transmitted as zero by the Session-Sender.

The Sender Discriminator (SD) field MUST contain an unsigned 32 bit integer generated by the Session-Sender. It is used by the Session-Reflector and/or Session-Sender to identify packets belonging to a test session. The Session-Sender MUST choose a nonzero discriminator value that is unique among all test sessions on its system. This field is present only if the Sender Discriminator Present bit is set to one.

The Last Seqno in Train 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 a test packet belongs to. The packets belonging to a train are determined by observing the test packet sequence number in relation to the Last Seqno for a train. The sequence number of a packet in a train MUST be lower than or equal to the Last Seqno for that train. The sequence number MUST also be larger than the Last Seqno for the previous train. This field is present only if the Last Seqno in Train Present bit is set to one.

The Desired Reverse Packet Interval (DRPI) MUST contain an unsigned 32 bit integer generated by the Session-Sender. It MUST indicate the desired inter-packet time 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 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 OWAMP [RFC4656]. This field is present only if the Desired Reverse Packet Interval Present bit is set to one.

The method by which the Sender Discriminator and Desired Reverse Packet Interval values are obtained is outside of the scope of this document.

6.2 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 is selected, the behavior of the Session-Reflector SHALL be as follows:

- The Session-Reflector MUST read the Version field. If Ver=1, the Session-Reflector MUST read the S, L and D flag bits. If Ver is not equal 1, the Session-Reflector MUST ignore the rest of the value-added padding octets and MUST follow the procedures and guidelines described in section 4.2 of [RFC5357]. The Session-Reflector SHOULD transmit the packet as quickly as possible including the test packets that are currently stored for the test session.

- If S=0, L=0 and D=0, the Session-Reflector MUST ignore the rest of the value-added padding octets and MUST follow the procedures and guidelines described in section 4.2 of [RFC5357]. The Session-Reflector SHOULD transmit the packet as quickly as possible including the test packets that are currently stored for the test session.

- If S=1, the Session-Reflector MUST continue reading and extracting the information from the Sender Discriminator field in the value-added padding octets.

- After reading and extracting the information from the Sender Discriminator field, the Session-Reflector SHOULD associate the test packets to the correct test session based on the value specified in the Sender Discriminator field and the source IP address specified in the IP header of the test packet. The actual method for demultiplexing the received test packets to the correct test session based on the Sender Discriminator and source IP address is outside the scope of this specification. The Session-Reflector MAY also use additional packet fields to demultiplex test packets to a test session.

- If L=1, the Session-Reflector MUST continue reading and extracting the information from the Last Seqno in Train field in the value-added padding octets.

- After reading and extracting the information from the Last Seqno in Train field, Last Seqno in Train field 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) and SHOULD transmit them as immediately as possible after the last packet of the train has been received. 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 the 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.

- If D=1, the Session-Reflector MUST continue reading and extracting the information from the Desired Reverse Packet Interval field in the value-added padding octets.

- After reading and extracting the information from the Desired Reverse Packet Interval field, 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 packets as quickly as possible.

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

6.2.1 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 re-use (reflect) the value-added padding octets (6, 10 or 14 octets) provided in the Sender's Packet Padding.

- o The Session-Reflector MAY re-use the rest of the padding octets in the Sender's Packet Padding.

When using the recommended truncation process [RFC5357], the Session-Reflector MUST truncate exactly 27 octets of padding in Unauthenticated mode,and exactly 56 octets in Authenticated and Encrypted modes.

6.3 Additional Considerations

It is not required to use the Sender Discriminator field for calculating the capacity metrics. The Sender Discriminator Present bit can be set to zero. However, the Session-Sender and Session-Reflector MUST implement a local policy to identify the test packets belonging to a specific test session. The method for demultiplexing the received test packets to the correct test session based on other packet fields (e.g. fields in the IP header) is outside the scope of this specification.

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 will only be capable to estimate or calculate the capacity metrics in the forward path direction of transmission. Capacity measurements in the reverse path direction requires the Session-Reflector to have knowledge of the session state and be capable to identify the test packets belonging to a specific test session. The method for creating a session state from the initial test packets on the TWAMP Light Session-Reflector is outside the scope of this specification.

7 Security Considerations

The value-added padding octets permit new attacks on the responder host communicating with core TWAMP [RFC5357]. The responder host MUST provide a mechanism to protect or limit the use of its local memory or buffer space.

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

8 IANA Considerations

This memo adds one mode to the IANA registry for the TWAMP Modes field, and describes behavior when the new modes are used. This field is a recognized extension mechanism for TWAMP.

8.1. Registry Specification

IANA has created a TWAMP-Modes registry (as requested in [RFC5618]). TWAMP-Modes are specified in TWAMP Server Greeting messages and Setup Response messages, as described in Section 3.1 of [RFC5357], consistent with Section 3.1 of [RFC4656], and extended by this memo. Modes are indicated by setting bits in the 32-bit Modes field that correspond to values in the Modes registry. For the TWAMP-Modes registry, we expect that new features will be assigned increasing registry values that correspond to single bit positions, unless there is a good reason to do otherwise (more complex encoding than single-bit positions may be used in the future to access the 2^32 value space).

8.2. Registry Contents

The TWAMP-Modes registry has been augmented as follows:

Value	Description	Semantics Definition
128	Valued-Added Octets Ver 1	This memo, Section 2 new bit position (7)

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 Ishac,J., "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., Ciavattone, L., TWAMP Reflect Octets and Symmetrical Size Features, RFC6038 , October 2010.

9.2 Informative References

[RRBNC] Ribeiro, V., Riedi, R., Baraniuk, R., Navratil, J., Cottrel, L., Pathchirp: Efficient available bandwidth estimation for network paths, Passive and Active Measurement Workshop, 2003.

[PDM] Dovrolis, C., Ramanathan, P., and Moore D., Packet Dispersion Techniques and a Capacity Estimation Methodology, IEEE/ACM Transactions on Networking, December 2004.

[ENHJMMB] Ekelin, S., Nilsson, M., Hartikainen, E., Johnsson, A., Mangs, J., Melander, B., Bjorkman, M., Real-time measurement of end-to-end available bandwidth using kalman filtering, Proceedings to the IEEE IFIP Network Operations and Management Symposium, 2006.

[SBW] Sommers, J., Barford, P., Willinger, W., Laboratory-based calibration of available bandwidth estimation tools, Microprocess Microsyst., 2007.

[Y1540] ITU-T Y.1540, Internet protocol data communication service - IP packet transfer and availability performance parameters, 2011.

[MRM] Morton, A., Ramachandran, G., Maguluri, G., Reporting Metrics Different Points of View, draft-ietf-ippm-reporting-metrics-03, June 2010.

Author's 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

Svante Ekelin
Ericsson
Farogatan 6
Stockholm, 164 80
Sweden
EMail: svante.ekelin@ericsson.com