IPPM Considerations for the IPv6 PDM Destination Option

Nalini Elkins – Inside Products, Inc.

We propose:

Requirement

- In basic IP transport
- Undisturbed by middle systems

Solution

- **Implementation** of existing extension header: Destination Options Header (DOH)
- Performance and Diagnostic Metrics (PDM) DOH

PDM

- Performance and Diagnostic Metrics Destination Option (PDM) contains the following fields: (by 5-tuple)

- PSNTP : Packet Sequence Number This Packet
- PSNLR : Packet Sequence Number Last Received
- DELTALR : Delta Last Received
- DELTALS : Delta Last Sent
- TIMEBASE : Base timer unit
- SCALEDL : Scale for Delta Last Received
- SCALEDS : Scale for Delta Last Sent

PDM Timing

- No time synchronization needed
- All times are in relation to self

Start Flow

- Packet 1 is sent from Host A to Host B. The time for Host A is 10:00AM.

- The time and packet sequence number are saved by Host A internally. The packet sequence number and delta times are sent in the packet.

Packet 1

Packet 1

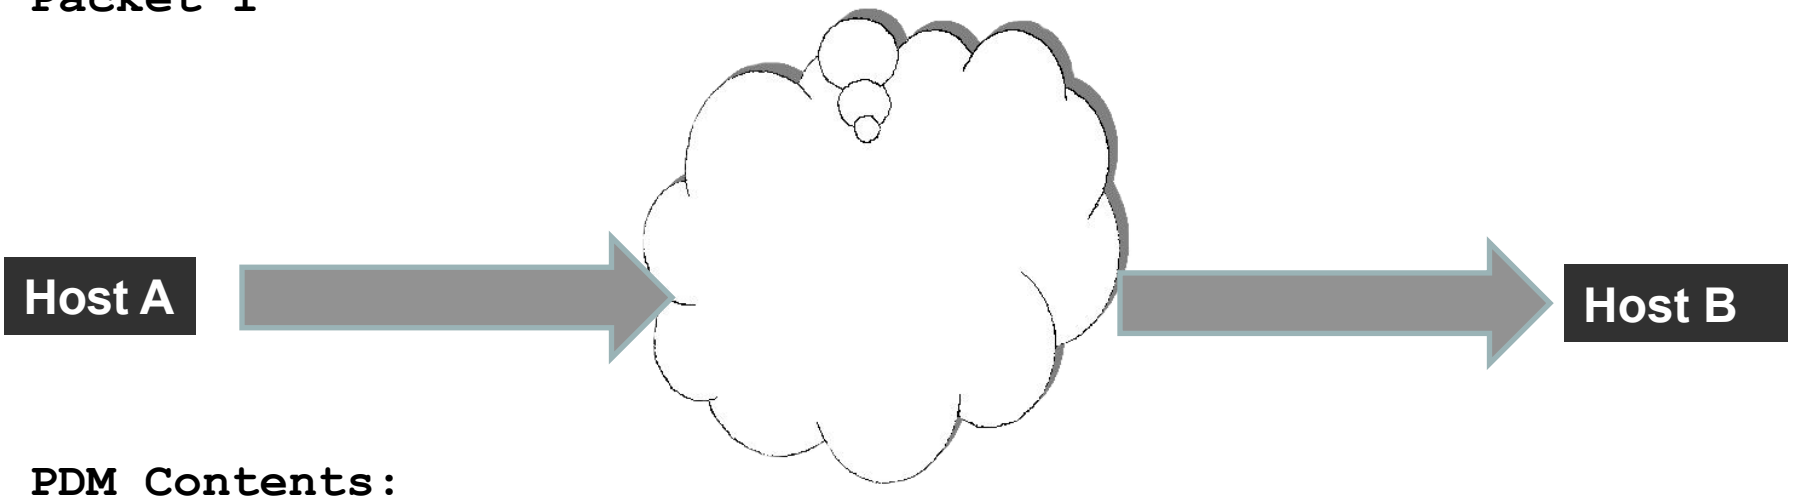

PDM Contents:

PSNTP	**:**	**Packet Sequence Number This Packet:**	**25**
PSNLR	**:**	**Packet Sequence Number Last Received:**	**-**
DELTALR	**:**	**Delta Last Received:**	**-**
DELTALS	**:**	**Delta Last Sent:**	**-**

Keep in Host A

- Internally, within the sender, Host A, it must keep:

- Packet Seq. Number of last packet sent: 25
- Time the last packet was sent: 10:00:00

Keep in Host B

- Packet 1 is received at Host B. Its time is set to one hour later than Host A. In this case, 11:00AM
- Internally, within the receiver, Host B, it must note:

- Packet Seq. Number of last packet received: 25
- Time the last packet was received : 11:00:03

Server Delay

- Host B processes packet 1 and creates a response (packet 2).
- Packet 2 is sent by Host B to Host A.
- This is the time taken by Host B or Server Delay
- Server Delay = Sending time (packet 2) - receive time (packet 1)

DeltaLR

- We will call the result of this calculation: Delta Last Received
- DELTALR = Sending time (packet 2) - receive time (packet 1)
- Note, both sending time and receive time are saved internally in Host B. They do not travel in the packet. Only the Delta is in the packet.

Host B Stats

- Within Host B is the following:

- Packet Sequence Number of the last packet received: 25
- Time the last packet was received: 11:00:03
- Packet Sequence Number of this packet: 12
- Time this packet is being sent: 11:00:07

- DELTALR = 4 seconds (11:00:07 - 11:00:03)
- DELTALR is Server Delay.

Packet 2

Packet 2

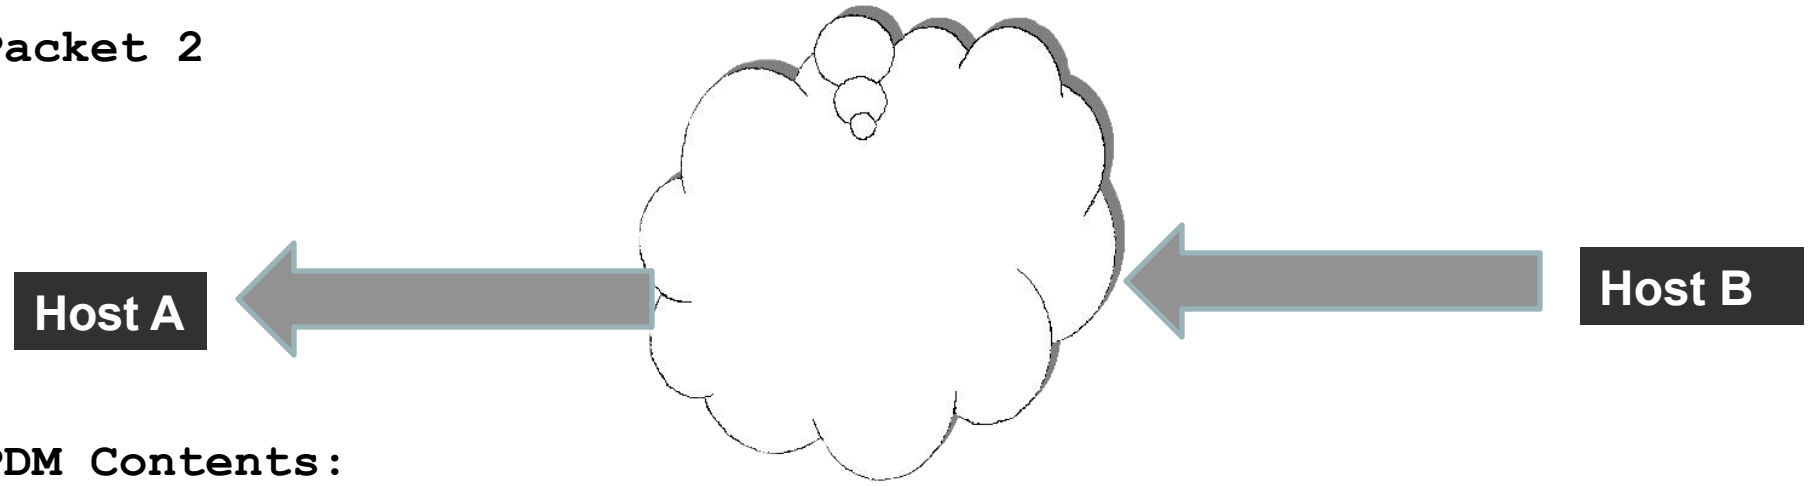

PDM Contents:

PSNTP	:	Packet Sequence Number This Packet:	12
PSNLR	:	Packet Sequence Number Last Received:	25
DELTALR	:	Delta Last Received:	4 seconds
DELTALS	:	Delta Last Sent:	-

Metrics Needed

- The metrics left to be calculated are end-to-end time and round-trip delay (network time).
- This will be calculated by Host A when it receives Packet 2.

Packet 2 Received

- Packet 2 is received at Host A. Remember, its time is set to one hour earlier than Host B. Internally, it must note:

- Packet Sequence Number of the last packet received: 12
- Time the last packet was received : 10:00:12

- Note, this timestamp is in Host A time. It has nothing whatsoever to do with Host B time.

End-to-End Time

- Now, Host A can calculate total end-to-end time.
- End-to-End Time = Time Last Received - Time Last Sent
- Packet 1 was sent by Host A at 10:00:00. Packet 2 was received by Host A at 10:00:12
- End-to-End time = 10:00:12 - 10:00:00 or 12

- This metric we will call DELTALS or Delta Last Sent

Network TIme

- We can now also calculate round trip delay (network time). The formula is:
- Round trip delay = DELTALS - DELTALR
- Or: End-to-end time – Server Delay
- Round trip delay = 12 - 4 or 8

How to Communicate?

- Now, the only problem is that at this point all metrics are in Host A only and not exposed in a packet.
- To do that, we need a third packet.

Packet 3

Packet 3

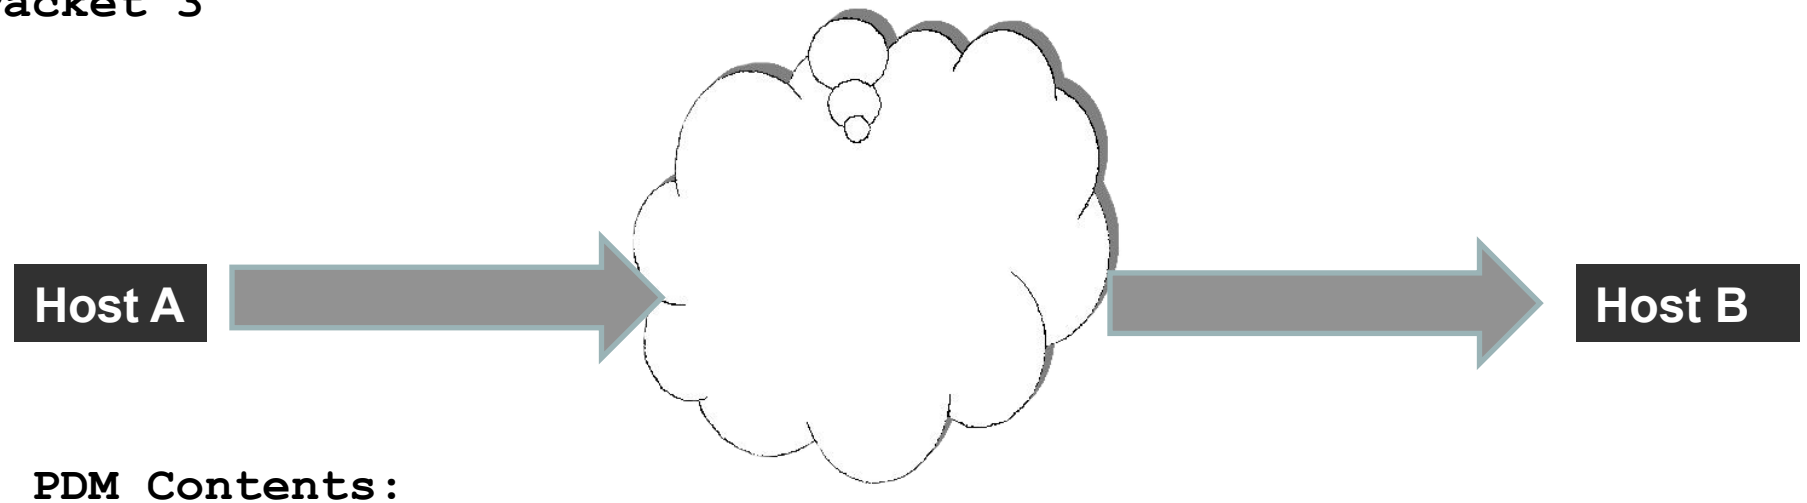

PDM Contents:

PSNTP	: Packet Sequence Number This Packet:	26	
PSNLR	: Packet Sequence Number Last Received:	12	
DELTALR	: Delta Last Received:	0	
DELTALS	: Delta Last Sent:	12	seconds

Show IPv6 PDM Type 16 Header
Using: Trace File:pdm16
Sort Order : Packet Number

-	-	Packet Number	Packet Date	Extension Header	Source Address	Destination Address	This Packet ID	Packet Last Received	Delta Last Received (Microseconds)	Delta Last Sent (Microseconds)
1		4	2014-01-10 13:30:22.857512	60	2001::2	2001::1	0	0	0	0
2		5	2014-01-10 13:30:22.860452	60	2001::1	2001::2	0	0	0	0
3		6	2014-01-10 13:30:23.865714	60	2001::2	2001::1	1	0	1006	983
4		7	2014-01-10 13:30:23.877588	60	2001::1	2001::2	1	0	1017	1017
5		8	2014-01-10 13:30:24.870476	60	2001::2	2001::1	2	1	1008	974
6		9	2014-01-10 13:30:24.871949	60	2001::1	2001::2	2	1	994	994
7		13	2014-01-10 13:30:25.879201	60	2001::2	2001::1	3	2	1005	995
8		14	2014-01-10 13:30:25.88565	60	2001::1	2001::2	3	2	1013	1013
9		17	2014-01-10 13:30:26.886962	60	2001::2	2001::1	4	3	1008	985
10		18	2014-01-10 13:30:26.897091	60	2001::1	2001::2	4	3	1011	1011
11		19	2014-01-10 13:30:27.891001	60	2001::2	2001::1	5	4	1007	974
12		20	2014-01-10 13:30:27.901722	60	2001::1	2001::2	5	4	1004	1004
13		27	2014-01-10 13:30:28.894605	60	2001::2	2001::1	6	5	1004	982
14		28	2014-01-10 13:30:28.905136	60	2001::1	2001::2	6	5	1003	1003

```
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: 2001::2 (2001::2), Dst: 2001::1 (2001::1)
    0110 .... = Version: 6
        [0110 .... = This field makes the filter "ip.version == 6" possible: 6]
    .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00000000
        .... 0000 00.. .... .... .... .... .... = Differentiated Services Field: Default (0x00000000)
        .... .... ..0. .... .... .... .... .... = ECN-Capable Transport (ECT): Not set
        .... .... ...0 .... .... .... .... .... = ECN-CE: Not set
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 71
    Next header: IPv6 destination option (60)
    Hop limit: 64
    Source: 2001::2 (2001::2)
    [Source Teredo Server IPv4: 0.0.0.0 (0.0.0.0)]
    [Source Teredo Port: 65535]
    [Source Teredo Client IPv4: 255.255.255.253 (255.255.255.253)]
    Destination: 2001::1 (2001::1)
    [Destination Teredo Server IPv4: 0.0.0.0 (0.0.0.0)]
    [Destination Teredo Port: 65535]
    [Destination Teredo Client IPv4: 255.255.255.254 (255.255.255.254)]
    [Source GeoIP: Unknown]
    [Destination GeoIP: Unknown]
    Destination Option
        Next header: TCP (6)
        Length: 2 (24 bytes)
        IPv6 Option (PadN)
            Type: PadN (1)
            Length: 6
            PadN: 000000000000
        IPv6 Option (Unknown 16)
            Type: Unknown (16)
            Length: 14
            Unknown Option Payload: 000103000003ee000003d700b326
Transmission Control Protocol, Src Port: 45862 (45862), Dst Port: 80 (80), Seq: 452248946, Len: 7
    Source port: 45862 (45862)
    Destination port: 80 (80)
    [Stream index: 2]
```

Breakout in WireShark

Timebase

- Possible values of Time Base:

- 00 - milliseconds
- 01 - microseconds
- 10 - nanoseconds
- 11 - picoseconds

Scale (DLR / DLS)

- 7-bit signed integer.
- Possible values from -64 to +63.
- Store most significant bits of timer value along with a scaling factor to indicate the magnitude.
- High-order 16 bits.