draft-ietf-ippm-6man-pdm-option-01
IPv6 PDM Destination Option

Nalini Elkins – Inside Products, Inc.
Mike Ackermann – BCBS Michigan
Rob Hamilton – Chemical Abstracts

Comments from Prague

We have added two new sections to address the comments from the WG session in Prague:

1. Comment: Indicate that PDM will not work when using IPv6 transition technologies

Action: Section 1.6 IPv6 Transition Technologies added.

2. Comment: Indicate that PDM must be placed BEFORE the ESP header.

Action: Section 3.3 "Header Placement" revised. New section 3.4 "Header Placement Using IPSec ESP Mode" added to further clarify header placement.

Testing

Used our implementation on FreeBSD

- Same subnet (cable) ✓
- Same administrative domain (TBD)
- Different administrative domains ✓
- Internet ✓

Implementation on Internet

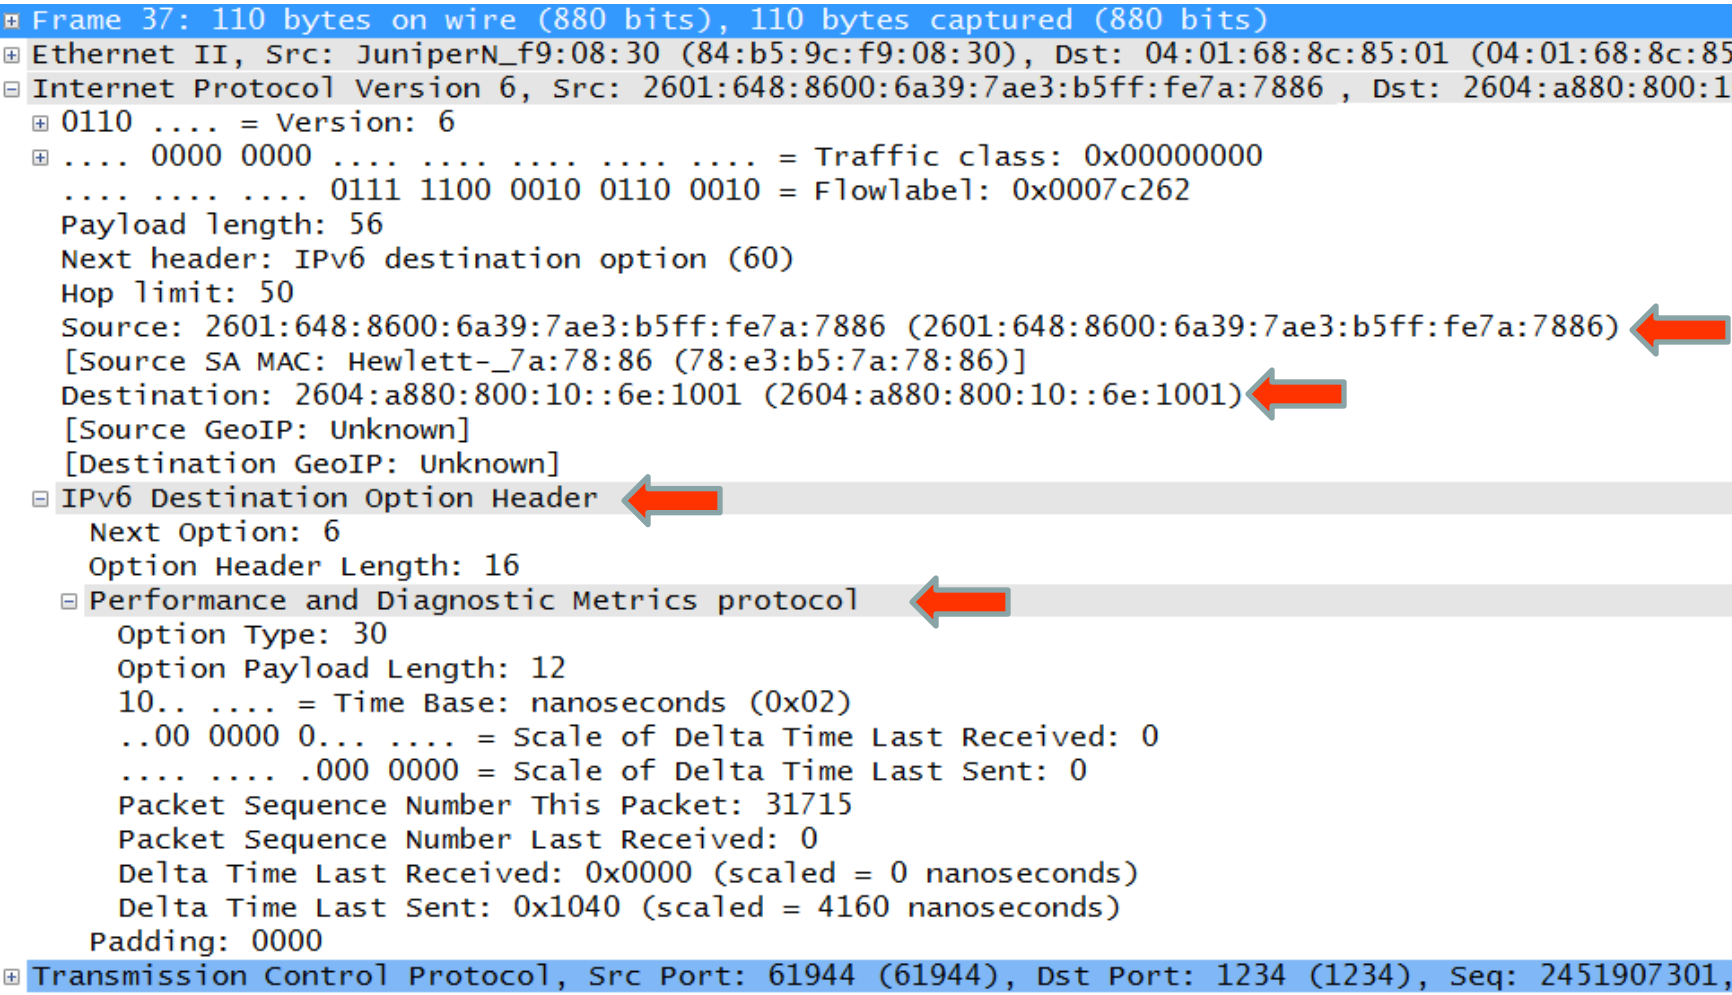

Geolocate Addresses

2601:648:8600:6A39:7AE3:B5FF:FE7A:7886	Comcast Cable	Martinez	California	United States
2604:A880:800:10::6E:1001	Digital Ocean	New York	New York	United States

- Obviously need more data points
- Working on that
- Fighting with VMs

Implementation on Stacks

- Request for Enhancement (RFE) submitted to IBM by large U.S. corporation
- Discussions held with IBM TCP/IP Chief Architect

Issues: Control Blocks

- What is in control blocks today (IP / TCP)
  - TCP CB do not know IP address
  - IP CB do not know other end IP address. Do not know port
  - Netstat commands have all info (see 5-tuple below) clearly that is in some control block

Output of Netstat –A

TCP	10.0.0.3:52987	67.217.64.244:https	TIME_WAIT
TCP	10.0.0.3:52988	54-249-66-39:https	TIME_WAIT
TCP	10.0.0.3:52989	67.217.64.244:https	TIME_WAIT

Issues : Seq Number Calculation

- How is sequence number for IPv4 (IPID) calculated?
  - Some do global counter
  - Some do counter per 5-tuple
  - For the stacks who do global counter, this will mean a big change

Issues: API

- Should (new) API be provided?
- Where does code to do PDM stats really belong?
- Our current proof-of-concept implementation intercepts each packet at interface

Issues: IPSec Diagnostics

- This is a big problem for users
- PDM may be a big help
- PDM Destination Option travels in the clear, even when using ESP mode

MTU Discussion: Need to Add

- Potential problem: “Packets become too large when adding the PDM header and results in <IPv6-fragementation-required> to the sending host” - Joachim

- Potential problem: “Size increase with PDM header makes stream exceed a network threshold and trigger channel capacity re-allocation” - Joachim

- Add caveat: When using hybrid modes, it becomes critical to not trigger such network events by careful implementation and planning. One thing that I have seen network operators do, when they know that they may have extra headers potentially added, is to "leave room". For example, send a packet with data payload of 1,430 rather than 1,480. With a packet that has a payload of 1,480 on a network with a 1,500 MTU, then just about anything you add is going to lead to fragmentation.

Comments?

- Thoughts?
- Issues?
- Questions?