draft-ietf-ippm-6man-pdm-option-00
IPv6 PDM Destination Option

Nalini Elkins – Inside Products, Inc.

Mike Ackermann – BCBS Michigan

Rob Hamilton – Chemical Abstracts

Agenda

1. Combined layout with IPPM considerations
2. WG adoption call on list
3. Implementation on FreeBSD / Linux
4. Security considerations
5. Wireshark dissector (LUA)
6. IPSec testing

Implementation

- NETMAP / VALE
  - part of standard FreeBSD distributions
  - single, non intrusive kernel module
  - http://info.iet.unipi.it/~luigi/netmap/

- framework for high speed packet I/O. Implemented as a kernel module for FreeBSD and Linux, it supports access to network cards (NICs), host stack, virtual ports (the "VALE" switch), and "netmap pipes". netmap can easily do line rate on 10G NICs (14.88 Mpps), moves over 20 Mpps on **VALE** ports, and over 100 Mpps on netmap pipes.

- Using 0x1E option type (experimental)

Implementation Findings

- Possible issues with checksum offload. Testing.
- Need to check on fragmentation issues.
- Have tested with local FreeBSD machines. Started testing across Internet.
- Security considerations
  - Resource consumption attacks (SYN flood protection)
  - MITM considerations

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

Wireshark Dissector

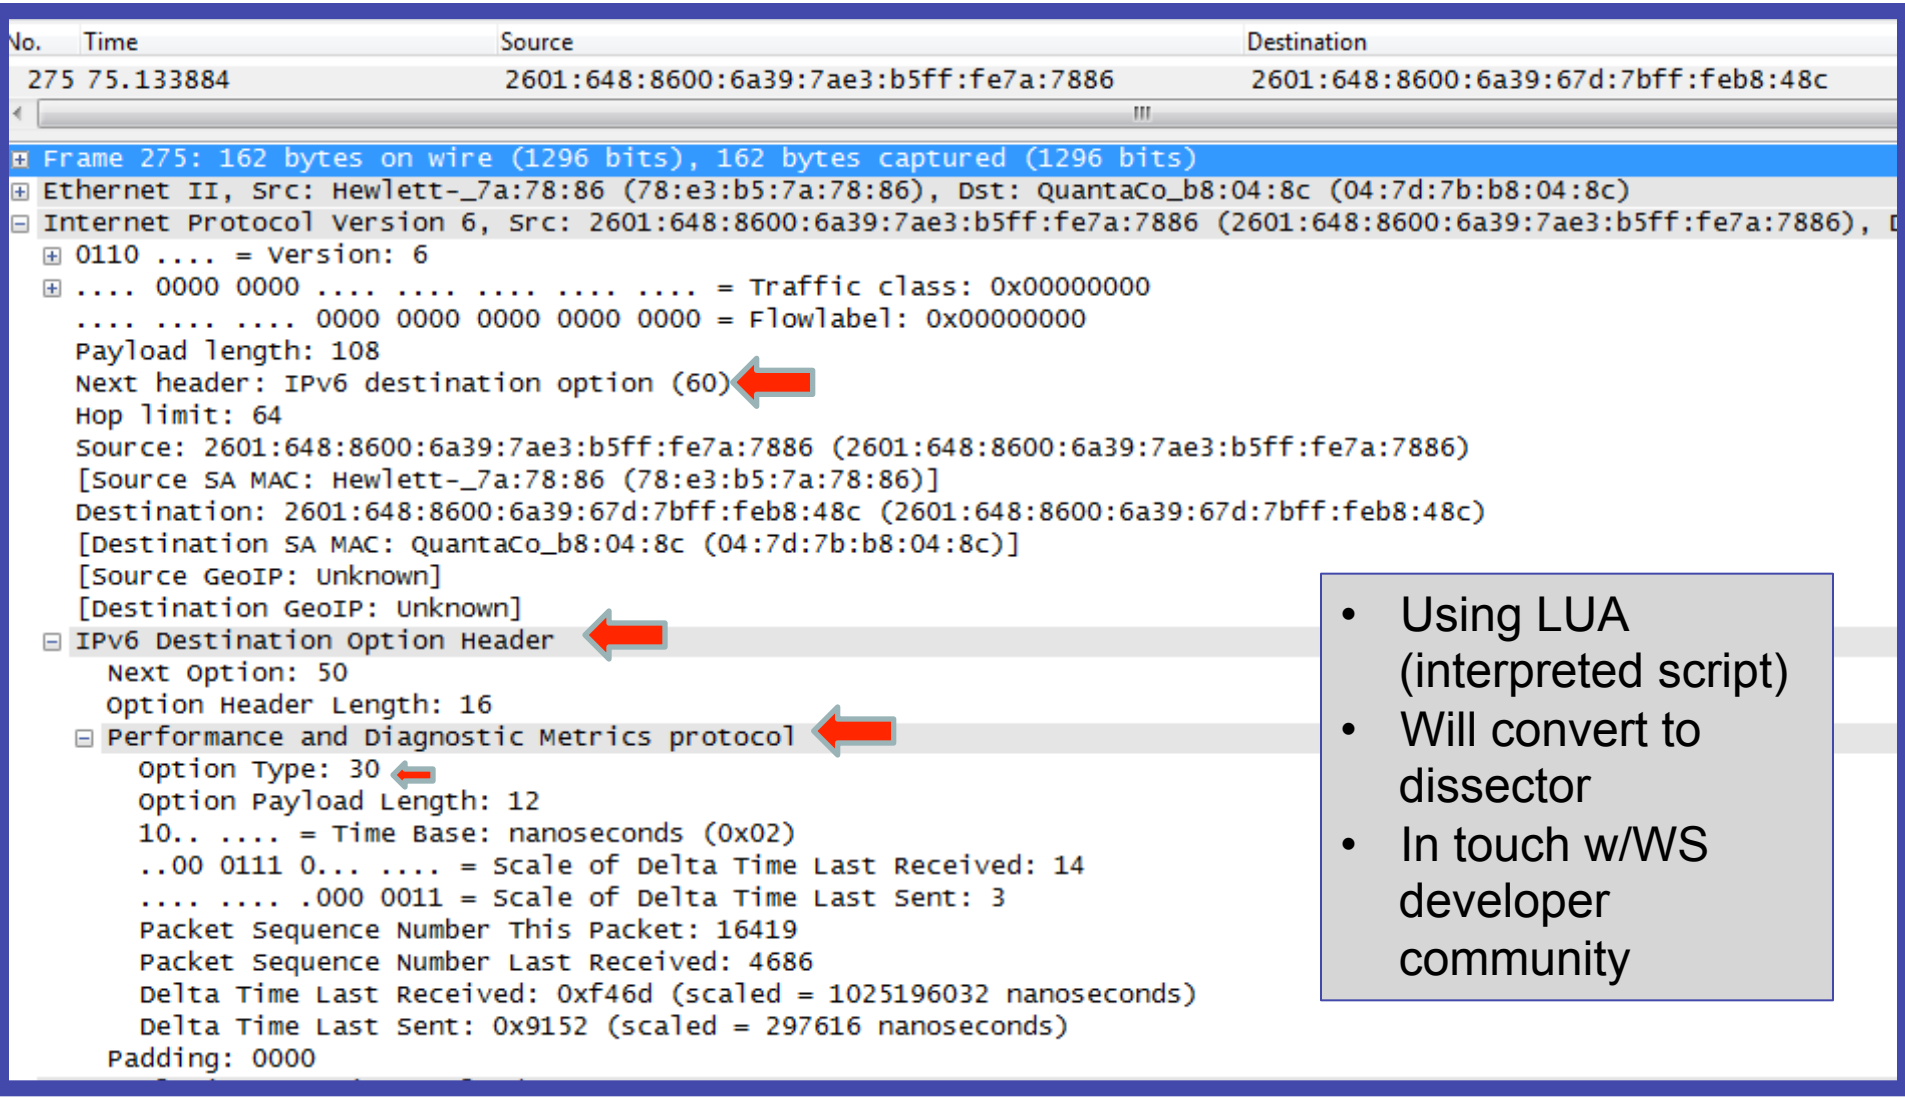

- Using LUA (interpreted script)
- Will convert to dissector
- In touch w/WS developer community

IPSec Testing

- Will PDM be available under ESP?
- From RFC4303 : IP Encapsulating Security Payload (ESP) section 3.1.1 Transport Mode processing

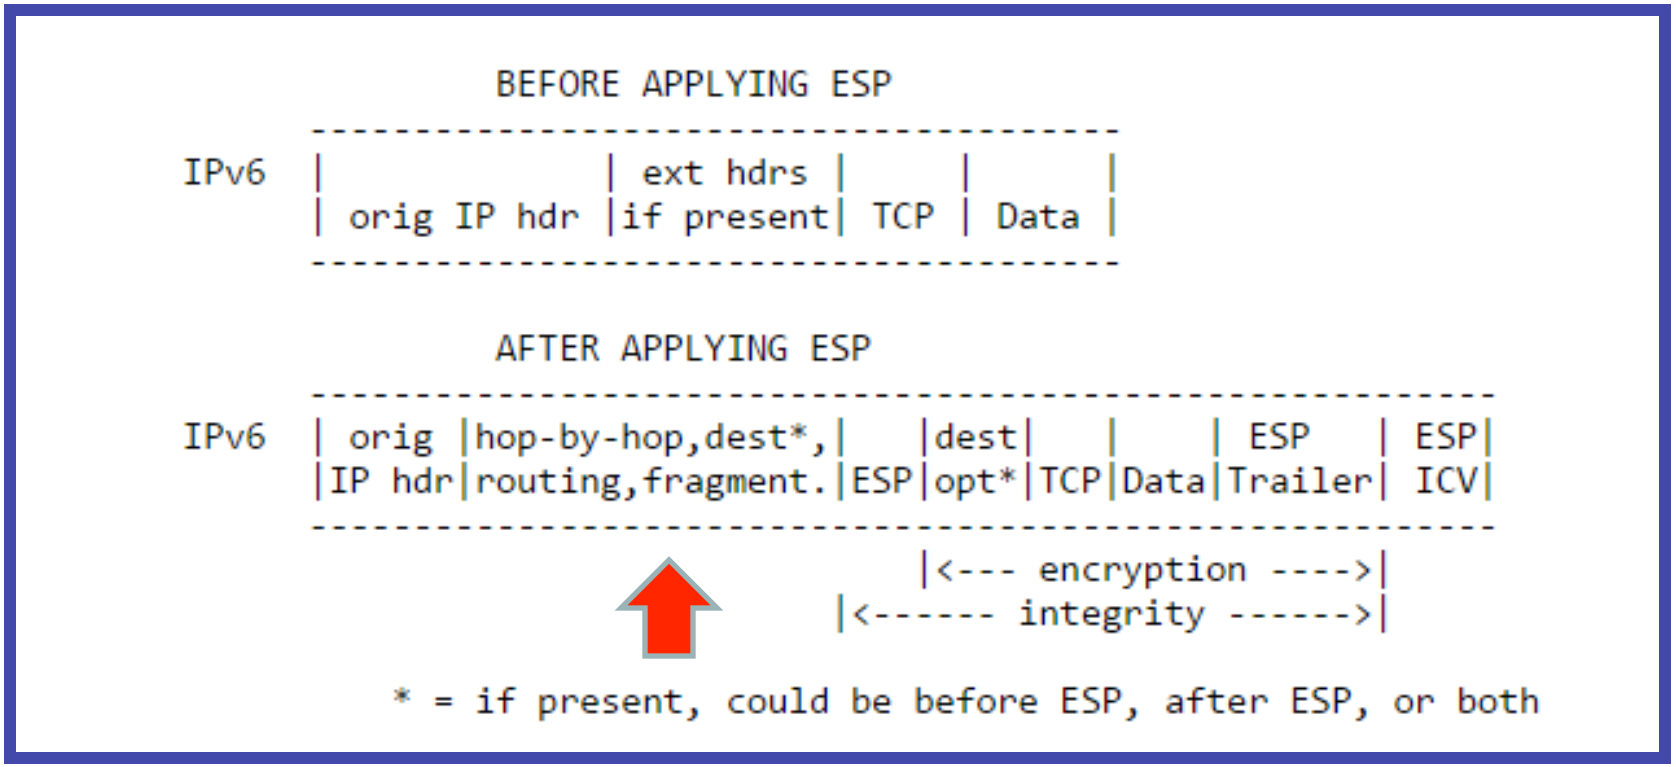

ESP Packet

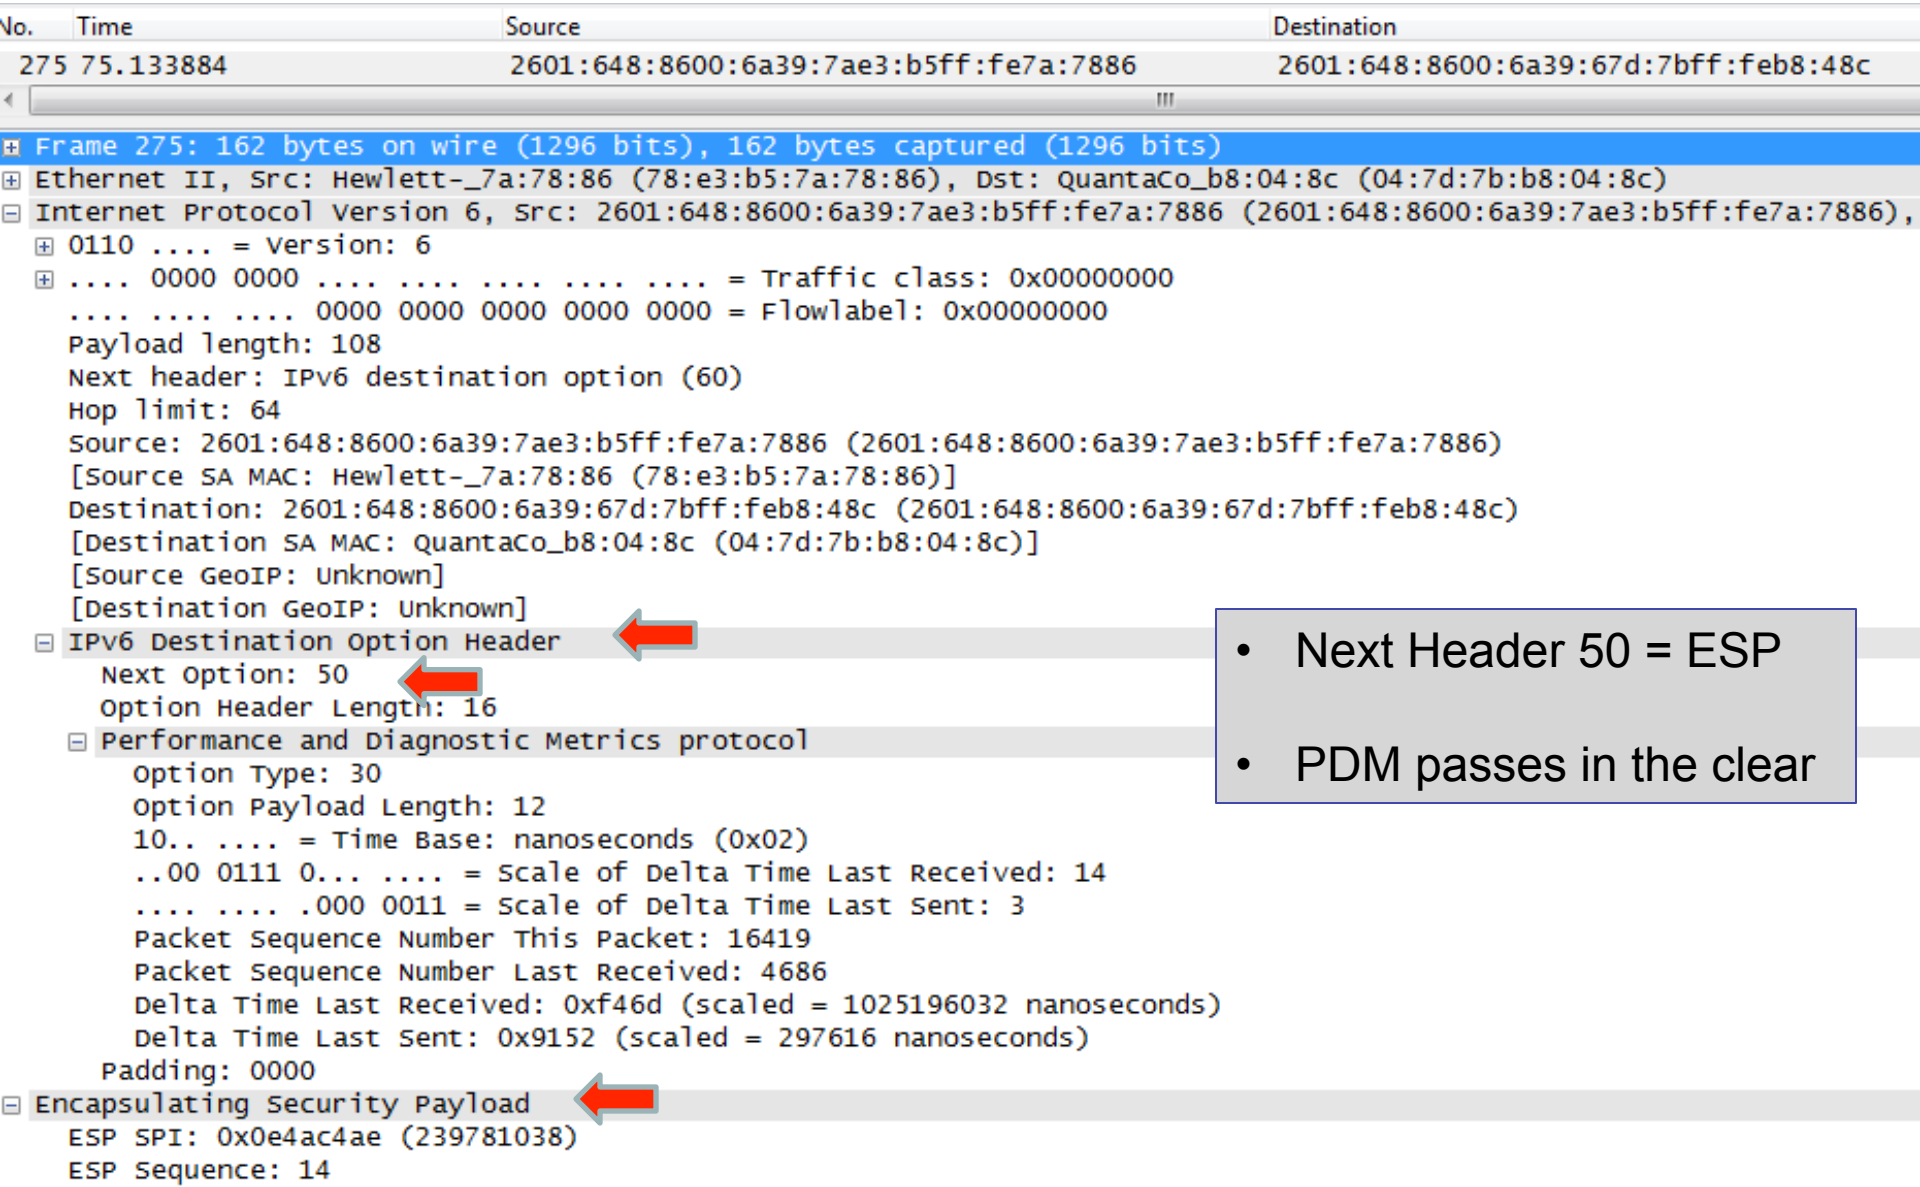

- Next Header 50 = ESP
- PDM passes in the clear

What does this buy us?

- Packet loss (is other side getting what we are sending)
- Out-of-order segments
- Hangs (can be deduced by exponential backoff)
- Timing

To Do

- Clean implementation
- Windows? IBM mainframe? Cell phones?
- Better analysis (pattern of flow)

Comments?

- Thoughts?
- Issues?
- Questions?