

THOTCON 0x1 – Where did those damn packets go?

Kevin A. Nassery <kevin@nassery.org>

0110110111
1001.0010.001



Today's Challenges

- Vanishing Perimeters
- Crazy Bandwidth
- Asymmetric and Redundant Paths
- Analyzer performance

How do Data Access Networks solve some of these issues?

0110110111

1001.0010.001



DAN

- Flexibility over traditional taps, port mirrors, aggregation and regeneration.
- Performance advantages of being able to filter traffic before it gets to your tools.
- Avoids the complexity of distributed sniffers and RSPAN.
- Easy to use, and to manage remotely.
- Aggregate multiple source ports
- Filter that input data
- Distribute the input data to output ports
- Filter that output data

1001.0010.001



Quick Example w/out DAN

Without DAN Switching

- Both switches must be configured to mirror traffic from uplink ports to sensor ports.

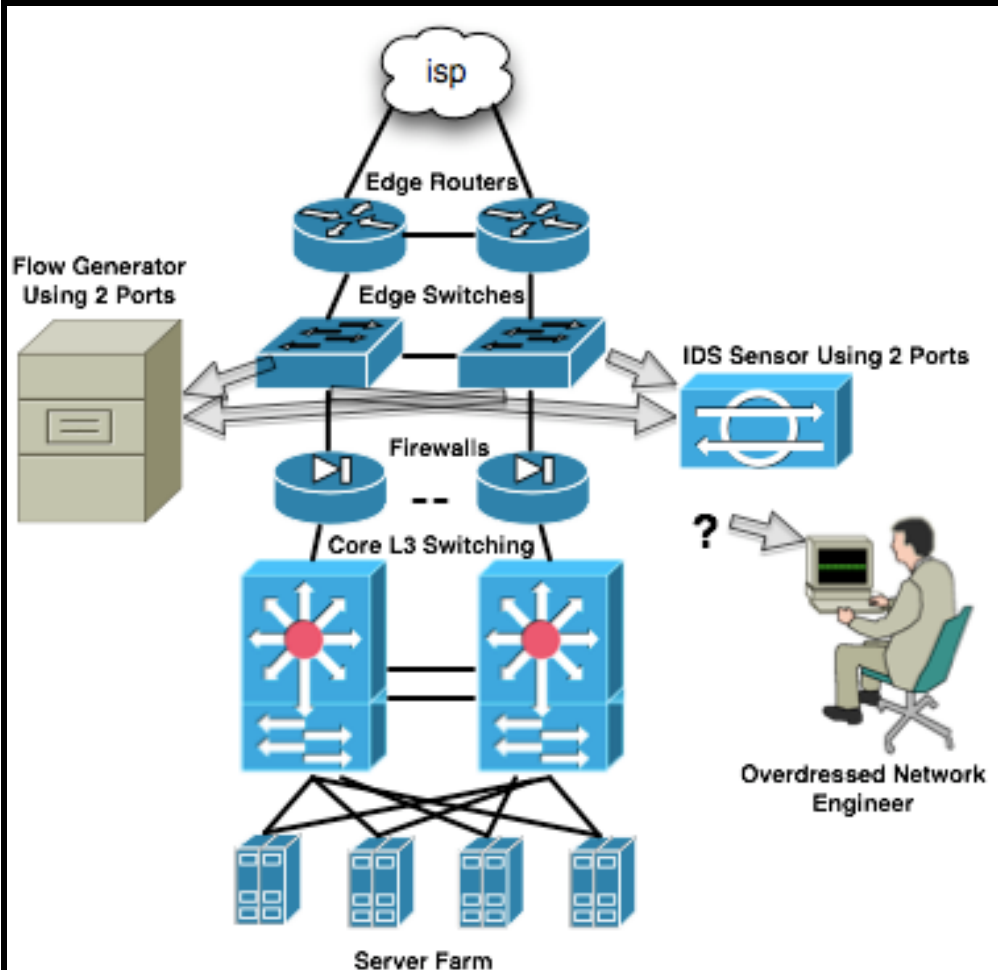
- IDS requires 2 ports 1 for each switch & must aggregate the data in software.

- Flow Generator and IDS are looking at the same traffic.

- Flow Generator also requires 2 ports & must aggregate data in software.

- 2 SPAN session limitation on switches means Network Engineer wishing to connect portable analyzer must disconnect an active analyzer.

- Network Engineer using portable analyzer can only see half of the external traffic assuming topology is load balanced.






w/DAN

With DAN Switching

- 1 SPAN port per switch is used.
- Switch SPAN ports are aggregated by DAN switch.
- DAN switch output to IDS sensor is 1 port.
- DAN switch sends the same output stream that the IDS sensor is using to the Flow Generator.
- Neteng now has choice of using free'd SPAN ports, or better yet, using a new output port on the DAN switch for his portable analyzer.
- If IDS or Flow generator are overloaded specific traffic can be excluded from stream (Ipsec VPN traffic perhaps).





The Bad #1: No Truncation

- Tools which only need headers get whole frames.
 - limits our ability to oversubscribe tool-ports.
- For example looking at headers of a 10gE link in production ~5% of information was headers.
 - Even with link saturation, truncated headers could be monitored with a gigE tool port.

0110110111

1001.0010.001



The Bad #2: No custom PDU offsets

- Filters can be written for basic protocol properties like TCP port, but cannot have filters on arbitrary offsets like TCP [0] to indicate the first byte of a TCP header.
- Typically we only get frame offsets which is too difficult to use consistently (for example dot1q variance, or IP options change TCP[0]).



The Bad #3: Limited ability to leverage 802.1q VLAN filters.

- Many switches strip VLAN tags off SPAN ports.
- This means DAN device must be inline dot1q links.
- Limitation of switch not DAN itself.

0110110111

1001.0010.001



The Bad #4: Source ports must be from single network layer.

- If src ports are combined from access, distribution, core, and perimeter networks packets are duplicated to the tool port at ever point they are seen confusing most tools.

0110110111
1001.0010.001



The Cure

- In-line frame truncation
- Tcpdump style PDU offsets for major protocols.
- Switch vendors need to support mirroring 802.1q VLAN tags to SPAN ports.
- IP TTL De-duplication (using TTL variances to separate distinct routing layers and eliminate duplicate frames).
- Input ports should be able to be labeled with arbitrary 802.1q tags so that tool ports can filter different access layers.
- Statistical sampling mode (send me 1/20 packets).