Deep Packet Inspection as a Service

Anat Bremler-Barr, Yotam Harchol, David Hay, Yaron Koral

Presented by: Han Zhang and Andrew Quinn

Deep Packet Inspection (DPI)

- Payload of packets is compared against *patterns*
- Used by middleboxes for all sorts of things:
  - Intrusion Detection (SNORT, BRO)
  - L7 Firewall (Linux L7-filter, ModSecurity)
  - L7 Load Balancing (F5, A10)
  - Network Analytics (Qosmos)
- Accounts for high per packet processing (2.9x slowdown)

Current Middlebox Architecture

- We often chain these service together in a pipeline…
- But each of these services do their own DPI!

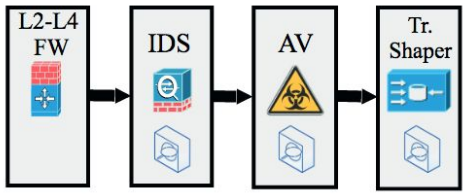

New Middlebox Approach

- DPI work in the beginning of pipeline
- Allow each middlebox to leverage the service

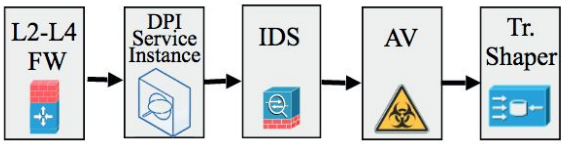

Major Benefit

Decouple DPI from Middlebox

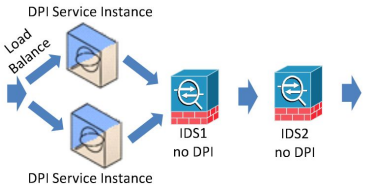

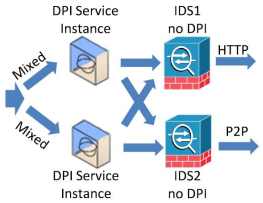

Outline

- Introduction
- Design
- Implementation
- Evaluation

DPI Background

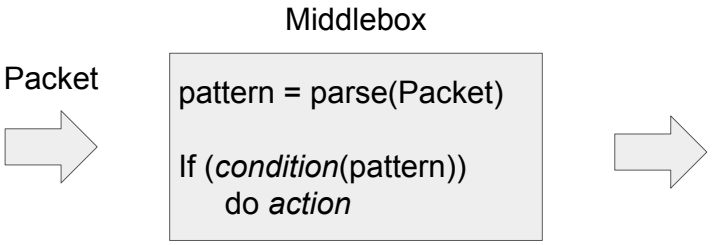

DPI-as-a-Service

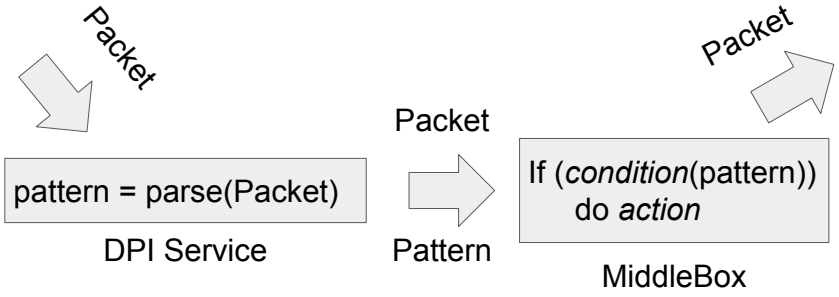

Architecture

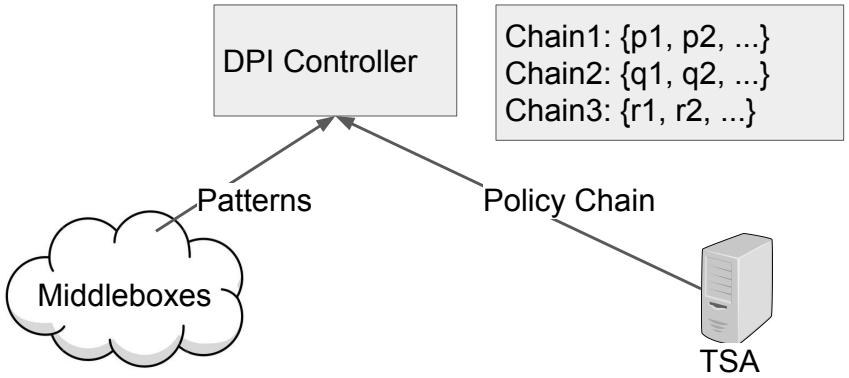

Instance Management

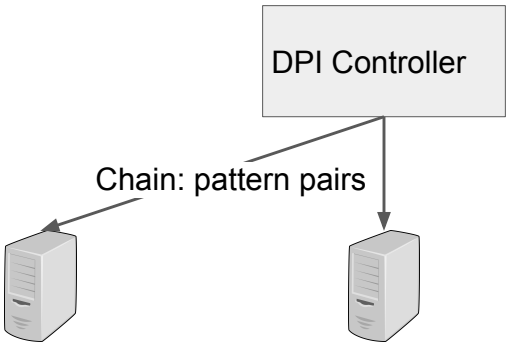

Instance Management

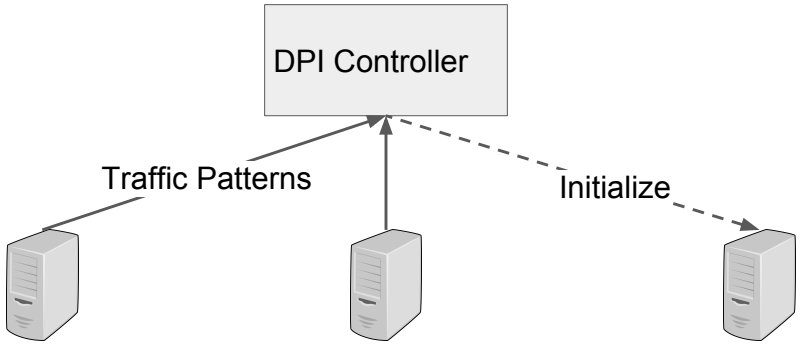

Outline

- Introduction
- Design
- Implementation
- Evaluation

DPI Service Instance

- Aggregate multiple pattern sets
- Scan incoming packets
- Generate *match-lists* of matching patterns
- Notify corresponding middleboxes if packets pattern matches

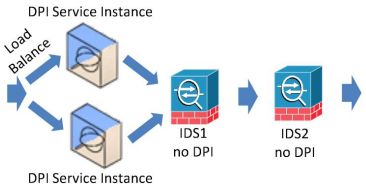

Implementation

- Build a system in Mininet with 2 user hosts, 2 middleboxes, and 1 DPI service instance
- Not used for system performance analysis due to Mininet overhead
- Instead test each component separately with custom input

Outline

- Introduction
- Design
- Implementation
- Evaluation

Virtualization Performance

Claim:

- Virtualization has minor impact
- The number of patterns has major impact

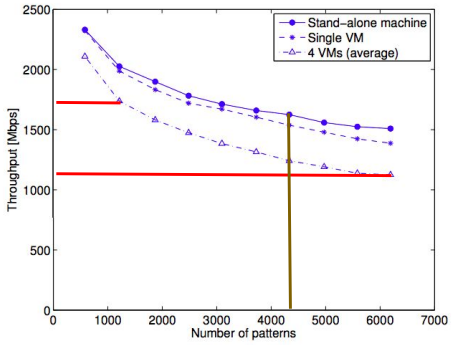

Gain from Virtual DPI

- Significantly higher throughput
- Faster DPI processing in general

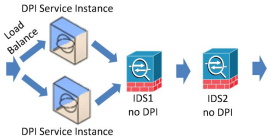

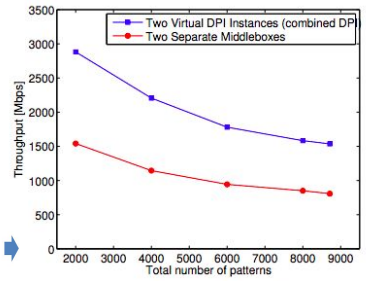

(a) *Snort1* and *Snort2*

Gain from Virtual DPI

- Two separate DPI services could go over 100% utilization, depending on the load

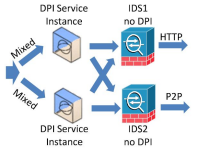

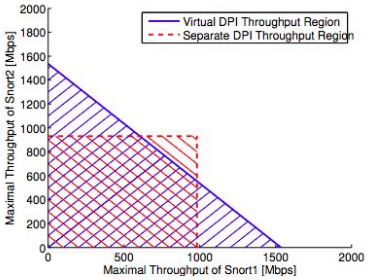

(a) *Snort1* and *Snort2*

Match Report Size

- Average 34 bytes
- 1% larger than 120 bytes
- More concerned about network delay

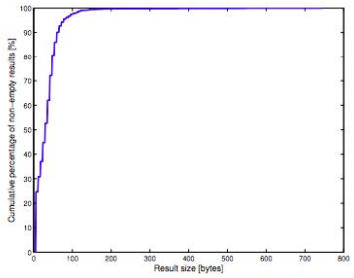

Conclusion

Insights:

- Network Function Virtualization (NFV) is important!
- Many common tasks in middleboxes

Limitations:

- System performance is tested in limited environment
- Simplistic middleboxes behavior
  - No consideration of middlebox performance without regards to DPI functions
  - Tradeoff between network delay vs hardware acceleration