IPv6 First Hop Security in Virtualized Environments

Christopher Werny cwerny@ernw.de

Who am I

- Network geek, working as security researcher for
- Germany based ERNW GmbH
  - Independent
  - Deep technical knowledge
  - Structured (assessment) approach
  - Business reasonable recommendations
  - We understand corporate

- Blog: www.insinuator.net

Agenda

- Introduction to IPv6 First-hop Security
- Lab Setup
- Test Cases
- Results
- Conclusion

Cisco First-Hop-Security

Cisco First-Hop-Security

- Cisco name for various security features in IPv6
- Rollout is/was planned in three stages
- Every Phase will release/released more IPv6 security features to achieve feature parity with the IPv4 world

Phase I

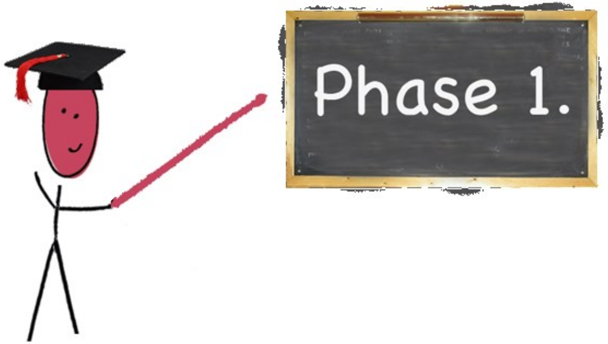

- Available since Summer 2010
- Introduced RA Guard and Port based IPv6 ACLs
- In the beginning, only supported on datacenter switches
  - Since 15.0(2) supported on C2960S and C3560/3750-X

RA Guard

- Implements *isolation* principle similar to other L2 protection mechanisms already deployed in v4 world.
- RFC 6105
- Works quite well against some flavors of problem.
  - On most platforms no logging or port deactivation can be implemented. RA packets are just dropped.

Phase II

- Available since end of 2011/ beginning of 2012 (depending on the platform)
- Introduced DHCPv6 Guard and NDP Snooping
  - The equivalent to DHCP Snooping and Dynamic ARP Inspection in the IPv4 World
- In the meantime good support on current access layer platforms

DHCPv6 Guard

- Similar functionality to DHCP Snooping in the IPv4 world
  - But more sophisticated
- Blocks reply and advertisement messages that originates from “malicious” DHCP servers and relay agents
- Provides finer level of granularity than DHCP Snooping.
- Messages can be filtered based on the address of the DHCP server or relay agent, and/or by the prefixes and address range in the reply message.

Cisco IPv6 Snooping

- IPv6 Snooping is the basis for several FHS security mechanisms

- When configured on a target (VLAN, Interface etc.), it redirects NDP and DHCP traffic to the switch integrated security module

IPv6 ND Inspection

- Learns and secures bindings for addresses in layer 2 neighbor tables.
- Builds a trusted binding table database based on the IPv6 Snooping feature
- IPv6 ND messages that do not have valid bindings are dropped.
- A message is considered valid if the MAC-to-IPv6 address is verifiable

FHS Availability - Cisco

Feature/Platform	Catalyst 6500 Series	Catalyst 4500 Series	Catalyst 2K/3K Series	ASR1000 Router	7600 Router	Catalyst 3850	Wireless LAN Controller (Flex 7500, 5508, 2500, WISM-2)	Nexus 3k/5k/6k/7k
RA Guard	15.0(1)SY	15.1(2)SG	15.0.(2)SE		15.2(4)S	15.0(1)EX	7.2	NX-OS 7.2
IPv6 Snooping	15.0(1)SY[1]	15.1(2)SG	15.0.(2)SE	XE 3.9.0S	15.2(4)S	15.0(1)EX	7.2	NX-OS 7.2
DHCPv6 Guard	15.2(1)SY	15.1(2)SG	15.0.(2)SE		15.2(4)S	15.0(1)EX	7.2	NX-OS 7.2
Source/Prefix Guard	15.2(1)SY	15.2(1)E	15.0.(2)SE[2]	XE 3.9.0S	15.3(1)S		7.2	NX-OS 7.2
Destination Guard	15.2(1)SY	15.1(2)SG	15.2(1)E	XE 3.9.0S	15.2(4)S			NX-OS 7.2
RA Throttler	15.2(1)SY	15.2(1)E	15.2(1)E			15.0(1)EX	7.2	
ND Multicast Suppress	15.2(1)SY	15.1(2)SG	15.2(1)E	XE 3.9.0S		15.0(1)EX	7.2	

Why is FHS so important in virtual environments?

- More and more systems get virtualized on different Hypervisor platforms.

- Private cloud environments will get more prevalent in the near future
  - And are already deployed by many environment

- Virtual Desktop Infrastructure gets more and more deployed

- Protecting those systems/assets from malicious client is paramount for the overall security of your environment

Why is FHS so important in virtual environments?

- Thinking about all that, your first line of defense (the access layer switch) also moves from the physical into the virtualized environment.

- While the support for FHS reaches a kind of “mature” state on several platforms (at least in the Cisco space) this might not necessarily be the case for virtual switches.
  - The reason for this talk ;)

Hypervisor Lab Setup

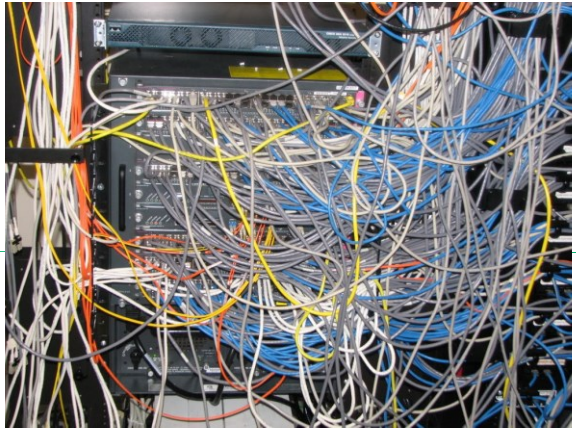

Lab Setup

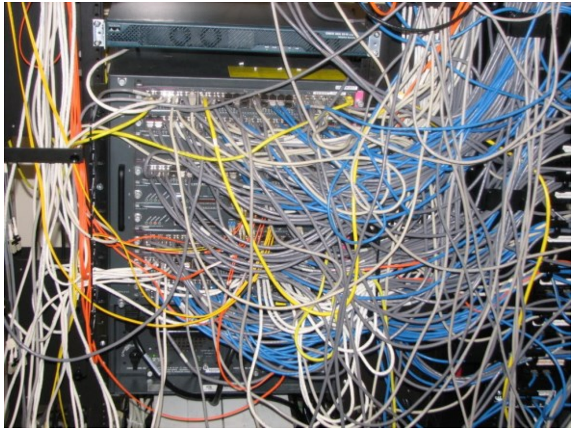

- Three different types of Hypervisors
  - Windows Server 2012 R2 Hyper-V
  - VMware ESXi 5.5
  - Kernel-based Virtual Machine (KVM)
- … with three different types of virtual switches
  - Hyper-V vSwitch
  - Cisco Nexus 1000V
  - Open vSwitch

Lab Setup

- Three different (fully patched as of 03.2015) operating systems
  - Windows 7
  - Windows 8
  - Kali
- Layer 2 adjacent residing on the same prefix/vlan

Lab Environment Overview

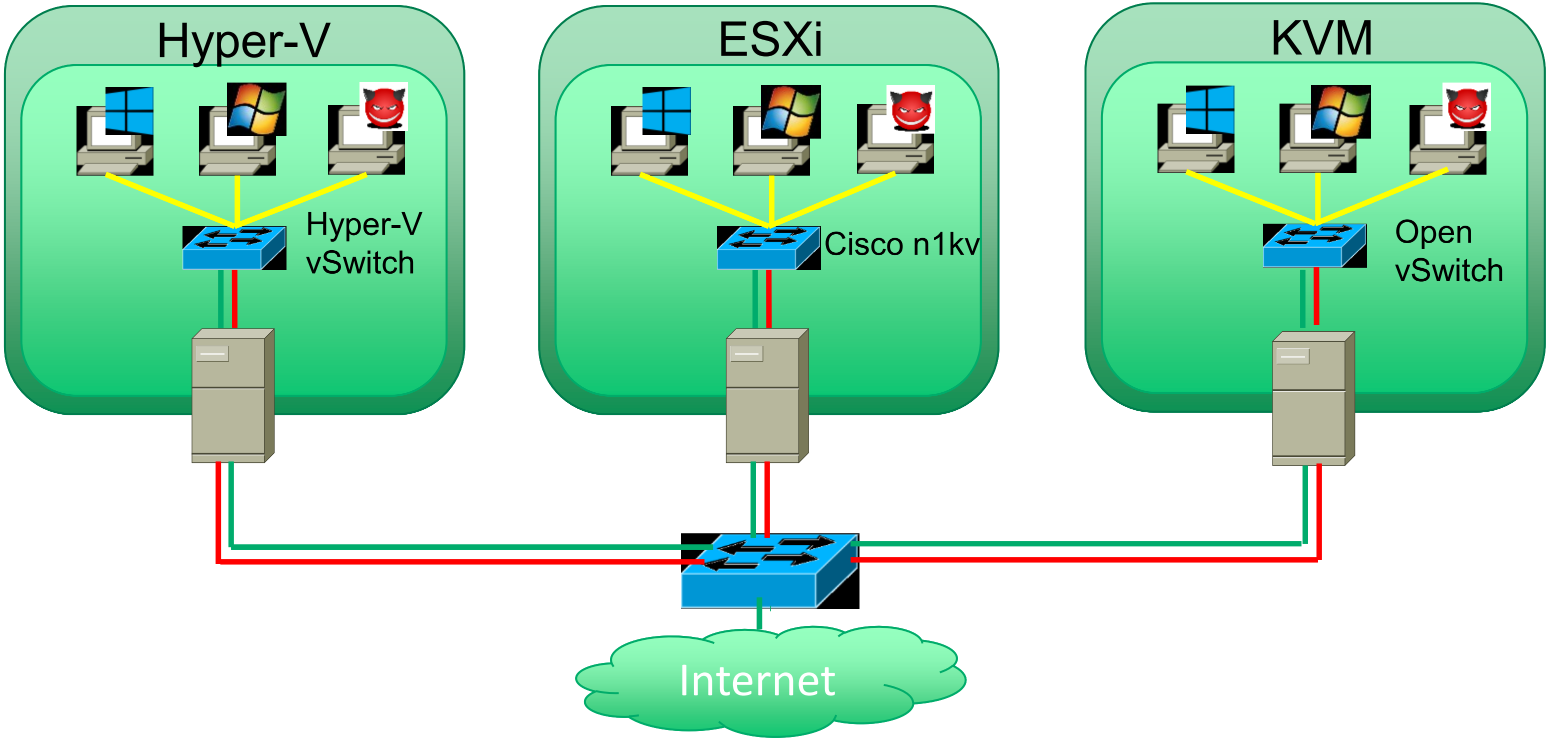

The Hypervisors and the virtual Switches

Hyper-V vSwitch

Hyper-V vSwitch

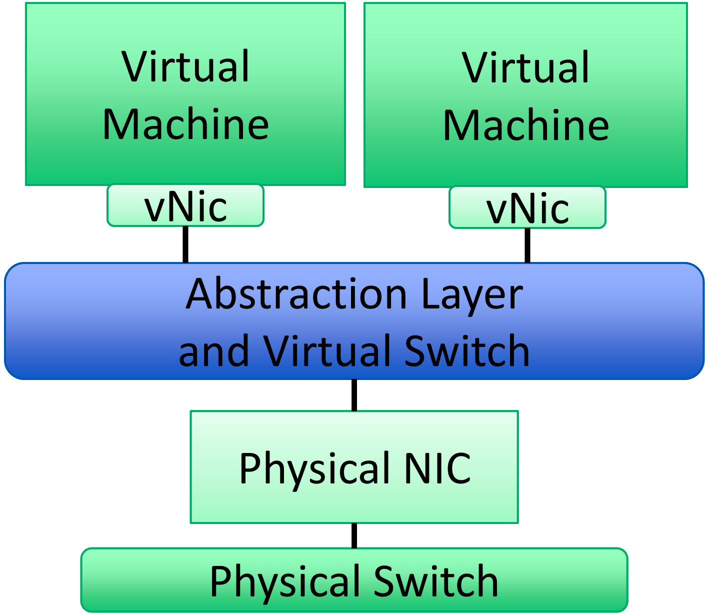

- The Hyper-V vSwitch supports:
  - RA-Guard
  - DHCPv6 Snooping
- In addition, since Server 2012:
  - Support for Extended/Stateful (IPv6) ACLs
  - Can only be configured via Powershell
  - No GUI element which shows the ACLs

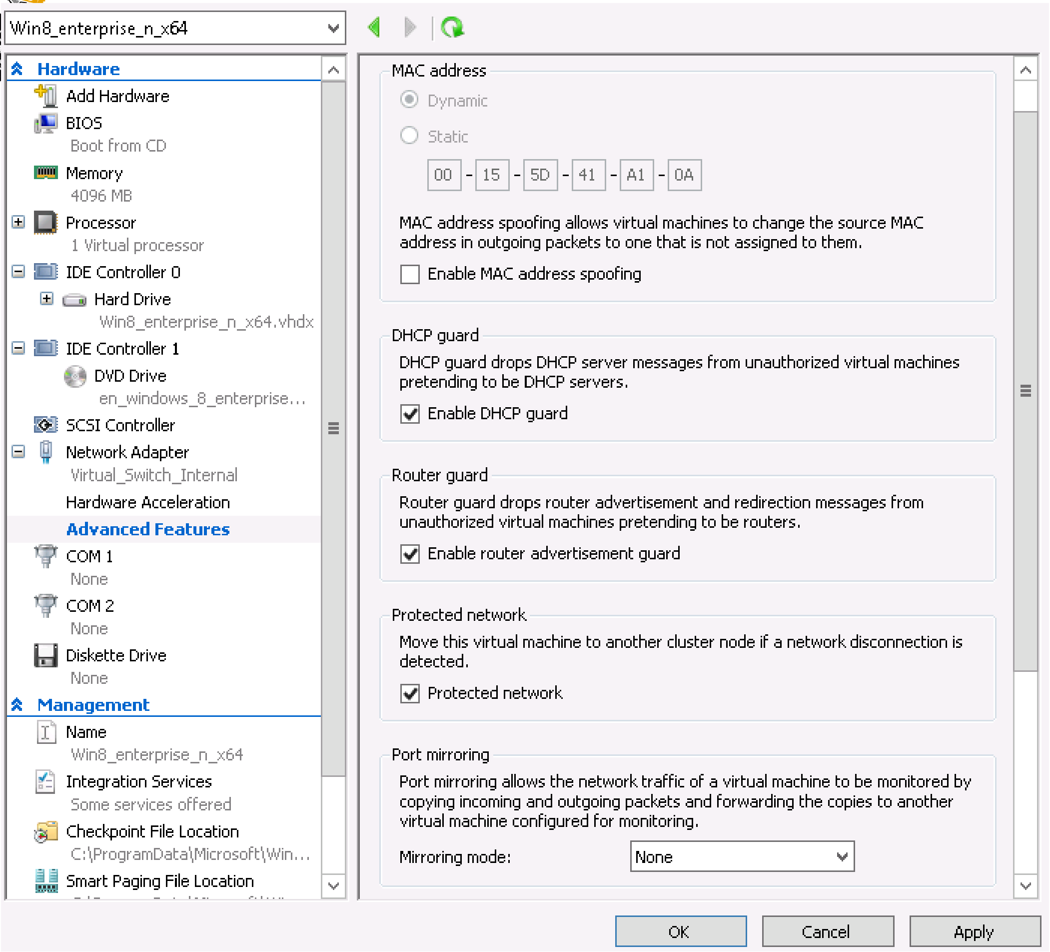

Hyper-V

Activate RA/DHCPv6 Guard

Testing Procedure

- First step:
  - vSwitch default configuration
  - Perform different (IPv6) attacks
  - Observe/document the results
- Second step:
  - Activate/configure the feature
  - Test again
  - Observe/document the results
- Third step:
  - Try to evade it ;)

Attacking tool “THC-IPv6 toolkit”

- Using THC-IPv6 toolkit
  - RA Guard tests
    - fake_router26 eth0 -A 2001:db8:dead:beef::/64
    - flood_router26 eth0
    - flood_router26 –H eth0
    - flood_router26 –F eth0
    - flood_router26 –D eth0
  - DHCPv6 Guard test
    - fake_dhcps6 eth0 2001:db8:dead:beef::/64

Results for Hyper-V

First Test Scenario

vSwitch default configuration

- Let‘s just say we had initially some unexpected results ;)
- Basic flooding of RAs didn‘t work at all.
- So i started to debug this behaviour

Some issues with THC-IPv6

- flood_router26 plain (without any EH)
  - RAs do _NOT_ get forwarded
  - No indication in any log file
  - They just disappeared

- After some digging:
  - Source MAC address in RAs is set to all zeros
  - Tested on various version of THC-toolkit (v2.3, 2.4, 2.7) and Hypervisors → all behave the same
  - Enabled “MAC address spoofing” on vSwitch so that the VM is allowed to send frames with different MAC addresses → No luck ☹ .
  - My assumption: the virtual switches treats the frames with a source mac address of all zeros as invalid and does not forward them

How to fix this issue

- A big thank you to Antonios who coded us the desired functionality into Chiron in just in a couple of hours
- I mean we could still have done it with Scapy, but having more features built into Chiron is a good thing ;)

Attacking tool „Chiron“

- List of tests (for your reference):
  - chiron_local_link.py eth0 –ra –rr
  - chiron_local_link.py eth0 –ra –rr –lfE 60 –nf 1
  - chiron_local_link.py eth0 -rr -ra -rand_ra -lfE 60
  - chiron_local_link.py eth0 -rr -ra -rand_ra -lfE 43
  - chiron_local_link.py eth0 -ra -rr -rand_ra -lfE 0,60
  - chiron_local_link.py eth0 -ra -rr -rand_ra -lfE5X60
  - chiron_local_link.py eth0 -ra -rr -rand_ra -nf 2
  - chiron_local_link.py eth0 -ra -rr -rand_ra -lfE 0
  - chiron_local_link.py eth0 -ra -rr -rand_ra-lfE0-nf 2
- – chiron_local_link.py eth0 -ra -rr -rand_ra -luE 0 -lfE 60 –nf 2
- – chiron_local_link.py eth0 -ra -pr 2001:db8:c001:cafe:: -lfE 60 –nf 2

Results for HyperV vSwitch

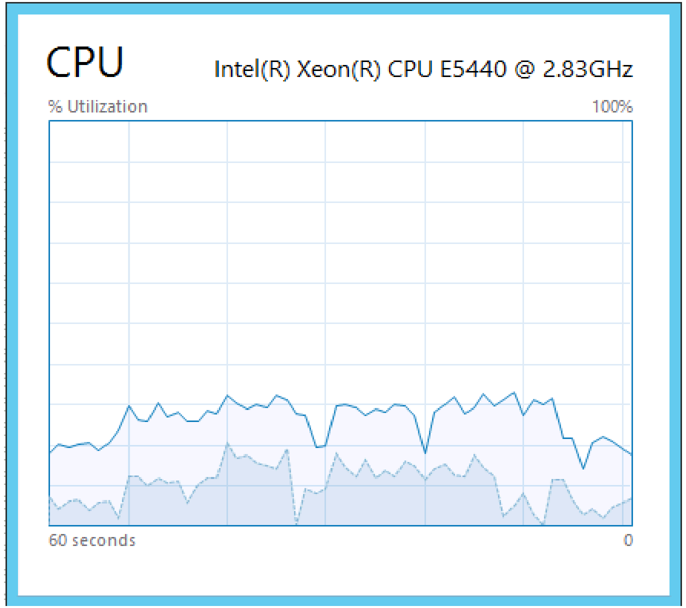

- Activating RA Guard on the vSwitch mitigated the attack
  - Even when using extension header
  - RA-Guard could not be evaded with use of Extension Headers
- DHCPv6 Guard works as well
- No CPU spikes on the HV during the attack
- Unfortunately we couldn’t find any log entries indicating that a VM is flooding RAs

Fragments only

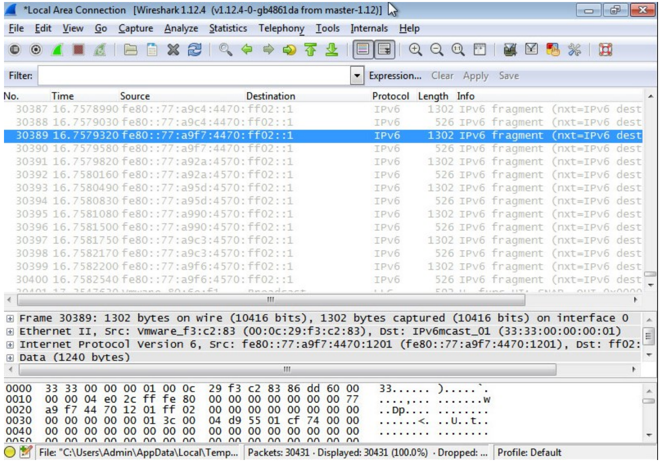

- Flood_router26 –D eth0
- Fragments are passing RA Guard
- The filter module is inspecting every fragment, and as soon as it sees the upper layer protocol information, this fragment is blocked

Cisco Nexus 1000v

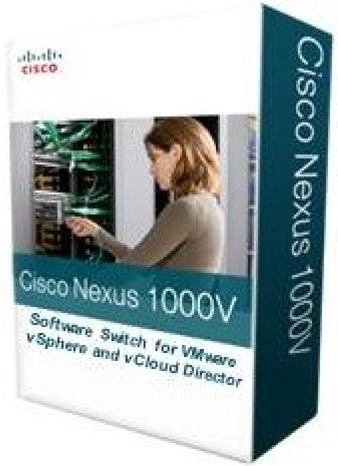

VMware ESXi

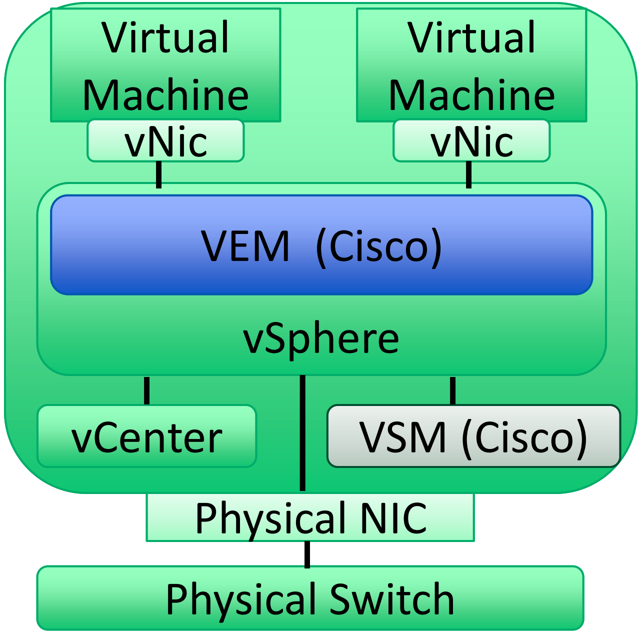

- ESXi 5.5.0
- vCenter Server 5.5.0
- Cisco Nexus 1000v
  - 5.2(1)SV3(1.2)

Nexus 1000V

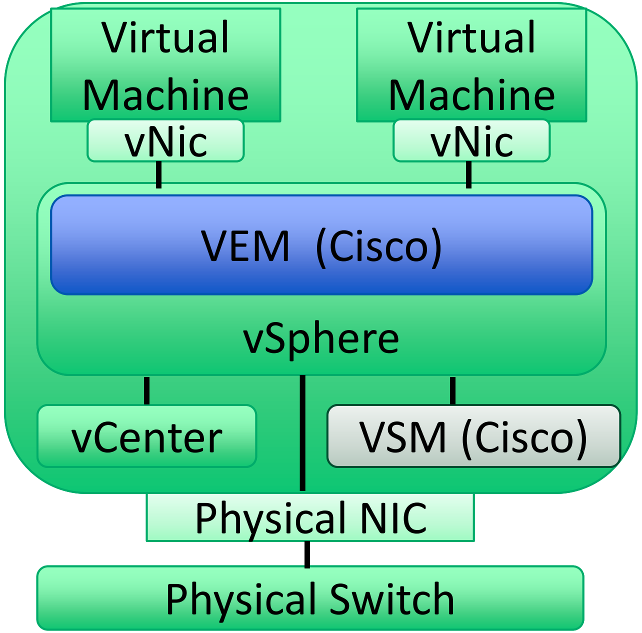

- Unfortunately, no IPv6 FHS features available on the Nexus 1000v
- The only option you have is using port based ACLs for filtering IPv6 traffic
- IPv6 ACLs were introduced in 5.2(1)SV3(1.1)
- I wasn't able to find any information whether FHS is on the roadmap for..
  - Quoting Ivan here: "Wave with your wallet ;)

Port-based ACLs

```
n1000V(config-ipv6-acl)# permit icmp any any ?
  <CR>
  <0-255>                  ICMPv6 message type
  beyond-scope             Destination beyond scope
  destination-unreachable  Destination address is unreachable
  dscp                     Match packets with given dscp value
  echo-reply               Echo reply
  echo-request             Echo request (ping)
  header                   Parameter header problems
  hop-limit                Hop limit exceeded in transit
  log                      Log matches against this entry
  mld-query                Multicast Listener Discovery Query
  mld-reduction            Multicast Listener Discovery Reduction
  mld-report               Multicast Listener Discovery Report
  nd-na                    Neighbor discovery neighbor advertisements
  nd-ns                    Neighbor discovery neighbor solicitations
  next-header              Parameter next header problems
  no-admin                 Administration prohibited destination
  no-route                 No route to destination
  packet-too-big           Packet too big
  parameter-option         Parameter option problems
  parameter-problem        All parameter problems
  port-unreachable         Port unreachable
  reassembly-timeout       Reassembly timeout
  redirect                 Neighbor redirect
  renum-command            Router renumbering command
  renum-result             Router renumbering result
  renum-seq-number         Router renumbering sequence number reset
  router-advertisement     Neighbor discovery router advertisements
  router-renumbering       All router renumbering
  router-solicitation      Neighbor discovery router solicitations
  time-exceeded            All time exceeded
  unreachable              All unreachable
```

- IPv6 ACLs are supported, but are kind of "limited"
  - No Extension Header support
  - No undetermined transport
- Configuration for RA ACL

n1kv(config)# ipv6 access-list RAGuardACL
n1kv(config-ipv6-acl)# deny icmp any any router-advertisement
n1kv(config-ipv6-acl)# permit ipv6 any any
n1kv(config-ipv6-acl)# Interface vethernet 1
n1kv(config-if)# Ipv6 port traffic-filter RAGuardACL in

Testing Procedure

- First step:
  - vSwitch default configuration
  - Perform different (IPv6) attacks
  - Observe/document the results
- Second step:
  - Activate/configure the feature
  - Test again
  - Observe/document the results
- Third step:
  - Try to evade it ;)

Results Nexus 1000v

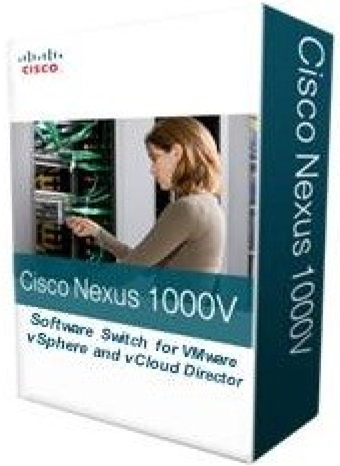

Results for Nexus 1000v

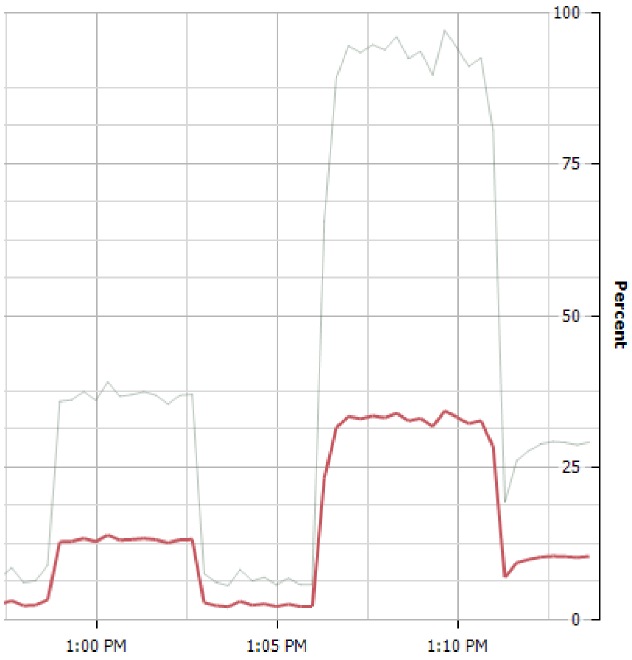

- Active RA ACL blocks all tried attacks
  - Basic attack without EH was blocked reliable
  - Could not be evaded with use of Extension Headers
  - Second fragment with the Upper Layer information is blocked
- flood_router eth0 causes 13% CPU load with one attacking machine while blocking the packets
- flood_router -D eth0 causes 34% CPU load with one attacking machine while blocking the packets

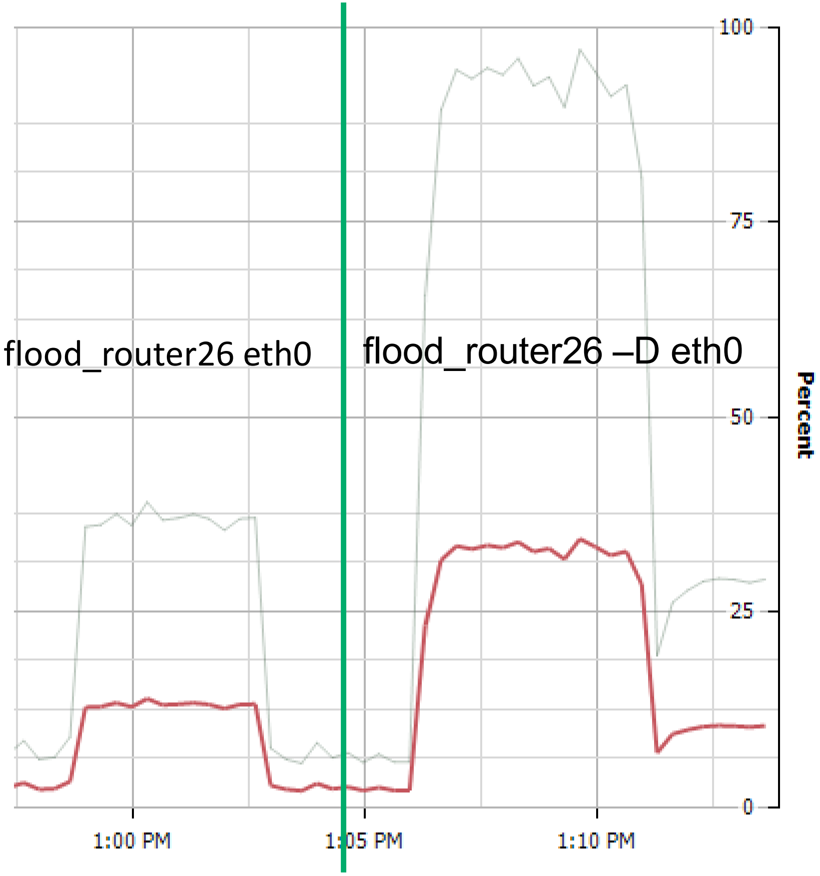

ESXi CPU load

flood_router26 eth0 and
flood_router26 –D eth0

Open vSwitch

KVM & Open vSwitch

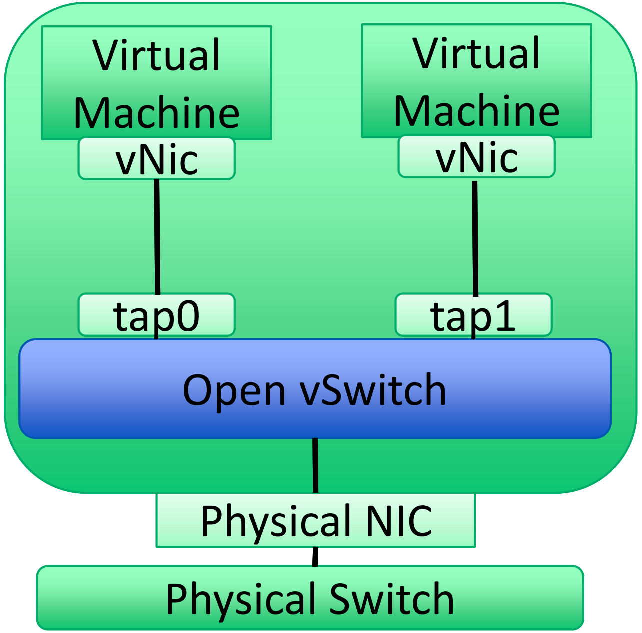

- Ubuntu 14.04.2 LTS
  - 3.13.0-32-generic
- QEMU
  - Version 2.0.0
- OpenFlow
  - 1.4
- Open vSwitch
  - 2.3.1

Open vSwitch

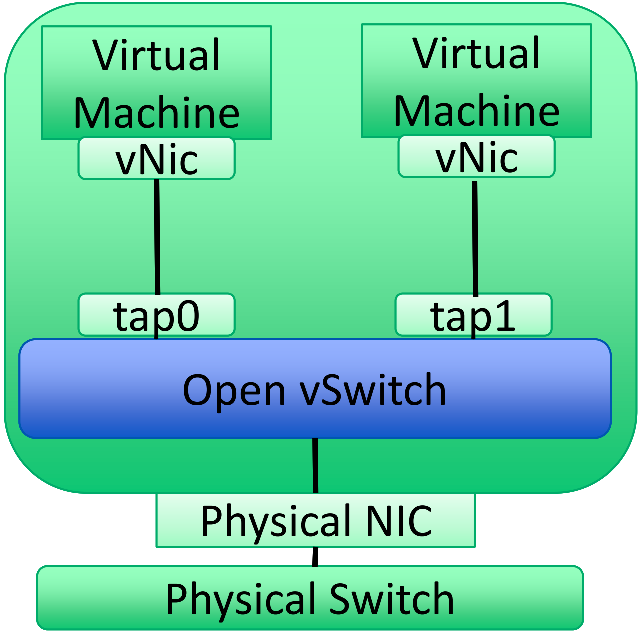

- Unfortunately no IPv6 FHS features available
- Only IPv6 ACL based behavior based on flow entries matching ICMPv6 type 134

OVS Main Components

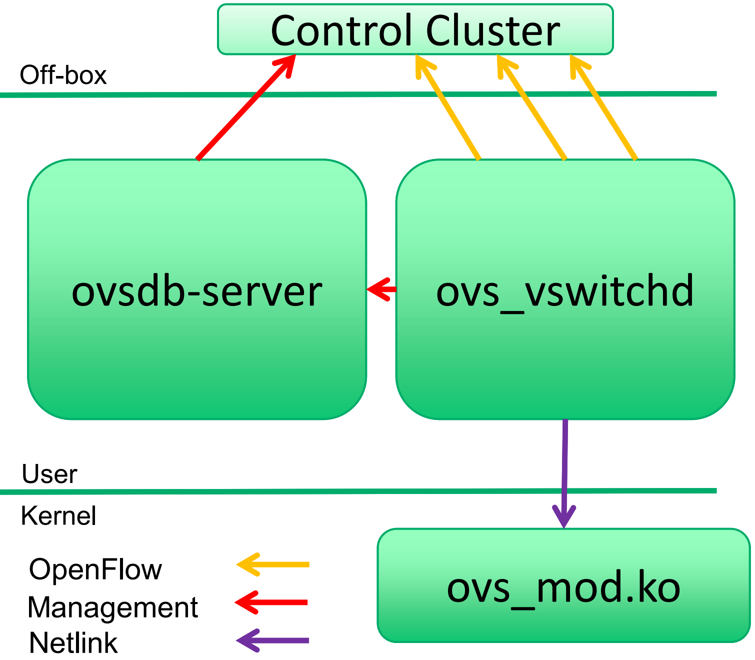

- ovsdb-server
  - Database that holds switch-level configuration
- ovs_mod.ko
  - Kernel module that handles switching
- ovs_vswitchd
  - Core component
    - Communicates with outside world using OpenFlow

OpenFlow

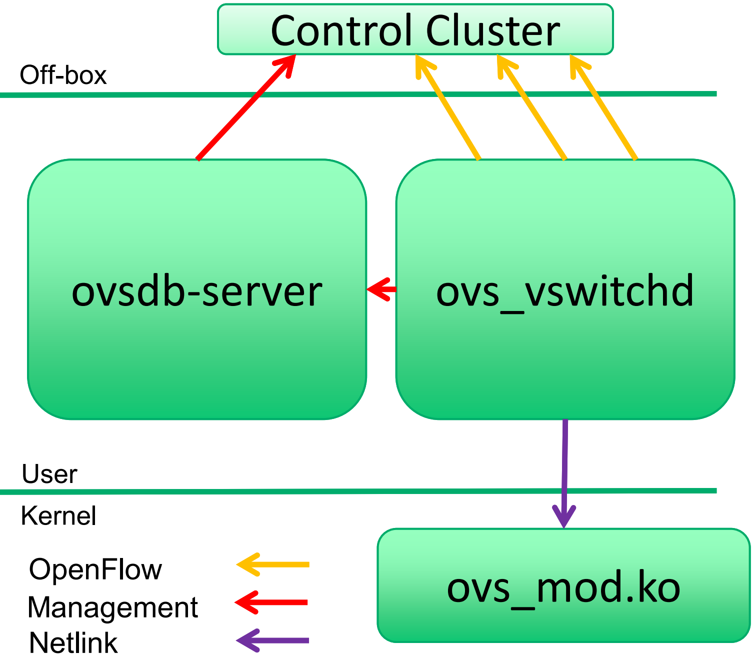

- OpenFlow is a communication protocol
- Centralized controller configures flow table
  - Lookup based on L2-L4
  - Supports full wildcarding and priorities
  - Flows associated with action: forward, drop, modify
  - Missed (might) flow go to controller
  - Extensible Match e.g. for IPv6 traffic
  - OpenFlow IPv6 support since version 1.2

Testing Procedure

- First step:
  - vSwitch default configuration
  - Perform different (IPv6) attacks
  - Observe/document the results
- Second step:
  - Activate/configure the feature
  - Test again
  - Observe/document the results
- Third step:
  - Try to evade it ;)

Open vSwitch Results

First Test Scenario

- Using basic flooding without configuration, no surprises here
- The victims get flooded and configures lot of (100) IPv6 addresses and CPU load spikes to the during the attack.
- We continued to configure the the flow entry for blocking the ICMPv6 type 134

Flow configuration

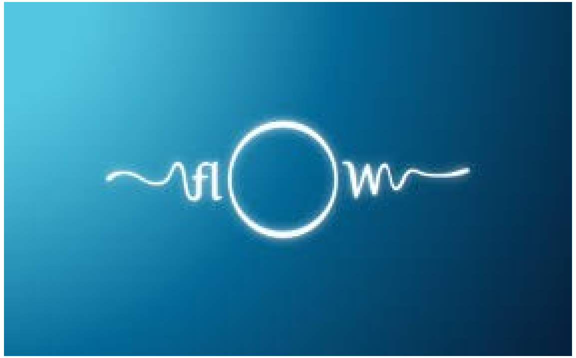

- ovs-ofctl add-flow bridge1
  "in_port=2, # Port of attacker
  dl_type=0x86dd, # IPv6
  nw_proto=58,   # ICMPv6
  icmp_type=134, # RA
  actions=drop" # drop packets

Results for Open vSwitch

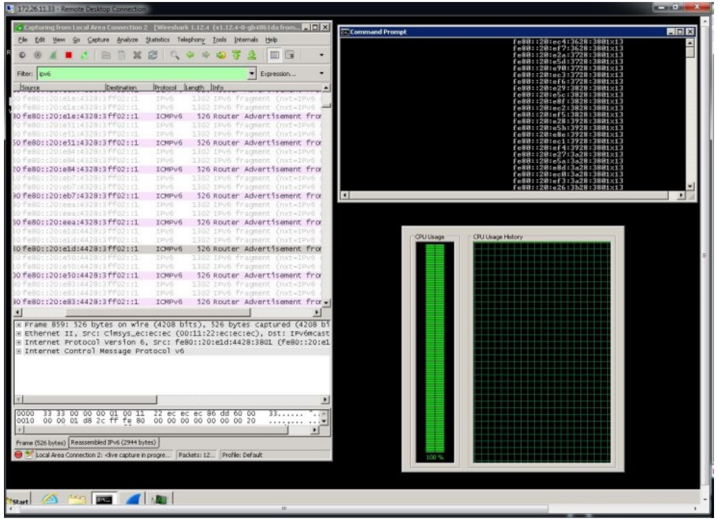

- Active OpenFlow ACL
  - Blocks most of the attacks, but…
    - flood_router26 –D eth0
    - chiron_local_link.py eth3 -ra -rr -rand_ra -nf 2
    - chiron_local_link.py eth3 -ra -rr -rand_ra -lfE 0 -nf 2
    - chiron_local_link.py eth3 -ra -rr -rand_ra -luE 0 -lfE 60 –nf 2
    - chiron_local_link.py eth3 -ra -pr 2001:db8:c001:cafe:: -lfE 60 –nf 2
  - … passed the ACL

What does this mean?

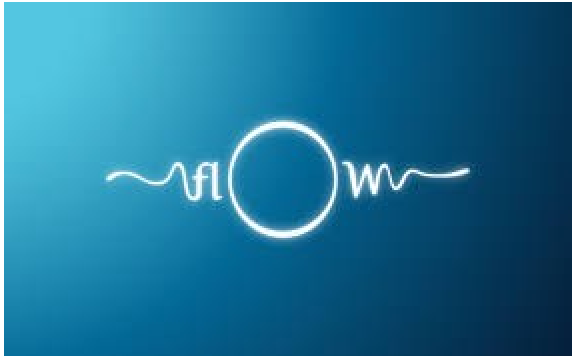

- It seems that only the first fragment gets checked against the configured ACL
- If it does not match, all subsequent fragments belonging to the IPv6 packet get forwarded without further checks.

Results for Open vSwitch

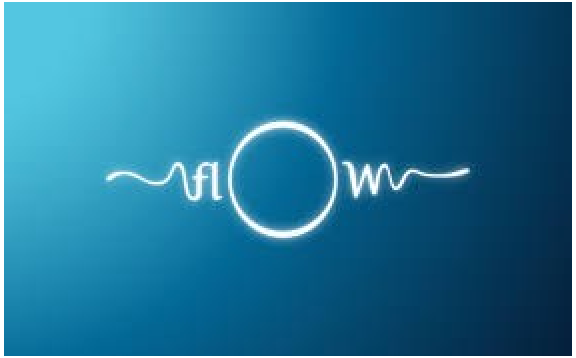

- Let's block fragments an test it again:
  ovs-ofctl bridge1
  "in_port=2, # Port of attacker
  dl_type=0x86dd, # IPv6
  ip_frag=yes, # Match fragments
  actions=drop"

- After configuring the fragment ACL
  - only flood_router26 –D fragments could pass the ACL, but the victim does NOT create an address or gateway

providing security.

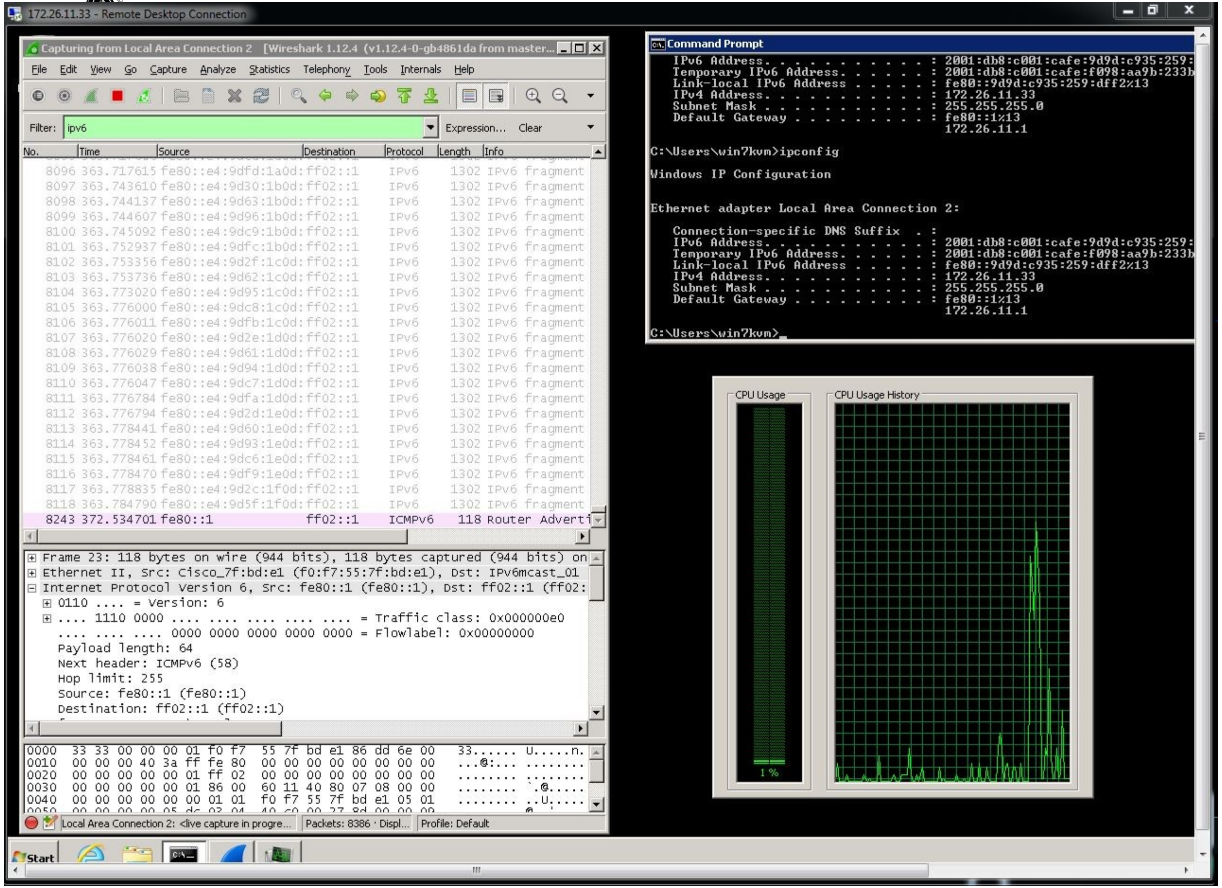

After deploying fragment ACL

flood_router26 –D eth0

Summary

First Hop Security Features	Hyper-V vSwitch	Nexus 1000v	Open vSwitch
RA Guard	Yes	No	No
DHCPv6 Guard	Yes	No	No
IPv6 ACLs	Yes	Yes	Yes
IPv6 Snooping	No	No	No
IPv6 Source Guard	No	No	No
IPv6 Prefix Guard	No	No	No
IPv6 Destination Guard	no	no	no

FHS availability

There is room for improvement ;)

Conclusion

- IPv6 First-hop Security features are _NOT_ wide spreaded in common virtual switches
- Hyper-V is as of right know the only one which supports a least few FHS features
- Thinking that it is 2015, that's quite an unfortunate state, and reminds me of the state of FHS on physical switches 4 years ago.
- Hopefully this will change in the near future.
- Again, wave with your wallet ;)

There’s never enough time…

THANK YOU…

...for yours!

Questions?