How to securely operate an IPv6 network

https://tools.ietf.org/html/draft-ietf-opsec-v6-06

LACNIC 23

Enrique Davila
enriqued@cisco.com
Released: May 2015

Agenda

- Management Plane
- Control Plane
  - Routing Information
  - Neighbor Discovery
  - Control Plane Protection
- Data Plane
  - Anti-spoofing
  - Access Control List
- Telemetry
- Summary

2

Management Plane

3

Management over IPv6

- SSH, syslog, SNMP, NetFlow all work over IPv6

- Dual-stack management plane
  - More resilient: works even if one IP version is down
  - More exposed: can be attacked over IPv4 and IPv6

- As usual, infrastructure ACL is your friend as well as out-of-band management

Control Plane: Routing Protocols

Preventing IPv6 Routing Attacks Protocol Authentication

- BGP, ISIS, EIGRP no change:
  - An MD5 authentication of the routing update
- OSPFv3 has changed and pulled MD5 authentication from the protocol and instead rely on transport mode IPsec (for authentication and confidentiality)
  - But see RFC ~~6506~~ 7166 *(not yet widely implemented)*
- IPv6 routing attack best practices
  - Use traditional authentication mechanisms on BGP and IS-IS
  - **Use IPsec** to secure protocols such as OSPFv3

BGP Route Filters

- Pretty obvious for customer links
- For peering, a relaxed one

```
ipv6 prefix-list RELAX deny 3ffe::/16 le 128
ipv6 prefix-list RELAX deny 2001:db8::/32 le 128
ipv6 prefix-list RELAX permit 2001::/32
ipv6 prefix-list RELAX deny 2001::/32 le 128
ipv6 prefix-list RELAX permit 2002::/16
ipv6 prefix-list RELAX deny 2002::/16 le 128
ipv6 prefix-list RELAX deny 0000::/8 le 128
ipv6 prefix-list RELAX deny fe00::/9 le 128
ipv6 prefix-list RELAX deny ff00::/8 le 128
ipv6 prefix-list RELAX permit 2000::/3 le 48
ipv6 prefix-list RELAX deny 0::/0 le 128
```

Source: http://www.space.net/~gert/RIPE/ipv6-filters.html

Link-Local Addresses vs. Global Addresses

- Link-Local addresses, fe80::/10, (LLA) are isolated
  - Cannot reach outside of the link
  - **Cannot be reached from outside of the link ☺**
- Could be used on the infrastructure interfaces
  - Routing protocols (inc BGP) work with LLA
  - **Benefit:** no remote attack against your infrastructure: implicit infrastructure ACL
  - Note: need to provision loopback for ICMP generation (notably traceroute and PMTUD)
  - *See also: RFC 7404*
  - LLA can be configured statically (not the EUI-64 default) to avoid changing neighbor statements when changing MAC

```
interface FastEthernet 0/0
   ipv6 address fe80::1/64 link-local
neighbor fe80::2%FastEthernet0/0
```

Control Plane: Neighbor Discovery

Scanning - Bad for CPU
Remote Neighbor Cache Exhaustion RFC 6583

- Potential router CPU/memory attacks if aggressive scanning
  - Router will do Neighbor Discovery... And waste CPU and memory
- **Local router** DoS with NS/RS/…

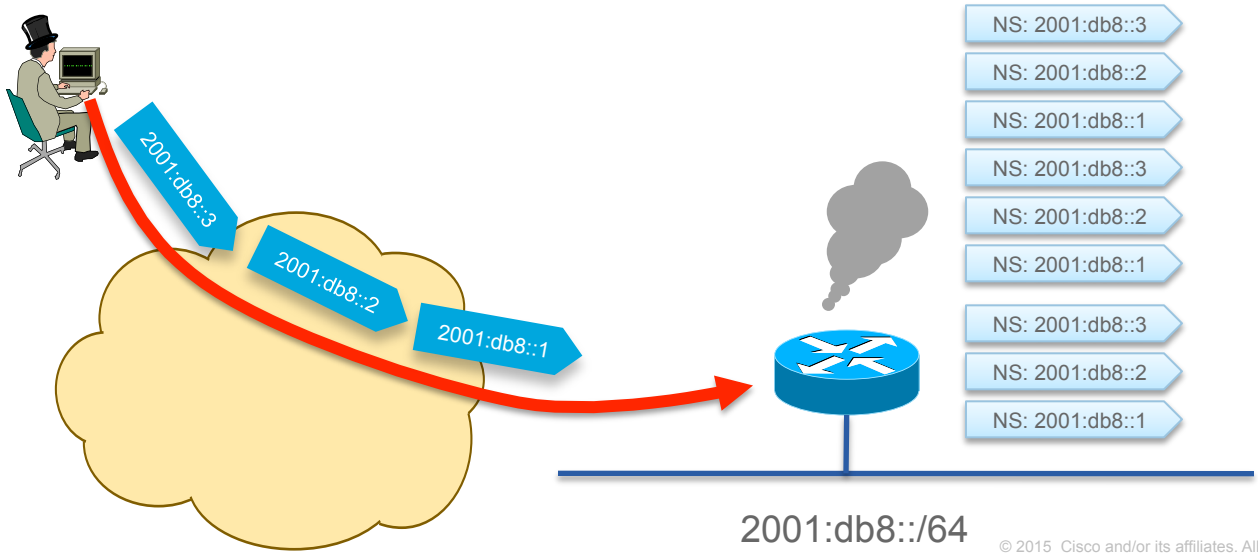

Mitigating Remote Neighbor Cache Exhaustion

- Built-in rate limiter with options to tune it
  - Since 15.1(3)T: `ipv6 nd cache interface-limit`
  - Or IOS-XE 2.6: `ipv6 nd resolution data limit`
  - Destination-guard is part of First Hop Security phase 3
  - Priority given to refresh existing entries vs. discovering new ones (RFC 6583)
- Using a /64 on **point-to-point links** => a lot of addresses to scan!
  - Using /127 could help (RFC 6164)
- **Internet edge/presence**: a target of choice
  - Ingress ACL permitting traffic to specific statically configured (virtual) IPv6 addresses only
- Using infrastructure ACL prevents this scanning
  - iACL: edge ACL denying packets addressed to your routers
  - Easy with IPv6 because new addressing scheme can be done ☺

http://www.insinuator.net/2013/03/ipv6-neighbor-cache-exhaustion-attacks-risk-assessment-mitigation-strategies-part-1

Simple Fix for Remote Neighbor Cache Exhaustion

- Ingress ACL allowing only valid destination and dropping the rest
- NDP cache & process are safe
- Requires DHCP or static configuration of hosts

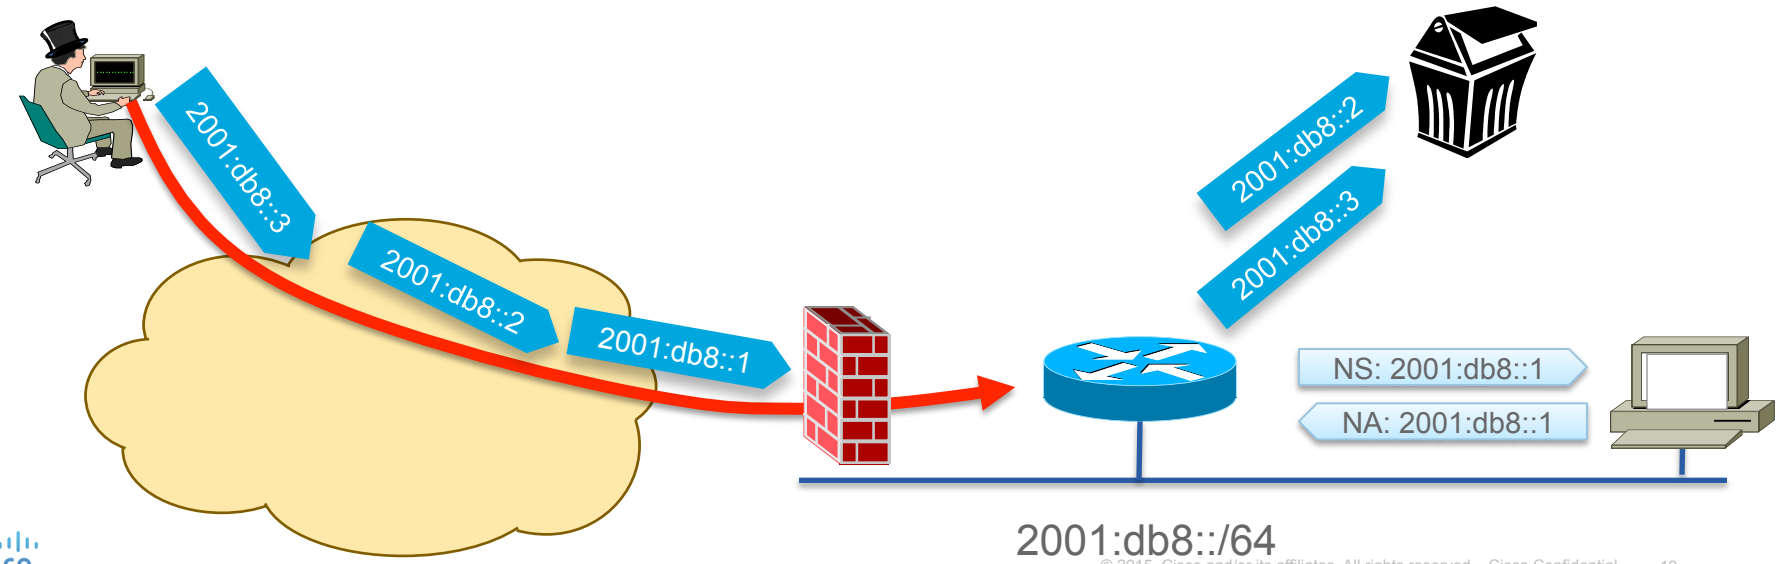

12

ARP Spoofing is now NDP Spoofing: Threats

- ARP is replaced by Neighbor Discovery Protocol
  - Nothing authenticated
  - Static entries overwritten by dynamic ones
- Stateless Address Autoconfiguration
  - rogue RA (malicious or not)
  - All nodes badly configured
    - DoS
    - Traffic interception (Man In the Middle Attack)
- Attack tools exist (from THC – The Hacker Choice)
  - Parasit6
  - Fakerouter6
  - ...

Mitigating Rogue RA: Host Isolation

- Prevent Node-Node Layer-2 communication by using:
  - Private VLANs (PVLAN) where nodes (isolated port) can only contact the official router (promiscuous port)
  - WLAN in ‘AP Isolation Mode’
  - 1 VLAN per host (SP access network with Broadband Network Gateway)
- Link-local multicast (RA, DHCP request, etc) sent only to the local official router: no harm
- Can break DAD
  - Advertise the SLAAC prefix without the on-link bit to force router to do ‘proxy-ND’

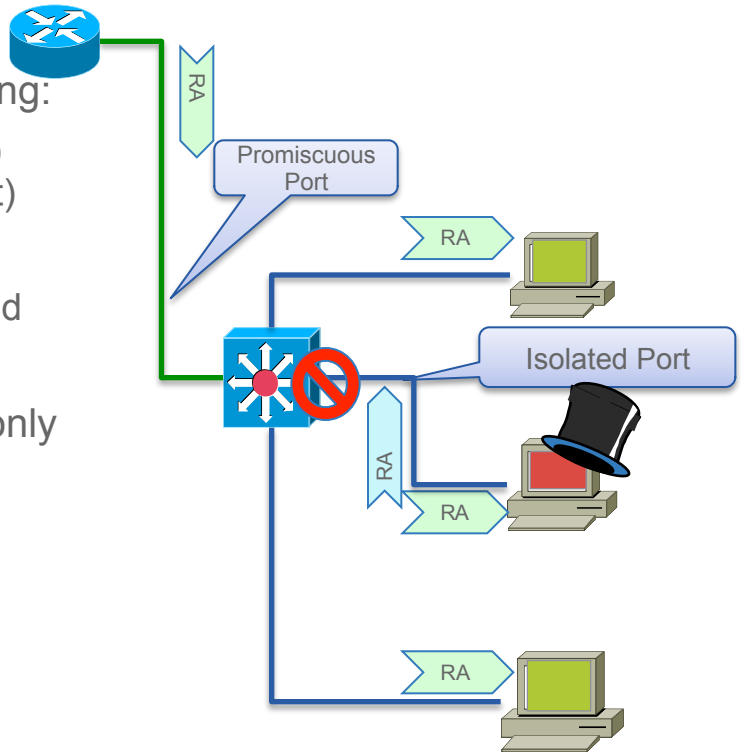

Secure Neighbor Discovery (SeND) RFC 3971

- **Cryptographically Generated Addresses** (CGA)
  - IPv6 addresses whose interface identifiers are cryptographically generated
- **RSA signature** option
  - Protect all messages relating to neighbor and router discovery
- Timestamp and nonce options
  - Prevent replay attacks
- Certification paths for authorized Routers
  - Anchored on trusted parties, expected to certify the authority of the routers on some prefixes
- Requires IOS 12.4(24)T
- **Not available on host OS** (Windows, OS/X, Android, iOS, …)

Securing Link Operations: First Hop Trusted Device

- Advantages
  - central administration, central operation
  - Complexity limited to first hop
  - Transitioning lot easier
  - Efficient for threats coming from the link
  - Efficient for threats coming from outside

- Disadvantages
  - Applicable only to certain topologies
  - Requires first-hop to learn about end-nodes
  - First-hop is a bottleneck and single-point of failure

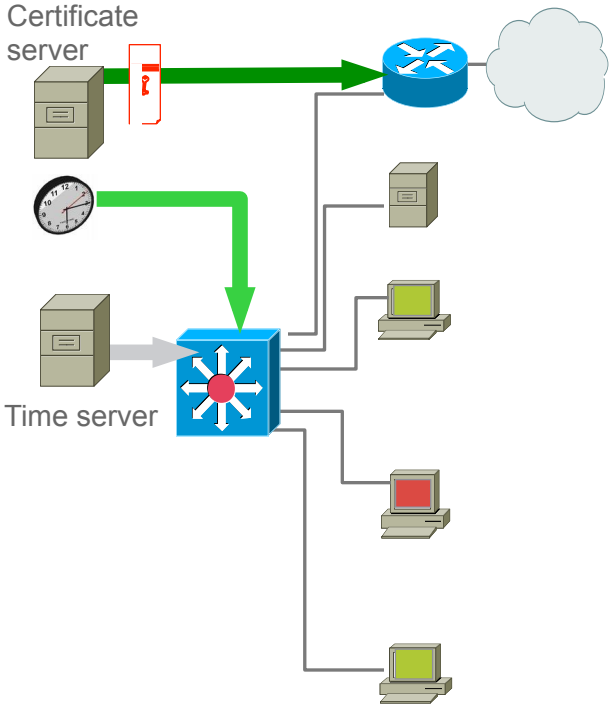

First Hop Security: RAguard since 2010 - RFC 6105

- **Port ACL** blocks all ICMPv6 RA from hosts

```
interface FastEthernet0/2
  ipv6 traffic-filter ACCESS_PORT in
  access-group mode prefer port
```

- **RA-guard lite** (12.2(33)SXI4 & 12.2(54)SG): also dropping all RA received on this port

```
interface FastEthernet0/2
  ipv6 nd raguard
  access-group mode prefer port
```

Can also enforce MTU, prefix, ... In RA

- **RA-guard** (12.2(50)SY, 15.0(2)SE)

```
ipv6 nd raguard policy HOST device-role host
ipv6 nd raguard policy ROUTER device-role router
ipv6 nd raguard attach-policy HOST vlan 100
interface FastEthernet0/0
  ipv6 nd raguard attach-policy ROUTER
```

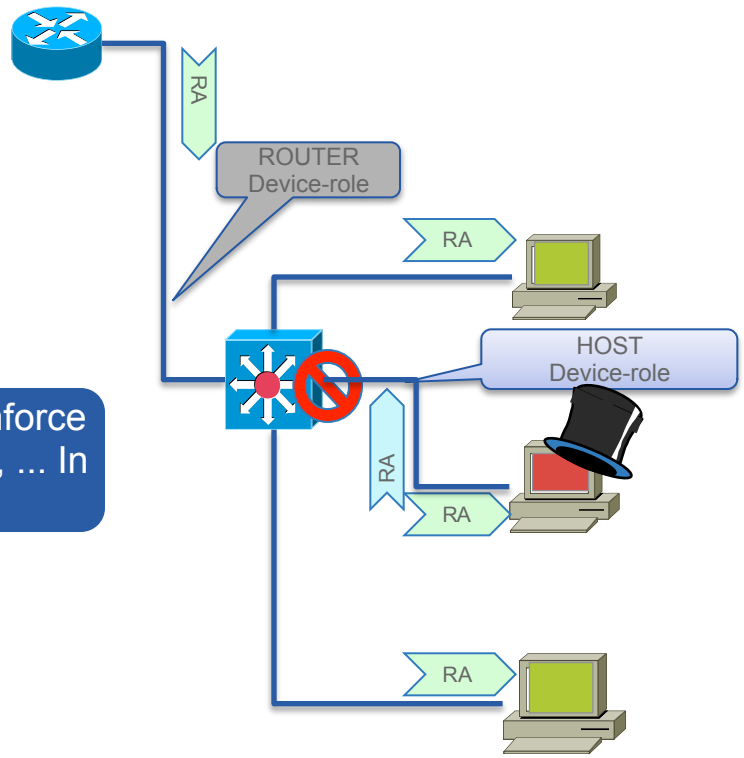

Control Plane Protection

Control Plane Policing for IPv6 Protecting the Router CPU

- Against DoS with NDP, Hop-by-Hop, Hop Limit Expiration...
- Software routers (ISR, 7200): works with CoPPr (CEF exceptions)
- See also RFC 6192
- Rate limiters

```
policy-map COPPr
 class ICMP6_CLASS
   police 8000
 class OSPF_CLASS
   police 200000
 class class-default
   police 8000
!
control-plane cef-exception
 service-policy input COPPr
```

Data Plane

DoS Example
Ping-Pong over Physical Point-to-Point

- Same as in IPv4, on real P2P without NDP, if not for me, then send it on the other side... Could produce looping traffic
- Classic IOS and IOS-XE platforms implement RFC 4443 **so this is not a threat**
  - Except on 76xx see CSCtg00387 (tunnels) and few others
  - IOS-XR see CSCsu62728
  - **Else use /127 on P2P link** (see also RFC 6164)
  - Or use infrastructure ACL or only link-local addresses

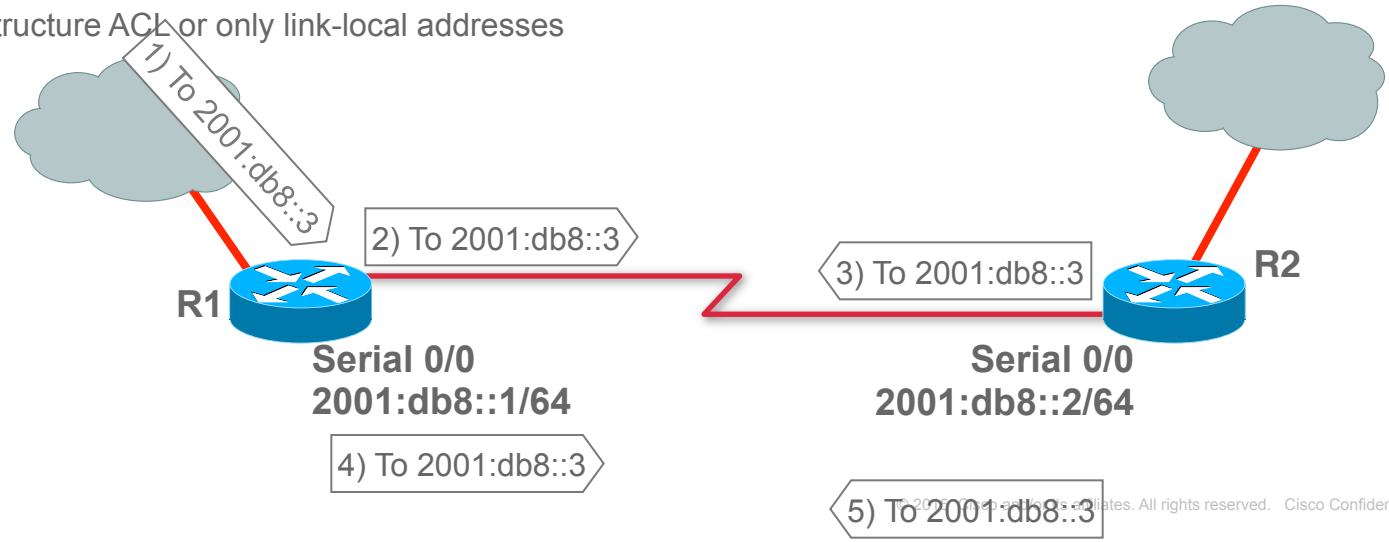

IPv6 Bogon and Anti-Spoofing Filtering

- IPv6 nowadays has its bogons:
  - http://www.team-cymru.org/Services/Bogons/fullbogons-ipv6.txt
- Every network should implement two forms of anti-spoofing protections:
  - Prevent spoofed addresses from entering the network
  - Prevent the origination of packets containing spoofed source addresses
- Anti-spoofing in IPv6 same as IPv4
  - => Same technique for single-homed edge= uRPF

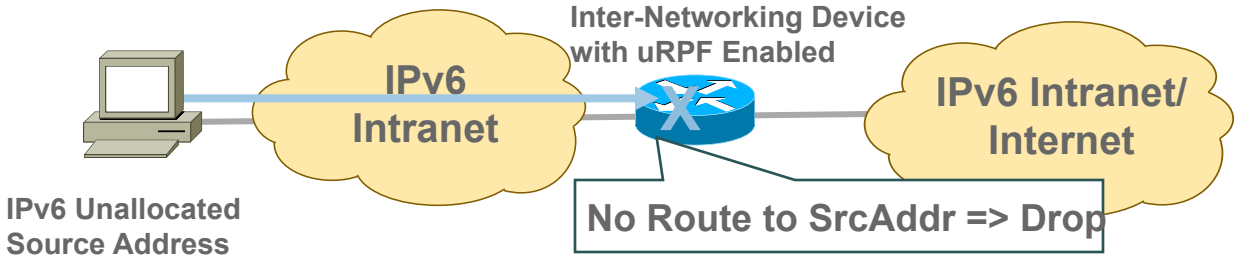

Remote Triggered Black Hole

- RFC 5635 RTBH is easy in IPv6 as in IPv4
- uRPF is also your friend for blackholing a source
- RFC 6666 has a specific discard prefix
  - 100::/64

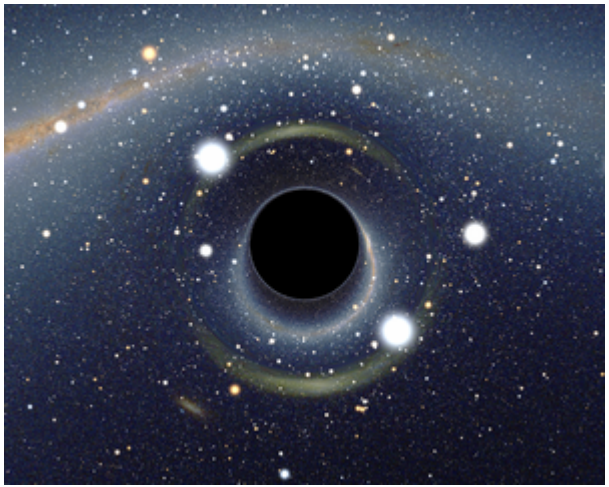

Source: Wikipedia Commons

- http://www.cisco.com/web/about/security/intelligence/ipv6_rtbh.html

Parsing the Extension Header Chain

- Finding the layer 4 information is not trivial in IPv6
  - Skip all known extension header
  - Until either known layer 4 header found => **MATCH**
  - Or unknown extension header/layer 4 header found... => **NO MATCH**

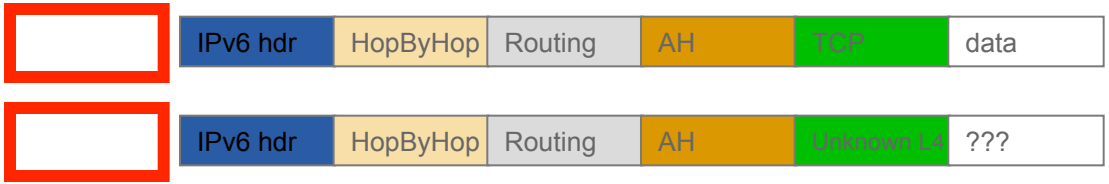

IOS IPv6 Extended ACL

- Can match on
  - Upper layers: TCP, UDP, SCTP port numbers, ICMPv6 code and type
  - TCP flags SYN, ACK, FIN, PUSH, URG, RST
  - Traffic class (only six bits/8) = DSCP, Flow label (0-0xFFFFF)
- IPv6 extension header
  - `routing` matches any RH, `routing-type` matches specific RH
  - `mobility` matches any MH, `mobility-type` matches specific MH
  - `dest-option` matches any destination options
  - `auth` matches AH
  - `hbh` matches hop-by-hop (since 15.2(3)T)
- `fragments` keyword matches
  - Non-initial fragments
- `undetermined-transport` keyword does not match if
  - TCP/UDP/SCTP and ports are in the fragment
  - ICMP and type and code are in the fragment
  - Everything else matches (including OSPFv3, …)
  - Only for deny ACE

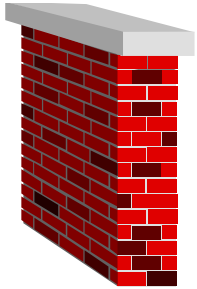

Check your platform & release as your mileage can vary…

Telemetry

Available Tools

- Usually IPv4 telemetry is available
- SNMP MIB
  - Not always available yet on Cisco gears
- Flexible Netflow for IPv6
  - Available in : 12.4(20)T, 12.2(33)SRE
  - Public domain tools: nfsen, nfdump, nfcpad…

Flexible Flow Record: IPv6 Key Fields (Version 9)

IPv6	
IP (Source or Destination)	Payload Size
Prefix (Source or Destination)	Packet Section (Header)
Mask (Source or Destination)	Packet Section (Payload)
Minimum-Mask (Source or Destination)	DSCP
Protocol	Extension
Traffic Class	Hop-Limit
Flow Label	Length
Option Header	Next-header
Header Length	Version
Payload Length	

Routing
Destination AS
Peer AS
Traffic Index
Forwarding Status
Is-Multicast
IGP Next Hop
BGP Next Hop

Flow
Sampler ID
Direction

Interface
Input
Output

Transport	
Destination Port	TCP Flag: ACK
Source Port	TCP Flag: CWR
ICMP Code	TCP Flag: ECE
ICMP Type	TCP Flag: FIN
IGMP Type	TCP Flag: PSH
TCP ACK Number	TCP Flag: RST
TCP Header Length	TCP Flag: SYN
TCP Sequence Number	TCP Flag: URG
TCP Window-Size	UDP Message Length
TCP Source Port	UDP Source Port
TCP Destination Port	UDP Destination Port
TCP Urgent Pointer	

Flexible Flow Record: IPv6 Extension Header Map

Bits 11-31	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Res	ESP	AH	PAY	DST	HOP	Res	UNK	FRA0	RH	FRA1	Res

- FRA1: Fragment header – not first fragment
- RH: Routing header
- FRA0: Fragment header – First fragment
- UNK: Unknown Layer 4 header (compressed, encrypted, not supported)
- HOP: Hop-by-hop extension header
- DST: Destination Options extension header
- PAY: Payload compression header
- AH: Authentication header
- ESP: Encapsulating Security Payload header
- Res: Reserved

29

Summary

Key Takeaway

- As expected IPv6 secure operations are quite similar to IPv4 (Main differences at layer 2)
- Management plane
  - Protect management plane with access-class
- Control plane
  - Authenticate IGP
  - Consider the use of link-local on P-P links?
  - Mitigate rogue-RA with RA-guard
  - Configure control plane policing
- Data plane
  - Beware of ping-pong on not /127 real P2P link
  - Apply anti-spoofing, anti-bogons
  - Use ACL where applicable, ACL must permit NDP
- Telemetry
  - SNMP MIB and Netflow v9 are your friends
  - Netflow can be used for inventory

TOMORROW starts here.