How to Securely Operate an IPv6 Network

Eric Vyncke, evyncke@cisco.com

@evyncke

```
OPSEC                                                K. Chittimaneni
Internet-Draft                                                Google
Intended status: Informational                               M. Kaeo
Expires: April 25, 2014                         Double Shot Security
                                                           E. Vyncke
                                                       Cisco Systems
                                                    October 22, 2013


        Operational Security Considerations for IPv6 Networks
                       draft-ietf-opsec-v6-04
```

Foreword

- All topics common to IPv4/IPv6 are unchanged:
  - Physical security
  - Role Base Access Control
  - ....
- I took the liberty to include Cisco configuration (as it may be useful for you) but I will not detail them

Agenda

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

Management Plane

Management over IPv6

- SSH, syslog, SNMP, NetFlow, RADIUS 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 (more to come) as well as out-of-band management

- So, protect all SNMP, SSH access from untrusted interfaces




Control Plane: Routing Protocols

Preventing IPv6 Routing Attacks
Protocol Authentication

- BGP, IS-IS, EIGRP no change:
  - An MD5 authentication of the routing update
- OSPFv3 originally has changed and pulled MD5 authentication from the protocol and instead rely on transport mode IPsec (for authentication and confidentiality)
  - But see RFC 6506 *(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

For Your Reference

- 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 ☺**
  - 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
```

```
OPsec Working Group                                  M. Behringer
Internet-Draft                                           E. Vyncke
Intended status: Informational                              Cisco
Expires: July 10, 2014                            January 6, 2014
```

Using Only Link-Local Addressing Inside an IPv6 Network
draft-ietf-opsec-lla-only-06

LLA-Only Pros and Cons

Benefits:

- no remote attack against your infrastructure links: implicit infrastructure ACL*
- Smaller routing table (links do not appear)
- Simpler configuration
- Easier to renumber

Cons:

- need to provision loopback for:
  - ICMP for Traceroute
  - ICMP for PMTUD
  - SNMP/NetFlow/syslog/ ...
- No interface ping

Special case for IXP:

- Usually a specific /64 which is not routed => uRPF will drop ICMP generated (PMTUd) by routers in the IXP
- LLA-only on the IXP interfaces => ICMP are generated from a non IXP interface

**: loopbacks are still routable/reachable*




Control Plane: Neighbor Discovery

Scanning Made 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/…

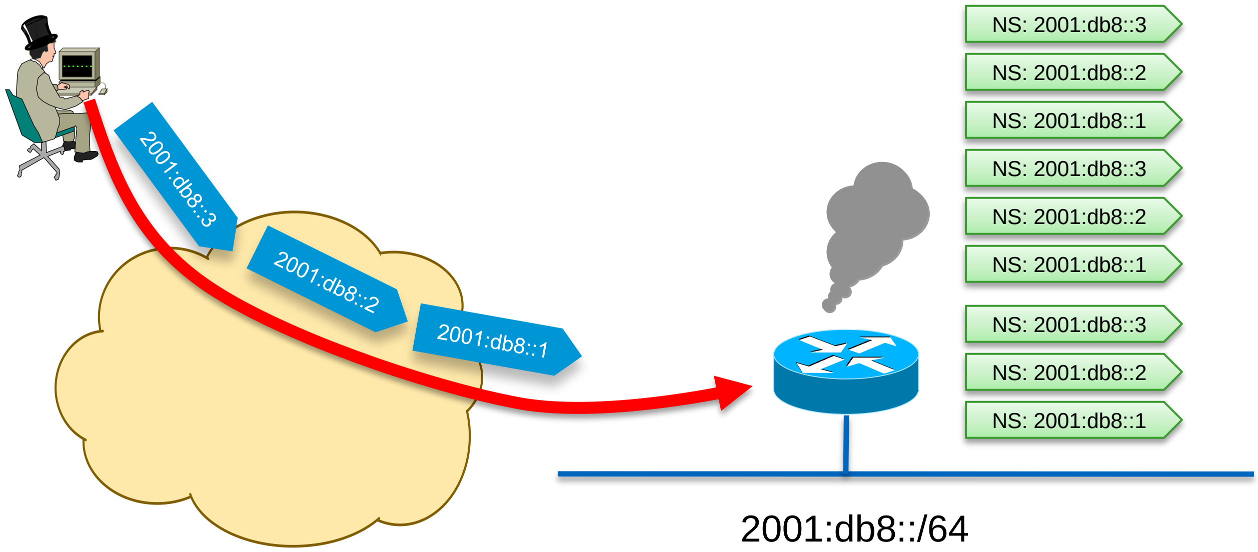

Mitigating Remote Neighbor Cache Exhaustion

- 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

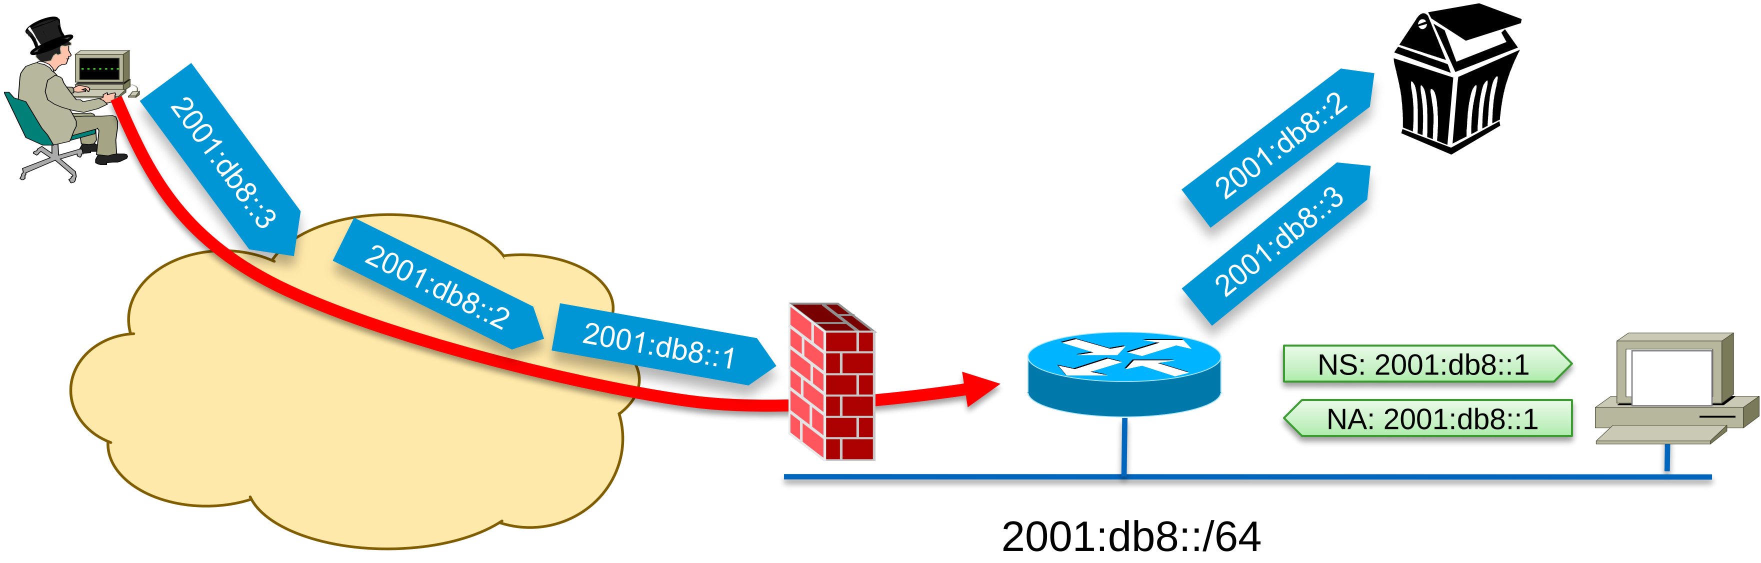

ARP Spoofing is now NDP Spoofing: Threats

- 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
  - ...

The Hacker's Choice

ARP Spoofing is now NDP Spoofing: Mitigation

- **GOOD NEWS**: dynamic ARP inspection for IPv6 is available
  - First phase (Port ACL & RA Guard) available since Summer 2010
  - Second phase (NDP & DHCP snooping) starting to be available since Summer 2011
  - http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-first_hop_security.html
- **(Kind of) GOOD NEWS**: Secure Neighbor Discovery
  - SeND = NDP + crypto
  - IOS 12.4(24)T But not in Windows Vista, 2008 and 7, Mac OS/X, iOS, Android
- Other **GOOD NEWS**:
  - Private VLAN works with IPv6
  - Port security works with IPv6
  - IEEE 801.X works with IPv6 (except downloadable ACL)

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'

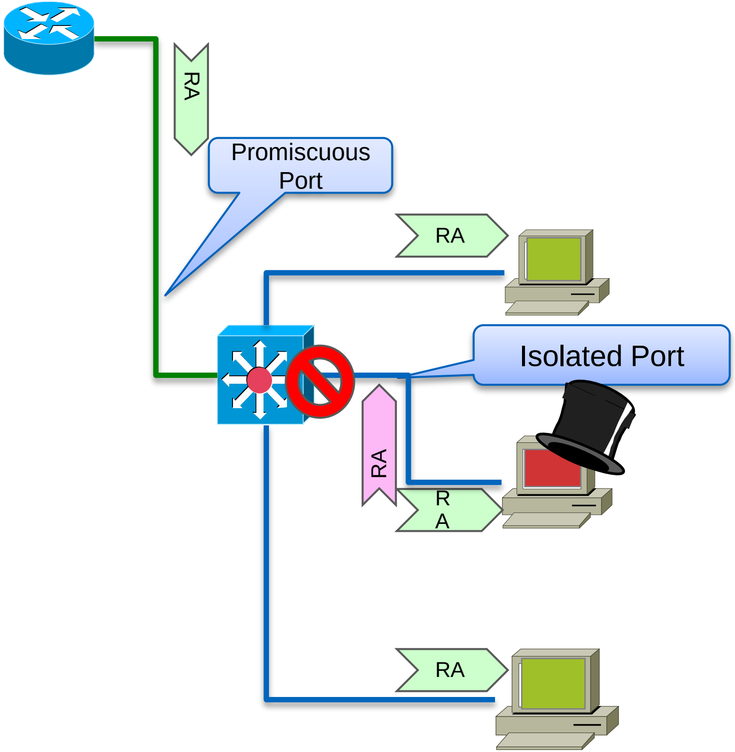

First Hop Security: RAguard since 2010 RFC 6105

- **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
```

- **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
```

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

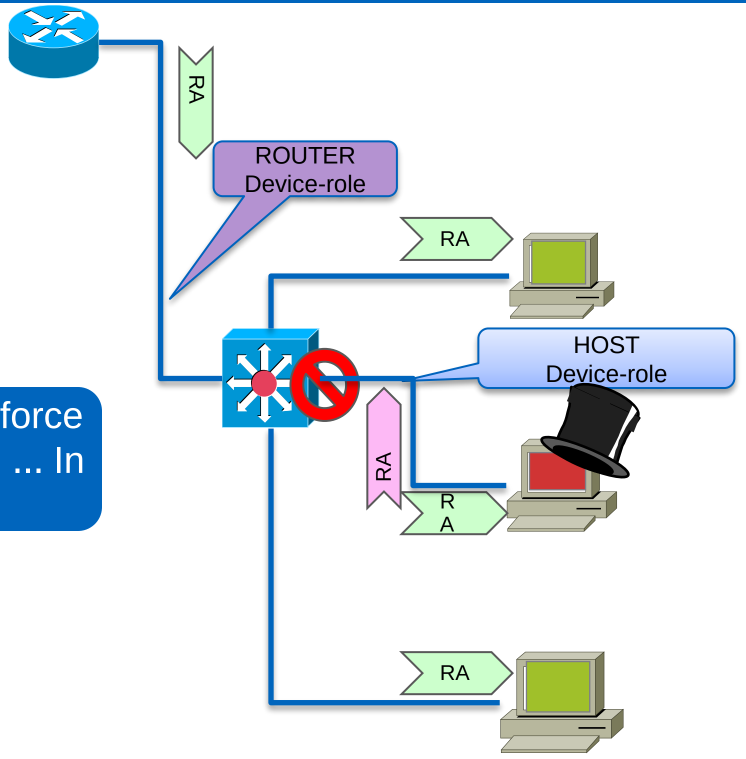




Control Plane Protection

Control Plane Policing for IPv6
Protecting the Router CPU

For Your Reference

- Against DoS with NDP, Hop-by-Hop, Hop Limit Expiration...
- See also RFC 6192

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

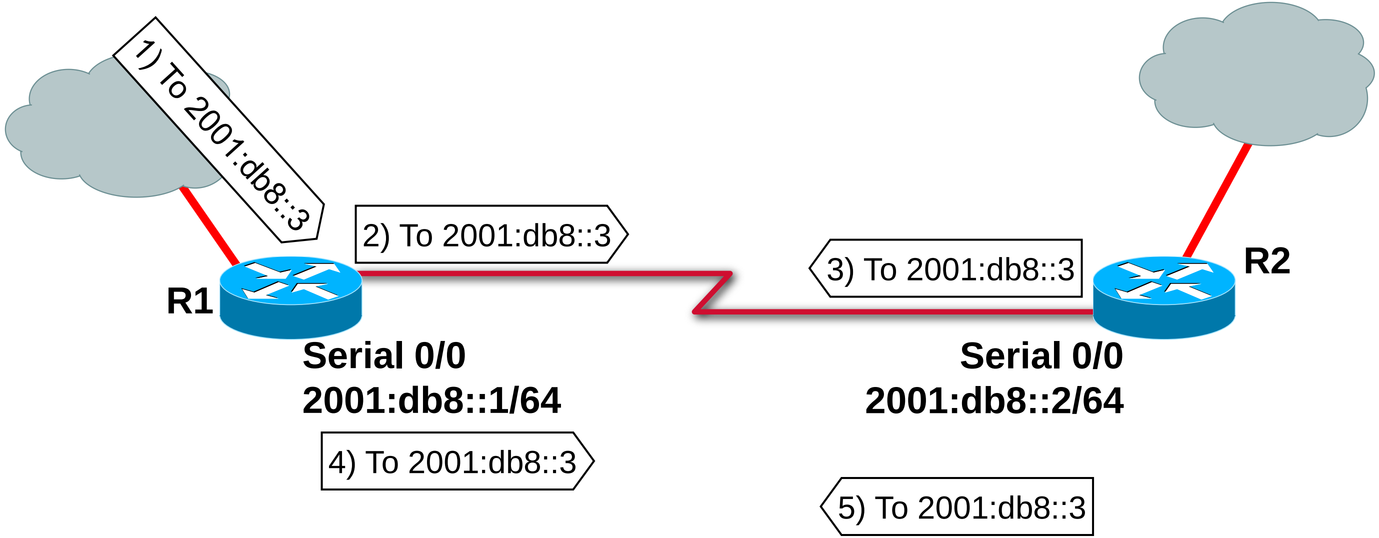

IPv6 Bogon and Anti-Spoofing Filtering

- 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

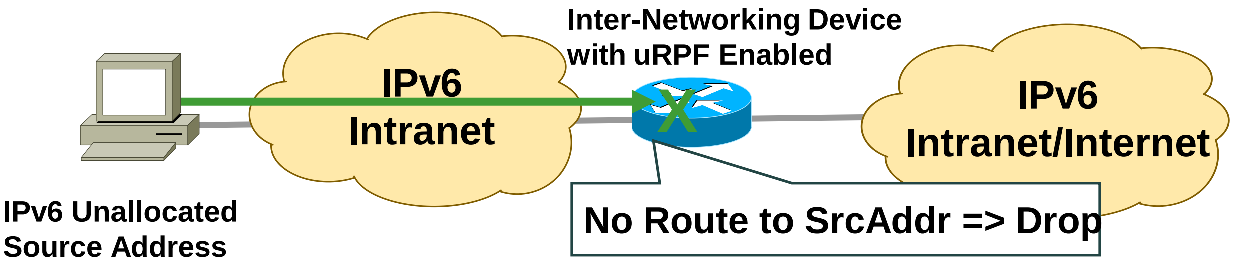

Bogons Filtering

- Detailed & updated list at:
  - http://www.team-cymru.org/Services/Bogons/fullbogons-ipv6.txt
- Or simpler but more relaxed

```
ipv6 access-list NO_BOGONS
    remark Always permit ICMP unreachable (Path MTU Discovery & co)
    permit icmp any any unreachable
    remark Permit only large prefix blocks from IANA
    permit ip 2001::/16 any
    permit ip 2002::/16 any
    permit ip 2003::/18 any
    permit ip 2400::/12 any
    permit ip 2600::/10 any
    permit ip 2800::/12 any
    permit ip 2a00::/12 any
    permit ip 2c00::/12 any
    Remark implicit deny at the end (but see later)
```

Source: http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xml

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

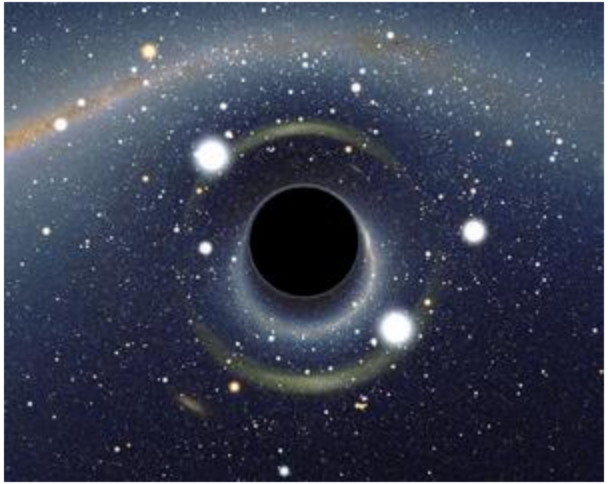
Source: Wikipedia Commons

- 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**

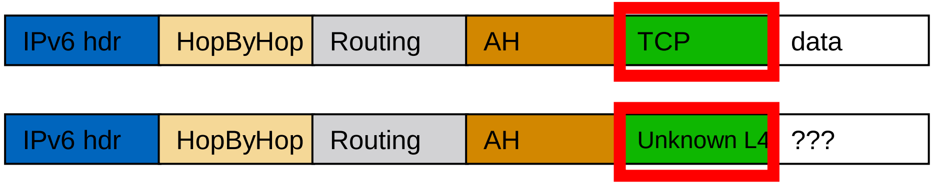

IOS IPv6 Extended ACL

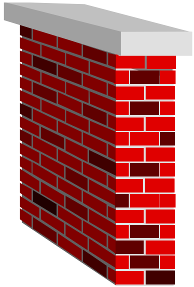

- Can match on
  - Upper layers: TCP, UDP, SCTP port numbers, ICMPv6 code and type
  - TCP flags SYN, ACK, FIN, PUSH, URG, RST
  - Traffic class (only six bits/8) = DSCP, Flow label (0-0xFFFFF)
- IPv6 extension header
  - `routing` matches any RH, `routing-type` matches specific RH
  - `mobility` matches 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

CRITICAL without this, there is a way to bypass STATELESS ACL!

Check your platform & release as your mileage can vary…

Looping Attack Between 6to4 and ISATAP (RFC 6324)

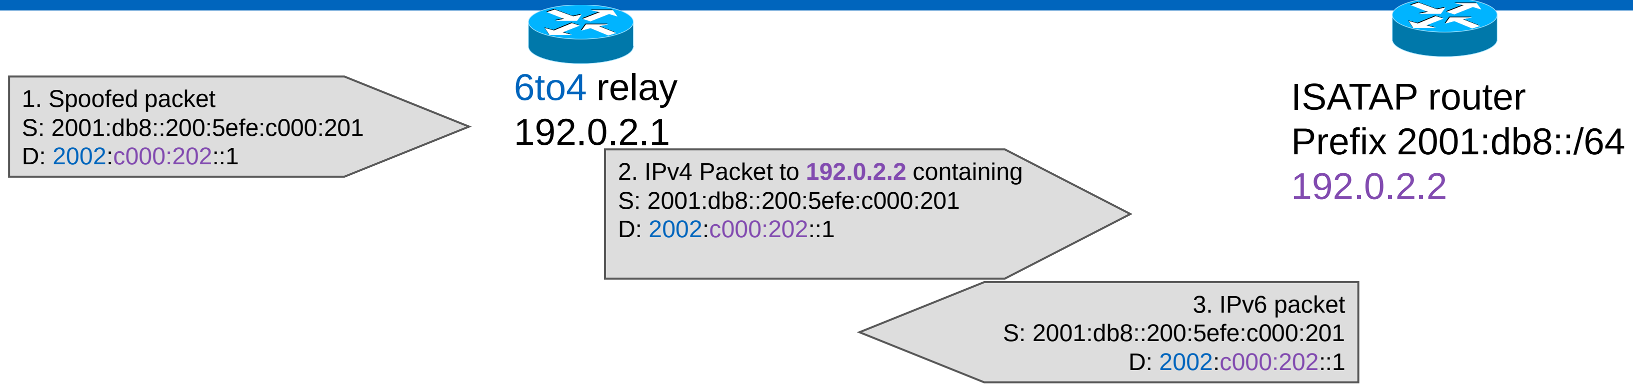

Repeat until Hop Limit == 0

- Root cause
  - Same IPv4 encapsulation (protocol 41)
  - Different ways to embed IPv4 address in the IPv6 address
- ISATAP router:
  - accepts 6to4 IPv4 packets
  - Can forward the inside IPv6 packet back to 6to4 relay
- Symmetric looping attack exists

Mitigation:
- Easy on ISATAP routers: deny packets whose IPv6 is its 6to4
- Less easy on 6to4 relay: block all ISATAP-like local address?
- Good news: not so many open ISATAP routers on the Internet
- Do not announce the 6to4 relay address outside of your AS and accepts protocol-41 packets only from your AS

6rd Relay Security Issues

- 6rd is more constrained than 6to4, hence more secure
- IPv4 ACL (or IPv4 routing) can limit the 6rd packets to the 6rd domain within the ISP
  - No more open relay
  - No more looping attacks

IPv6 security is similar to IPv4 security
No excuse to operate an insecure IPv6 network




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…

IPv6 MIB Implementation

For Your Reference

	IP FWD (ROUTES)	IP	ICMP	TCP	UDP
Original IPv4 only	2096	2011		2012	2013
IPv6 only	2465		2466	2452	2454
Protocol Version Independent (PVI)	rfc2096-update = 4292	rfc2011-update = 4293 = IP-MIB			
				rfc2012-update = 4022	rfc2013-update = 4113

- IPv4/IPv6 stats can be monitored from CLI “show interface accounting” on most platforms
- RFC 4292 and 4293 – Interface Stats table are added, also required HW support
- Tunnel MIB (RFC 4087)

Using SNMP to Read Interfaces Traffic

For Your Reference

```
evyncke@charly:~$ snmpwalk -c secret -v 1 udp6:[2001:db8::1] -Cw 70 -m IP-MIB
ipNetToPhysicalPhysAddress

SNMP table: IP-MIB::ipIfStatsTable

  index ipIfStatsInReceives ipIfStatsHCInReceives ipIfStatsInOctets
 ipv4.1              683929                     ?          55054803
 ipv4.2             1123281                     ?         107467461
 ipv6.1              152612                     ?          17261398
 ipv6.2            15083935                     ?        2131680450


evyncke@charly:~$ snmpwalk -c secret -v 1 udp6:[2001:db8::1] -Cw 70 ifTabl

SNMP table: IF-MIB::ifTable

 index ifIndex          ifDescr            ifType ifMtu    ifSpeed
     1       1 FastEthernet0/0   ethernetCsmacd  1500  100000000
     2       2 FastEthernet0/1   ethernetCsmacd  1500  100000000
```

Using SNMP to Read IPv4/IPv6 Neighbors Cache

For Your Reference

```
evyncke@charly:~$ snmpwalk -c secret -v 1 udp6:[2001:db8::1] -m IP-MIB
ipNetToPhysicalPhysAddress

IP-MIB::ipNetToPhysicalPhysAddress.1.ipv4."192.168.0.2" = STRING: 0:13:c4:43:cf:e

IP-MIB::ipNetToPhysicalPhysAddress.1.ipv4."192.168.0.3" = STRING: 0:23:48:2f:93:24

IP-MIB::ipNetToPhysicalPhysAddress.1.ipv4."192.168.0.4" = STRING: 0:80:c8:e0:d4:be

...

IP-
MIB::ipNetToPhysicalPhysAddress.2.ipv6."2a:02:05:78:85:00:01:01:02:07:e9:ff:fe:f2:a0:c6
" = STRING: 0:7:e9:f2:a0:c6

IP-
MIB::ipNetToPhysicalPhysAddress.2.ipv6."2a:02:05:78:85:00:01:01:02:20:4a:ff:fe:bf:ff:5f
" = STRING: 0:20:4a:bf:ff:5f

IP-
MIB::ipNetToPhysicalPhysAddress.2.ipv6."2a:02:05:78:85:00:01:01:30:56:da:9d:23:91:5e:ea
" = STRING: 78:ca:39:e2:43:3

...
```

```
evyncke@charly:~$ snmptable -c secret -v 1 udp6:[2001:db8::1] -Ci -m IP-MIB
ipNetToPhysicalTable
```

Flexible Flow Record: IPv6 Key Fields

For Your Reference

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

For Your Reference

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

Netflow Reverse Usage

- Scanning an IPv6 network is impossible (address space too large)
- **How can we run a security audit?**
- Easy
  - Get all IPv6 addresses from Netflow
  - Note: scanning link-local addresses requires layer-2 adjacency, i.e.
    - Ping6 ff02::1

Vulnerability Scanning in a Dual-Stack World

- Finding all hosts:
  - Address enumeration does not work for IPv6
  - Need to rely on DNS or NDP caches or NetFlow
- Vulnerability scanning
  - IPv4 global address, IPv6 global address(es) (if any), IPv6 link-local address
  - Some services are single stack only (currently mostly IPv4 but who knows...)
  - Personal firewall rules could be different between IPv4/IPv6
- **IPv6 vulnerability scanning MUST be done for IPv4 & IPv6 even in an IPv4-only network**
  - IPv6 link-local addresses are active by default




Forensic

Multiple Facets to IPv6 Addresses

- Every host can have multiple IPv6 addresses simultaneously
  – Need to do correlation!
  – Alas, no Security Information and Event Management (SIEM) supports IPv6
  – Usually, a customer is identified by its /48 ☺
- Every IPv6 address can be written in multiple ways
  – 2001:0DB8:0BAD::0DAD
  – 2001:DB8:BAD:0:0:0:0:DAD
  – 2001:db8:bad::dad (this is the canonical RFC 5952 format)
  – => Grep cannot be used anymore to sieve log files…

Perl Grep6

For Your Reference

```
#!/usr/bin/perl -w
use strict ;
use Socket ;
use Socket6 ;

my (@words, $word, $binary_address, $address) ;

$address = inet_pton AF_INET6, $ARGV[0] ;
if (! $address) { die "Wrong IPv6 address passed as argument" ; }

## go through the file one line at a time
while (my $line = <STDIN>) {
          @words = split /[ \n\(\)\[\]]/, $line ;
          foreach $word (@words) {
                    $binary_address = inet_pton AF_INET6, $word ;
                    if ($binary_address and $binary_address eq $address) {
                              print $line ;
                              next ;
                    }
          }
}
```

How to Find the MAC Address of an IPv6 Address?

- Easy if EUI-64 format as MAC is embedded
  - 2001:db8::0226:bbff:fe4e:9434
    - *(need to toggle bit 0x20 in the first MAC byte = U/L)*
  - Is 00:26:bb:4e:94:34

How to Find the MAC Address of an IPv6 Address?

- DHCPv6 address or prefix… the client DHCP Unique ID (DUID) can be
  - MAC address: trivial
  - Time + MAC address: simply take the last 6 bytes
  - Vendor number + any number: no luck… next slide can help
  - No guarantee of course that DUID includes the real MAC address.

```
# show ipv6 dhcp binding
Client: FE80::225:9CFF:FEDC:7548
  DUID: 000100010000000A00259CDC7548
  Username : unassigned
  Interface : FastEthernet0/0
  IA PD: IA ID 0x0000007B, T1 302400, T2 483840
    Prefix: 2001:DB8:612::/48
            preferred lifetime 3600, valid lifetime 3600
            expires at Nov 26 2010 01:22 PM (369)
```


DHCPv6 in Real Live…

- Not so attractive ☹
- Only supported in Windows Vista, and Windows 7, Max OS/X Lion
  - Not in Linux (default installation), …
- Windows Vista does not place the used MAC address in DUID but any MAC address of the PC

```
# show ipv6 dhcp binding
Client: FE80::FDFA:CB28:10A9:6DD0
  DUID: 0001000110DB0EA6001E33814DEE
  Username : unassigned
  IA NA: IA ID 0x1000225F, T1 300, T2 480
    Address: 2001:DB8::D09A:95CA:6918:967
            preferred lifetime 600, valid lifetime
600
            expires at Oct 27 2010 05:02 PM (554
seconds)
```

Actual MAC address:
0022.5f43.6522

How to Find the MAC Address of an IPv6 Address?

- Last resort… look in the live NDP cache (CLI or SNMP)

```
#show ipv6 neighbors 2001:DB8::6DD0
IPv6 Address            Age Link-layer Addr State Interface

2001:DB8::6DD0            8 0022.5f43.6522  STALE Fa0/1
```

- If no more in cache, then you should have scanned and saved the cache…
- EEM can be your friend
- First-Hop Security phase II can generate a syslog event on each new binding
  - **`ipv6 neighbor binding logging`**




Summary

Our journey...

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

Key Takeaway /1

- 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
  - Disable source routing
  - Use ACL where applicable
    - ACL must permit NDP

Key Takeaway /2

- Telemetry
  - SNMP MIB and Netflow v9 are your friends
  - Netflow can be used for inventory
- Forensic
  - Multiple addresses per node, multiple ways to write an IPV6 address
  - Finding MAC address from IPv6:
    - EUI-64,
    - DHCPv6 (not so trivial)
    - else periodic NDP cache dumps...
- Lawful Interception
  - implemented, missing mediation device




Questions and Answers?

cisco TM