ERNW
providing security.

Basic Attacks and Mitigation Strategies

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
- Conference: www.troopers.de

Agenda

- Refresher about Link-Layer Behavior of IPv6
- Basic IPv6 Security Assumptions
- Attacks inside a Layer 2 Domain
- Mitigation Strategies
- Attacks from Outside
- Mitigation Strategies
- Conclusion

ICMPv6

- ICMPv6 is the new version of ICMP.
  It was first specified in RFC 2462, latest in RFC 4443.

- ICMPv6 includes “traditional" ICMP functions, functionalities of IGMP (RFC 1112), IGMPv2 (RFC 2236) and extensions of the type "Multicast Listener Discovery” (MLD) for IPv6.

- Additionally ICMPv6 includes the Neighbor Discovery Protocol (RFC 2461, updated by RFC 4861).

- ICMPv6 is an integral part of every IPv6 implementation;
  every IPv6 stack must include ICMPv6.

- ICMPv6 has the next-header value 58.

Overall ICMPv6 Header Format

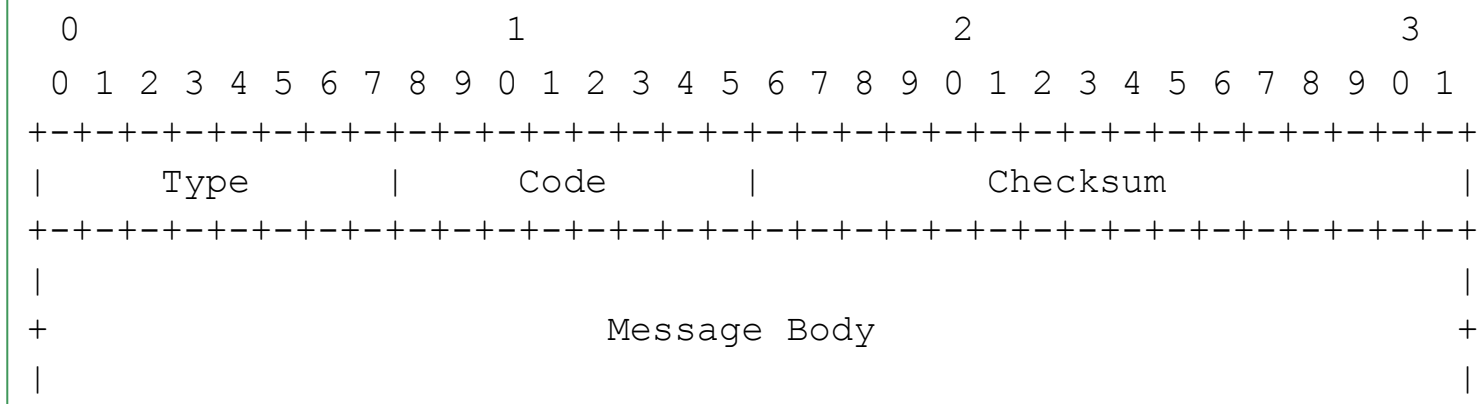

- Type: specifies the type of the message and also describes the format of the message.
- Code: Using this field, whose interpretation depends on the type, other subtypes are defined.
- Checksum: used to detect errors during transmission.
- Message Body: Depends on the type and code of the header.

(Main) ICMPv6 Types

Type(Value)	Description
1	Destination Unreachable (with codes 0,1,2,4)
2	Packet too big (Code 0)
3	Time Exceeded (Code 0,1)
4	Parameter Problem (Code 0,1,2)
128	Echo Request (Code 0)
129	Echo Reply (Code 0)
130	Multicast Listener Query
131	Multicast Listener Report
132	Multicast Listener Done
133	Router Solicitation
134	Router Advertisement
135	Neighbor Solitication
136	Neighbor Advertisement
137	Redirect

Neighbor Discovery Protocol

- *Neighbor Discovery* (ND) provides mechanisms for the following tasks:
  1. Neighbor Discovery / Address Resolution
  2. Router Discovery
  3. Prefix Discovery
  4. Parameter Discovery
  5. Address Autoconfiguration
  6. Next-Hop Determination
  7. Neighbor Unreachability Detection
  8. Duplicate Address Detection
  9. Redirect

Neighbor Discovery

- The address resolution is the exchange of *neighbor solicitation* and *neighbor advertisement* messages to the link-layer address, for example, to resolve the next hop.
  - Multicast Neighbor Solicitation Message
  - Unicast Neighbor Advertisement Message
- Both nodes involved update their *Neighbor Cache*.
- Once this is done successfully, the nodes can communicate with each other via unicast.
- Replaces the ARP (Address Resolution Protocol) in IPv4.

Neighbor Solicitation

Ethernet Header
- Dest.-MAC: 33-33-FF-03-04-05

IPv6 Header
- Source-IP: 2001::cafe:201:2FF:FE03:406
- Dest.-IP: FF02::1:FF03:405
- Hop limit: 255

Neighbor Solicitation Header
- Dest. Address is 2001::cafe:201:2FF:FE03:405

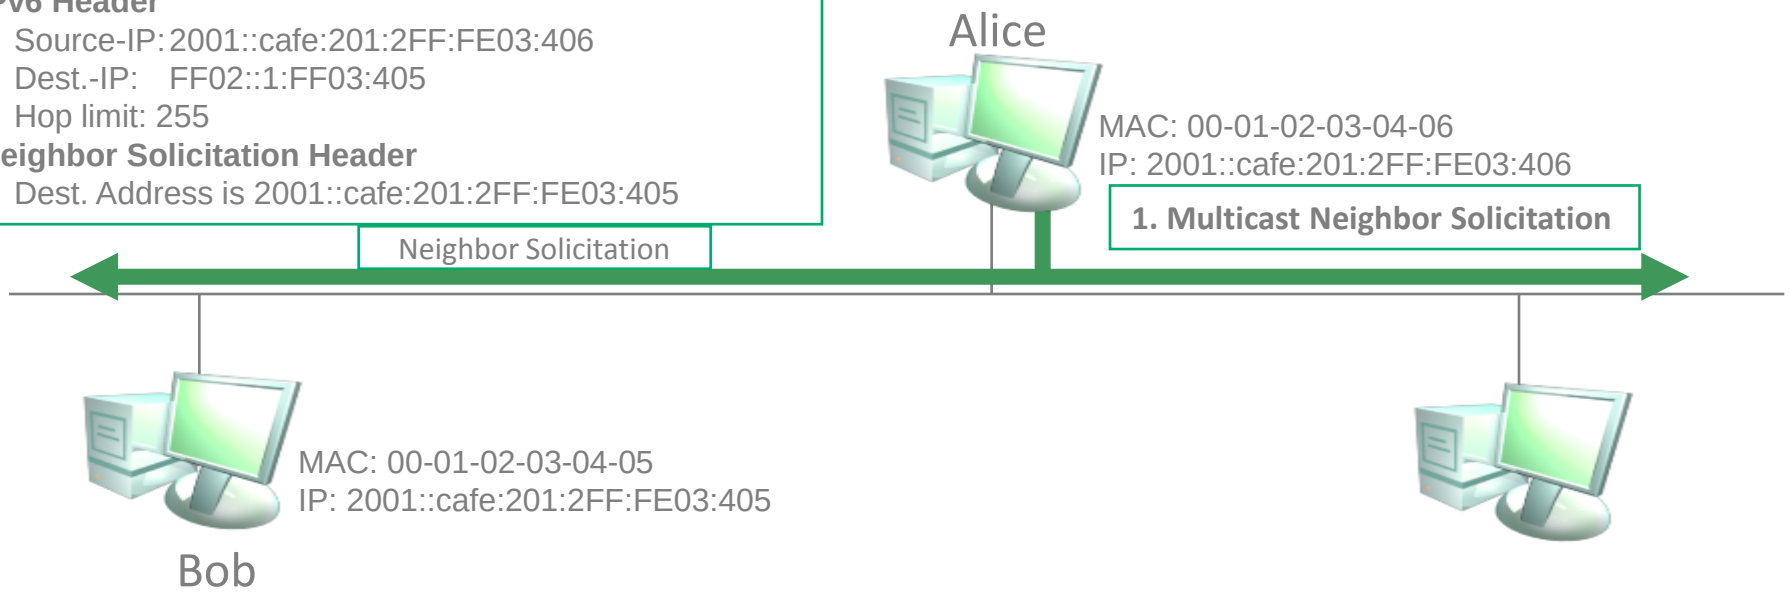

Neighbor Advertisement

Ethernet Header
Dest.-MAC: 00-01-02-03-04-06
IPv6 Header
Source-IP: 2001::cafe:201:2FF:FE03:405
Dest.-IP: 2001::cafe:201:2FF:FE03:406
Hop limit: 255
Neighbor Advertisement Header
Source Address is 2001::cafe:201:2FF:FE03:405
Neighbor Discovery Option
Source Link-Layer Address (00-01-02-03-04-05)

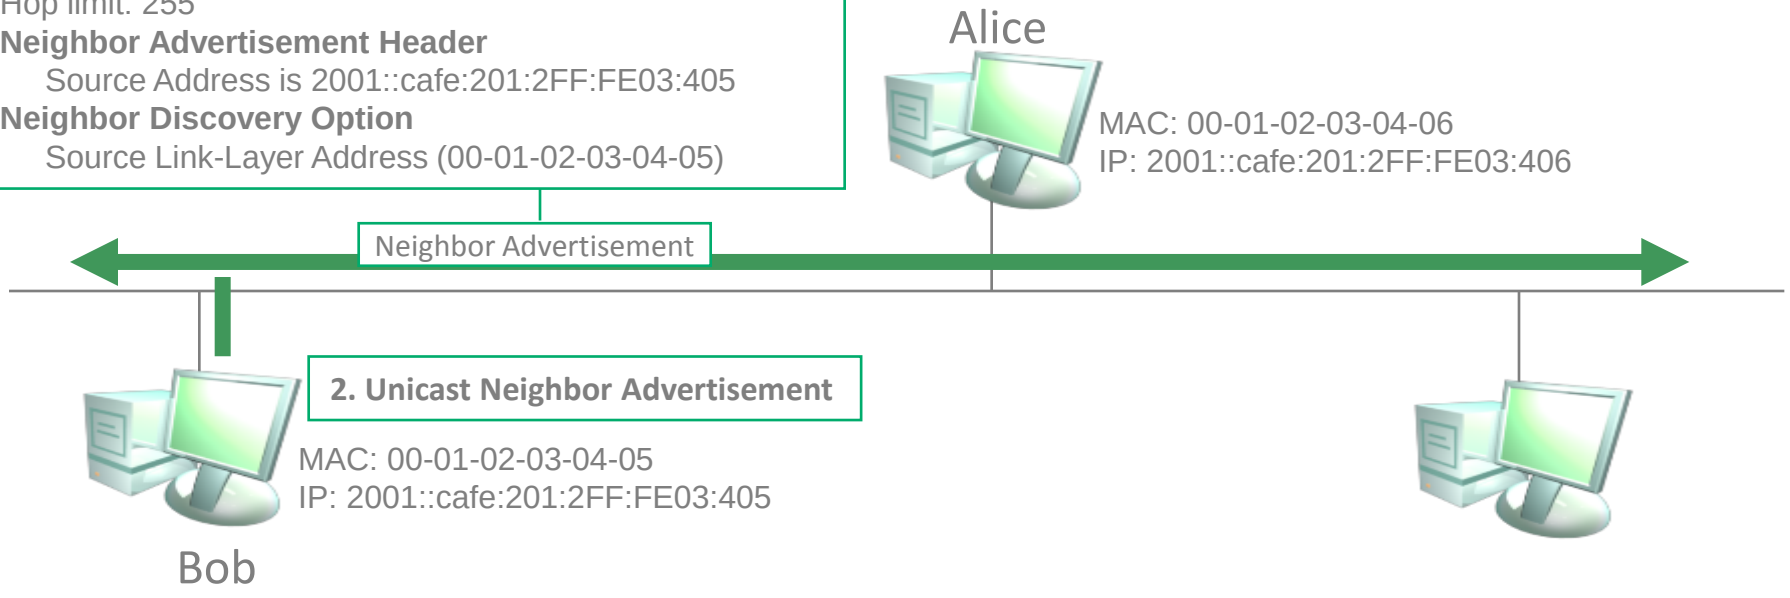

Duplicate Address Detection (DAD)

- To detect duplicate addresses neighbor solicitations are used.
  - The destination address in the neighbor solicitation message is the *Solicited-Node* IPv6 Multicast address which corresponds to the address that should be checked for uniqueness.
  - As the source address, the unspecified address "::" is used.

- If the address is already in use, the host which already uses this address, responds with a multicast *Neighbor Advertisement*.
  - The destination address is the link-local multicast address for all nodes (FF02::1 “all-nodes“).
  - Corresponds to *gratuitous ARP* functionality of IPv4.

Duplicate Address Detection

Ethernet Header
- Dest.-MAC: 33-33-FF-03-04-05

IPv6 Header
- Source-IP: ::
- Dest.-IP: FF02::1:FF03:405
- Hop limit: 255

Neighbor Solicitation Header
- Dest. Address is 2001::cafe:201:2FF:FE03:405

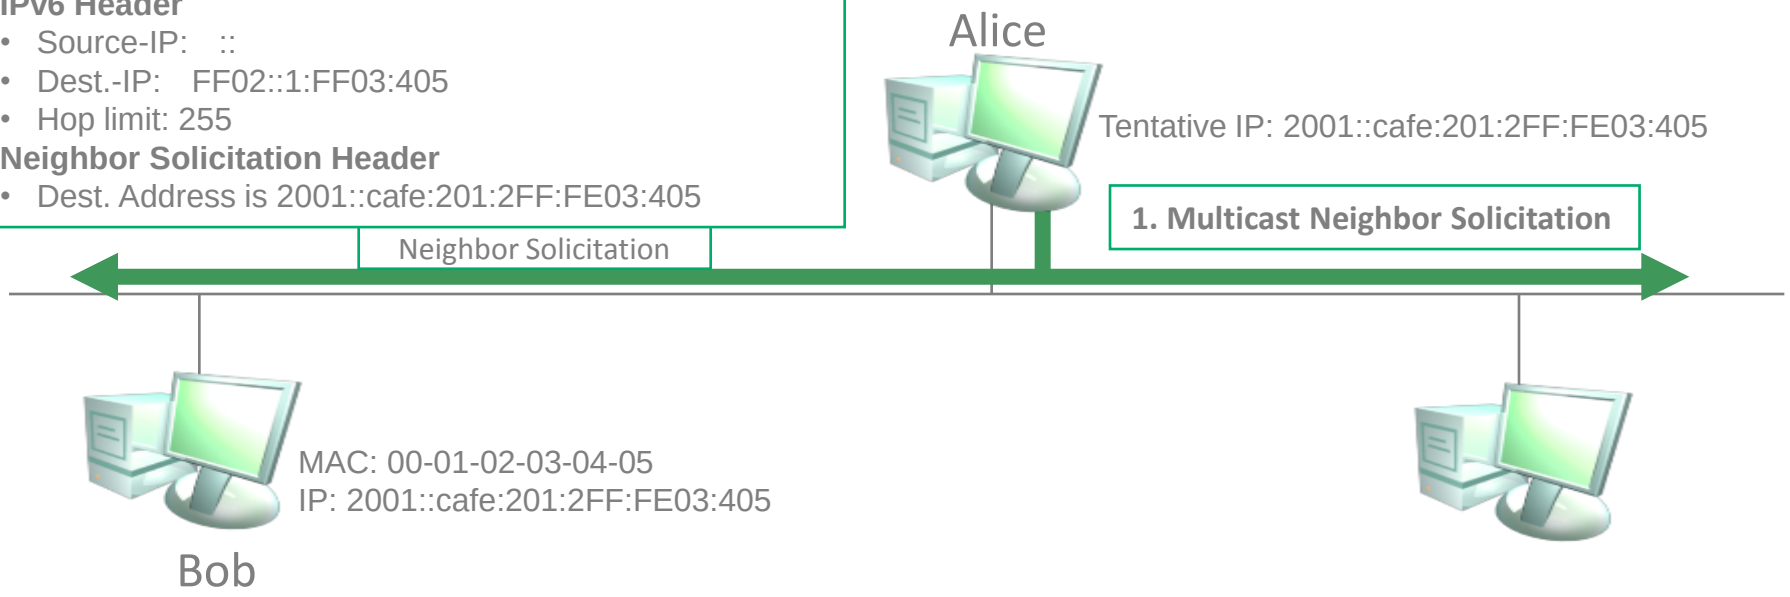

Duplicate Address Detection

Ethernet Header
- Dest.-MAC: 33-33-00-00-00-01

IPv6 Header
- Source.-IP: 2001::cafe:201:2FF:FE03:405
- Dest.-IP: FF02::1
- Hop limit: 255

Neighbor Advertisement Header
- Source Address is 2001::cafe:201:2FF:FE03:405

Neighbor Discovery Option
- Source Link-Layer Address

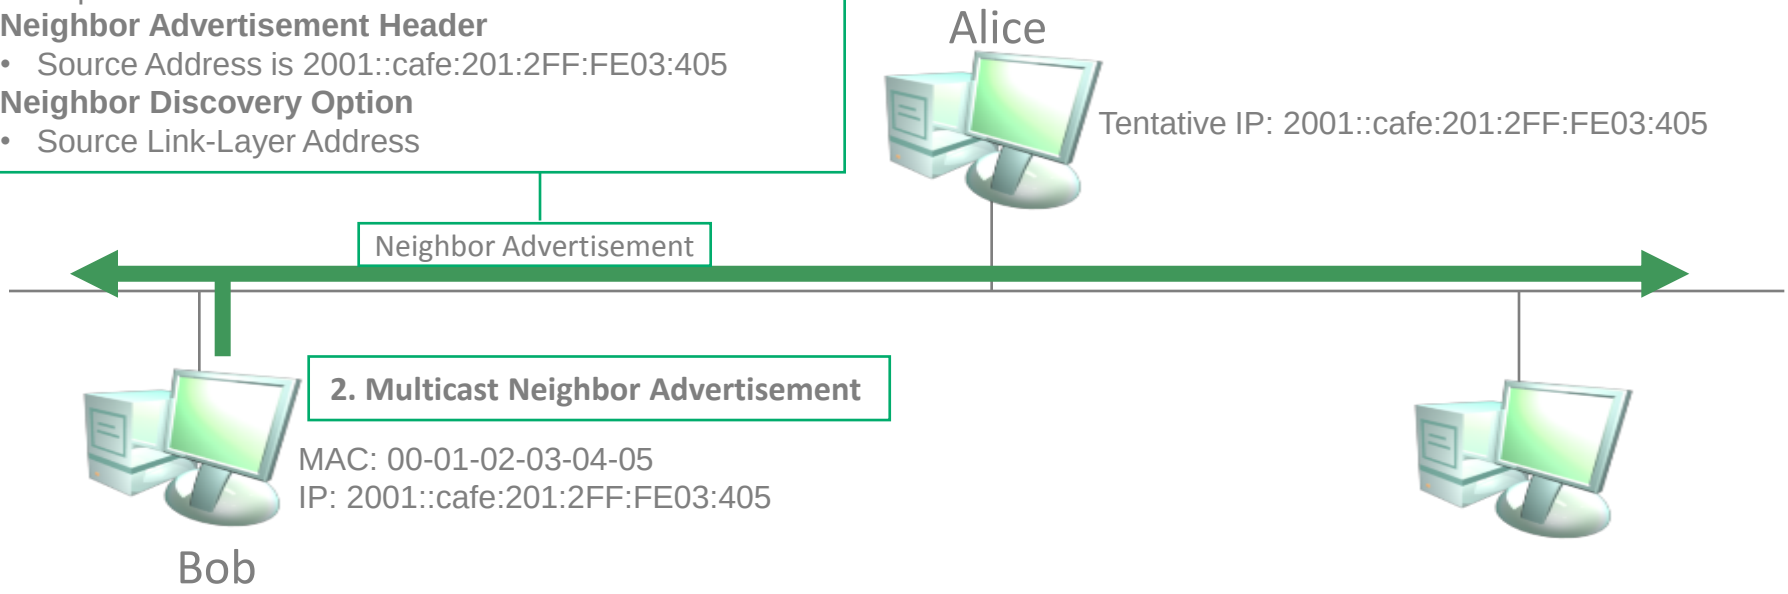

Router Discovery

- Used to detect routers that are connected to the local network.

- IPv6 router discovery also provides the following information:
  - Default value for the "Hop Limit" field
  - Whether any "stateful address protocol” (DHCPv6) should be used.
  - Settings for the “Retransmission Timer”
  - The network prefix for the local network
  - The MTU of the network
  - Mobile IPv6 Information
  - Routing Information

Multicast Router Solicitation Message

Ethernet Header
- Dest.-MAC: 33-33-00-00-00-02

IPv6 Header
- Source-IP: fe80::10
- Dest.-IP: FF02::2
- Hop limit: 255

Router Solicitation

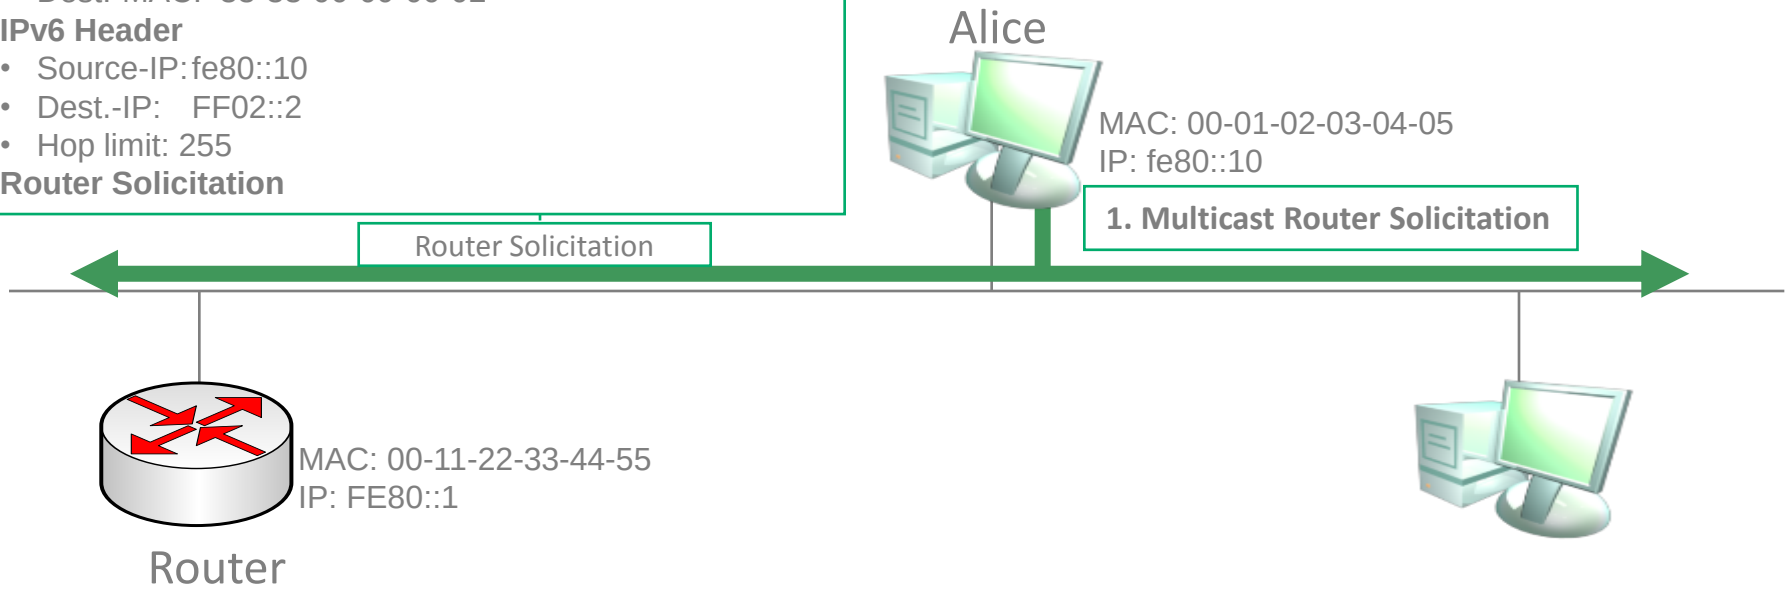

Router Advertisement Message

Ethernet Header
- Dest.-MAC: 33-33-00-00-00-01

IPv6 Header
- Source-IP: FE80::1
- Dest.-IP: FF02::1
- Hop limit: 255

Router Advertisement Header
- Current Hop Limit, Flags, Router Lifetime, Reachable and Retransmission Timers

Neighbor Discovery Options
- Source Link-Layer Address
- Prefix-Informationen

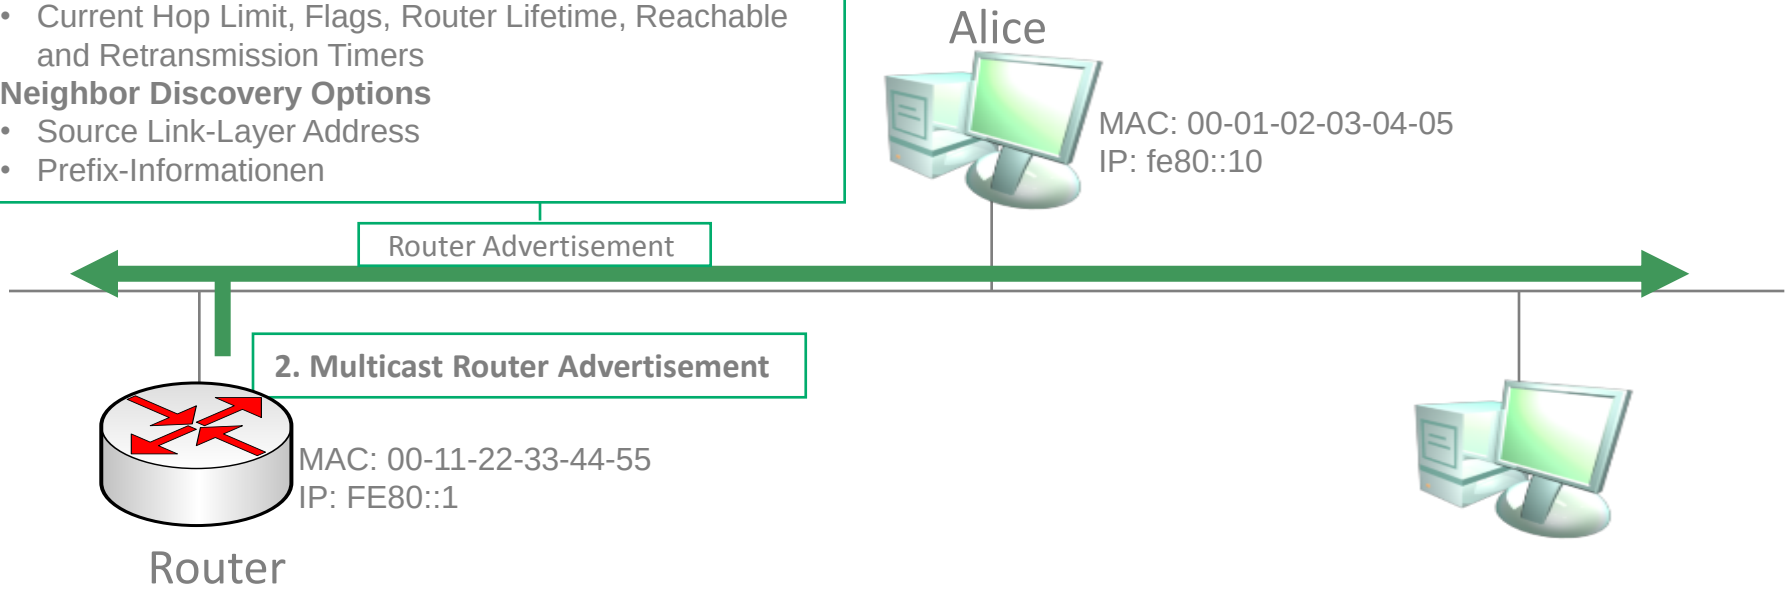

IPv6 basic security assumptions

- “The local network is trustworthy.”
  - Ha Ha!
  - And: what’s “local” anyway?
- “If there are threats, we can address them with crypto.”
  - Ever heard of “operational overhead” (caused by crypto)?
- “Everybody will use IPsec (that’s why we’ve built it in).”
  - As of RFC 6434, it is not a mandatory for an IPv6 node to implement IPSec
  - Well, IPsec is a success story in itself, isn’t it? ;-)
  - Unfortunately this has lead to a number of debatable protocol decisions.
    - OSPFv3 without MD5 (or SHA-x for that matter) is … crazy…

Attacks from within $SEGMENT

Attacks against Router Discovery

Router Advertisement Spoofing/Flooding

The Rogue RA Problem Statement

- Router advertisements (as part of autoconfig approach) fundamental part of “IPv6 DNA”.
  - Modifying this behavior (e.g. by deactivating their processing on the host level) is a severe “deviation from default” and as such “operationally expensive”.
  - Such an approach might be hard to maintain through a system’s lifecycle as well.
    - Think service packs in MS world, kernel updates, installation of libs/tools/apps.

- By default, local link regarded trustworthy in IPv6 world
  - All ND related stuff (which includes RAs) unauthenticated, by default.

Rogue Router Advertisements

- Some RA-generating entity accidentally active in your network
  - IPv6 capable SOHO device connected by user.
  - Windows system with ICS enabled
    - No longer valid, see http://support.microsoft.com/kb/2750841/en-us.
  - Virtual machine running sth. emitting RAs…
- Attacker interferes with router discovery

Impact

- Traffic redirection through malicious default gateway configuration on the host.
- Denial-of-Service against IPv6 connectivity in case of accidentally transmitted router advertisements.

Impact – fake_router26

```
C:\>ipconfig

Windows IP Configuration


Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . :
   IPv6 Address. . . . . . . . . . . : 2001:db8:cafe:1234:c906:1f8:57a9:a974
   IPv6 Address. . . . . . . . . . . : 2001:db8:dead:beef:c906:1f8:57a9:a974
   Link-local IPv6 Address . . . . . : fe80::c906:1f8:57a9:a974%13
   Autoconfiguration IPv4 Address. . : 169.254.169.116
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : fe80::21a:a1ff:fec1:6311%13
                                       fe80::216:36ff:fe12:3bc6%13

Wireless LAN adapter Wireless Network Connection:
```

Router Advertisement Flooding

- Flooding the local network with router advertisements, IPv6 hosts and router update the IPv6 configuration
- Consuming all available resources
  - -> Denial of Service against targets

Impact - flood_router26

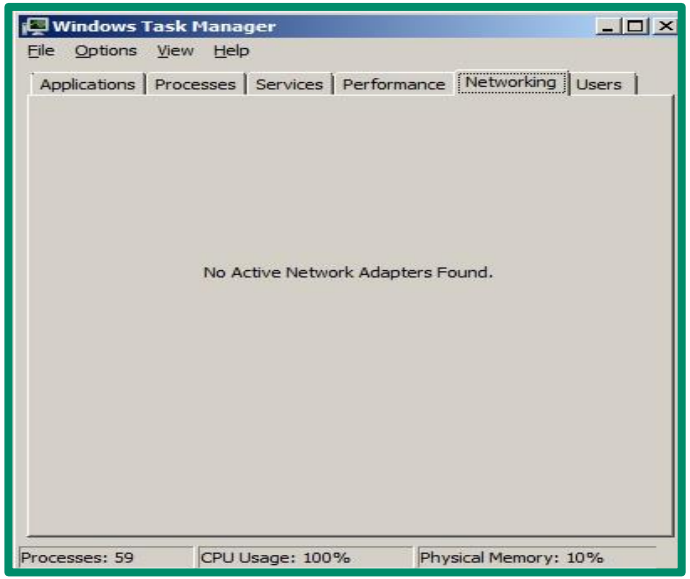

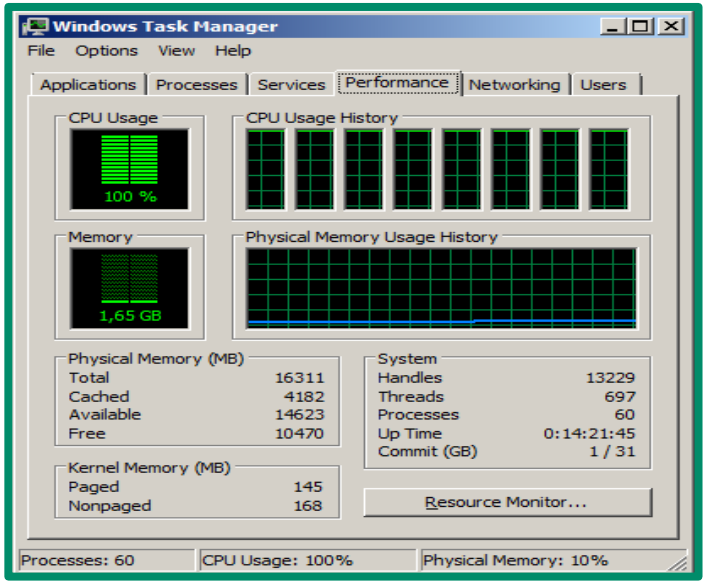

Spoofed RA protection as of RFC 6104

- Manual configuration
- RA Snooping (RA Guard)
- Using ACLs
  - See also: http://rmv6tf.org/wp-content/uploads/2013/04/5-IPv6-Attacks-and-Countermeasures-v1.2.pdf
- SEcure Neighbor Discovery (SeND)
- Router Preference
- Relying on Layer 2 Admission Control
- Host-Based Packet Filters
- Using an “Intelligent” Deprecation Tool
  - E.g. NDPMon
- Using Layer 2 Partitioning

Default Router Preference

- In RFC 4191 an additional flag was introduced within RA messages to indicate the preference of a default router in case more than one are present on the local link.

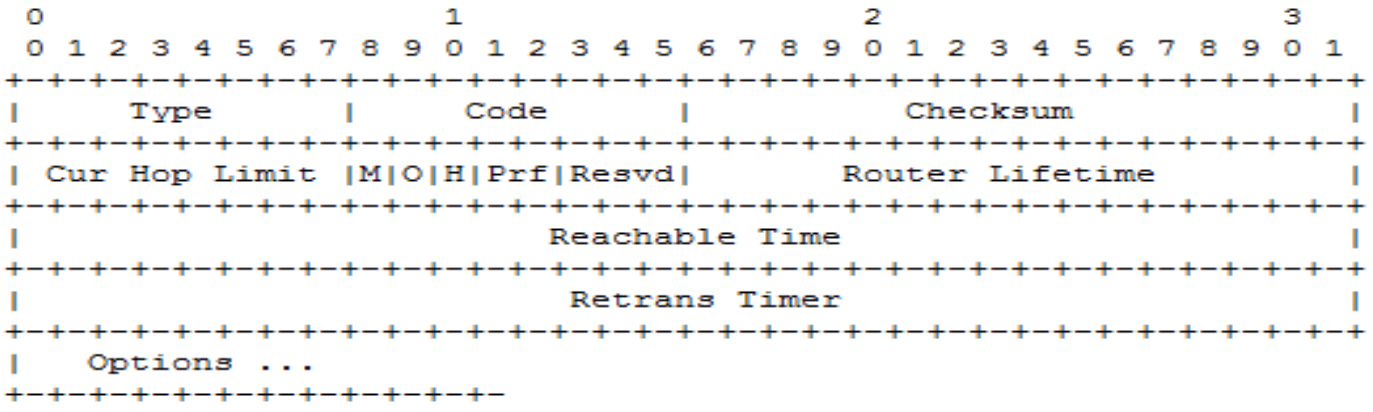

Router Preference Values

- The *preference* values are encoded as a two-bit signed integer with the following values:
  - 01 High
  - 00 Medium (default)
  - 11 Low
  - 10 Reserved

RA Messages

- When the *preference* is set, the RA messages look like:

```
Internet Control Message Protocol v6
   Type: 134 (Router advertisement)
   Code: 0
   Checksum: 0xded0 [correct]
   Cur hop limit: 64
 Flags: 0x08
     0... .... = Not managed
     .0.. .... = Not other
     ..0. .... = Not Home Agent
     ...0 1... = Router preference: High
   Router lifetime: 1800
   Reachable time: 0
   Retrans timer: 0
 ICMPv6 Option (Source link-layer address)
 ICMPv6 Option (MTU)
 ICMPv6 Option (Prefix information)
```

```
Internet Control Message Protocol v6
   Type: 134 (Router advertisement)
   Code: 0
   Checksum: 0xcdc6 [correct]
   Cur hop limit: 64
 Flags: 0x00
     0... .... = Not managed
     .0.. .... = Not other
     ..0. .... = Not Home Agent
     ...0 0... = Router preference: Medium
   Router lifetime: 1800
   Reachable time: 0
   Retrans timer: 0
 ICMPv6 Option (Source link-layer address)
 ICMPv6 Option (MTU)
 ICMPv6 Option (Prefix information)
```

Configuration (Cisco)

- The configuration of the preference is done with the following command:

```
Router(config)# interface f0/1
Router(config-if)# ipv6 nd router-preference {high | medium | low}
```

- If the command is not configured, the default value of medium will be used in the RA messages.

- Command available since IOS Version 12.4(2)T

Further Mitigating Controls

- Prevent Node-Node Layer-2 communication by using
  - Private VLANs where nodes can only speak to the legitimate router
- WLAN with isolated clients
- One VLAN per Host (e.g. SP access network)

RA Guard – Host Mode

- Implements *isolation* principle similar to other L2 protection mechanisms already deployed in v4 world.
- RFC 6105
- Works quite well against some attacks.
  - But it seems currently no logging or port deactivation can be implemented. RA packets are just dropped.
- Can be easily circumvented

Cisco First-Hop-Security

- Cisco name for various security features in IPv6
- Staged in three phases
- Every Phase will release/released more IPv6 security features to achieve feature parity with the IPv4 world

Phase I

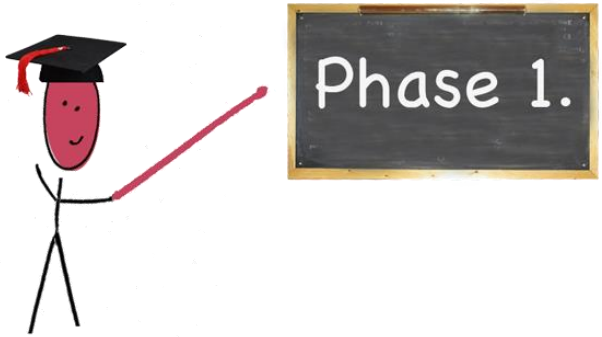

- Available since Summer 2010
- Introduced RA Guard and Port based IPv6 ACLs
- In the beginning, only supported on datacenter switches
  - Since 15.0(2) supported on C2960S and C3560/3750-X

General Principles on FH Command Interface[1]

Each FH feature provides a configuration mode to create and populate policies (+ one implicit “default” policy)

```
ipv6 nd raguard policy MYHOST
 device-role host
```

Each FH feature provides commands to attach policies to targets: box,vlan, port

```
vlan configuration 100
 ipv6 nd raguard attach-policy MYHOST
 ipv6 snooping
interface e0/0
 ipv6 nd raguard attach-policy MYROUTER
```

Packets are processed by the lowest-level matching policy for each feature

Packets received on e0/0 are processed by policy ra-guard “MYROUTER” AND policy snooping “default”

Packets received on any other port of vlan 100 are processed by policy ra-guard “MYHOST” AND policy snooping “default”

RA Guard – Host Mode


```
Router(config-if)#ipv6 nd ?
  raguard  RA_Guard Configuration Command
Router(config-if)#ipv6 nd raguard ?
   <cr>
Router(config-if)#switchport mode access
Router(config-if)#ipv6 nd raguard
Router(config-if)#exit
Router(config)#exit

Router# show version
Cisco IOS Software, s3223_rp Software (s3223_rp-
IPBASEK9-M), Version 12.2(33)SXI5, RELEASE SOFTWARE
(fc2)
```

Fragmentation of RAs

- Issue
  - RA Guard works like a stateless ACL filter of ICMP Type 134
  - Can be circumvented by “pushing” ULP (RA) into second fragment
- Possible Solutions
  - Block all fragments sent to ff02::1
  - Drop packet if first fragment does not have ULP header
    - `deny ipv6 any any undetermined-transport`

Port-based ACLs


```
4948E(config)#ipv6 access-list IPv6
4948E(config-ipv6-acl)#deny ipv6 any any undetermined-transport
4948E(config-ipv6-acl)#deny icmp any any router-advertisement
4948E(config-ipv6-acl)#permit ipv6 any any
4948E(config)#interface g1/19
4948E(config-if)#ipv6 traffic-filter IPv6 in
```

Block Forwarding of RAs on Infrastructure Level

- RA Guard or ACLs
  - _Or_!
- RA Guard currently (Mar 2014) not a bullet-proof solution.
  - -DF switch in THC's `fakerouter6` does the trick.
    - See also http://www.insinuator.net/2011/05/yet-another-update-on-ipv6-security-some-notes-from-the-ipv6-kongress-in-frankfurt/
- ACLs might be operationally expensive.
  - Probably port based ACLs not part of your current ops model, right?
  - HW support needed
    - http://docwiki.cisco.com/wiki/Cisco_IOS_IPv6_Feature_Mapping#IPv6_Features
  - Still, currently best protection approach that's available
    - See also http://www.insinuator.net/2012/03/the-story-continues-another-ipv6-update/
- RA Guard will (hopefully) evolve
  - Some IETF drafts out there to address evasion problem
    - http://tools.ietf.org/html/draft-ietf-v6ops-ra-guard-implementation-07

Evaluation of RFC 6104 Controls

Control	Sec Benefit	Operational Feasibility
Manual configuration	4	1
RA Snooping (RA Guard)	4	4
Using ACLs	5	3
SEcure Neighbor Discovery (SEND)	5	1
Router Preference	2	5
Relying on Layer 2 Admission Control	5	2
Host-Based Packet Filters	3	1
Using an “Intelligent” Deprecation Tool	2	1
Using Layer 2 Partitioning	4	3

RFC 6980 – Sec Implications of NDP Fragmentation

- Standards Track RFC (August 2013) from Fernando Gont
  - Who is also here at the Security Summit ;)
- Completely forbids fragmentation of NDP packets.

Neighbor Discovery Threats

- **IPv6 Neighbor Spoofing:**
  - An Attacker can send malicious Neighbor Advertisement to poison the neighbor cache of the victim (similar to ARP spoofing in IPv4).

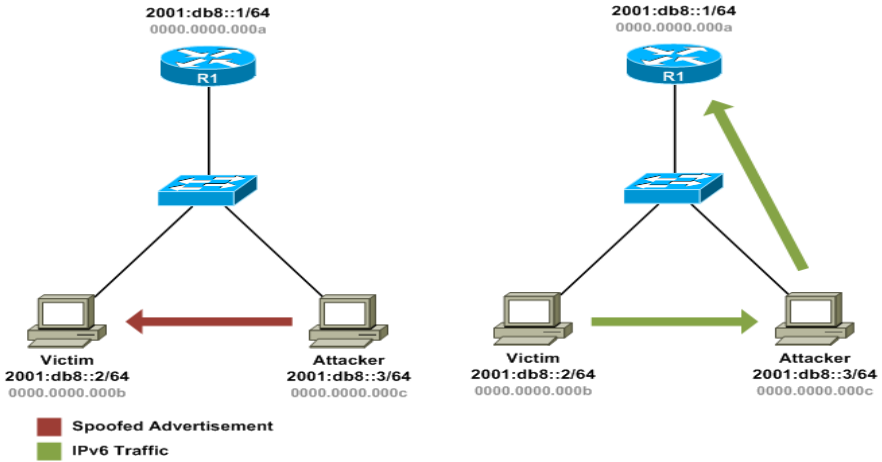

DoS of Address Assignment

- DoS against IPv6 nodes by always answering to neighbor solicitation for duplicate address detection

- Can happen in IPv4 and IPv6.
  - → Pretty much the same overall risk.
  - Attack not expected to be seen very often (as was the case in v4 netw.).
  - Probably much more efficient DoS scenarios available to attacker.

Cisco IPv6 Snooping

- IPv6 Snooping is the basis for several FHS security mechanisms
  - Including ND Inspection and address glean
- 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

Attacks against DHCPv6

DHCPv6

- Defined in RFC 3315
- Uses UDP Ports 546 (Clients) and 547 (Server/Relays)
- DHCPv6 can be used to replace or complement SLAAC
- Also called *stateful configuration* because the DHCP Server has to keep state (of the addresses leased to the DHCP clients).

DHCP Message Exchange

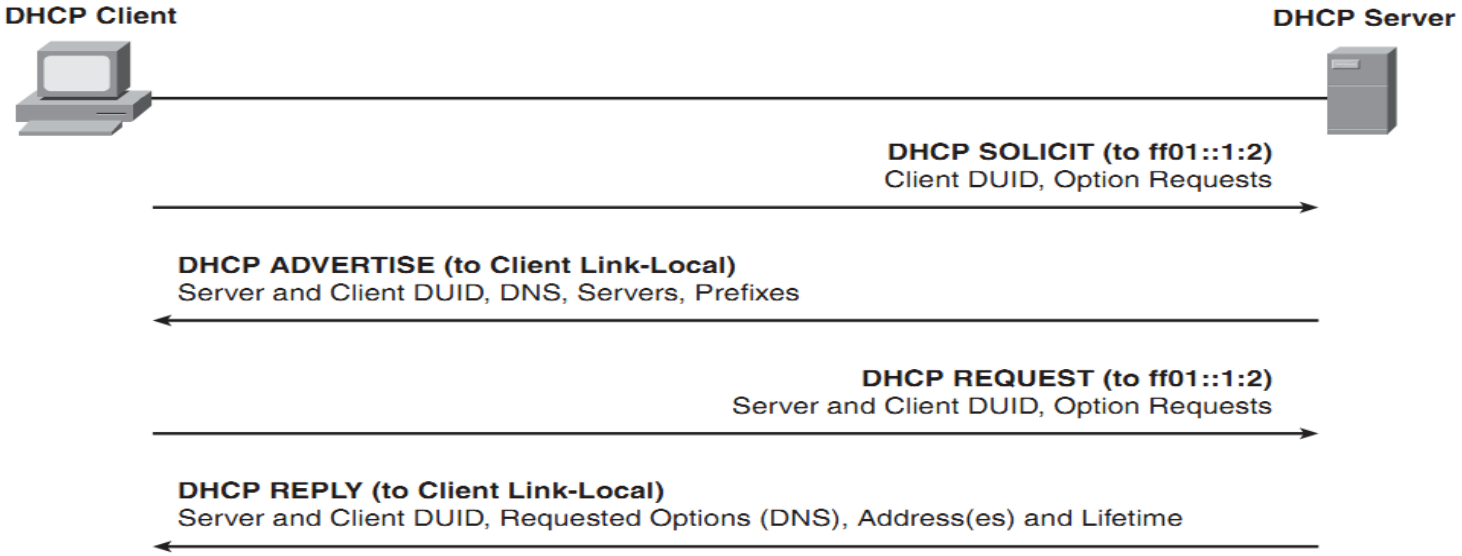

Differences to DHCPv4

- All message exchanges include a 24-bit transaction identifier.

- Messages can optionally be authenticated with the help of a hash-based message authentication code which is defined in RFC 3118.
  - We do not see this to be widely deployed in the near future.

Threats against DHCPv6

- Starvation
  - The attacker plays the role of many DHCPv6 clients and requests too many addresses, which depletes the pool of IPv6 addresses.

- Denial of service (DoS)
  - The Attacker sends a huge amount of SOLICIT messages to the servers, forcing them to install a state for a while and potentially causing a huge load on the servers' CPU and file systems, up to the point that legitimate clients can no longer be served.

Threats against DHCPv6

- Scanning
  - If leased addresses are generated sequentially, the usual network scanning can be reused to detect potential targets.

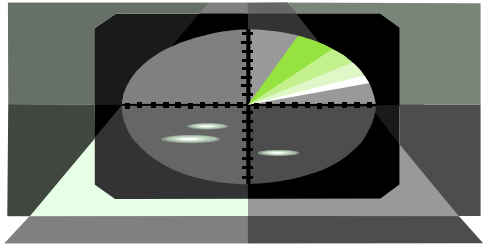

- Rogue DHCPv6 server
  - The Attacker sends malicious ADVERTISE and REPLY messages to legitimate clients. These messages contain falsified information about the default gateway, DNS servers, and so on that could be used to redirect the traffic.
  - An Attacker can join the site-local Multicast Group (FF05::1:3) and will receive a copy of all SOLICIT Messages destined to this group.

Mitigating Controls

- The authentication mechanism (RFC 3118) could address various threats but…
  - The pre-shared keys must be configured on all clients and DHCP servers.
  - Change of PSK involves a lot of work ;)
  - Very few DHCPv6 Implementations support the Authentication Mechanism (Windows Server 2008 does not support it)

- *DHCP snooping* for IPv6.
  - aka DHCP Guard

1-Slide Sec Discussion of DHCPv6

- As in v4 *rogue DHCP servers* can cause harm.
  - Nothing new here.
  - We'll probably release a *LOKI* module for this soon…
- Overall risk pretty much the same as in v4.
  - Might be even smaller as link-local address not deployed by DHCP.
- Same mitigation techniques (if any) will apply.
  - In case some "DHCP snooping successor" available for $YOUR_PLATFORM.

Phase II

- Available since end of 2011/ beginning of 2012 (depending on the platform)
- Introduced DHCPv6 Guard and NDP Snooping
  - DHCP Snooping and Dynamic ARP Inspection in the IPv4 World
- As of march 2014, available on 2960S/3560/3750-X
  - And on Cat 4500, Cat 4948 (E/F) and 7600 Routers

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.

DHCPv6 Guard


```
ipv6 access-list acl1
 permit host FE80::1 any
ipv6 prefix-list abc permit
2001:0DB8:c001:babe::/64 le 128

ipv6 dhcp guard policy pol1
 device-role server
 match server access-list acl1
 match reply prefix-list abc
 trusted-port <optional>

vlan configuration 10
ipv6 dhcp guard attach-policy pol1

show ipv6 dhcp guard policy pol1
```

Attacks from the Outside World

Neighbor Cache Entries

State	Description
INCOMPLETE	Neighbor Solicitation has been sent, but no Neighbor Advertisement has been retrieved.
REACHABLE	Positive confirmation was received within the last *ReachableTime* milliseconds, no special actions necessary.
STALE	ReachableTime milliseconds have elapsed, no actions takes place. This is entered upon receiving an unsolicited Neighbor Discovery message → entry must actually be used.
DELAY	ReachableTime milliseconds have elapsed and a packet was sent within the last *DELAY_FIRST_PROBE_TIME* seconds. If no message was sent → change state to PROBE.
PROBE	A reachability confirmation is actively sought by retransmitting Neighbor Solicitations every *RetransTimer* milliseconds until reachability confirmation is received.

Neighbor Cache Exhaustion [this slide stolen from Eric Vyncke]

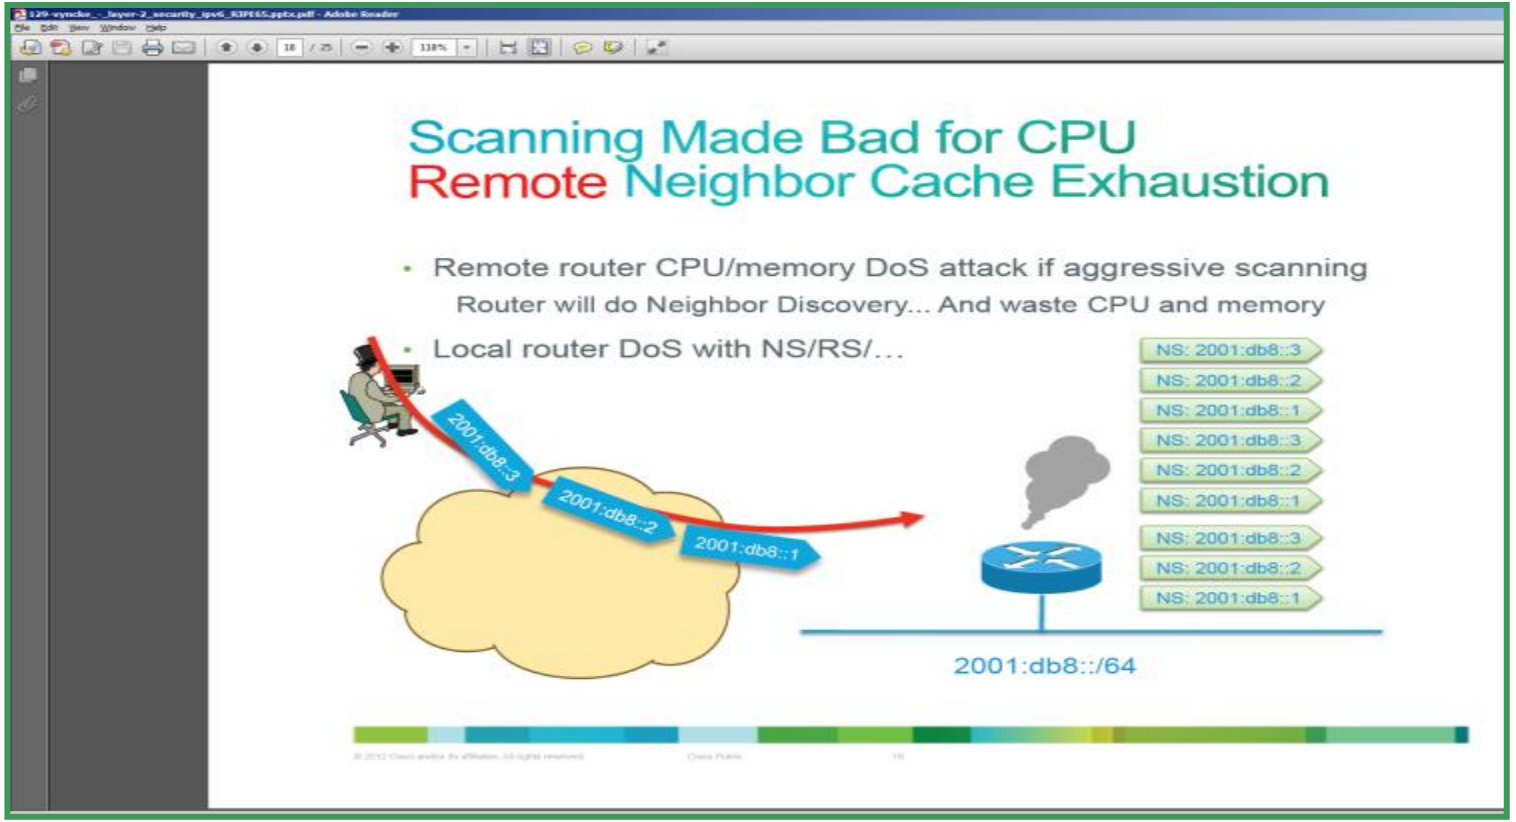

RFC 6583

Internet Engineering Task Force (IETF)
Request for Comments: 6583
Category: Informational
ISSN: 2070-1721

I. Gashinsky
Yahoo!
J. Jaeggli
Zynga
W. Kumari
Google, Inc.
March 2012

Operational Neighbor Discovery Problems

Abstract

In IPv4, subnets are generally small, made just large enough to cover the actual number of machines on the subnet. In contrast, the default IPv6 subnet size is a /64, a number so large it covers trillions of addresses, the overwhelming number of which will be unassigned. Consequently, simplistic implementations of Neighbor Discovery (ND) can be vulnerable to deliberate or accidental denial of service (DoS), whereby they attempt to perform address resolution for large numbers of unassigned addresses. Such denial-of-service attacks can be launched intentionally (by an attacker) or result from legitimate operational tools or accident conditions. As a result of these vulnerabilities, new devices may not be able to "join" a network, it may be impossible to establish new IPv6 flows, and existing IPv6 transported flows may be interrupted.

This document describes the potential for DoS in detail and suggests possible implementation improvements as well as operational mitigation techniques that can, in some cases, be used to protect against or at least alleviate the impact of such attacks.

RFC 6583, Potential Controls

- Filtering of Unused Address Space
  - RFC 6583: “it is fully understood that this is ugly (and difficult to manage); but failing other options, it may be a useful technique especially when responding to an attack.“

- Obviously this requires static addressing.
- If you do this, use *stateless* filtering.
  - ACLs might be your friend.
  - Do *not* induce additional state by stateful filtering!
    - The more overall state maintained, the higher the overall vulnerability for DoS.

RFC 6583, Potential Controls

- *Minimal Subnet Sizing*
  - RFC 6583: “this approach is not suitable for use with hosts that are not statically configured.”
- Well, this violates the /64 paradigm.
  - Doesn’t RFC 6164 “allow” this violation anyway?
  - Still, this is about leaving “a standard path”. Be careful!
    - “Organization’s culture” may play a role here.
  - Yes, we are aware of sect. 3 of RFC 5375.
    - We don’t regard this as relevant here though.
- Overall this approach might have quite good *operational feasibility*. Provided nothing breaks due to deviation f. /64.
- If you do this, still assign full /64, but configure /120 or sth.
  - So you can revert to /64 in case of problems or once better solutions are available (see below).

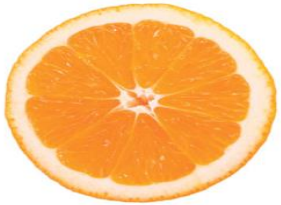

RFC 6583, Potential Controls

- Routing Mitigation
  - “For obvious reasons, host participation in the IGP makes many operators uncomfortable, but it can be a very powerful technique if used in a disciplined and controlled manner. One method to help address these concerns is to have the hosts participate in a different IGP (or difference instance of the same IGP) and carefully redistribute into the main IGP.”

- Honestly, this approach is so ridiculous both from an architecture and operations perspective, that we’ll not discuss this further.
  - Anybody remembers the days of `routed` on some Unix systems… and how happy we were to get rid of it?

RFC 6583, Potential Controls

- Tuning of the NDP Queue Rate Limit
  - “It is worth noting that this technique is worth investigating only if the device has separate queues for resolution of unknown addresses and the maintenance of existing entries.”
- We expect this to become “the main approach”
  - Vendors already start to implement this. (see below)
- In Cisco land:
  - `ipv6 nd cache interface-limit`
    - See also http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/command/ipv6-i3.html#GUID-FC37F82B-5AAC-4298-BB6C-851FB7A06D88
    - This one provides some logging, too. Might come in handy for attack detection.
      - `Mar 10 15:11:51.719: %IPV6_ND-4-INTFLIMIT: Attempt to exceed interface limit on GigabitEthernet0/1 for 2001:DB8:0:900D::2:329A` (So use it in any case!)
  - on IOS-XE 2.6: `ipv6 nd resolution data limit`
    - Thanks to Jim Small for this hint.Might address another problem though.
- Another suggestion: lowering retrans-timer to a sub-second value (suggestion f. Benedikt Stockebrand. Thx!)

NCE, Some Notes from the Lab

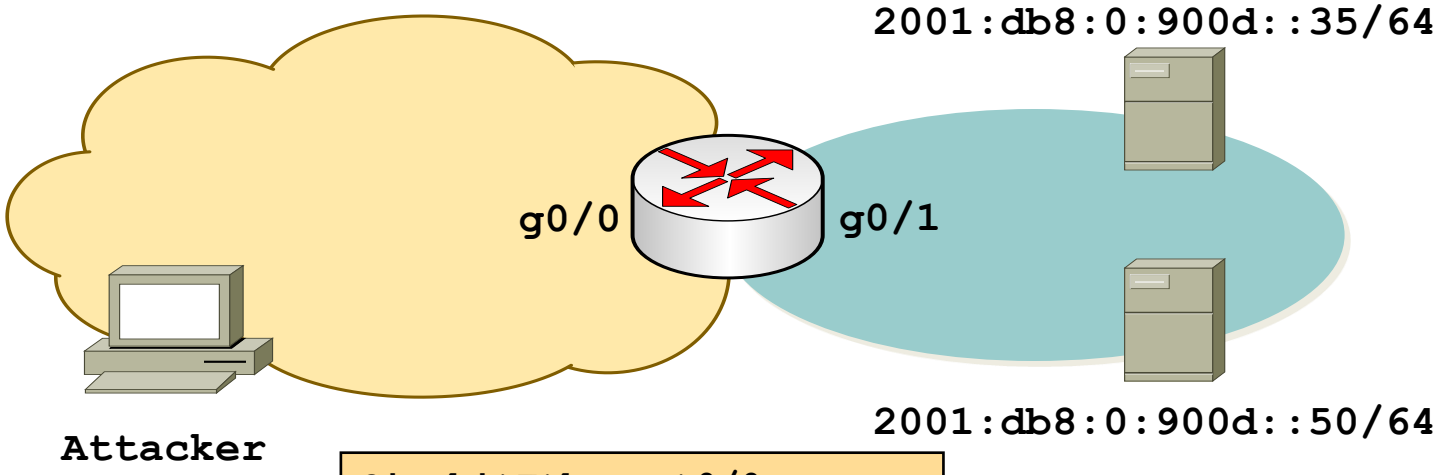


```
GigabitEthernet0/0
     FE80::BAAD:1
     2001:DB8:0:BAAD::1/64
GigabitEthernet0/1
     FE80::900D:1
     2001:DB8:0:900D::1/64
```

NCE, Conclusions from the Lab

- All tested Cisco devices do not store more than 512 INCOMPLETE entries in neighbor cache, at any given time.
  - Four different IOS-based medium-end devices tested.
- Furthermore reading RFC 4861 sect. 7.2.2 indicates INCMP entries will be deleted after three seconds anyway.
- So NCE *seems* not to be a major problem here (C land).
  - Various sources told us that Juniper space actually *is* susceptible to (NCE) problems.
  - We'll do some lab testing with an M7i and keep you posted.
    - Right now we can't comment on this further.
- Details of testing to be found here
  - http://www.insinuator.net/2013/03/ipv6-neighbor-cache-exhaustion-attacks-risk-assessment-mitigation-strategies-part-1/

FHS - Phase 3

- Available since December 2012
- Introduced Destination-Guard
  - To mitigate Neighbor Cache Exhaustion attack
- As of marc 2014, currently only supported on catalyst 4948E

IPv6 Destination Guard

Overview

- Blocks and filters traffic from an unknown source and filters IPv6 traffic based on the destination address.
- Uses „first-hop security binding table“
  - populates all active destinations into it and blocks data traffic when the destination is not identified.

IPv6 Destination Guard

Requirements

- Implemented in Cisco 7600 and Cisco Catalyst 4500/4900
- Requires 15.3S, 15.2S or 15.1SG

IPv6 Destination Guard

Example Configuration


```
Router(config)# vlan configuration 300
Router(config-vlan-config)# ipv6 destination-guard attach-
policy destination
% Warning - 'ipv6 snooping' should be configured before
destination-guard


Router(config-vlan-config)# ipv6 snooping attach-policy ND
Router(config)# vlan configuration 300
Router(config-vlan-config)# ipv6 destination-guard attach-
policy destination
Router(config-vlan-config)#

Router# show ipv6 destination-guard policy destination
Destination guard policy Destination:
  enforcement always
        Target: vlan 300
```

References

- [1] IPv6 First Hop Security: Eric Vyncke