

SLAACers

IPv6 Accountability without DHCPv6

Alexander Clouter <ac56@soas.ac.uk>

Library and Information Services
School of Oriental and African Studies
London

Networkshop 39, 2011

The Problem

Objective is to tie in accountability¹ (`abuse@`). IPv4 to MAC address mapping comes from DHCP:

- typically manual sync config to topology and changes
- failover support can be non-trivial
- tedious to extract mappings from raw log files

difficulty

n. an effort that is inconvenient [syn: {trouble}]

opinion

DHCP is a latent effort - SLP, DNS SRV, multicast NTP, etc

¹L1 → L2 map from RADIUS 802.1X/mac-auth

Situation Worse for IPv6

IPv6 assignment can be via static, DHCP or SLAAC².

- same problems as DHCP for IPv4
- local-link (fe80::/64)
- multiple IP assignment encouraged
 - EUI-64 (subnet + MAC address)
 - privacy/temporary/randomised
 - static
 - Mobile IPv6 (non-local)
- lack of OS support (eg. pre Mac OS X 10.7)

Question

Can we add accountability to SLAAC?

²StateLess Address Auto Configuration

Situation Worse for IPv6

IPv6 assignment can be via static, DHCP or SLAAC².

- same problems as DHCP for IPv4
- local-link ($fe80::/64$)
- multiple IP assignment encouraged
 - EUI-64 (subnet + MAC address)
 - privacy/temporary/randomised
 - static
 - Mobile IPv6 (non-local)
- lack of OS support (eg. pre Mac OS X 10.7)

Question

Can we add accountability to SLAAC?

²StateLess Address Auto Configuration

Situation Worse for IPv6

IPv6 assignment can be via static, DHCP or SLAAC².

- same problems as DHCP for IPv4
- local-link ($fe80::/64$)
- multiple IP assignment encouraged
 - EUI-64 (subnet + MAC address)
 - privacy/temporary/randomised
 - static
 - Mobile IPv6 (non-local)
- lack of OS support (eg. pre Mac OS X 10.7)

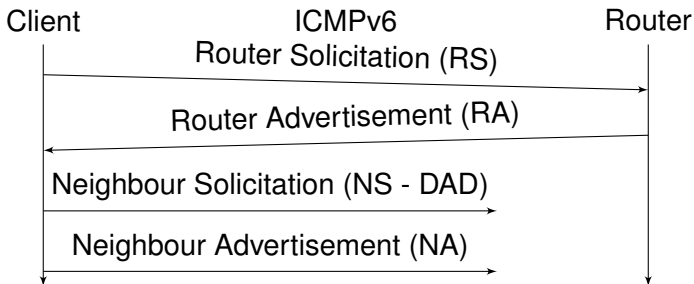
Question

Can we add accountability to SLAAC?

²StateLess Address Auto Configuration

Polling Routers

ARP provides IPv4→MAC mapping, IPv6 uses Neighbour Discovery Protocol. NDP table queriable via the CLI and SNMP.



Unfortunately this only gives a snapshot of state. Polling is compromise between granularity and load (not event driven³).

³polling typically considered a Bad Idea™

Router Logging Facilities

Some routers can print this information ('debug ipv6 nd'):

- CPU bound (C3750 - PowerPC 405 125MHz)
- interaction with router can become impossible
- syslog files can become very big very quickly
- logs might not be suitable - seems not to be on a C3750


```
Apr  6 21:09:37: ICMPv6-ND: Sending NS for  
FE80::D0E0:1C9B:4007:C24F on Vlan76
```

```
Apr  6 21:09:37: ICMPv6-ND: Received NA for  
FE80::D0E0:1C9B:4007:C24F on Vlan76  
from FE80::D0E0:1C9B:4007:C24F
```

```
Apr  6 21:10:48: ICMPv6-ND:  
Neighbour 2001:DB8:54:4:1C6D:17DF:570E:382D  
on Vlan76 : LLA 0011.2233.4455
```

Directly Parse ICMPv6 NS/NA Traffic

- what needs collection
 - Neighbour Advertisement (NA)
 - optionally Neighbour Solicitation (NS - DAD) - multicast
- how to collect - port mirroring (remote)
 - Cisco (R)SPAN
 - similar systems exist for HP and Juniper (probably others)



SLAACer

Packet capturing⁴ based Perl daemon that dissects NS/NA traffic and generates machine (and human) readable syslog messages:

```
2011-04-03T01:02:10+01:00: 00-11-22-33-44-55
 2001:db8:....:1234 [SOLICIT,dad]
2011-04-03T01:02:11+01:00: 00-11-22-33-44-55
 fe80::1234::4321 [SOLICIT]
2011-04-03T01:02:14+01:00: 00-11-22-33-44-55
 2001:db8:....:1234 [ADVERT,solicited,override]
2011-04-03T01:02:15+01:00: 00-11-22-33-44-55
 fe80::1234::4321 [ADVERT,solicited,override]
```

⁴PCAP - tcpdump/wireshark

SLAACer

Raw syslog still makes for awkward mapping extraction.
syslog-ng can push this information into an SQL database:

```
syslog → syslog-ng → parser → 'pipe' → (p|my)sql → SQL
```

```
SELECT record_addr(ts, mac, inet)
```

But wait! We can also do this for IPv4 ISC DHCP syslog too!

```
MAC addr => IP addr => first seen => last seen
```

```
00-11-22-33-44-55 => 2001:db8:::1234
```

```
=> Apr 4 08:28:22+01 => Apr 4 16:48:15+01
```

```
00-11-22-33-44-55 => fe80:::4321
```

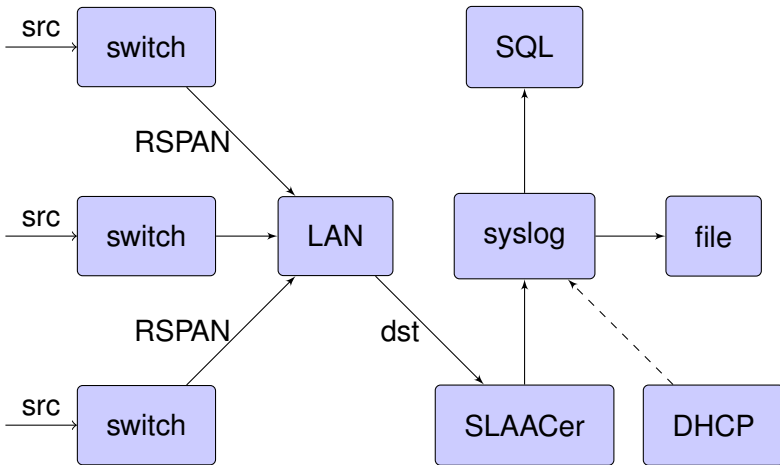
```
=> Feb 23 10:06:23+00 => Apr 4 16:48:04+01
```

```
00-11-22-33-44-55 => 10.20.30.40
```

```
=> Mar 7 08:23:02+00 => Apr 4 07:21:30+01
```

SLAACer

Realtime Event Driven (DHCP) IPv4 and (monitored) IPv6 Mappings



Caveats

SLAACer is not without it's problems:

- a lot of traffic is being sent to SLAACer
 - only capture ingress traffic
 - filters (VACLs) help immensely (eg. trim IPv4)
 - 3750 will not VACL RSPAN sources
 - 3750/6500 will not VACL with IPv6 ACLs
 - 3750 will not VACL IPv6 Ethernet frames with MAC ACLs
 - 6500/DFC3A will not VACL Ethernet frames with MAC ACLs
 - compromise, rely on just multicast NS-DAD packets
 - CSCtd72626 - RSPAN misses 33:33:ww:xx:yy:zz
 - keep an eye on bandwidth graphs - use port-channels
- RSPAN VLAN floods everywhere - MAC learning disabled
 - unique per switch-stack RSPAN
- duplicate MAC at different locations
 - L1→L2 problem - solve with RADIUS Simultaneous-Use

Summary

- SLAACer records realtime (event driven) IPv6 usage
- no need for HA DHCPv6 - router is now point of failure
- sysadmins can now choose either DHCPv6 and/or SLAAC
- SLAACer is not 'the' solution but will tie us over till something better comes along. GPL means you can 'fix' it.
- Outlook
 - record RA announcements from misconfigured clients
 - experiment with SPF loopbacks - VACL at RSPAN source
 - wish for something like SNMP traps for IPv6→MAC events



Alexander Clouter (me!)

SLAACer - Accountability with IPv6

<http://www.digriz.org.uk/slaacer>