# IPv6 Prefix Intelligence

## Measurement and Analysis for Protocols Proposed Research Group (maprg)

Buenos Aires, April 4, 2016

Dave Plonka <plonka@akamai.com>

# Agenda

- The Problem
  - To mitigate abuse from IPv6 hosts, what prefix length (mask) should be applied to the source address(es) for rate-limiting and in access control lists?
- State of the Internet that affects IPv6 abuse mitigation
  - IPv6 addressing heterogeneity
- How does this map to IETF-defined protocols?
  - Any IPv6 traffic destined for services hosted on IPv6
- Proposed solution: IPv6 [/64] prefix intelligence
  - Measurement & Analysis technique: Temporal & Spatial IPv6 Address Classification
  - Differentiated prefix length (specificity) for the active IPv6 [client] address space

## The problem

- We *know* that the strawman for a service to rate-limit or block a single ostensibly abusive IPv6 source by using a /64 prefix netmask will sometimes adversely "nearby" many IPv6 hosts, *i.e.,* those within the same /64 prefix as the source address of the abuse.

- RFC 7136 ("Significance of IPv6 Interface Identifiers") says to treat an IPv6 address' Interface ID (low 64 bits) as semantically opaque.

  This presents a significant challenge to fine-grained address-based abuse mitigation and will likely to be ignored when defending an IPv6 service from network abuse.
  If the /64 is opaque, then access control will use at most a /64 prefix for filtering, leading to collateral damage by affecting traffic involving other hosts whose address addresses are covered by that /64 prefix.

# Current State of the Internet: IPv6 Address Heterogeneity

## State of the Internet: IPv6 address heterogeneity

IPv6 addresses in presentation format:

```
2001:db8:0:1cdf:21e:c2ff:fec0:11db
2001:db8:10:1::103
2001:db8:167:1109::10:901
2001:db8:4137:9e76:3031:f3fd:bbdd:2c2a
```

## State of the Internet: IPv6 address heterogeneity

IPv6 addresses in presentation format:

```
2001:db8:0:1cdf:21e:c2ff:fec0:11db
2001:db8:10:1::103
2001:db8:167:1109::10:901
2001:db8:4137:9e76:3031:f3fd:bbdd:2c2a
```

Consider 16-bit (4 character) and 4-bit (1 character) segments:

```
2001:0db8:0000:1cdf:021e:c2ff:fec0:11db
2001:0db8:0010:0001:0000:0000:0000:0103
2001:0db8:0167:1109:0000:0000:0010:0901
2001:0db8:4137:9e76:3031:f3fd:bbdd:2c2a
```

## Stateless address classification: `addr6 -s -i`

```
** IPv6 General Address Analysis **
Total IPv6 addresses: 4
Unicast:         4 (100.00%)         Multicast:         0 (0.00%)
Unspec.:         0 (0.00%)

** IPv6 Unicast Addresses **
Loopback:              0 (0.00%)     IPv4-mapped:           0 (0.00%)
IPv4-compat.:          0 (0.00%)     Link-local:            0 (0.00%)
Site-local:            0 (0.00%)     Unique-local:          0 (0.00%)
6to4:                  0 (0.00%)     Teredo:                0 (0.00%)
Global:                4 (100.00%)

** IPv6 Interface IDs **
Total IIDs analyzed: 4
IEEE-based:          1 (25.00%)      Low-byte:                2 (50.00%)
Embed-IPv4:          0 (0.00%)       Embed-IPv4 (64):         0 (0.00%)
Embed-port:          0 (0.00%)       Embed-port (r):          0 (0.00%)
ISATAP:              0 (0.00%)       Byte-pattern:            0 (0.00%)
Randomized:          1 (25.00%)
```

# Example 1: IPv6 hosts using a subnet prefix length of /64 (SLAAC)

```
20010db8000e0000 00172cd5fa4bd6b1 75 0d
20010db8000e0000 002ae748ea083efb 75 0d
20010db8000e0000 005d58e18441347a 79 1d
20010db8000e0000 005f1dd3864f2d03 79 0d
20010db8000e0000 00872ce4d7e0d16c 76 0d
... (1594 more addresses) ...
20010db8000e0000 fdbefa6dce8d096c 80 1d
20010db8000e0000 fdbf6e62e74a33a4 80 1d
20010db8000e0000 fdd4f4f54264cc52 75 0d
20010db8000e0000 fdf73310ae0043da 75 2d
20010db8000e0000 feedfacedeadbabe 71 3d
```

# Example 2: IPv6 hosts using a subnet prefix length > /64 (DHCPv6)

```
20010db802000013000600000000010 128 4d
20010db802000013000600000000011 128 4d
20010db802000013000600000000012 128 4d
20010db802000013000600000000013 128 4d
20010db802000013000600000000014 128 4d
... (70 more addresses) ...
20010db802000013002600000000011 128 4d
20010db802000013002600000000012 128 4d
20010db802000013002600000000013 128 4d
20010db802000013002600000000014 128 4d
20010db802000013002600000000015 128 1d
```

# Example 2: DHCPv6

```
20010db802000013000600000000010 128 4d
20010db802000013000600000000011 128 4d
20010db802000013000600000000012 128 4d
20010db802000013000600000000013 128 4d
20010db802000013000600000000014 128 4d
... (70 more addresses) ...
20010db802000013002600000000011 128 4d
20010db802000013002600000000012 128 4d
20010db802000013002600000000013 128 4d
20010db802000013002600000000014 128 4d
20010db802000013002600000000015 128 1d
```

/64 prefix

IID

Discriminating Prefix Length (DPL)

Days Stable

# Example 2: DHCPv6

```
20010db802000013000600000000010 128 4d
200                             011 128 4d
200                             012 128 4d
200                             013 128 4d
200                             014 128 4d
…
200                             011 128 4d
200                             012 128 4d
20010db802000013002               000013 128 4d
20010db802000013002600000         0014 128 4d
20010db802000013002600000000      5 128 1d
2001:0db8:0200:0013::/63 80
20010db802000013 80 SLAAC-probable: 0.00% stabl
                 perDay{ave=63.8 median=67
          80 stableDays{median=4.0 sttd
                 dPL{median=128.0 sttdev=0 CV=0}
           0 randomizedDPL{median=0.0 sttdev=UNDEF C
          0 SLAACstableDays{median=0.0 sttdev=UNDEF CV=UNDEF}
          55 stableStableDays{median=4.0 sttdev=0.413 CV=0.109}
```

addr6 reports all having Interface IDs (IIDs) that are neither IEEE-based nor randomized.

Note that discriminating prefix lengths (DPL) of these 80 addresses are very high, *e.g.*, 128 (bits).

We'd expect it to take more than 1 million random IIDs for a DPL to reach even 110 (bits).

# Example 2: DHCPv6

```
20010db802000013000600000000010 128 4d
20010db802000013000600000000011 128 4d
20010db802000013000600000000012 128 4d
20010db802000013000600000000013 128 4d
20010db802000013000600000000014 128 4d
… (70 more addresses) ...
20010db802000013002600000000011 128 4d
20010db802000013002600000000012 128 4d
20010db802000013002600000000013 128 4d
20010db802000013002600000000014 128 4d
20010db802000013002600000000015 1
2001:0db8:0200:0013::/63 80
20010db802000013 80 SLAAC-probable: 0.00% stable: 68.75%
                 perDay{ave=63.8 median=67.0 sttdev=11.72 CV=0.183}
          80 stableDays{median=4.0 sttdev=1.197 CV=0.390}
                 dPL{median=128.0 sttdev=0 CV=0}
           0 randomizedDPL{median=0.0 sttdev=UNDEF CV=UNDEF}
          0 SLAACstableDays{median=0.0 sttdev=UNDEF CV=UNDEF}
          55 stableStableDays{median=4.0 sttdev=0.413 CV=0.109}
```

**Therefore, classify this as a non-SLAAC prefix, *e.g*, use /128 in ACLs or other filters meaning to affect only individual clients.**

# Example 1: SLAAC

```
20010db8000e0000 00172cd5fa4bd6b1 75 0d
20010db8000e0000 002ae748ea083efb 75 0d
20010db8000e0000 005d58e18441347a 79 1d
20010db8000e0000 005f1dd3864f2d03 79 0d
20010db8000e0000 00872ce4d7e0d16c 76 0d
... (1594 more addresses) ...
20010db8000e0000 fdbefa6dce8d096c 80 1d
20010db8000e0000 fdbf6e62e74a33a4 80 1d
20010db8000e0000 fdd4f4f54264cc52 75 0d
20010db8000e0000 fdf73310ae0043da 75 2d
20010db8000e0000 feedfacedeadbabe 71 3d
```

/64 prefix

IID

DPL

Days Stable

# Example 1: SLAAC

```
20010db8000e000000172cd5fa4bd6b1 75 0d
20010db8000e0000002ae748ea083efb 75 0d
20010db8000e0000005d58e18441347a 79 1d
20010db8000e0000005f1dd3864f2d03 79 0d
20010db8000e000000872ce4d7e0d16c 76 0d
... (1594 more addresses) ...
20010db8000e0000fdbefa6dce8d096c 80 1d
20010db8000e0000fdbf6e62e74a33a4 80 1d
20010db8000e0000fdd4f4f54264cc52 75 0d
20010db8000e0000fdf73310ae0043da 75 2d
20010db8000e0000feedfacedeadbabe 71 3d

$ addr6 -a 20010db8000e0000feedfacedeadbabe
unicast=global=global=randomized=unspecified
```

# Example 1: SLAAC

```
20010db8000e0000 00172cd5fa4bd6b1 75 0d
20010db8000e0000 002ae748ea083efb 75 0d
20010db8000e0000 005d58e18441347a 79 1d
20010db8000e0000 005f1dd3864f2d03 79 0d
20010db8000e0000 00872ce4d7e0d16c 76 0d
... (1594 more addresses) ...
20010db8000e0000 fdbefa6dce8d096c 80 1d
20010db8000e0000 fdbf6e62e74a33a4 80 1d
20010db8000e0000 fdd4f4f54264cc52 75 0d
20010db8000e0000 fdf73310ae0043da 75 2d
20010db8000e0000 feedfacedeadbabe 71 3d

20010db8000e0000 1604 SLAAC-probable: 90.52% stable: 9.48%
                 perDay{ave=556.4 median=613.0 sttdev=124.6 CV=0.223}
          1604 stableDays{median=0.0 sttdev=1.068 CV=1.431}
                 dPL{median=77.0 sttdev=2.204 CV=0.028}
           1594 randomizedDPL{median=77.0 sttdev=1.861 CV=0.024}
          1452 SLAACstableDays{median=0.0 sttdev=0.622 CV=1.345}
          152 stableStableDays{median=3.0 sttdev=0.497 CV=0.144}
```

# Example 1: SLAAC

```
20010db8000e000000172cd5fa4bd6b1 75 0d
20010db8000e0000002ae748ea083efb 75 0d
20010db8000e0000005d58e18441347a
20010db8000e0000005f1dd3864f2d0
20010db8000e000000872ce4d7e0d16
... (1594 more addresses) ...
20010db8000e0000fdbefa6dce8d096
20010db8000e0000fdbf6e62e74a33a
20010db8000e0000fdd4f4f54264cc5
20010db8000e0000fdf73310ae0043d
20010db8000e0000feedfacedeadba

20010db8000e0000 1604 SLAAC
                 perDay{av
          1604 stableDays{
                 dPL{median=77.0 sttdev=2.204 CV=0.028}
           1594 randomizedDPL{median=77.0 sttdev=1.861 CV=0.024}
          1452 SLAACstableDays{median=0.0 sttdev=0.622 CV=1.345}
          152 stableStableDays{median=3.0 sttdev=0.497 CV=0.144}
```

There is an expected maximum Discriminating Prefix Length (DPL) for a set, size *n*, of IPv6 addresses with random IIDs.

At probability of 0.99 (99%), *e.g.*, *n*=1643 such addresses have expected max. DPL <= 91 (bits).

Here, where *n*=1604, the observed max. DPL was 86 (bits); thus, they have plausibly random IIDs.

Address *Dendrachronology:* from dendra (tree limbs); khronos (time)

## Related Work? Questions? Comments?

- **"Temporal and Spatial Classification of Active IPv6 Addresses"** (IMC 2015) **http://www.akamai.com/technical-publications/**

- Should there be a protocol by which network operators can communicate their suggested prefix length or mask to indicate how to filter or block hosts in their network?
  - How is this related to other host reputation solutions?

Dave Plonka <plonka@akamai.com>