# Rotating Prefixes in xDSL networks

Lutz Donnerhacke

IKS Service GmbH

# IPv6 – the privacy killer

- Standards leaks
  - EUI-64 maps MAC into public IPs
  - Allows tracking of devices across networks
  - Privacy *extensions* are optional
- Provider accidentally leaks
  - Routing aggregation groups dynamic prefixes
  - Prefix not changed without reconnection (long time)
  - Allows tracking by prefix even with privacy extensions
  - Prefix change prevents server operation

# Multiple prefixes

- Hand out three prefixes
  - A dynamic prefix D1 for new connections
  - A dynamic prefix D2 for still open connections
  - A stable prefix S for server operation
- Preferred vs. valid lifetime
  - New connections use preferred lifetime > 0
  - Only D1 has positive preferred lifetime
  - Connection can stay open until valid lifetime is reached
  - D2 is the most recent D1 to keep open connections
  - S has a really long valid lifetime

# Rotating prefixes

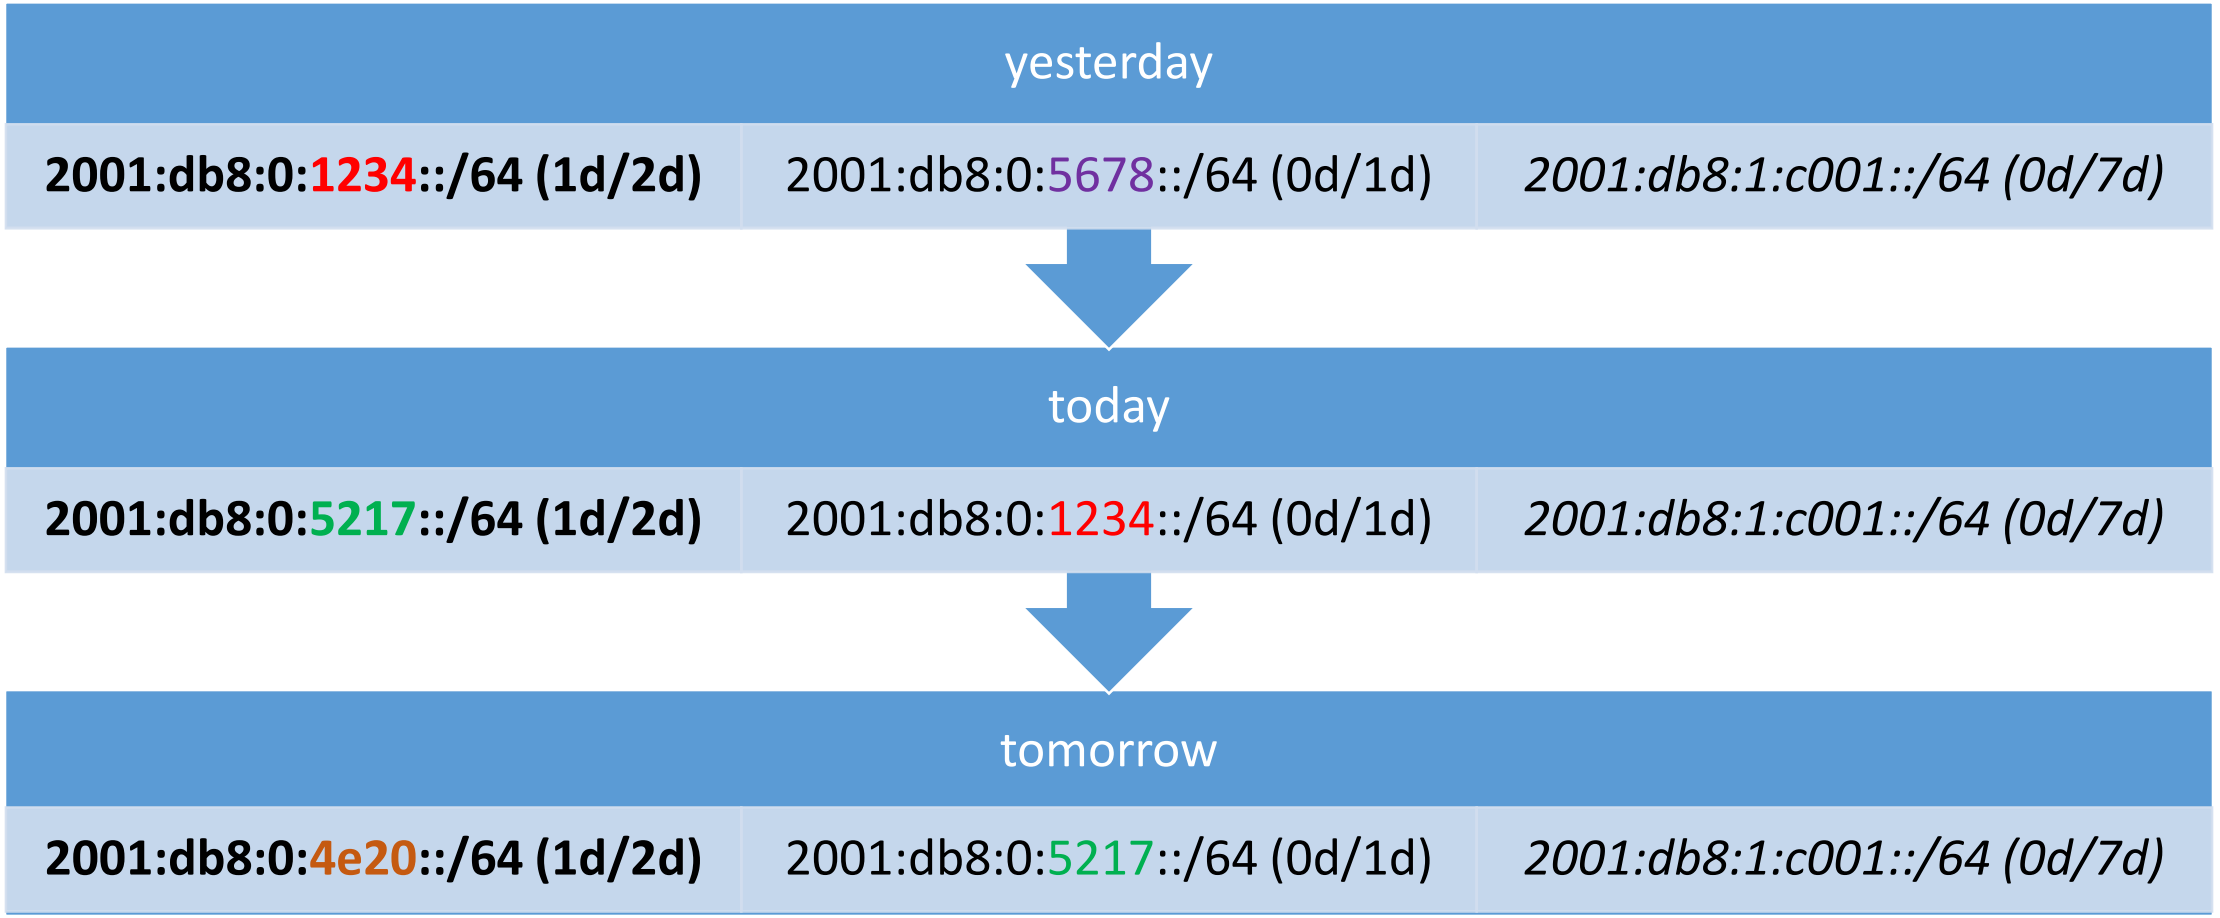

# Practical considerations

- Prefix size
  - AVM needs *at least* /60 for LAN, Guest, and Uplink
  - AVM behind AVM (extender) requires /59 for PD
- Rotation time
  - Experiments: hourly rotation, 1h preferred, 2h valid
  - CPE can request used prefixes to be extended
  - Different pools per hour (4 bit)
- Pool size
  - For randomness three times the customer count
  - 50k customers * /60 -> /44 * 16h -> /40 * 3 -> /38 minimum
  - Add large enough pool for static assignment, aggregateable

# Where to act

- PPP
  - L2TP termination via improved mpd (other talk)
  - Divert DHCP using ipfw on ng*
  - Perl-Script to assign prefixes and adding routes to kernel
  - Works since years
- DHCP
  - Routes need to be added on Router -> DHCP relay
  - Central DHCP server managing prefixes
  - Did not work for years

# IPv6 in L2 switched xDSL-networks

- Carriers filter aggressively
  - No broadcast to the CPE (beside ARP to the learned IP)
  - No multicast to the CPE
  - No unknown multicast in the network
- No multicast -> no ICMPv6 -> no IPv6
  - AVM tries SLAAC (fails), then blindly send DHCPv6
  - We see DHCPv6 requests from CPE
  - Router does not forward DHCPv6 responses to CPE
  - Router miss ICMPv6 ND responses (multicasted)

# Fixing IPv6 in L2 switched xDSL-networks

- Obtain DLSAM
  - Project for tests with carrier
  - Technician only meetings
  - Get console access and documentation
  - Try almost all IPv6 related option: **icmpv6-sec-fltr**
  - Adding security opens the multicast channels
- Outside the lab
  - Option will not be rolled out
  - Fake the ND responses, like we fake ARP since years

# Faking ND in production

- Extending parpd
  - Derive MAC from EUI-64 link-local IP
  - Learn and respond to normal ND traffic
  - Monitor unmatched ND traffic

```
interface lagg140
  timeout  1.164
  rule     0.0.0.0/0   0.0.0.0/0   ignore
  rule6    fe80::/96   fe80::/64   eui
  rule6    fe80::/64   fe80::/64   tell
  rule6    ::/0        ::/0        verbose ignore
end
```

# Prefixes in production

- show ipv6 dhcp relay binding
  Total number of Relay bindings = 261
  Total number of IAPD bindings = 261
  Total number of IANA bindings = 0
  Total number of Relay bindings added by Bulk lease = 0
- show ipv6 route static
  S 2A01:75C0:1140:E200::/60 [1/0]
  via FE80::464E:6DFF:FE20:7F73, Vlan140
  S 2A01:75C0:1140:E280::/60 [1/0]
  via FE80::464E:6DFF:FE5D:B7D9, Vlan140
  S 2A01:75C0:1140:E4E0::/60 [1/0]
  via FE80::E228:6DFF:FE28:3C32, Vlan140
  S 2A01:75C0:1140:E5A0::/60 [1/0]

# IPv6 is really used

00:07:16.429150 fe80::51:140 > fe80::ca0e:14ff:fe29:e316: [icmp6 sum ok] icmp6: neighbor sol: who has fe80::ca0e:14ff:fe29:e316(src lladdr: 5c:83:8f:3b:6f:3f) [class 0xe0] (len 32, hlim 255)

*Unicasted* ND query from the router for a CPE. (not seen by parpd)

00:07:16.446221 fe80::ca0e:14ff:fe29:e316 > fe80::51:140: [icmp6 sum ok] icmp6: neighbor adv: tgt is fe80::ca0e:14ff:fe29:e316(SO)(tgt lladdr: c8:0e:14:29:e3:16) (len 32, hlim 255)

*Real* ND response from the CPE.

00:07:16.537675 2a01:75c0:1140:ff90:219:99ff:fe99:2b8e.33270 > 2a01:4f8:202:34a2:2:100:0:111.443: P [tcp sum ok] 2818149766:2818150359(593) ack 2150523048 win 1432 <nop,nop,timestamp 352457737 615010237> [flowlabel 0xb79fc] (len 625, hlim 63)

00:07:16.567626 2a01:4f8:202:34a2:2:100:0:111.443 > 2a01:75c0:1140:ff90:219:99ff:fe99:2b8e.33270: P [tcp sum ok] 1:209(208) ack 593 win 899 <nop,nop,timestamp 615010987 352457737> [flowlabel 0x92afe] (len 240, hlim 55)

00:07:16.581155 2a01:75c0:1140:ff90:219:99ff:fe99:2b8e.33270 > 2a01:4f8:202:34a2:2:100:0:111.443: . [tcp sum ok] 593:593(0) ack 209 win 1432 <nop,nop,timestamp 352457780 615010987> [flowlabel 0xb79fc] (len 32, hlim 63)