

Stateless IPv4 Network Address Translation

draft-tsou-stateless-nat44-02

Tina Tsou <tina.tsou.zouting@huawei.com>

Will Liu <liushucheng@huawei.com>

Simon Perreault <simon.perreault@viagenie.ca>

Reinaldo Penno <repenno@cisco.com>

Maoke Chen <fibrib@gmail.com>

Sunset4

IETF 85, Atlanta

2012-11-05

Problem

- Stateful NAT44 in CGN
 - Fragile
 - Complex
 - Hard to scale up
 - Hard to log mappings
 - Doesn't like asymmetric routing
 - Etc.

Solution overview

- Each subscriber gets part of an external address (port set)
- The external address and port set are encoded in the internal address that is assigned to the CPE.
- The CPE restricts itself to its allowed port set.
- The stateless NAT44 only translates addresses, extracting the needed information from the address itself.

Topology and bridge mode

- Stateless NAT44 topology




Figure 1: Stateless NAT44 topology

- Bridge mode - CPE is configured working as a transparent bridge




Figure 2: Stateless NAT44 topology: CPE as bridge

Address formats

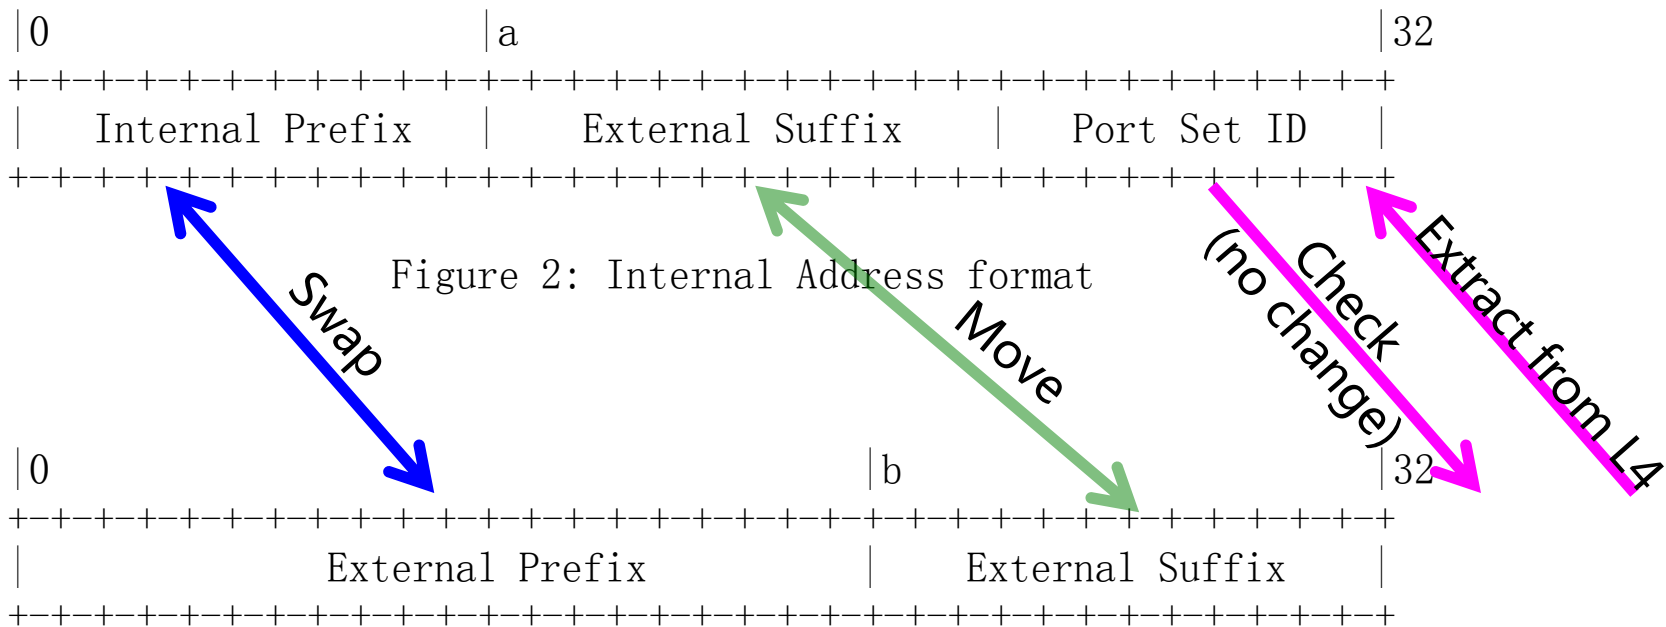


Figure 3: External Address format

Fragment handling

- Address translation depends on L4 port number, therefore fragments need to be reassembled
- Same considerations as NAT64, text adapted from RFC 6146

Non-contiguous port sets

- Optionally, non-contiguous port sets can be used.
- Requires provisioning a port set **mask** to the CPE.

```

|0          |5
+++++++
|1 1 1 0 1| Port Set ID (length n = 5 bits)
+++++++
& & & &
+++++++
|1 1 1 1 1| Port Set Mask
+++++++
| | | | |
V V V V V
+++++++
|1 1 1 0 1 x x x x x x x x x x x| Port Set = 59392-61439
+++++++
|0          |16

|0          |8
+++++++
|0 0 1 0 1 1 1 1| Port Set ID (length n = 8 bits)
+++++++
& & & & & & &
+++++++
|0 0 1 1 1 1 1 1| Port Set Mask
+++++++
| | | | | | | |
V V V V V V V V
+++++++
|x x 1 0 1 1 1 1 x x x x x x x x x| Port Set = 12032-12287, 28416-28671,
+++++++
                                         44800-45055, 61184-61439
|0          |16

```

Advantages

- Stateless CGN
 - No logging
 - Robust, scalable, etc.
 - Allows asymmetric routing (careful with fragments)
- Minimal modifications to CPE
 - A vanilla Linux home router can do this today.
- Fits into existing infrastructure and operational practices
- Very flexible port set definition

How is this different from...

- SD-NAT44 [draft-penno-softwire-sdnat-01]
 - SD-NAT has per-customer state.
SLNAT44 has no state at all.
 - SLNAT44 supports non-contiguous port sets.
 - SLNAT44 has no signaling between CGN and CPE.
 - SD-NAT44 has been dropped from draft -02.
- MAP, 4rd
 - SLNAT44 is IPv4-only, no impact on IPv6.

Next steps

- Is Sunset4 interested in this draft?
- What would be the next steps to progress this draft?