

# Source Address Dependent Route Information Option for Router Advertisements



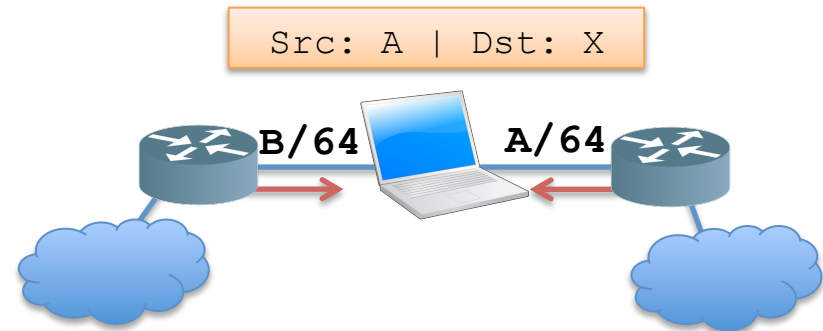
draft-pfister-6man-sadr-ra-00

Speaker: Pierre Pfister

# The Problem

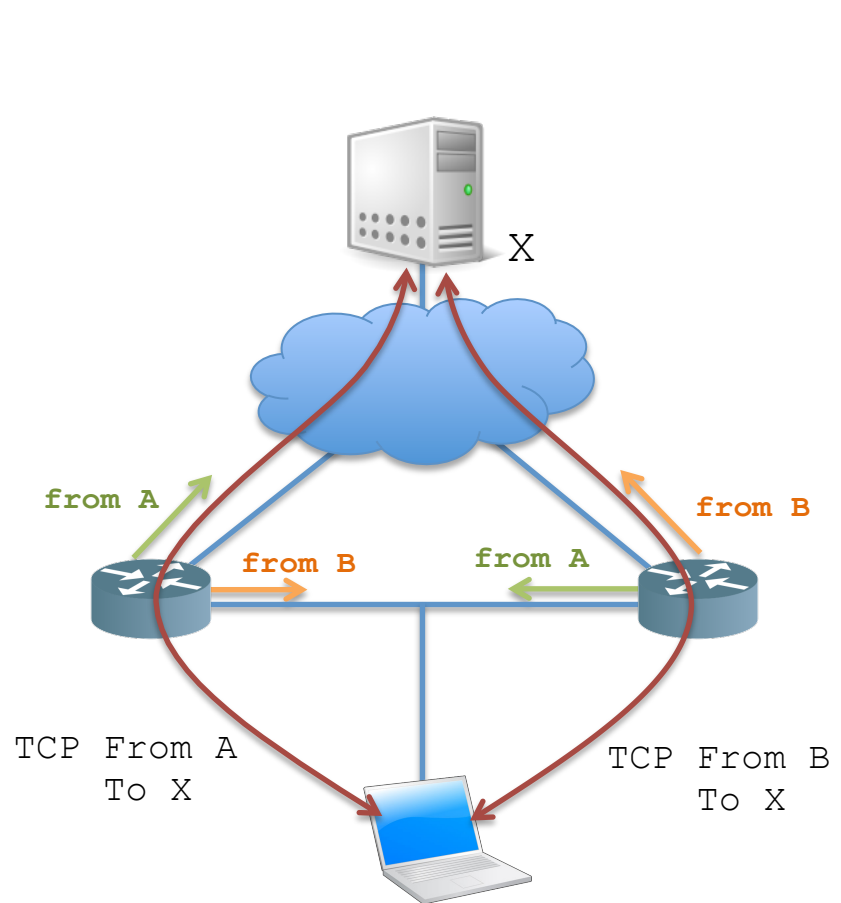
draft-sarikaya-6man-sadr-overview

- Routers do SADR and/or BCP38.
- Hosts will be connected to such routers.
- RFC6724 rule 5.5 may help but...
  - Application may pick the address.  
(RFC6724 does not apply then)
  - PIO presence does not imply best next-hop.
- PIO associated with an address can be advertised by multiple or zero routers on the link.



# The Consequences

BCP 38: Drop packet OR Inefficient routing OR Redirect ping-pong !



homenet

Redirect ping-pong in a nutshell:

**Step 1:** Routers are SADR aware and know each other.

**Step 2:** Host opens two TCP connections with same dst. with different src.

**Step 3:** Flows get constantly redirected.

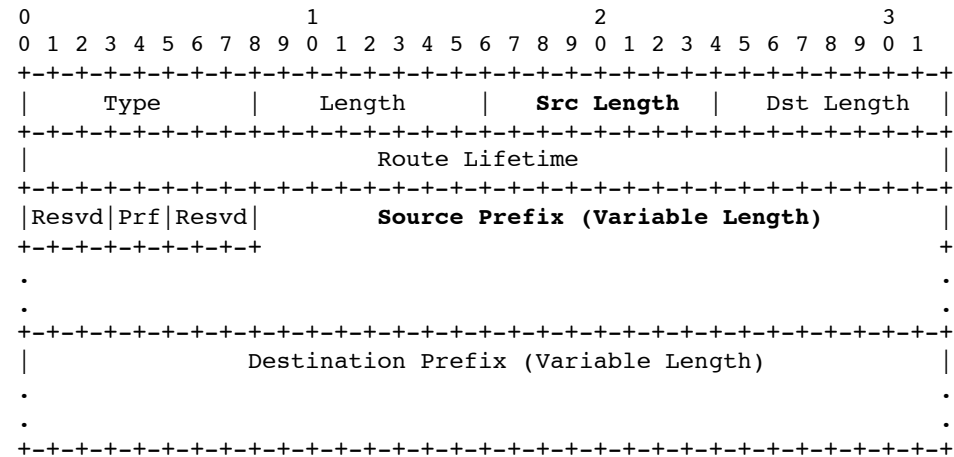
# Solution Overview 1/2

## 1. One new RA Option

Source Address Dependent Routing Information Option - SADRIO

RFC4191 RIO  
+  
Source Prefix

=




- All RIO fields are kept.
- Processing is the same (But with SADR routes).

# Solution Overview 2/2

## 2. One bit in the classic RIO (RFC4191)

Ignore bit for SADRIO capable hosts.



For backward compatibility with RFC4191.

**I bit is set:** Ignore the TLV.

**I bit is not set:** Consider it as a SADRIO with `::/0` source prefix.

Allows independent configuration of SADRIO aware and unaware host.

# Hosts Requirements

- Hosts have a set of SADR entries identified by:
  - Source Prefix
  - Destination Prefix
  - Router (link-local) address
  - Interface on which the SADRIO is received
- When parsing a RA:
  - Consider RIOs with **I** not set as SADRIO with `::/0` src prefix.
  - *add* the route if not present (and non-0 lifetime)
  - *update* lifetime + preference if present.
  - *delete* if present and 0 lifetime
- When sending a packet prefer entry with (in this order):
  1. Longest dst match length
  2. Longest src match length
  3. Greater router preference value

# Routers Requirements

- Do not send multiple SADRIOs with same src and dst prefix.
- Do not send multiple RIOs with same dst prefix (RFC4191).
- Do not send RIO with **I** bit not set and SADRIO with same destination prefix and `::/0` src prefix.
- Deprecate route (0 lifetime) when removed.

SADRIO aware hosts and non-aware hosts can be independently configured. e.g:

- **I** bit always set.
- Use SADRIOs with `::/0` src. prefix

# Why this design ?

1. Why not just using PIOs ?
  - No router priority.
  - Only ::/0 destinations.
  - Decoupling Configuration Vs Routing.
  
2. Why not ignoring a RIO when a SADRIO is present ?
  - Would require more TLVs per packet.
  - Multi-pass RA parsing (Not stateless).
  
3. TLV alignment is awkward.
  - 32bits alignment => 3 to 6 wasted bytes.
  - 64 bits alignment => 7 to 14 wasted bytes.
  - IP header is often not even aligned in memory...



# Thanks

## References:

draft-sarikaya-6man-sadr-overview  
draft-sarikaya-6man-sadr-ra

## Linux Kernel Patch (with different TLV format):

```
include/net/ip6_route.h | 11  
include/net/ndisc.h     |  3  
net/ipv6/ndisc.c        | 39  
net/ipv6/route.c        | 69  
4 files changed, 92 insertions(+), 30 deletions(-)
```

# Questions ?