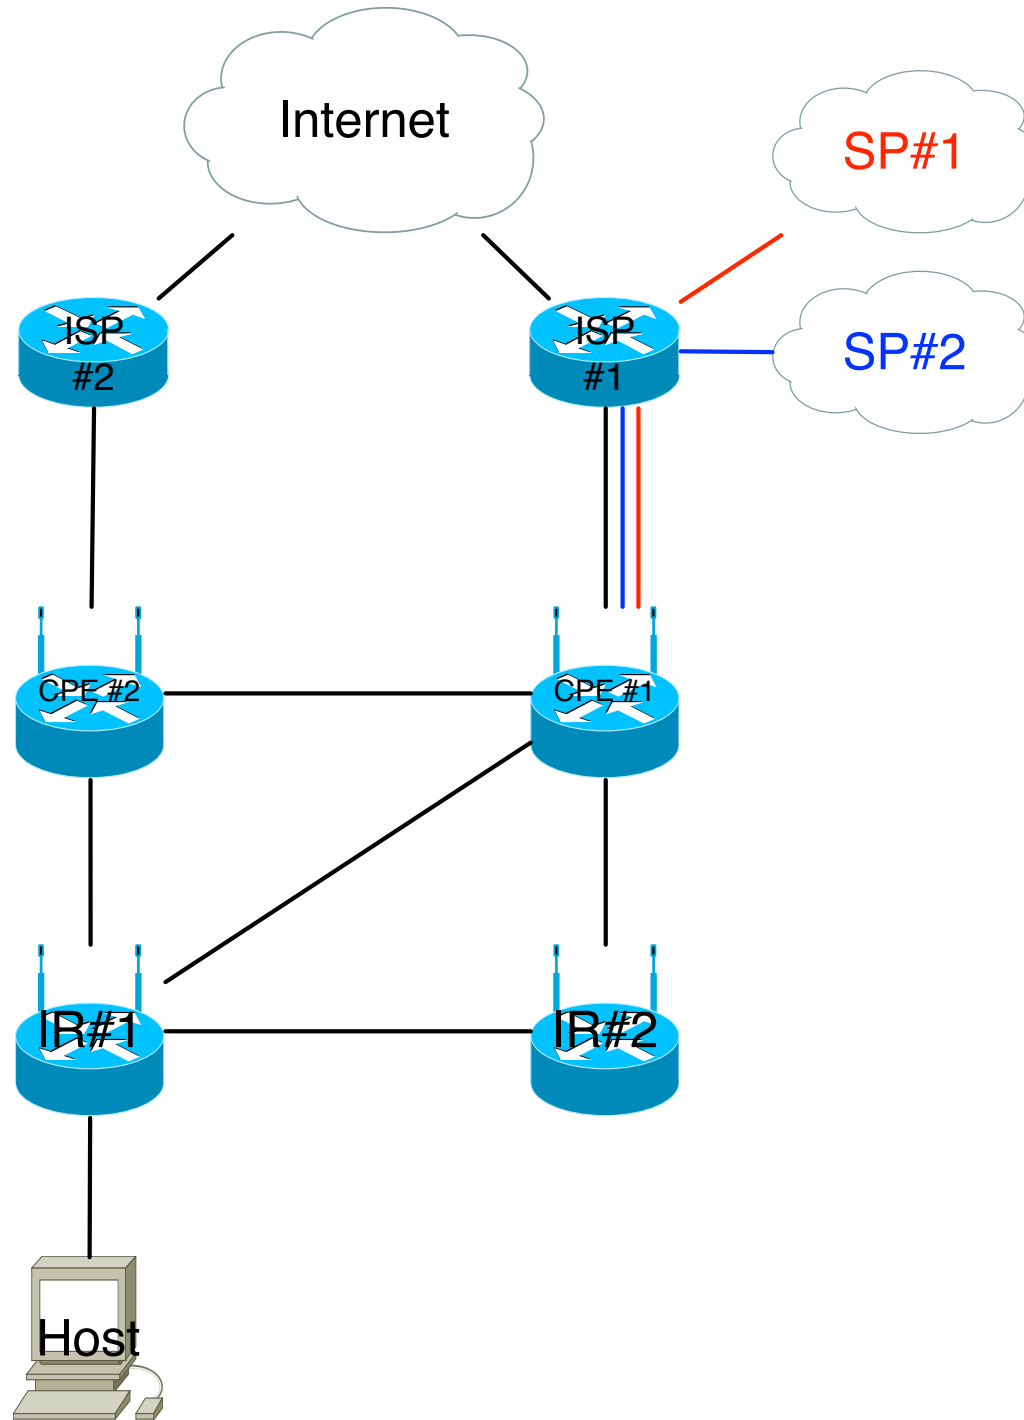


# draft-troan-homenet-sadr

Ole Trøan & Lorenzo Colitti

# Multihoming model

- External networks provide aggregates
  - E.g., a /56 via DHCPv6 PD
- Internal routers autoconfigure /64s from all aggregates on all interfaces
- All hosts get one address per aggregate
- Host chooses source address, network uses SADR to send packets to correct external network



# Forwarding - Notation

- (Source, Destination) / (S, D) routes
- (\*, 2::/64) - normal destination based route
  - Matches packets to 2::/64 from anywhere
- (1::/56, ::/0) – source specific default route
  - Matches packets to anywhere from 1::/56
- (1::/56, 2::/64) – more specific S,D route
  - Matches packets from 1::/56 to 2::/64
  
- Forwarding must use two lookup keys

# Forwarding - Which route wins?

- SA: 1::1, DA:: 2::1
- RIB:
  1. (\*, 2::1/128)
  2. (1::1/128, ::/0)
- Probably #1
  - The goal is to work with ingress filtering

# Forwarding

- Document says do longest match on DA first, then do longest match on the source constrained paths. If no match backtrack.
- Implementation can also do SA first.
  - Requires all (\*,D) routes replicated in all SA tables.
- All routers must follow same rules, otherwise network will loop
  - A non-SADR router in the middle can cause a loop!

# Populating routing tables

- Explicit (S,D) routes in routing protocol.
  - See \*baker\*-00
  - Ensures all routers support SADR (capability)
- Implicitly (hack alert!) Any external routes advertised by edge routers that also advertises usable prefixes is a S,D route.
  - Accept micro loops, since usable prefix and AS externals might not be inserted in the table exactly at the same time.
  - How do deal with legacy routers?

# Host - Multi-homing issues

- ISP connection down
  - (S,\*) disappears
  - Host will instantly get ICMP destination unreachable with code 0 (No route to destination) or 5 (Source address failed ingress/egress policy)
  - This doesn't help, RFC 1122 says TCP MUST NOT abort connection on code 0 unreachable
- Deprecating addresses will fix problem for new connections
  - Delay for dampening?
  - Leave last source prefix alone?