

Conditional Router Advertisements for Enterprise PA Multihoming

draft-ietf-v6ops-conditional-ras

Jen Linkova,
ENOG15, Moscow, June 2018

Enterprise Multihoming: Requirements

- Using Provider-aggregatable address space
- No BGP
- No NAT
- No changes on hosts

Problems with IPv6 PA Multihoming

Q: How to send packets to the correct uplink?

Q: How to implement policies?

Q: How to react to links failure/recovery?

Selecting the Uplink

Various Solutions Available/Being Developed

- *Good Old Policy Based Routing*
- *Source-Address Dependent Routing*
- *IPv6 Segment Routing*


etc...

Mutihomed Hosts: Work Being Done

- Multiple Provisioning Domains (mPVD)
 - Work in progress, takes a while to deploy
- MultiPath Transport (MPTCP, QUIC etc)
 - Other protocols?
 - Systems w/o multipath transport support?

Limiting the Scope

- Two uplinks used for Internet access (primary/backup or active/active)
 - No “walled gardens” etc
- Simple network topologies
 - *so even policy based routing would work...*



Problems to Solve: Signalling Events to Hosts

Uplink failure:

Addresses from that ISP prefix SHOULD NOT be used

Uplink recovery:

Addresses from that ISP prefix can be used again

Primary/Backup uplinks:

Addresses from the backup ISP SHOULD NOT be used if the primary uplink is up.

Existing Mechanism: SLAAC

"Addresses in the prefix can be used"



Preferred addresses (preferred lifetime > 0)

"Addresses in the prefix should not be used"



Deprecated addresses (preferred lifetime == 0)

Solution

Uplink failure:

Send RA with PIO preferred_lifetime = 0

Uplink recovery:

Send RA with PIO preferred_lifetime > 0

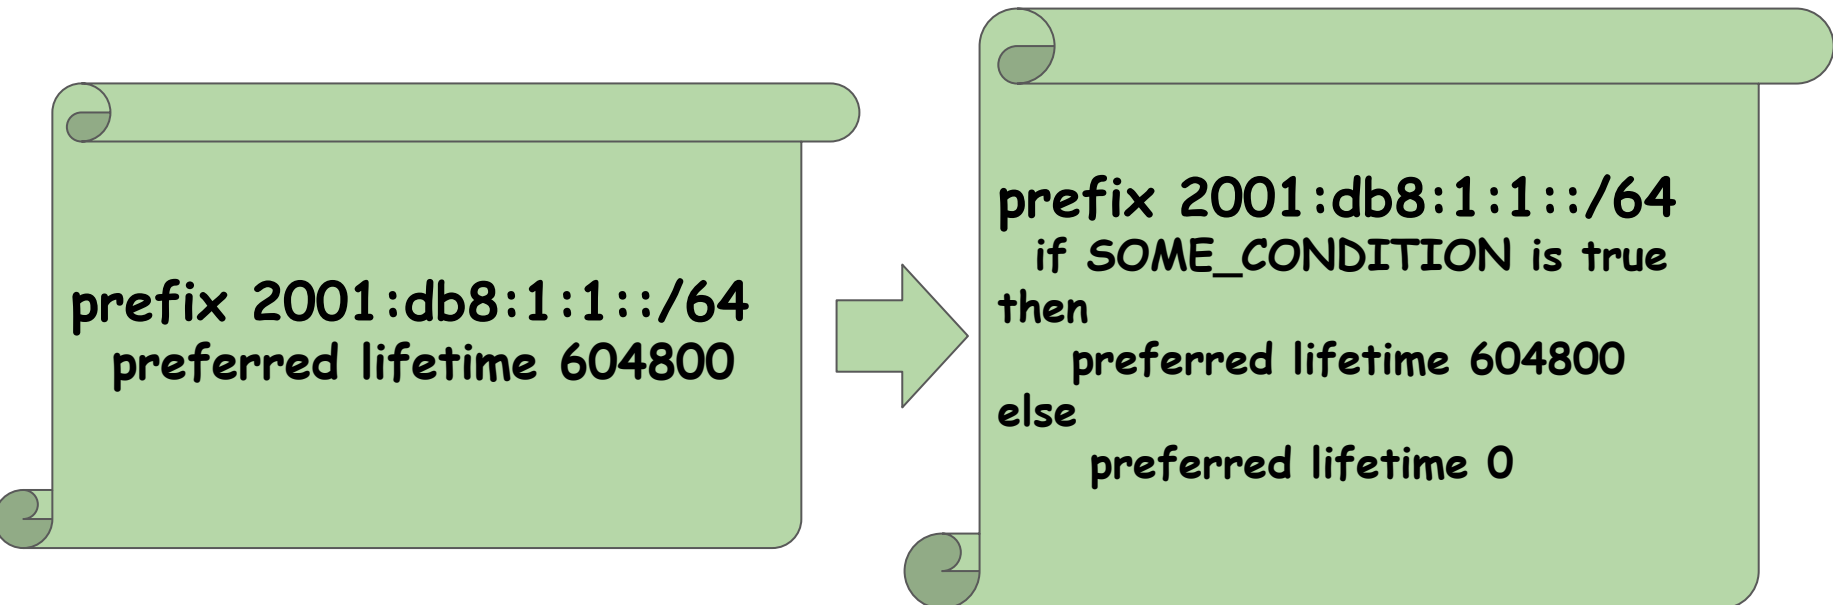
Primary/Backup uplinks:

Backup prefix preferred_lifetime = 0 if the primary uplink is up

Backup prefix preferred_lifetime > 0 if the primary uplink is down

Proposed Approach

- Router Advertisement fields values set conditionally
- Network events trigger new RAs being sent



prefix 2001:db8:1:1::/64
preferred lifetime 604800

prefix 2001:db8:1:1::/64
if SOME_CONDITION is true
then
preferred lifetime 604800
else
preferred lifetime 0

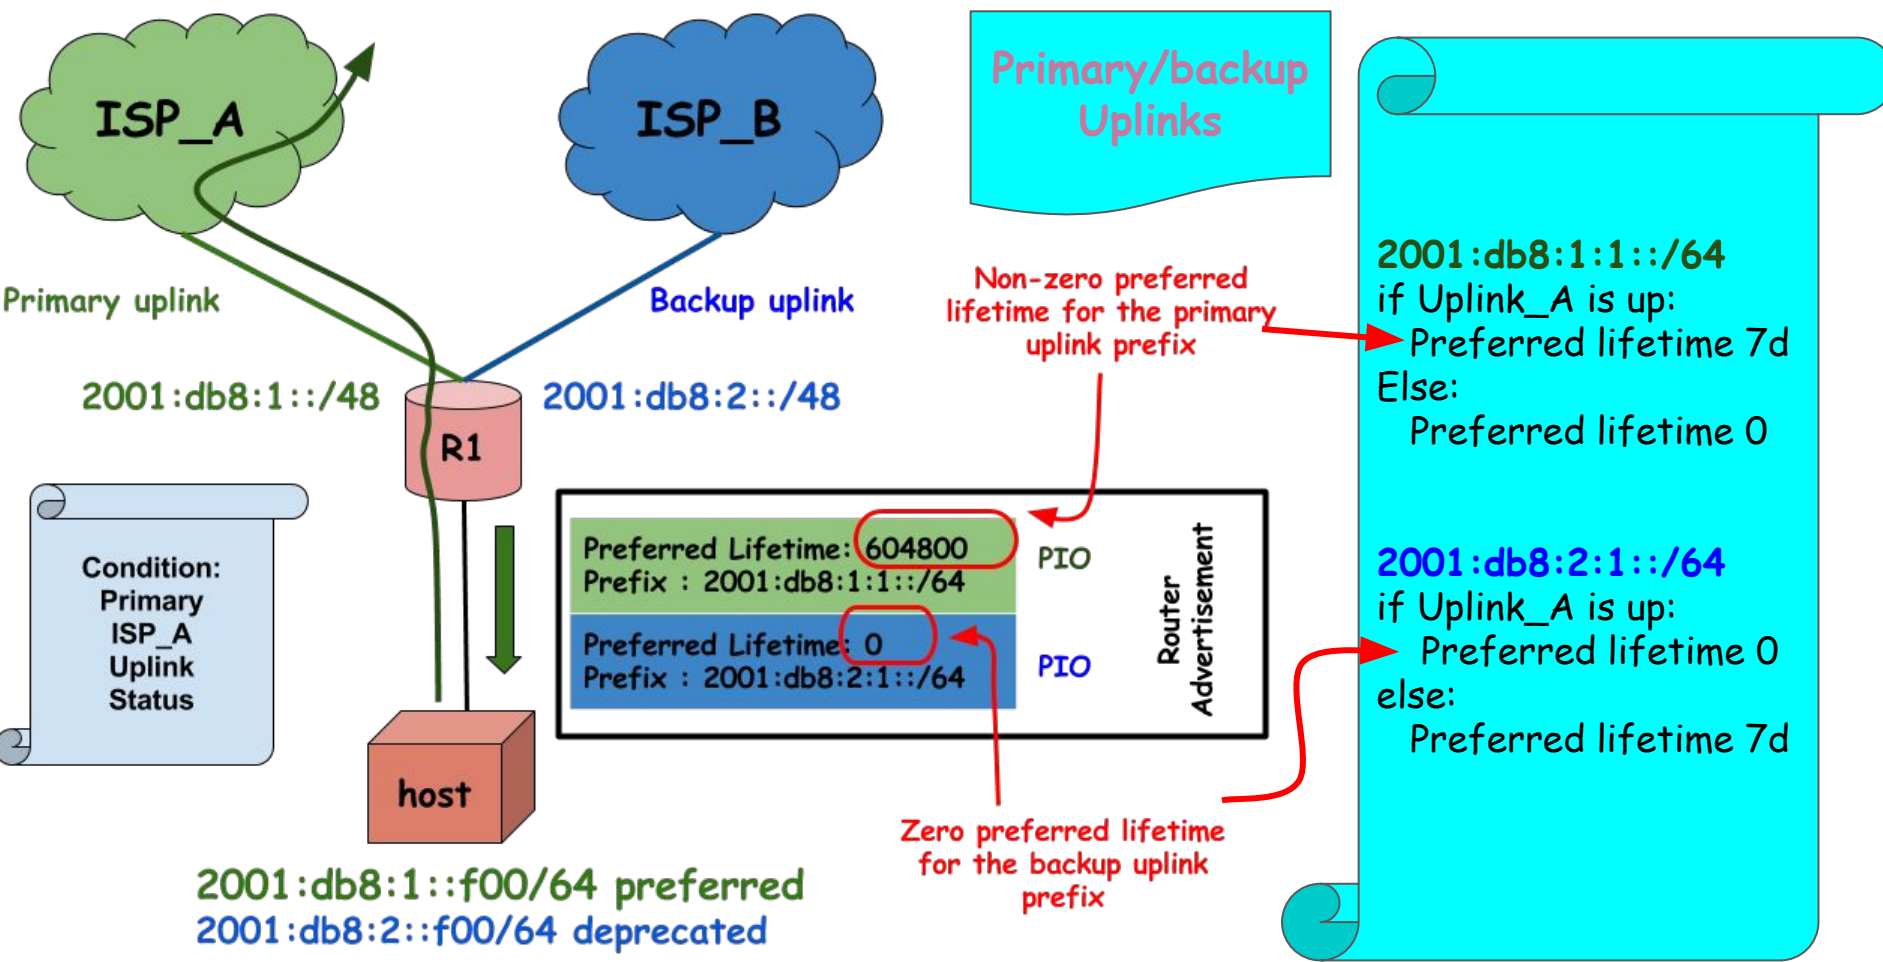
Potential Triggers

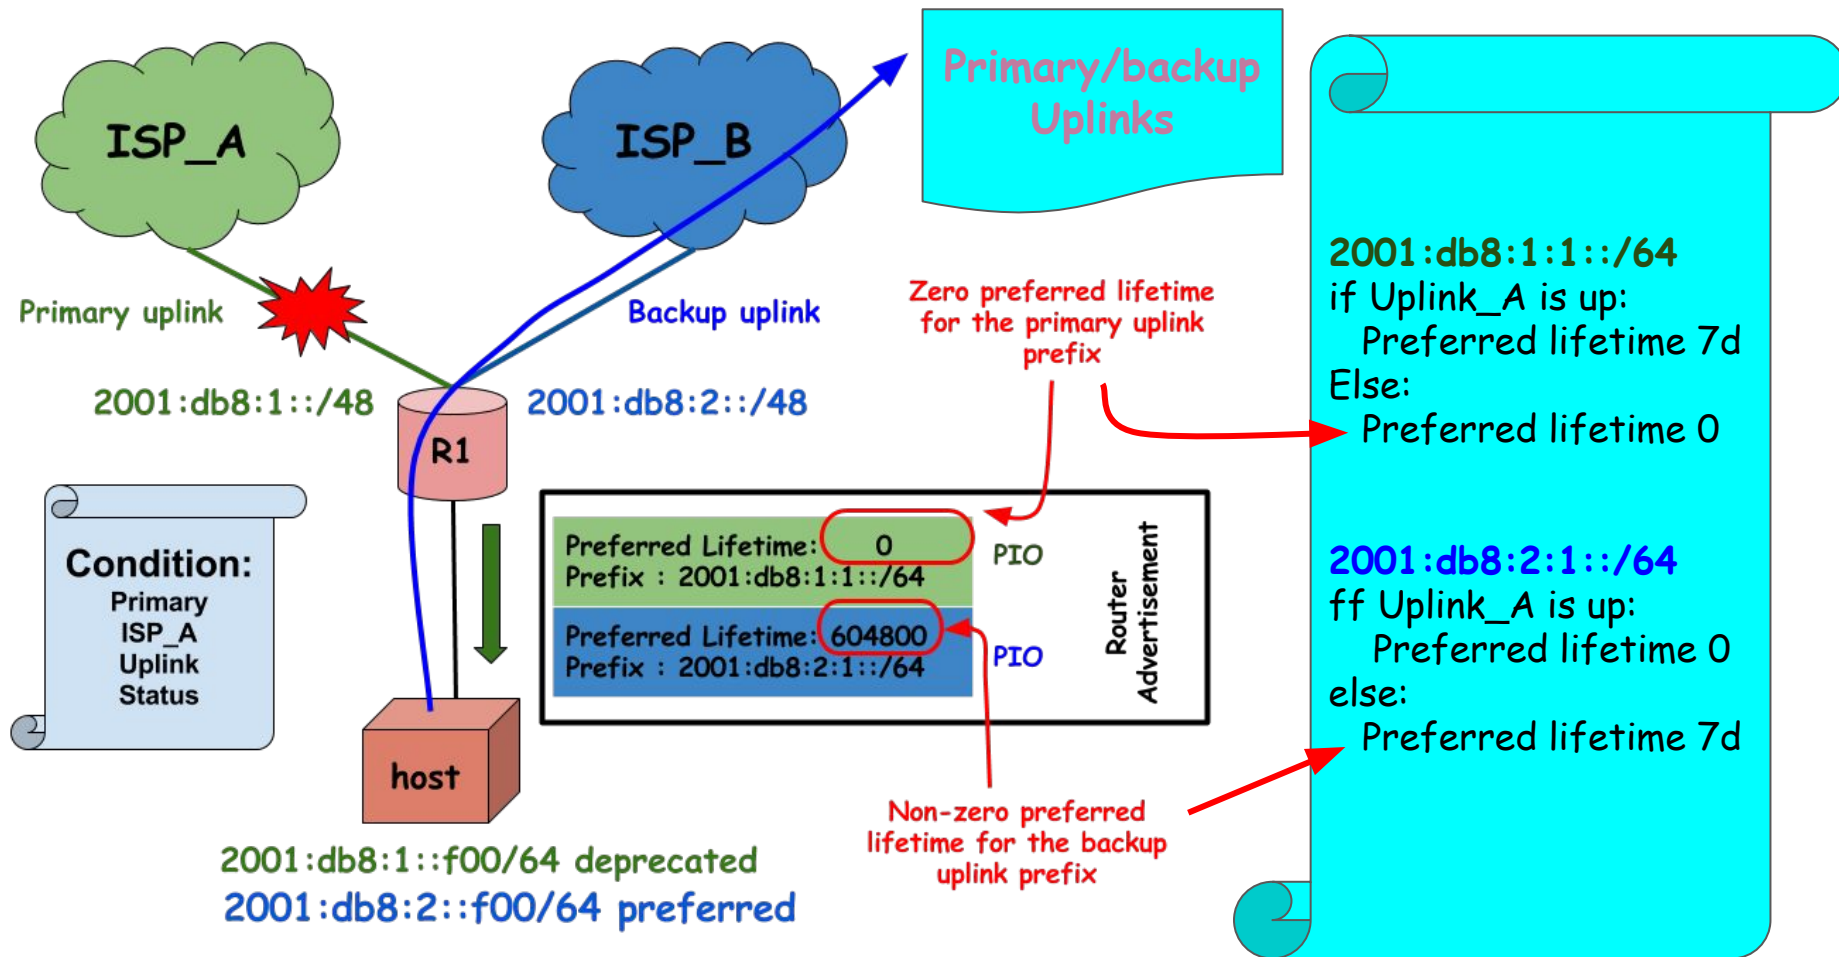
- Interface state
- Route presence
- Smth else...

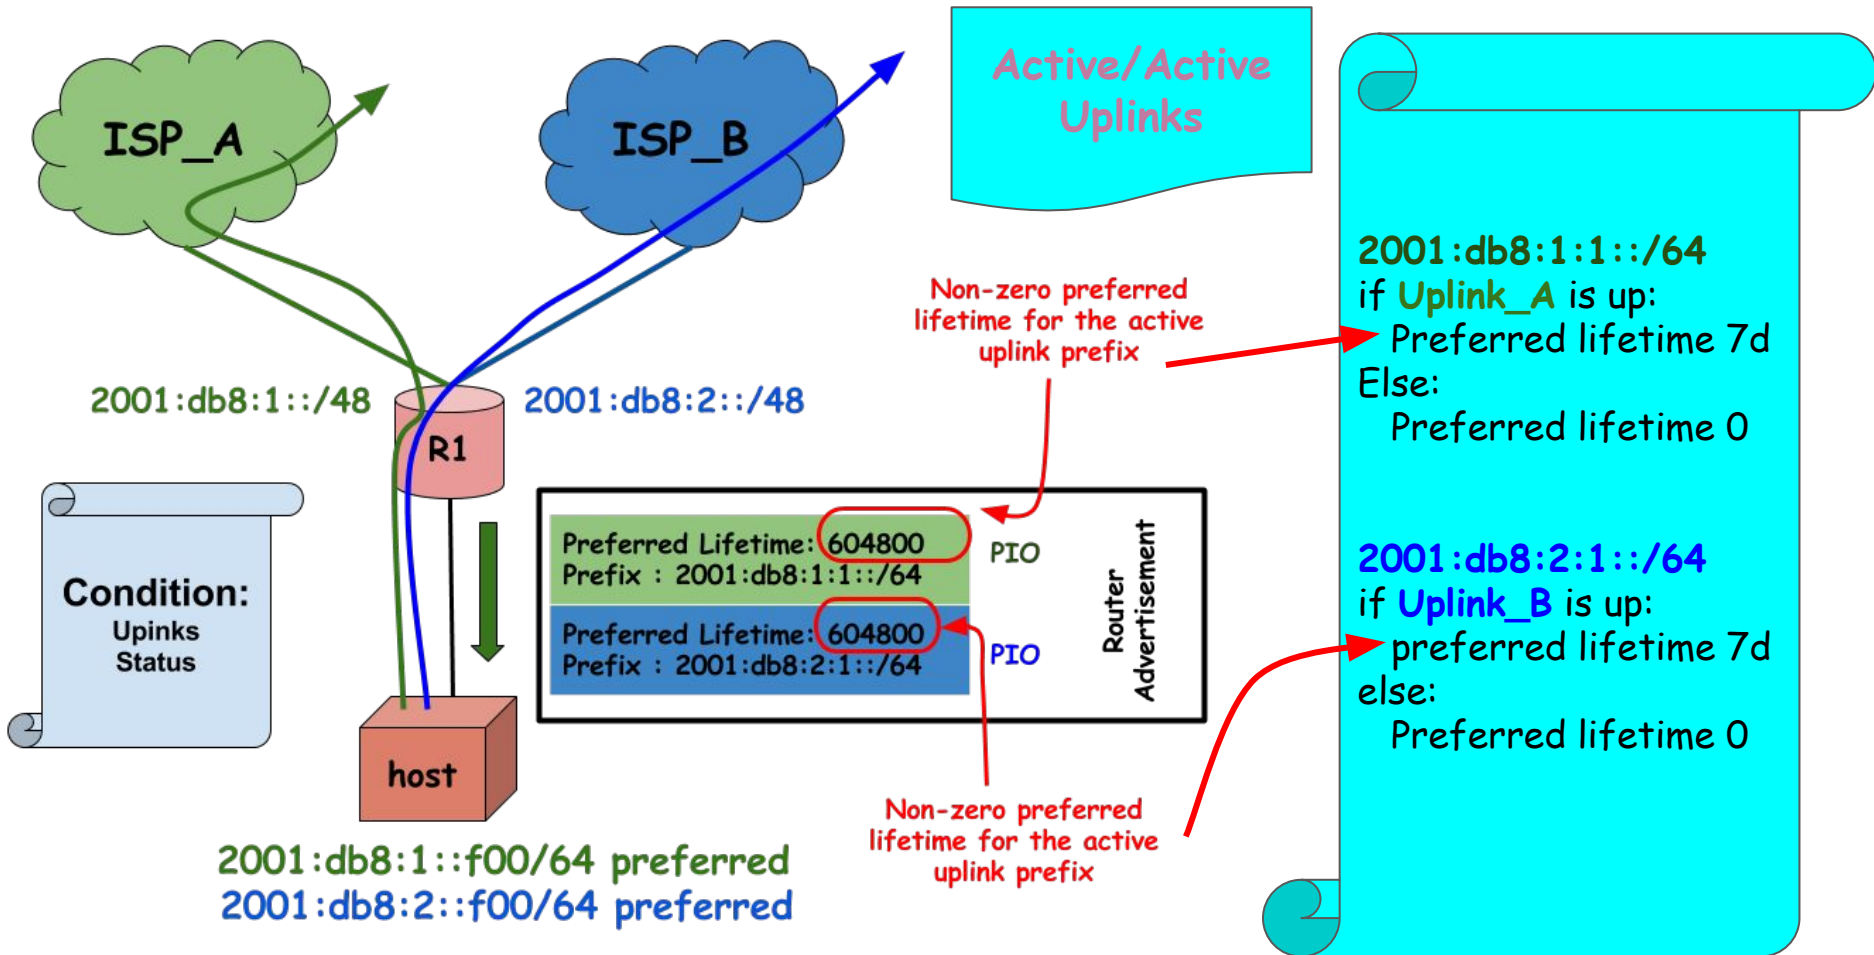
Fields to Be Updated

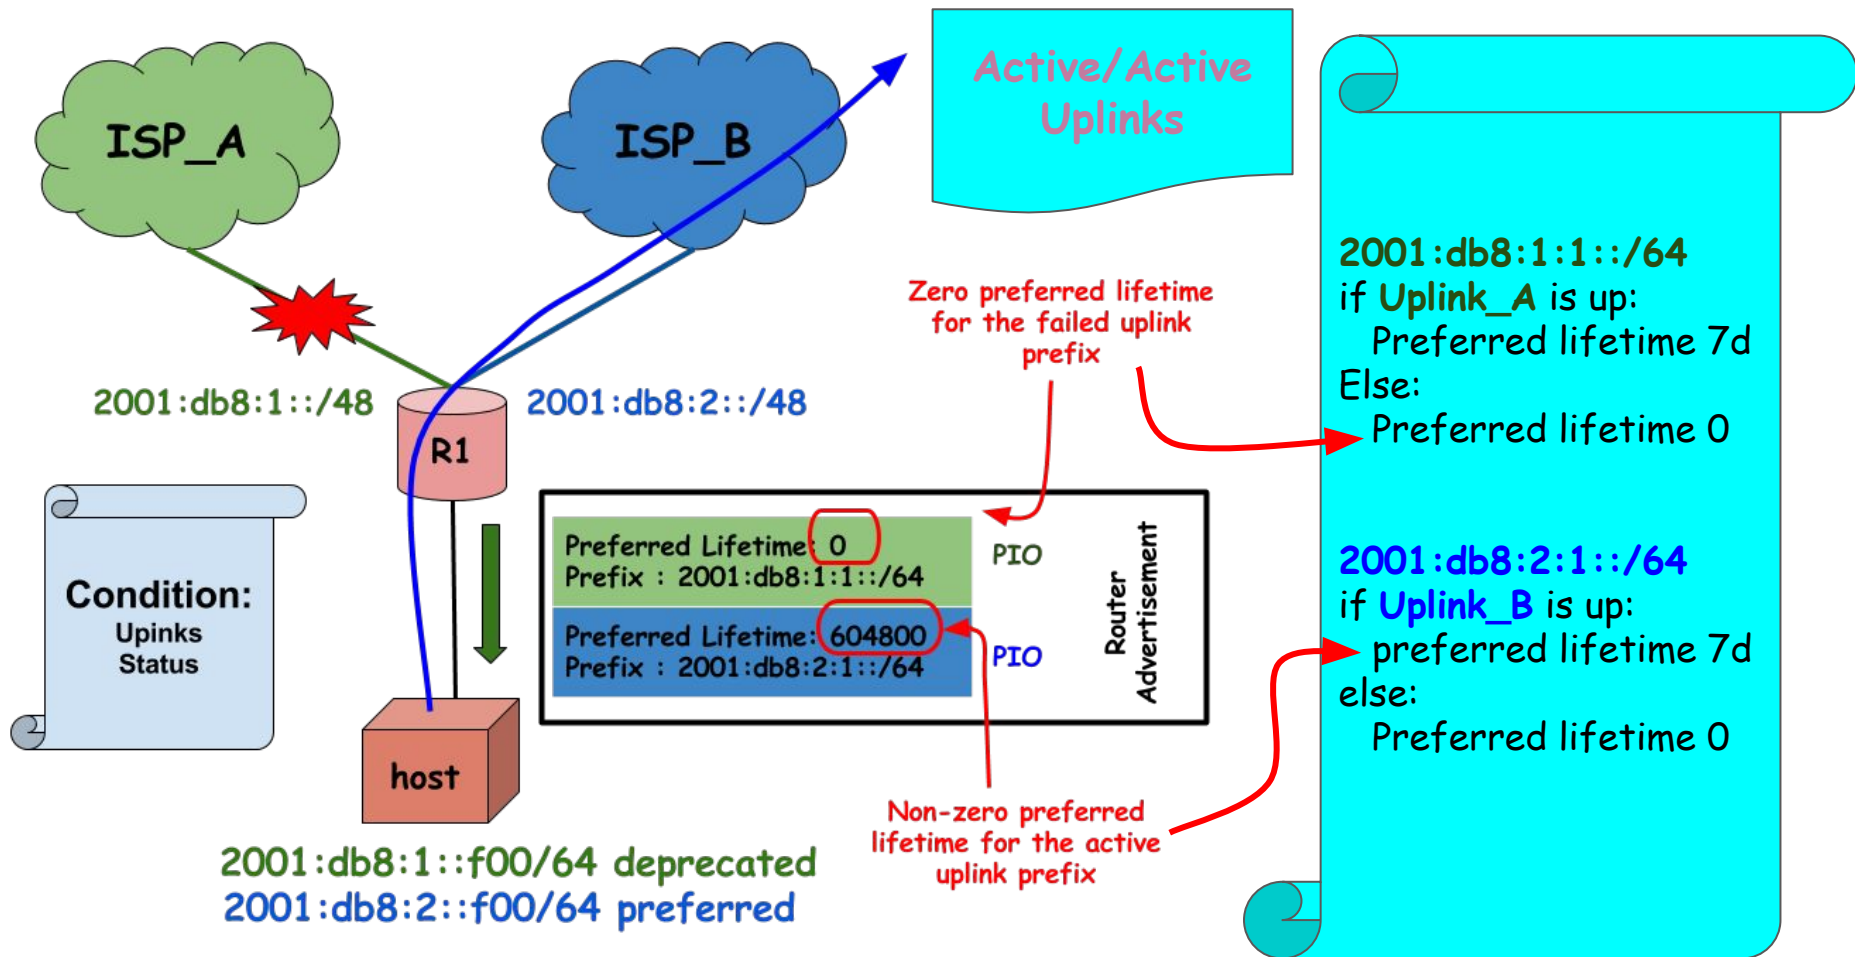
- PIO Preferred Lifetime
- RDNSS Lifetime
- Router Lifetime

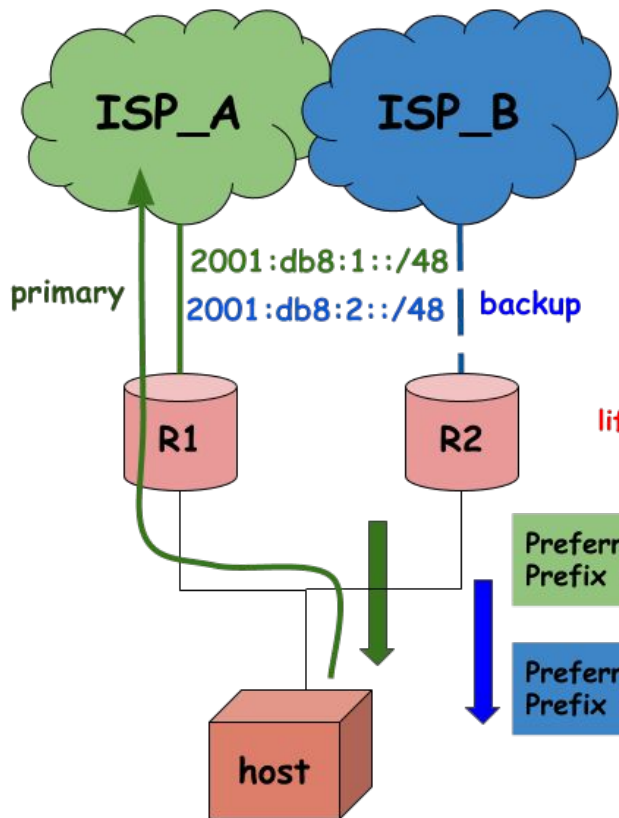
Example Scenarios











2001:db8:1:1::f00/64 preferred
2001:db8:2:1::f00/64 deprecated

Primary/Backup
Uplinks,
Two Routers

Non-Zero preferred
lifetime for the primary
uplink prefix

Preferred Lifetime: 604800
Prefix : 2001:db8:1:1::/64 R1 PIO

Preferred Lifetime: 0
Prefix : 2001:db8:2:1::/64 R2 PIO

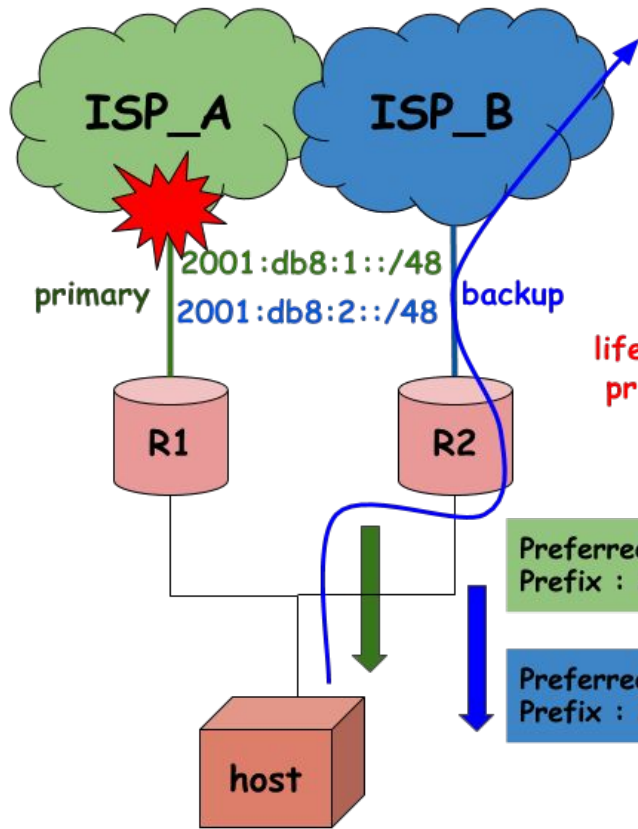
Zero preferred lifetime
for the backup uplink
prefix

R1 Policy:

2001:db8:1:1::/64
if Uplink_A is up:
Preferred lifetime 7d
Else:
Preferred lifetime 0

R2 Policy:

2001:db8:2:1::/64
if Uplink_A_route::
Preferred lifetime 0
else:
Preferred lifetime 7d



Primary/Backup
Uplinks,
Two Routers

Zero preferred
lifetime for the failed
primary uplink prefix

Preferred Lifetime: 0
Prefix : 2001:db8:1:1::/64

R1 PIO

Preferred Lifetime: 604800
Prefix : 2001:db8:2:1::/64

R2 PIO

Non-Zero preferred
lifetime for the backup
uplink prefix

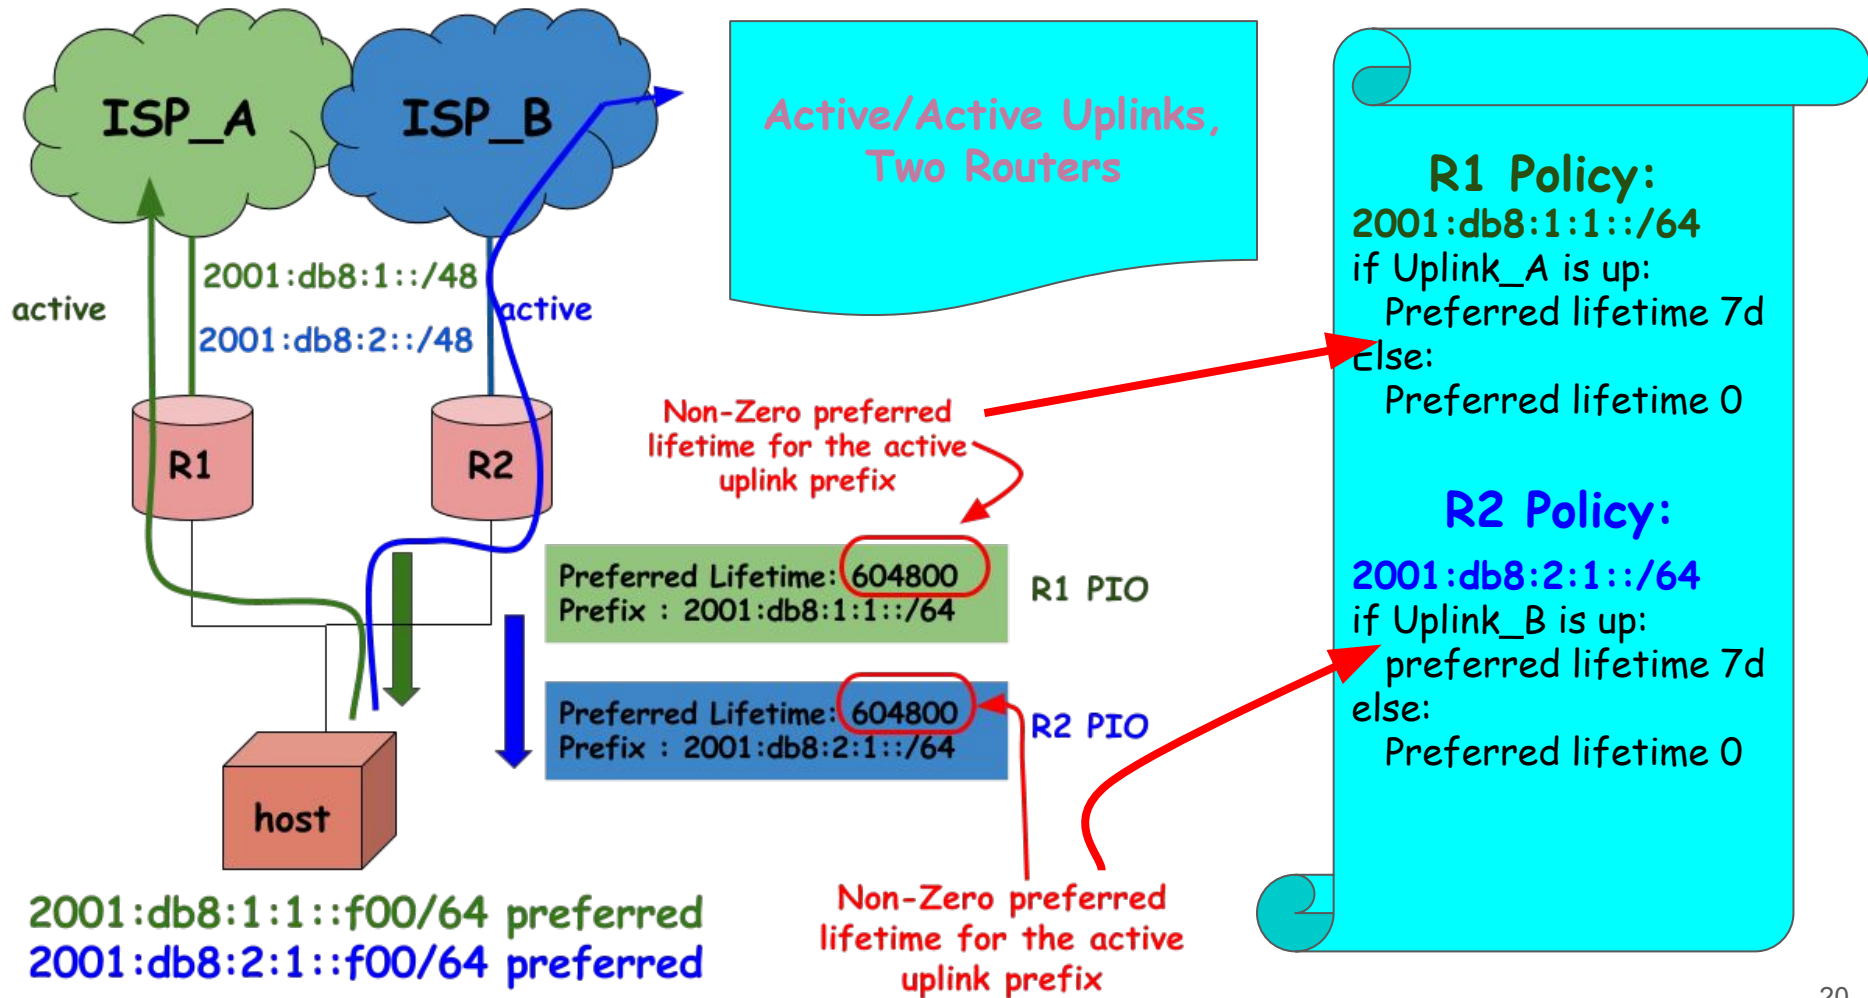
2001:db8:1:1::f00/64 deprecated
2001:db8:2:1::f00/64 preferred


R1 Policy:

2001:db8:1:1::/64
if Uplink_A is up:
Preferred lifetime 7d
Else:
Preferred lifetime 0

R2 Policy:

2001:db8:2:1::/64
if Uplink_A_route:
Preferred lifetime 0
else:
Preferred lifetime 7d





2001:db8:1:1::f00/64 deprecated
2001:db8:2:1::f00/64 preferred

Active/Active Uplinks,
Two Routers

Zero preferred lifetime
for the failed uplink
prefix

Preferred Lifetime: 0
Prefix : 2001:db8:1:1::/64
R1 PIO

Preferred Lifetime: 604800
Prefix : 2001:db8:2:1::/64
R2 PIO


Non-Zero preferred
lifetime for the active
uplink prefix

R1 Policy:

2001:db8:1:1::/64
if Uplink_A is up:
Preferred lifetime 7d
Else:
Preferred lifetime 0

R2 Policy:

2001:db8:2:1::/64
if Uplink_B is up:
preferred lifetime 7d
else:
Preferred lifetime 0



Active/Active Uplinks,
Dedicated Border
Routers

Non-Zero preferred
lifetime for the active
uplink prefix

Preferred Lifetime: 604800	PIO	Router Advertisement
Prefix : 2001:db8:1:1::/64		
Preferred Lifetime: 604800	PIO	
Prefix : 2001:db8:2:1::/64		

Non-Zero preferred lifetime for
the active uplink prefix

R3 Policy:

2001:db8:1:1::/64
if Uplink_A_route:
Preferred lifetime 7d
Else:
Preferred lifetime 0

2001:db8:2:1::/64
if Uplink_B_route:
preferred lifetime 7d
else:
Preferred lifetime 0

2001:db8:1:1::f00/64 preferred
2001:db8:2:1::f00/64 preferred

et cetera, et cetera...

Connection Preservation:

- Uplink failure:
 - connections interrupted (like IPv4 NAT)
- Uplink recovery:
 - connections are NOT interrupted (unlike IPv4 NAT)

Not Something New

IPv6 CPE L-13 requirement (RFC7084):

- Explicit prefix invalidation
- Homenet routers deprecate prefixes

Deploying Right Now

While we are awaiting for vendors to implement it...

We can still use it!

Example: JunOS event policies/event scripts

QUESTIONS?