Effect of IPv6 Prefix with Referrals

X. Li, C. Bao

2009-07-26

Outline

- Introduction
- Referral model
- Stateless translation and address representation
- Prefix selection
- Examples and discussions
- Remarks
- Recommendations

Introduction

- This presentation discusses the affect of IPv6 prefix with referrals, it is based on draft-bcx-behave-learn-address-00 and draft-xli-behave-v4v6-prefix-00.
- It is for scenarios 1 and 2, “an IPv6 network to the IPv4 Internet” and “the IPv4 Internet to an IPv6 network”.
- We discuss the stateless translation (e.g. IVI).

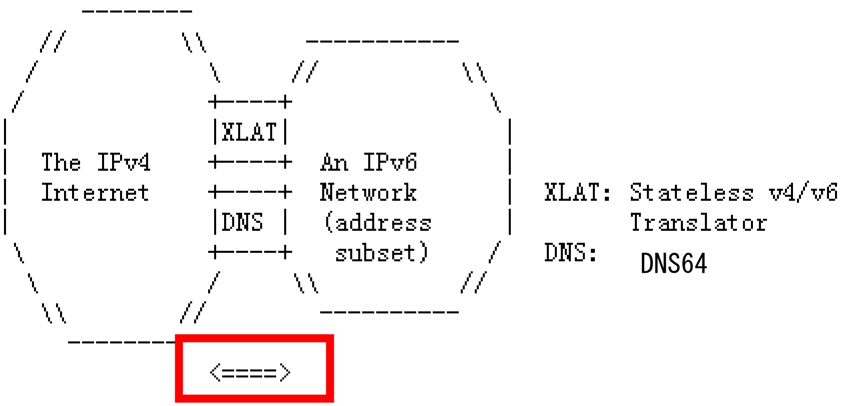

XLAT: Stateless v4/v6 Translator

DNS: DNS64

Figure 9: Stateless translator for Scenarios 1 and 2

Referral model

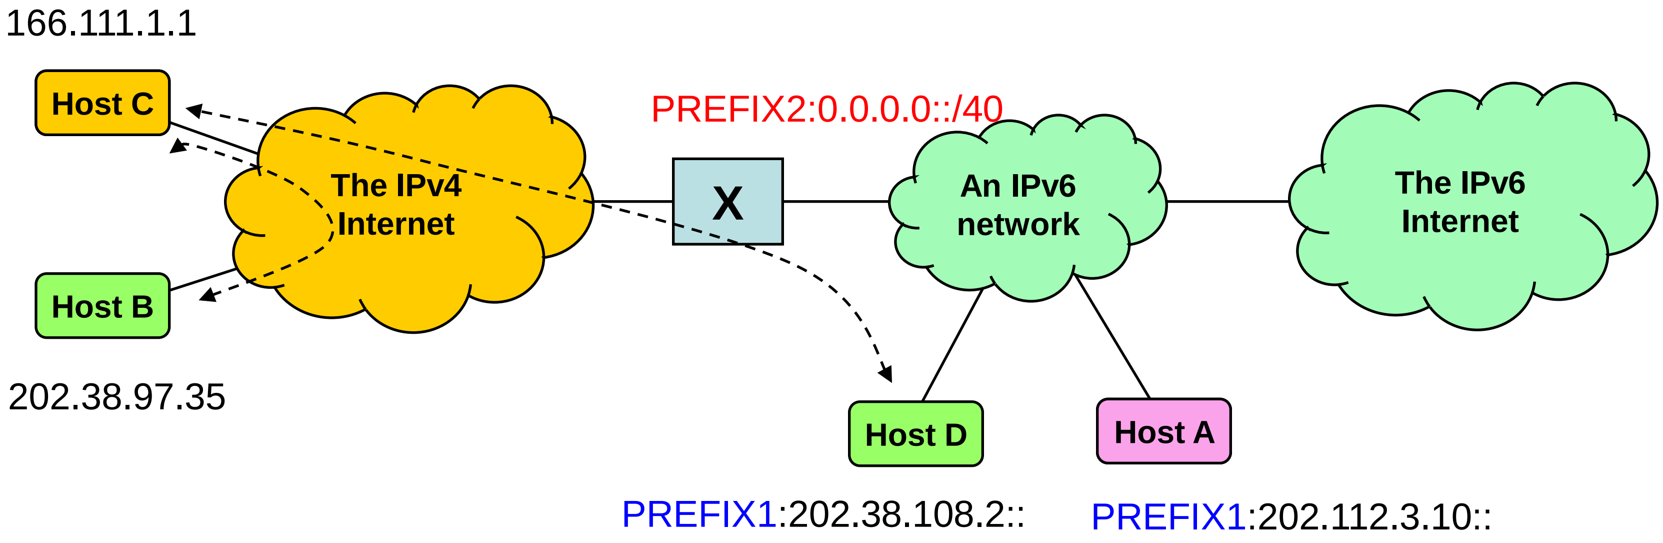

- Host A and Host D are IPv6 hosts using IPv4 embedded IPv6 addresses in an IPv6 network
- Host B and Host C are IPv4 hosts in the IPv4 Internet.
- Host C communicates with Host B in IPv4 Internet and communicates with Host D via translator.
- **Host A gets Host B and Host D’s IPv4 addresses (literals) from Host C via referral, i.e. Host A (PREFIX1:202.112.3.10) gets the IPv4 literals**
  - **Host B=202.38.97.35**
  - **Host D=202.38.108.2**

Stateless translation and address representation (1)

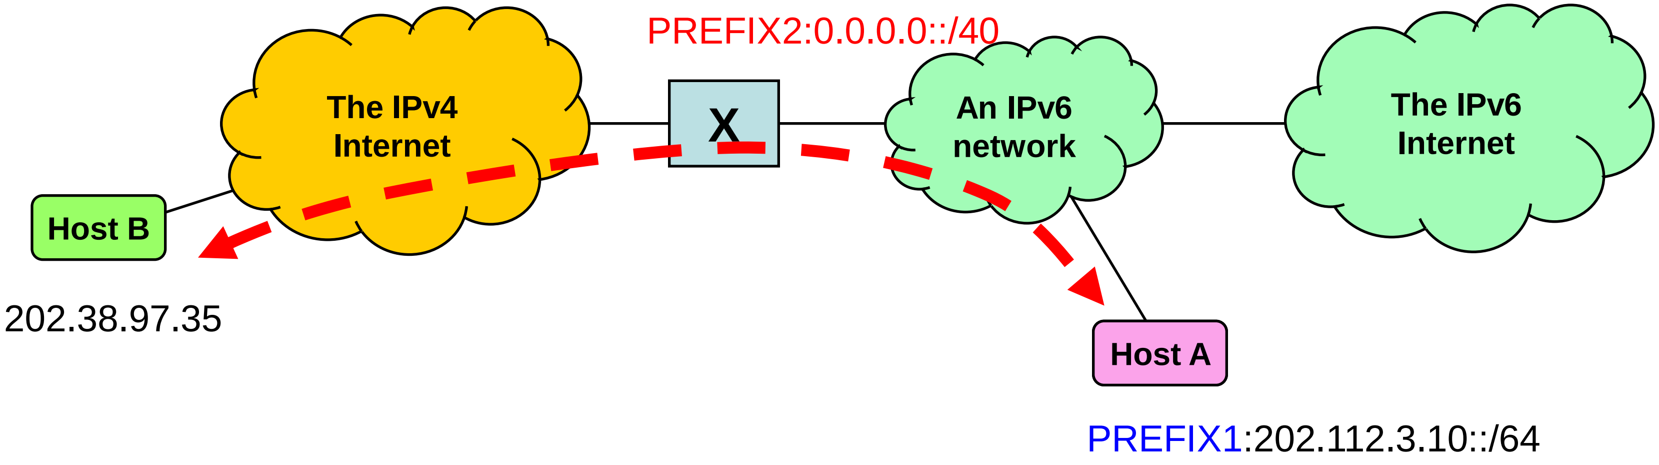

Represent IPv4 address in IPv6 using PREFIX2

In an IPv4 network, IPv4 addresses are assigned to IPv4 hosts (Host B), and these IPv4 addresses need to be translated to IPv6 addresses that can be reached by IPv6 hosts (Host A).

Represent IPv6 addresses in IPv4 using PREFIX1

In an IPv6 network, IPv6 addresses are assigned to IPv6 hosts (Host A), and these IPv6 addresses need to be translated to IPv4 addresses that can be reached by IPv4 hosts (Host B).

Stateless translation and address representation (2)

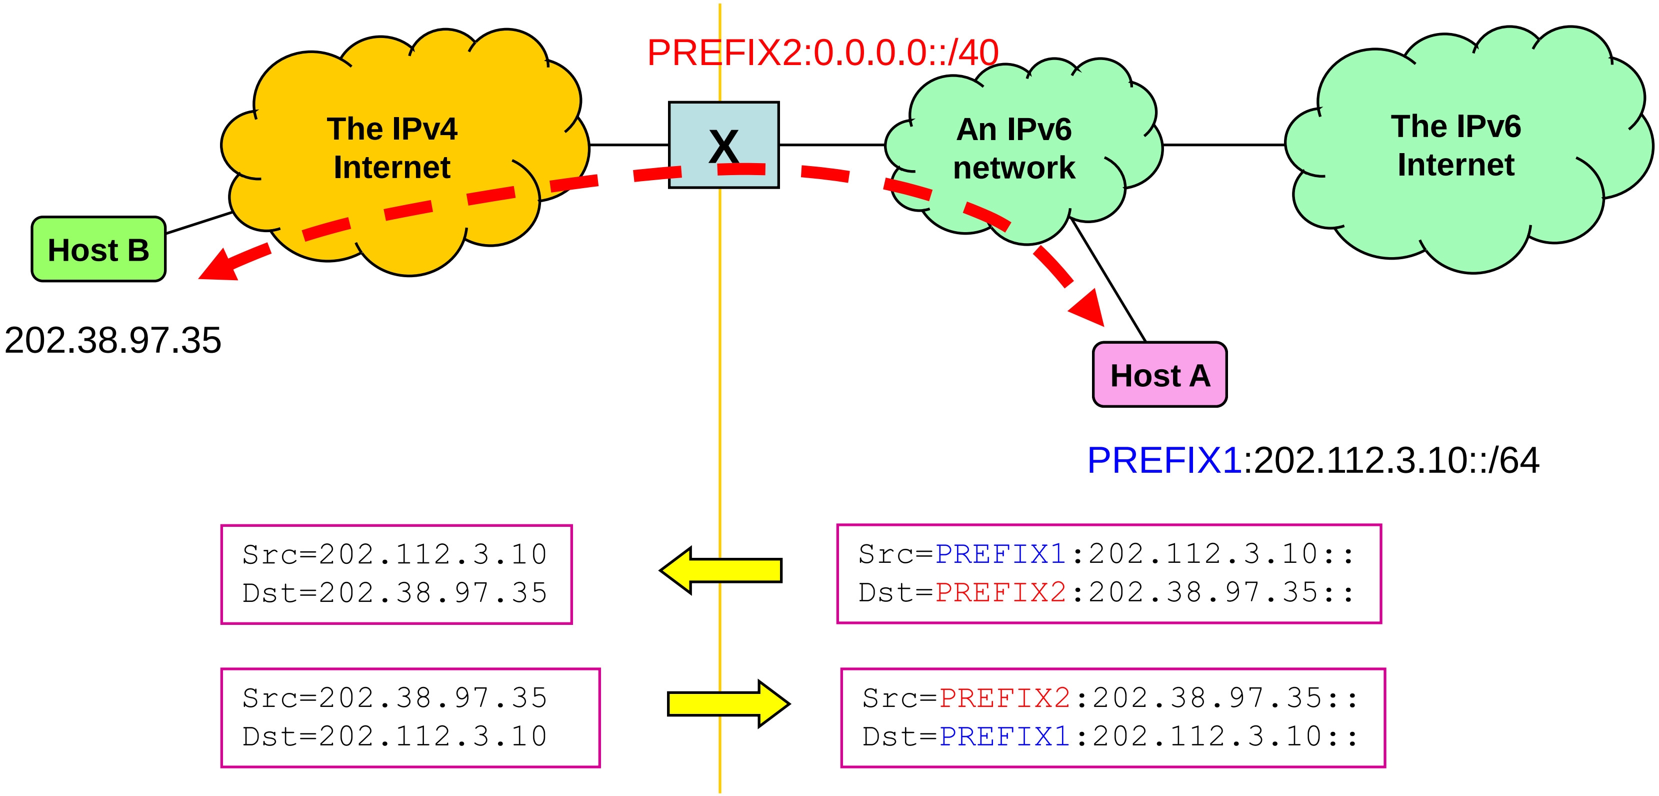

```
Src=202.112.3.10
Dst=202.38.97.35
```

```
Src=PREFIX1:202.112.3.10::
Dst=PREFIX2:202.38.97.35::
```

```
Src=202.38.97.35
Dst=202.112.3.10
```

```
Src=PREFIX2:202.38.97.35::
Dst=PREFIX1:202.112.3.10::
```

Stateless translation and address representation (3)

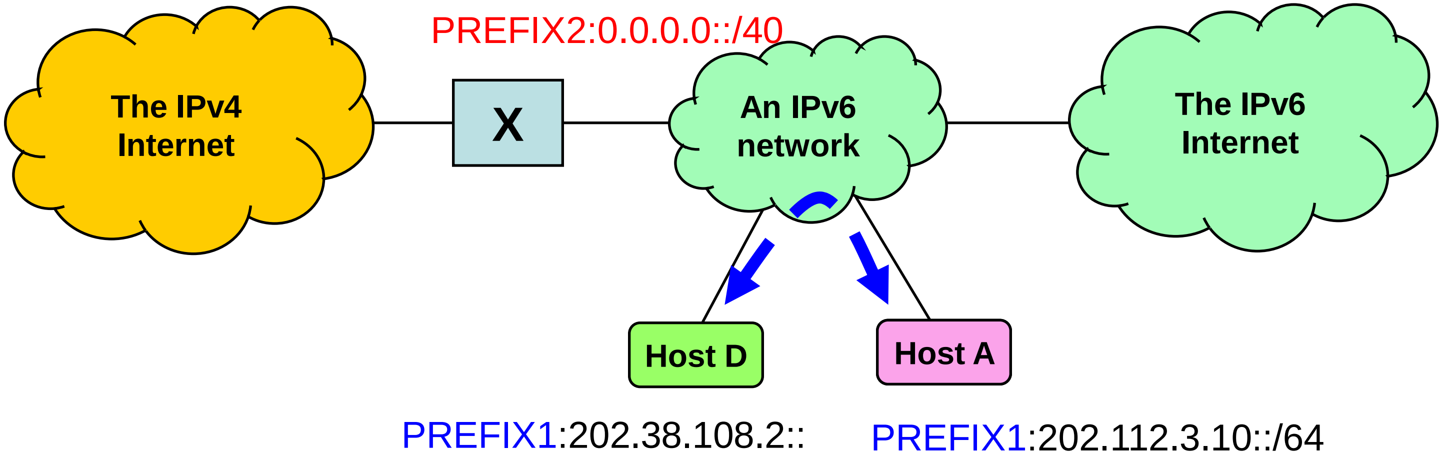

```
Src=PREFIX1:202.112.3.10::
Dst=PREFIX1:202.38.108.2::
```

```
Src=PREFIX1:202.38.108.2::
Dst=PREFIX1:202.112.3.10::
```

PREFIX1 and PREFIX2 selection

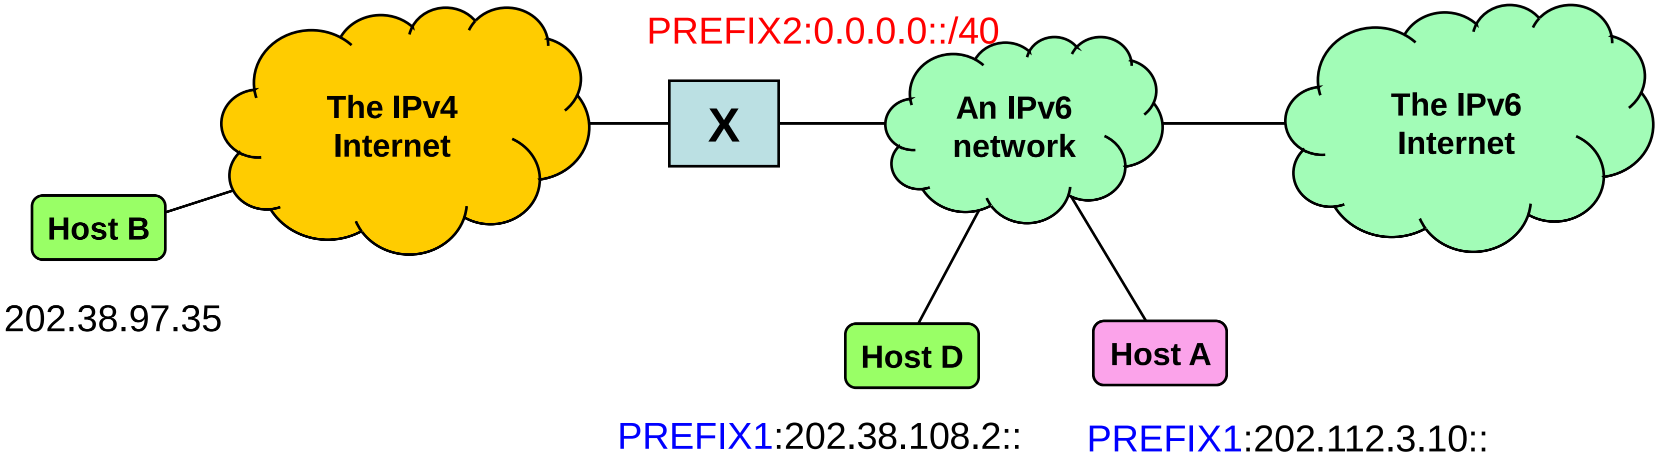

- How to choose PREFIX1 (host prefix) and PREFIX2 (translator prefix)
  1. PREFIX1=LIR, PREFIX2=LIR
  2. PREFIX1=WKP, PREFIX2=WKP
  3. PREFIX1=WKP, PREFIX2=LIR
  4. PREFIX1=LIR, PREFIX2=WKP

PREFIX1 considerations

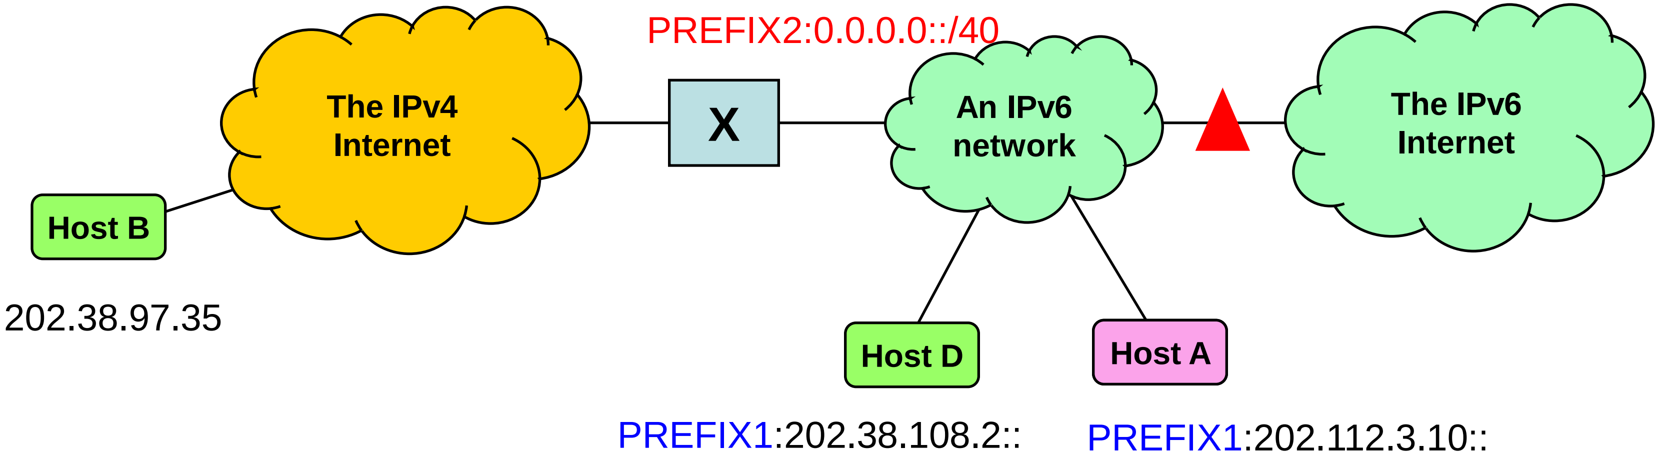

PREFIX2:0.0.0.0::/40

The IPv4 Internet

X

An IPv6 network

The IPv6 Internet

Host B

202.38.97.35

Host D

Host A

PREFIX1:202.38.108.2::

PREFIX1:202.112.3.10::

- Considering

  2. PREFIX1=WKP, PREFIX2=WKP (SIIT)
  3. PREFIX1=WKP, PREFIX2=LIR

  There is routing leakage problem for PREFIX1=WKP (see draft-xli-behave-v4v6-prefix-00). The workaround is to use multiple addresses for each host. However, there is no routing leakage problem for PREFIX1=LIR

  Therefore, for simplicity, PREFIX1 MUST be LIR

- Then we only need to evaluate

  1. PREFIX1=LIR, PREFIX2=LIR
  4. PREFIX1=LIR, PREFIX2=WKP

1. PREFIX1=LIR and PREFIX2=LIR

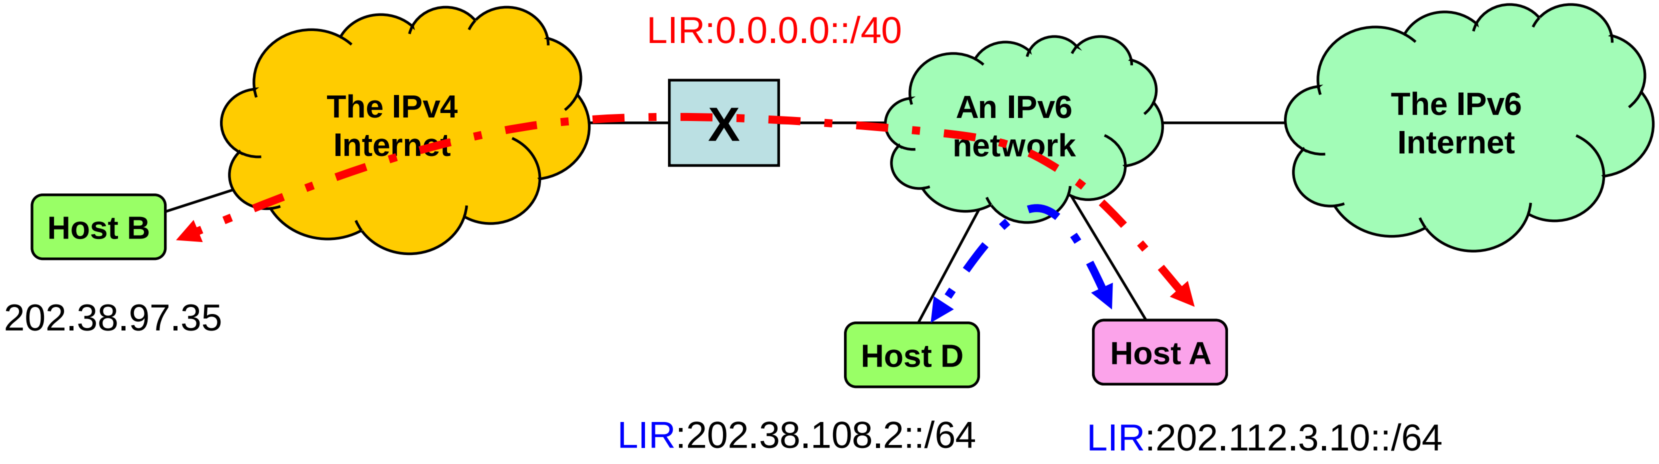

```
Host A to Host B → OK
Src=LIR:202.112.3.10::
Dst=LIR:202.38.97.35::
```

```
Host A to Host D → OK
Src=LIR:202.112.3.10::
Dst=LIR:202.38.108.2::
```

- Based on the “more specific win” routing principle, Host A automatically reaches the IPv4 literals of Host B and/or Host D, no matter it is used by IPv4 host or IPv6 host.
  - Host B’s routing entry: LIR:0.0.0.0::/40
  - Host D’s routing entry: LIR:202.38.108.0::/64
- Simple and clean solution.

4. PREFIX1=LIR and PREFIX2=WKP

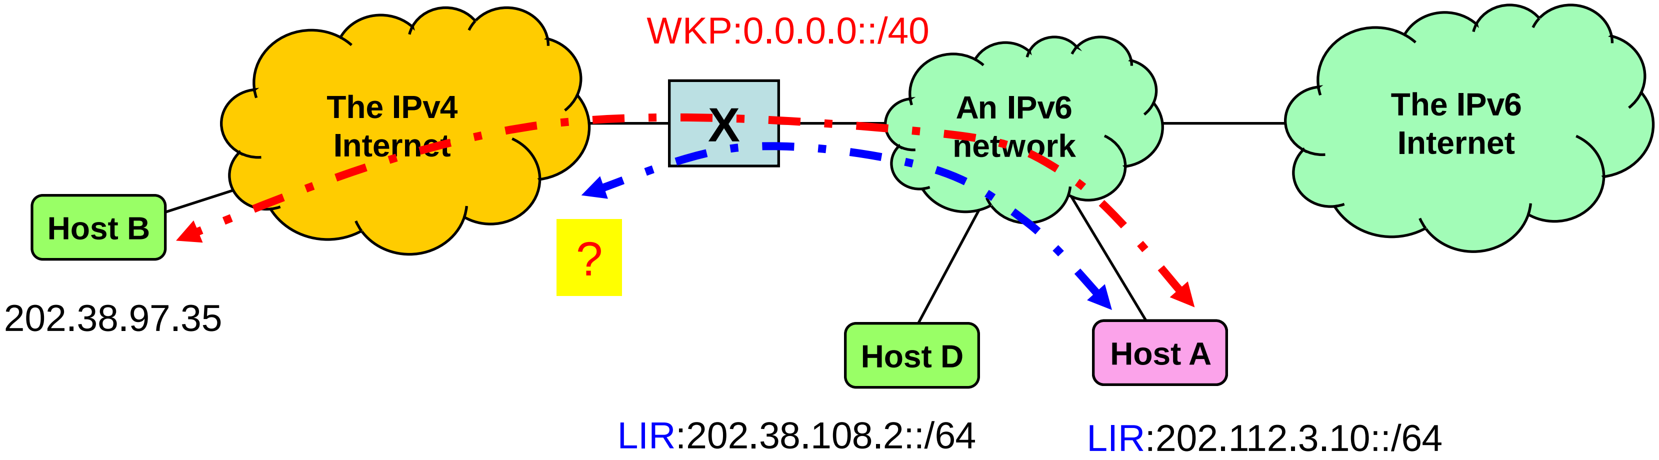

```
Host A to Host B ➔ OK
Src=LIR:202.112.3.10::
Dst=WKP:202.38.97.35::
```

```
Host A to Host D ➔ lost
Src=LIR:202.112.3.10::
Dst=WKP:202.38.108.2::
```

- Host A can communicate with Host B, but host A cannot communicate with Host D.
  - Host B’s routing entry: WKP:0.0.0.0::/40
  - Host D’s routing entry: LIR:202.38.108.2::/64
- If Host A needs to communicate with an IPv4 literals, Host A needs to know the location of the host with this IPv4 literals to create the IPv6 representation with the right IPv6 PREFIX, etc.

Work around

- Design a protocol to download the IPv4 address database to host A
  - Not simple.
- Configure the DNS64 based on the IPv4 address database
  - Not simple, cannot support Referrals if IPv4 address is literally used.
- Add the Hairpin function to the stateless translator based on the IPv4 address database
  - Non-optimal routing.
  - It is related to hairpin issue of the stateful translation, but not the same.

Design a protocol to download the IPv4 address database to host A

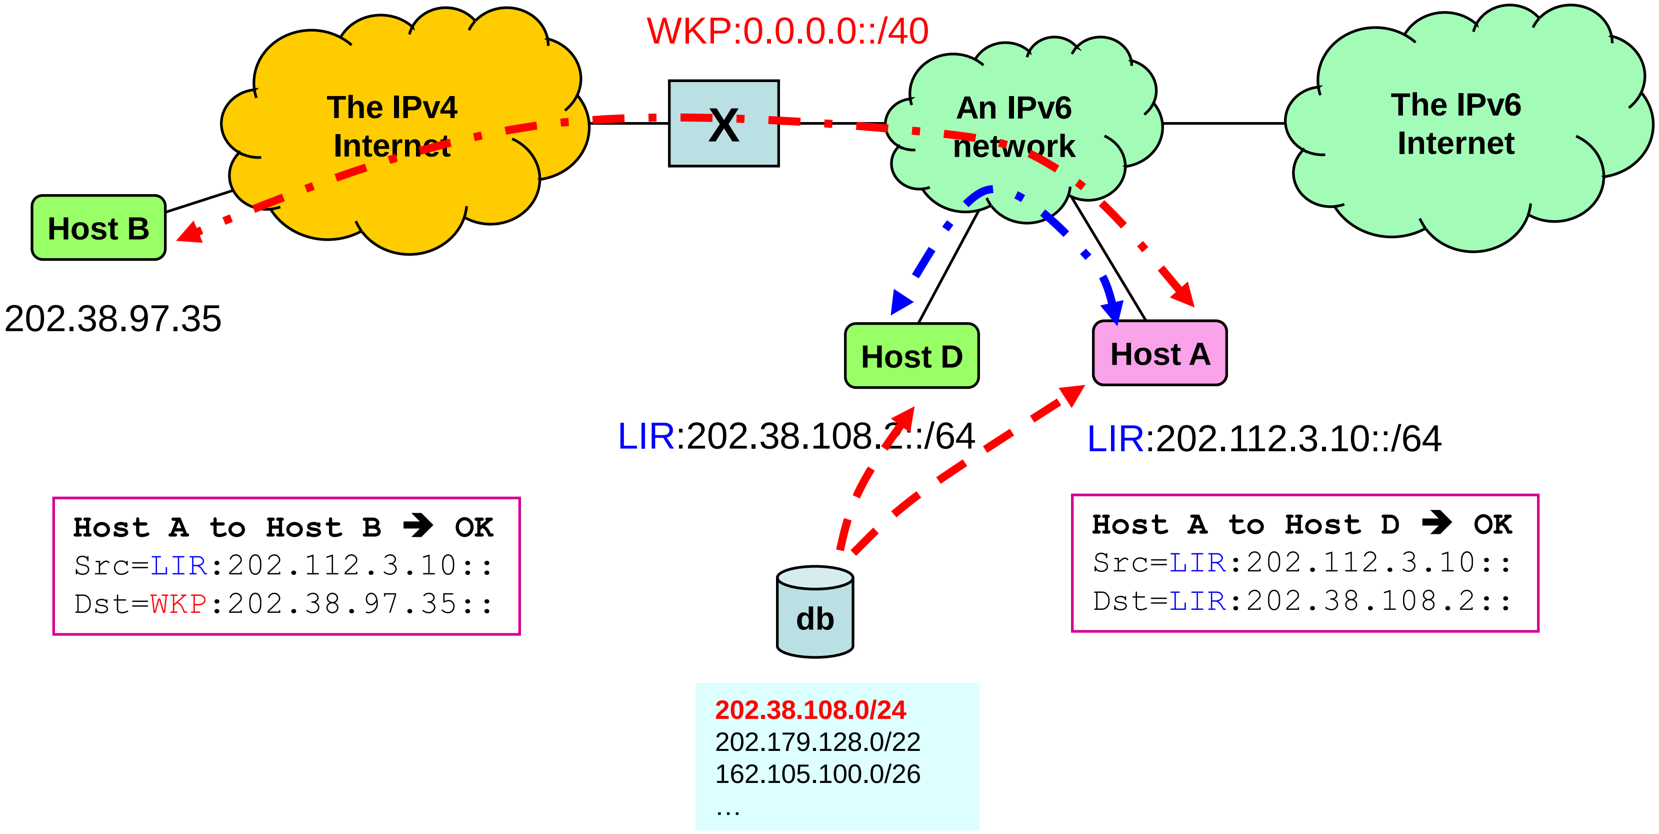

```
Host A to Host B → OK
Src=LIR:202.112.3.10::
Dst=WKP:202.38.97.35::
```

```
Host A to Host D → OK
Src=LIR:202.112.3.10::
Dst=LIR:202.38.108.2::
```

202.38.108.0/24
202.179.128.0/22
162.105.100.0/26
…

- Not a simple solution

Configure the DNS64 based on the IPv4 address database

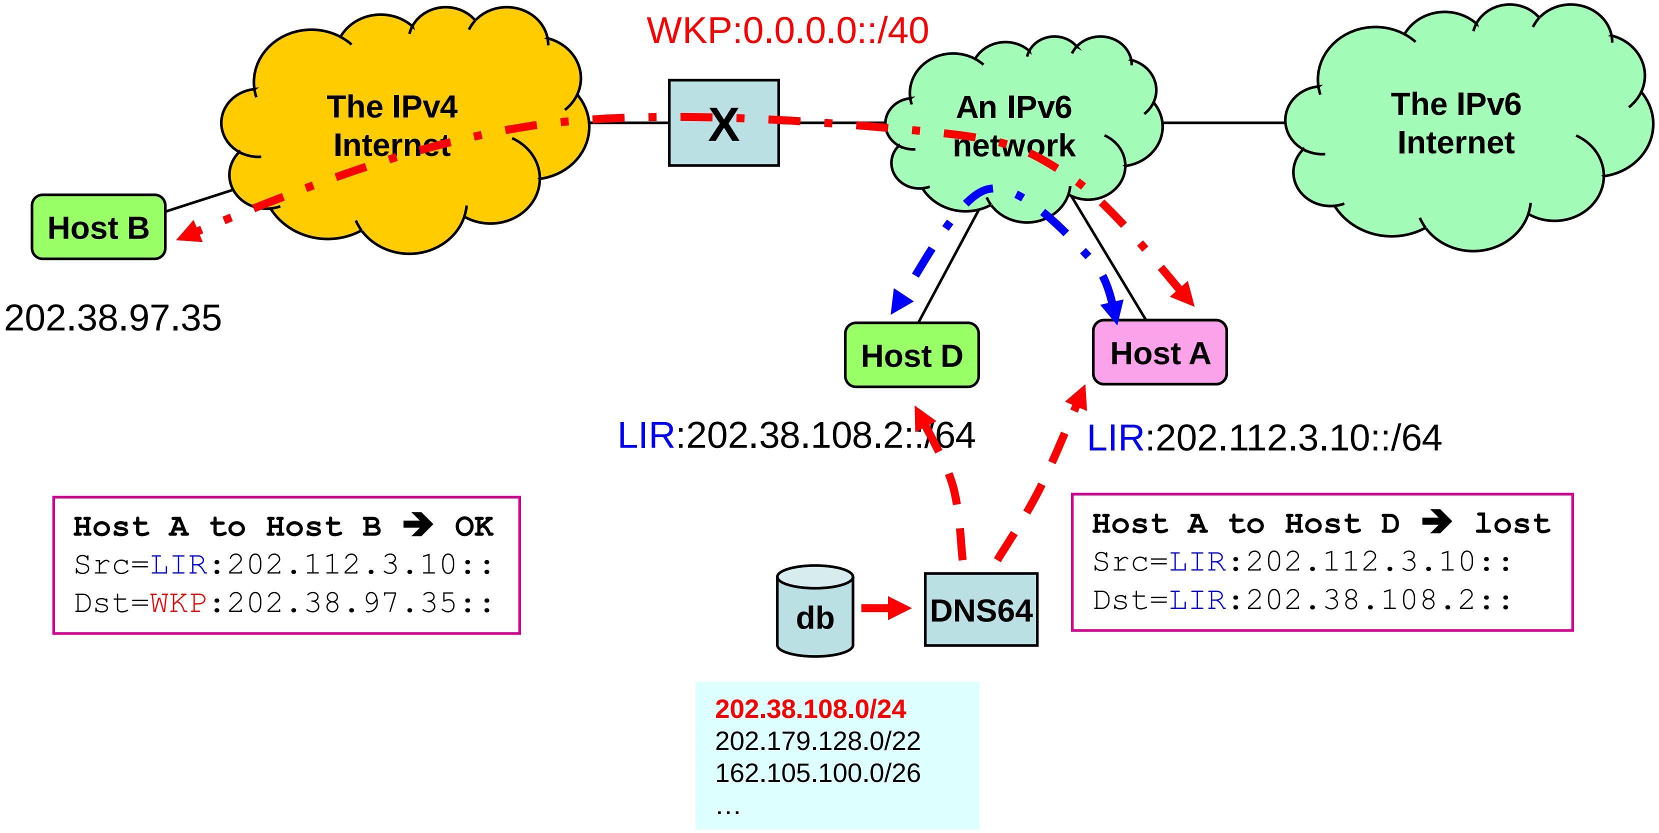

```
Host A to Host B → OK
Src=LIR:202.112.3.10::
Dst=WKP:202.38.97.35::
```

```
Host A to Host D → lost
Src=LIR:202.112.3.10::
Dst=LIR:202.38.108.2::
```

202.38.108.0/24
202.179.128.0/22
162.105.100.0/26
…

- Not a simple solution, this cannot work without using DNS as the referral

Add the hairpin function to the stateless translator based on the IPv4 address database

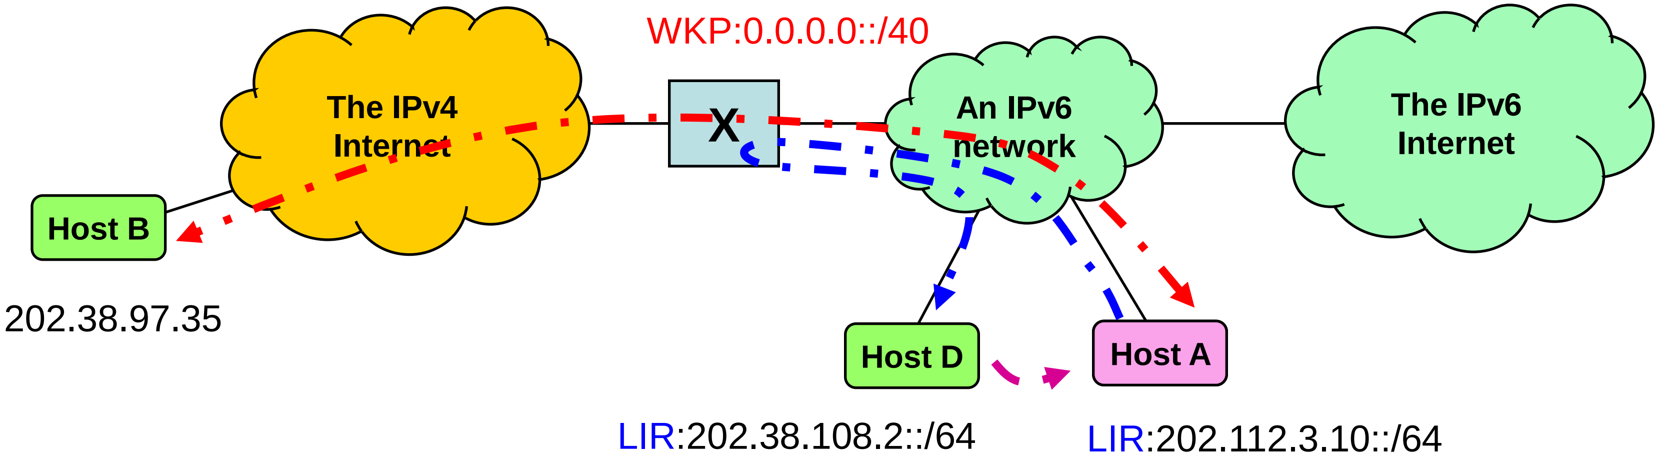

WKP:0.0.0.0::/40

The IPv4 Internet

X

An IPv6 network

The IPv6 Internet

Host B

202.38.97.35

Host D

Host A

LIR:202.38.108.2::/64

LIR:202.112.3.10::/64

```
Host A to Host B ➔ OK
Src=LIR:202.112.3.10::
Dst=WKP:202.38.97.35::
```

```
Host A to Host D ➔ OK
Src=LIR:202.112.3.10::
Dst=WKP:202.38.108.2::
```

- Non-optimal routing

Remarks

- For stateless translation (e.g. IVI) in scenarios 1 and 2.
  - For routing scalability and operation simplicity, the PREFIX1 must be LIR.
  - For operation simplicity, the PREFIX2 and PREFIX1 must be the same, and PREIFX2 must be LIR.
- For stateful translation (e.g. NAT64) in scenarios 1 and 3.
  - There is no PREFIX1, since stateful translator uses the state table to represent IPv6 hosts in IPv4.
  - There is not need to advertise PREFIX2 in the IPv6 Internet.
  - Therefore, either WKP or LIR can be used.

PREFIX Recommendation (1)

- Scenario 1: an IPv6 network to the IPv4 Internet
  - Stateless solution (IVI implementation)
    - If IPv6 host configured with only one site prefix, it MUST be LIR. If IPv6 host configured with multiple site prefixes, the site prefix with global IPv6 connectivity MUST be LIR, while the site prefix without global IPv6 connectivity CAN be WKP.
    - The translator prefix must be the same as host prefix and must be LIR.
  - Stateful solution (NAT64)
    - Either LIR or WKP CAN be used as the translator prefix.
- Scenario 2: the IPv4 Internet to an IPv6 network
  - Stateless solution (IVI implementation)
    - If IPv6 host configured with only one site prefix, it MUST be LIR. If IPv6 host configured with multiple site prefixes, the site prefix with global IPv6 connectivity MUST be LIR, while the site prefix without global IPv6 connectivity CAN be WKP.
    - The translator prefix must be the same as host prefix and must be LIR.

PREFIX Recommendation (2)

- Scenario 3: the IPv6 Internet to an IPv4 network
  - Stateful solution (NAT64)
    - LIR MUST be used.
- Scenario 4: an IPv4 network to the IPv6 Internet
  - No solution.