

IVI operational considerations and practice

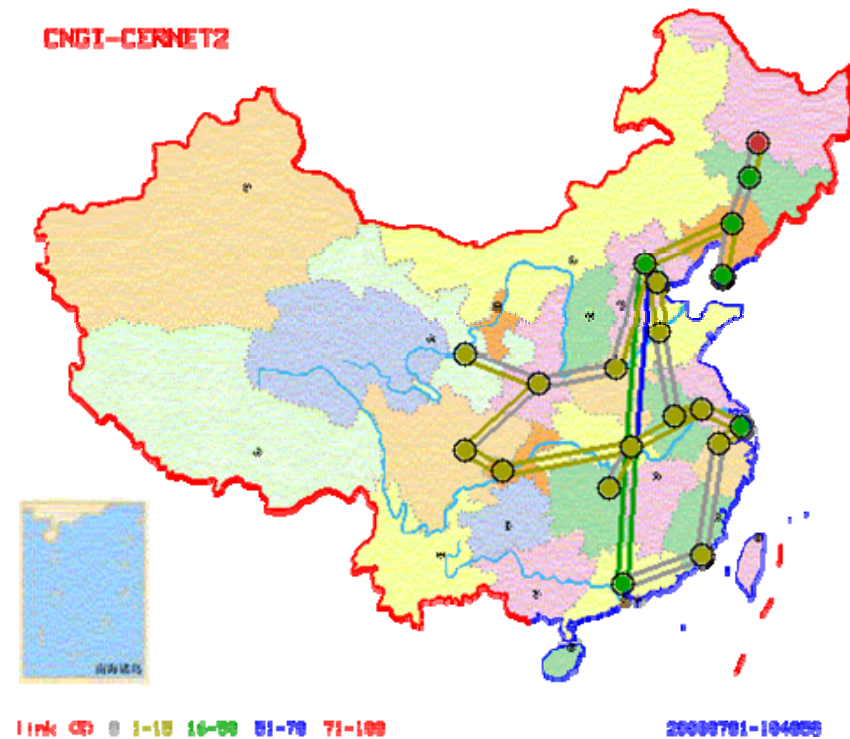
Xing Li

IETF-72, Dublin, ops area mtg, 29 July 2008

Outline

- Background
- IVI Scheme
- Operation practice
- Transition
- Remarks

CNGI-CERNET2




- CNGI-CERNET is an IPv6 single stack network.
- The original promotion concept
 - It is free and it is light loaded.
 - The users need to export the application into IPv6.
- But this concept did not work well.
 - The connectivity is the most important issue.
- So we developed IVI
 - IV means 4
 - VI means 6
 - IVI means 4|6 coexistence and transition
 - IVI is symmetric and both v6 and v4 initiated communication are supported

The Overview of the IVI Mechanism

- The IVI is a prefix-specific and explicit address mapping scheme.
 - Embed global IPv4 addresses into a subset of each ISP's IPv6 address block.
 - Based on this mapping rule, each ISP can borrow a portion of its IPv4 addresses and use it in IPv6.
- The SIIT stateless translation is implemented in the IVI gateway.
- The IPv4 address multiplexing techniques can be used.
- Ref:
 - <http://www.ietf.org/internet-drafts/draft-xli-behave-ivi-00.txt>

Address Mapping (1)



Mapping Rule: IPv4 addresses are embedded from bit 40 to bit 72 of the IPv6 addresses of a specific /32.

Example:

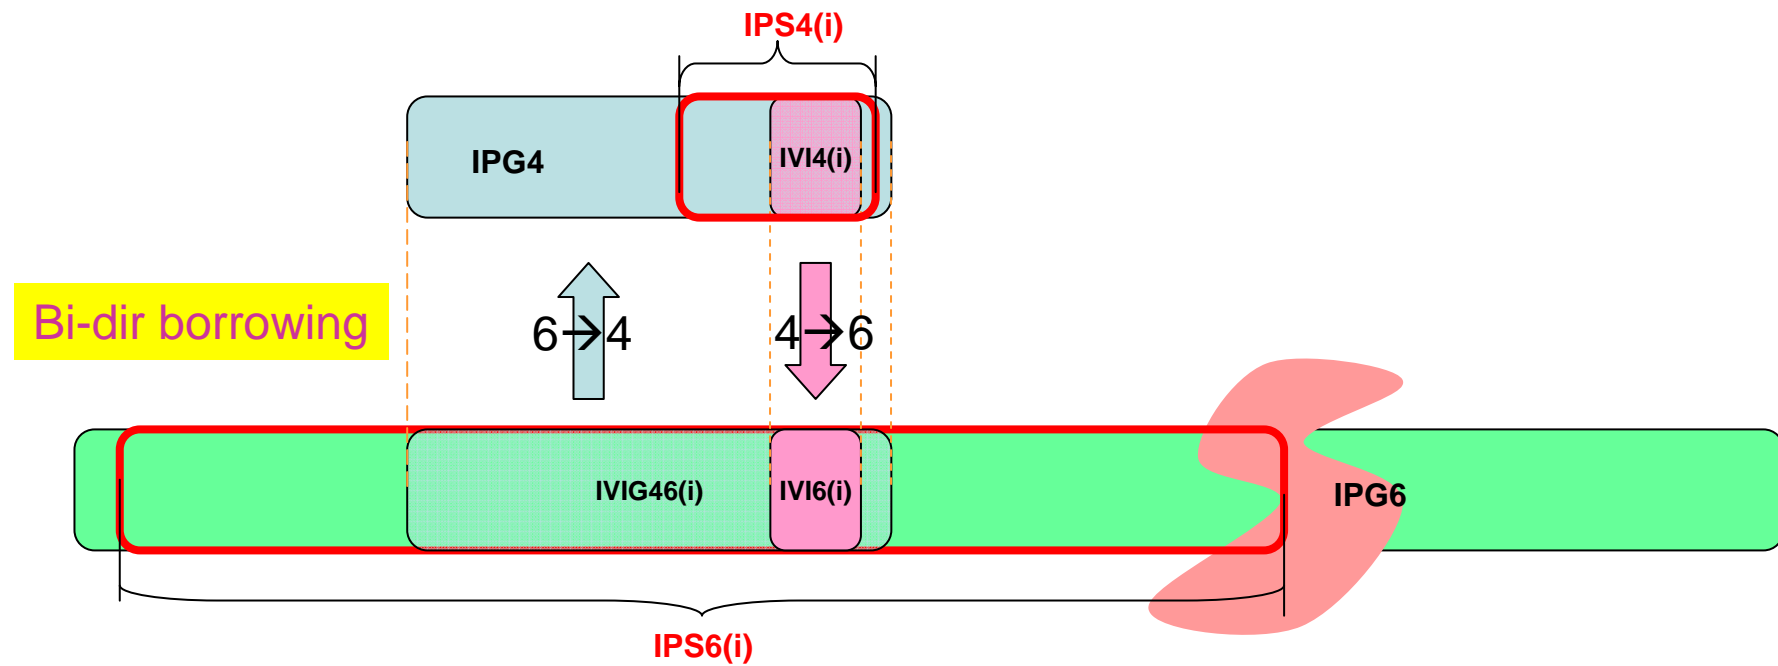
ISP's IPv6 /32	2001:250::/32
borrowed IPv4 address (IVI4):	202.38.108.0/24
mapped IVI IPv6 address (IVI6):	2001:250:ffca:266c::/64

The general IVI address mapping

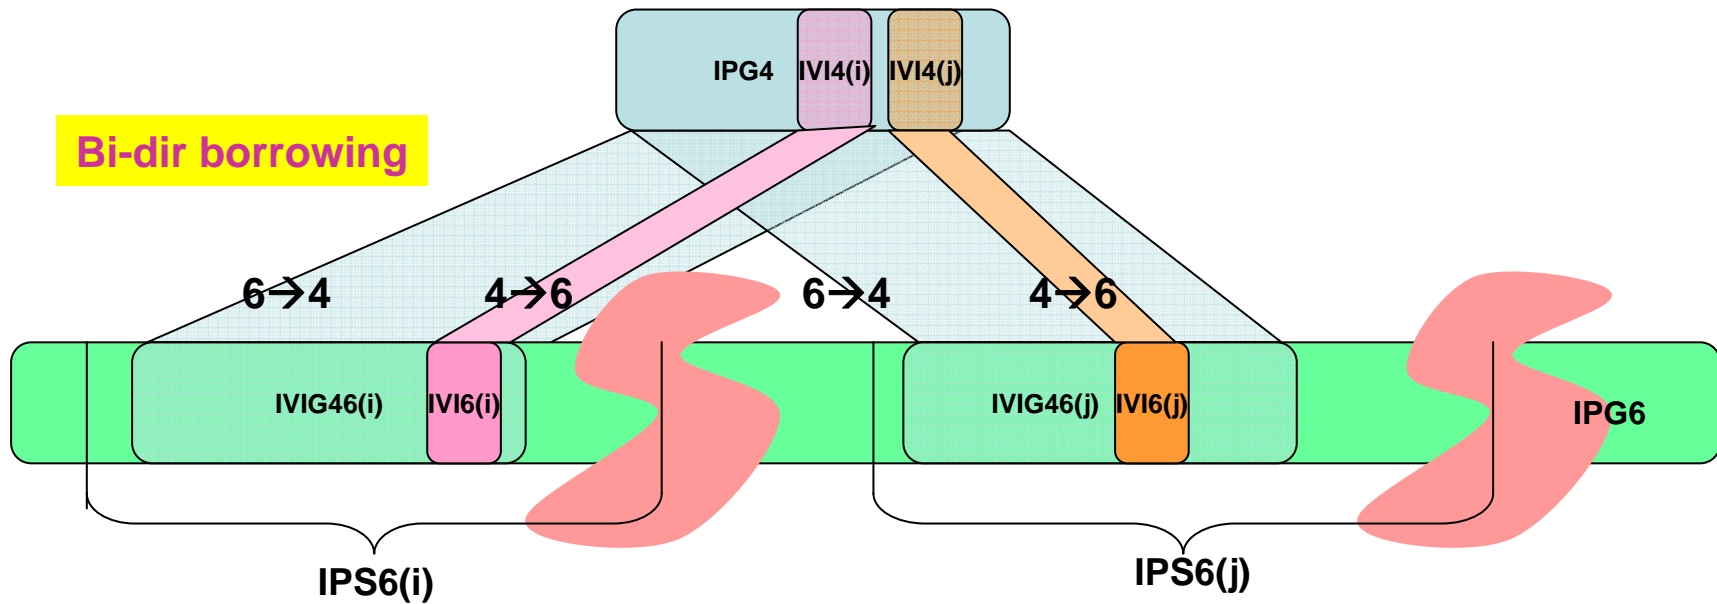
2001:DB8:FF00::/40
2001:DB8:FFFF::/48,
2001:DB8:ABCD:FF00::/56
2001:DB8:ABCD:FFFF::/64
2001:DB8:XXXX:XXXX:XXXX:XXXX::/96

backbone scope (implemented)
site scope
sub-site scope
subnet scope
IPv4 mapped alike scope

Address mapping (2)

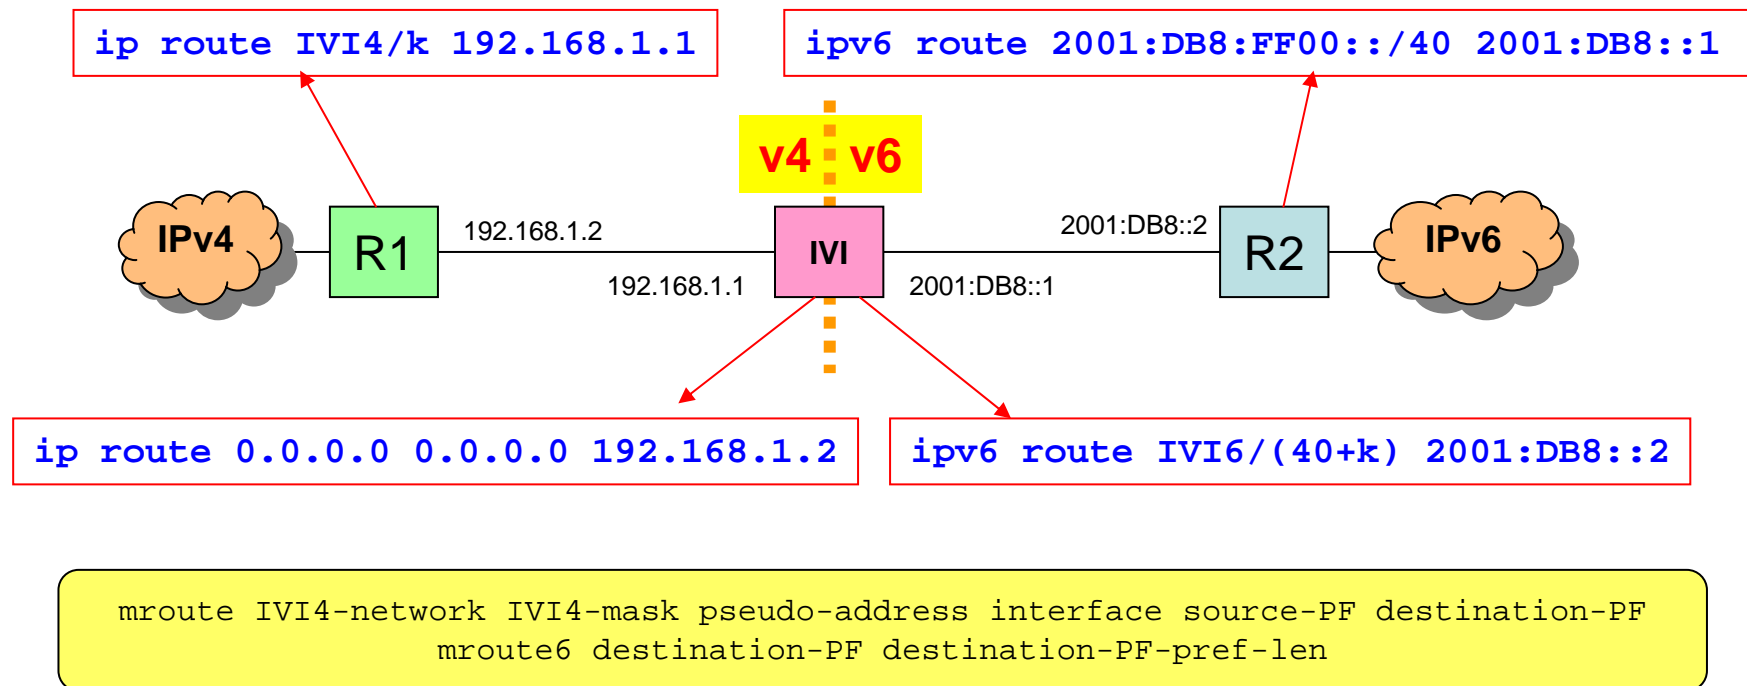


Address mapping (3)




Routing and Forwarding

Routing and mapping configuration example

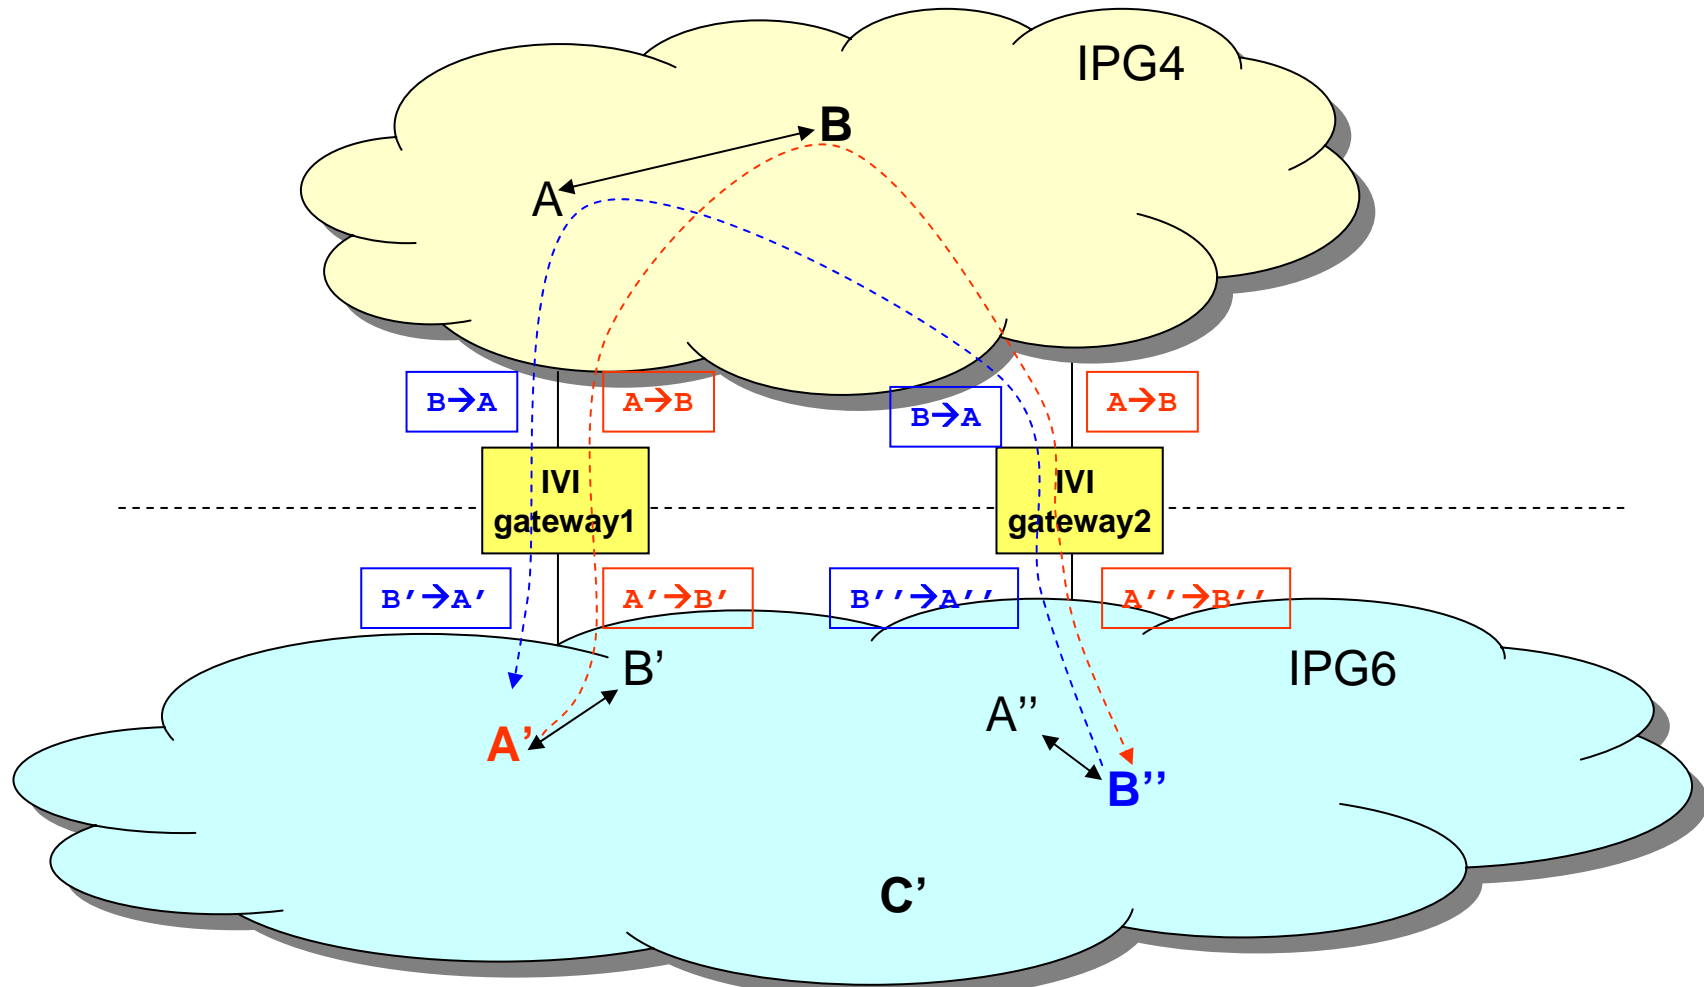


IVI Communication Scenarios (1)



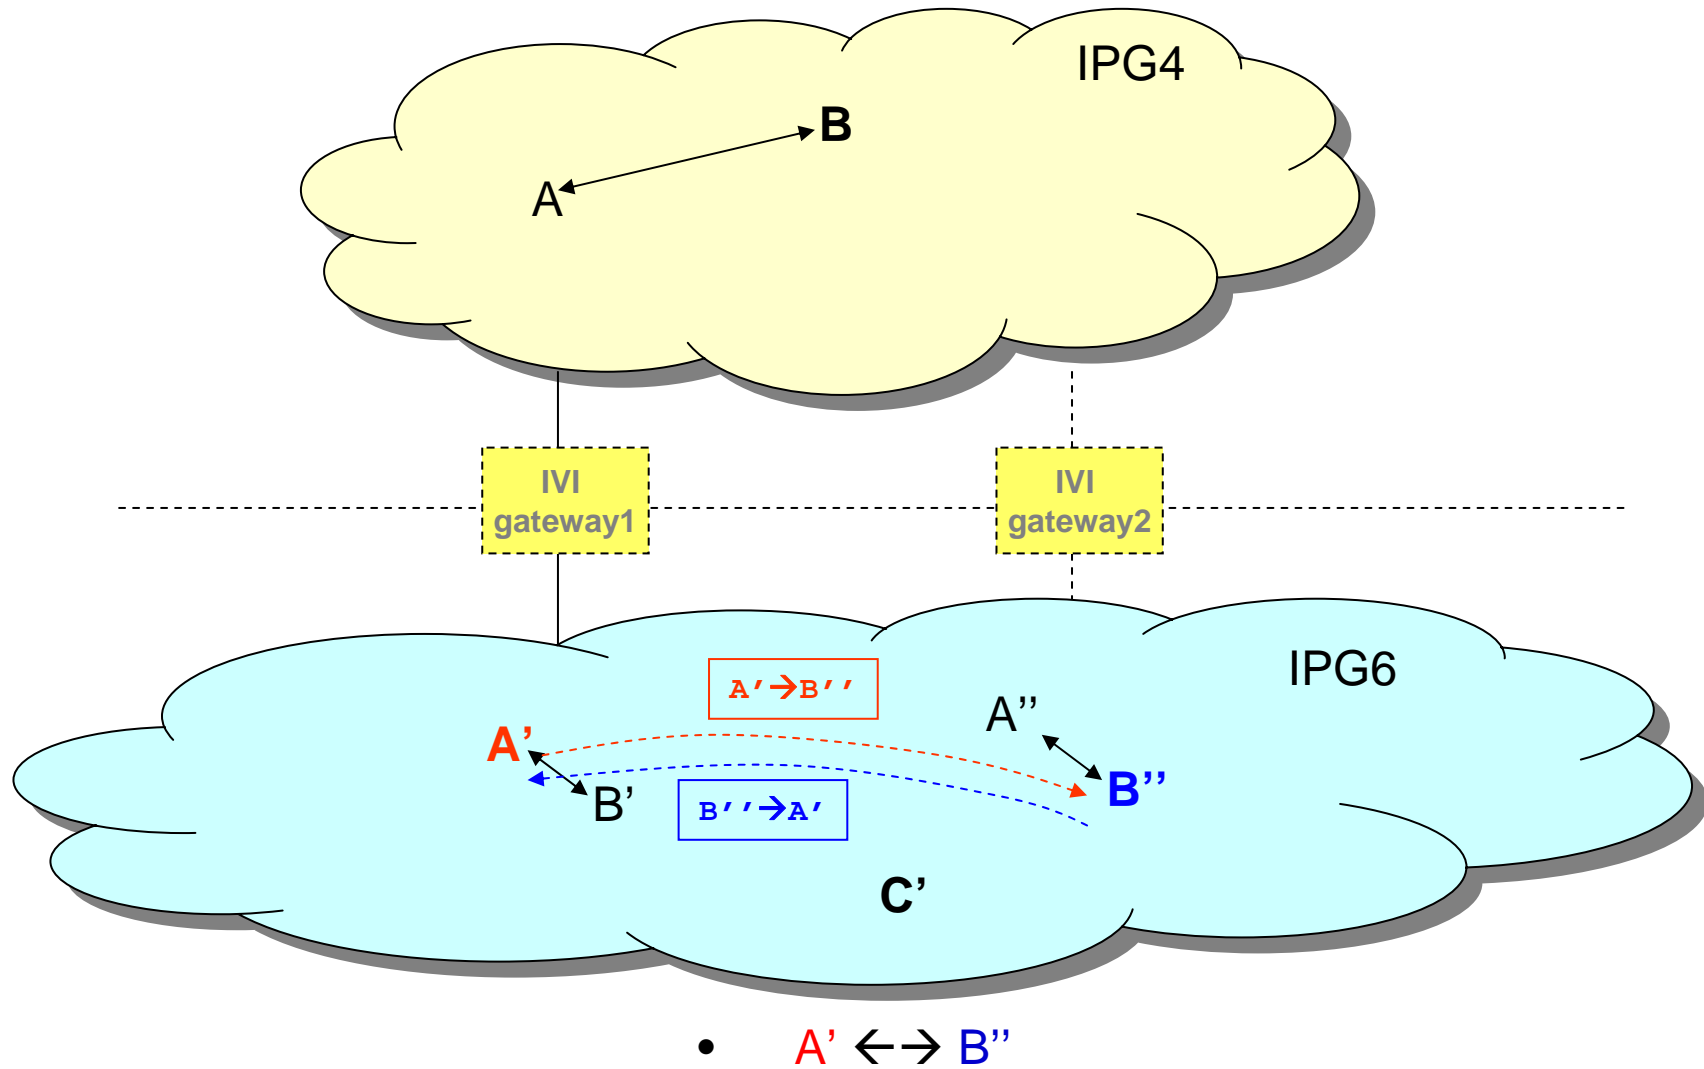
- $A' \leftrightarrow B$
- $A' \leftrightarrow C'$

IVI Communication Scenarios (2)

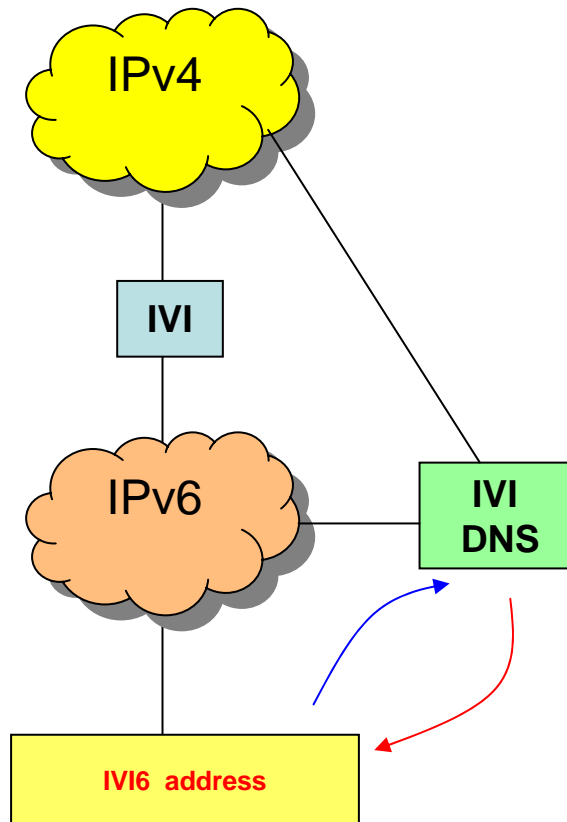


- $A' \leftarrow B \& A \rightarrow B''$

IVI Communication Scenarios (3)



DNS Configuration and Mapping




- For providing primary DNS service for IVI4(i) and IVI6(i), each host will have both A and AAAA records
- Authoritative DNS server
 - Example
 - www.ivi2.org A 202.38.108.2
 - www.ivi2.org AAAA 2001:250:ffca:266c:200::
- For resolving IVI46(i) for IVI6(i), use IVI DNS to do the dynamic mapping based on the IVI rule.
- Caching DNS server
 - Example
 - www.mit.edu A 18.7.22.83
 - www.mit.edu AAAA 2001:250:ff12:0716:5300::


Multiplexing of the Global IPv4 Addresses

- Temporal Multiplexing
 - Dynamic assignment of IVI6(i)
- Port Multiplexing
 - Combine address with the port number
- Spatial Multiplexing
 - Server 1:1 mapping
 - Client 1:N mapping
- Multiplexing using IPv4 NAT-PT
 - Cascade IPv4 NAT-PT and IVI (1:1 mapping)

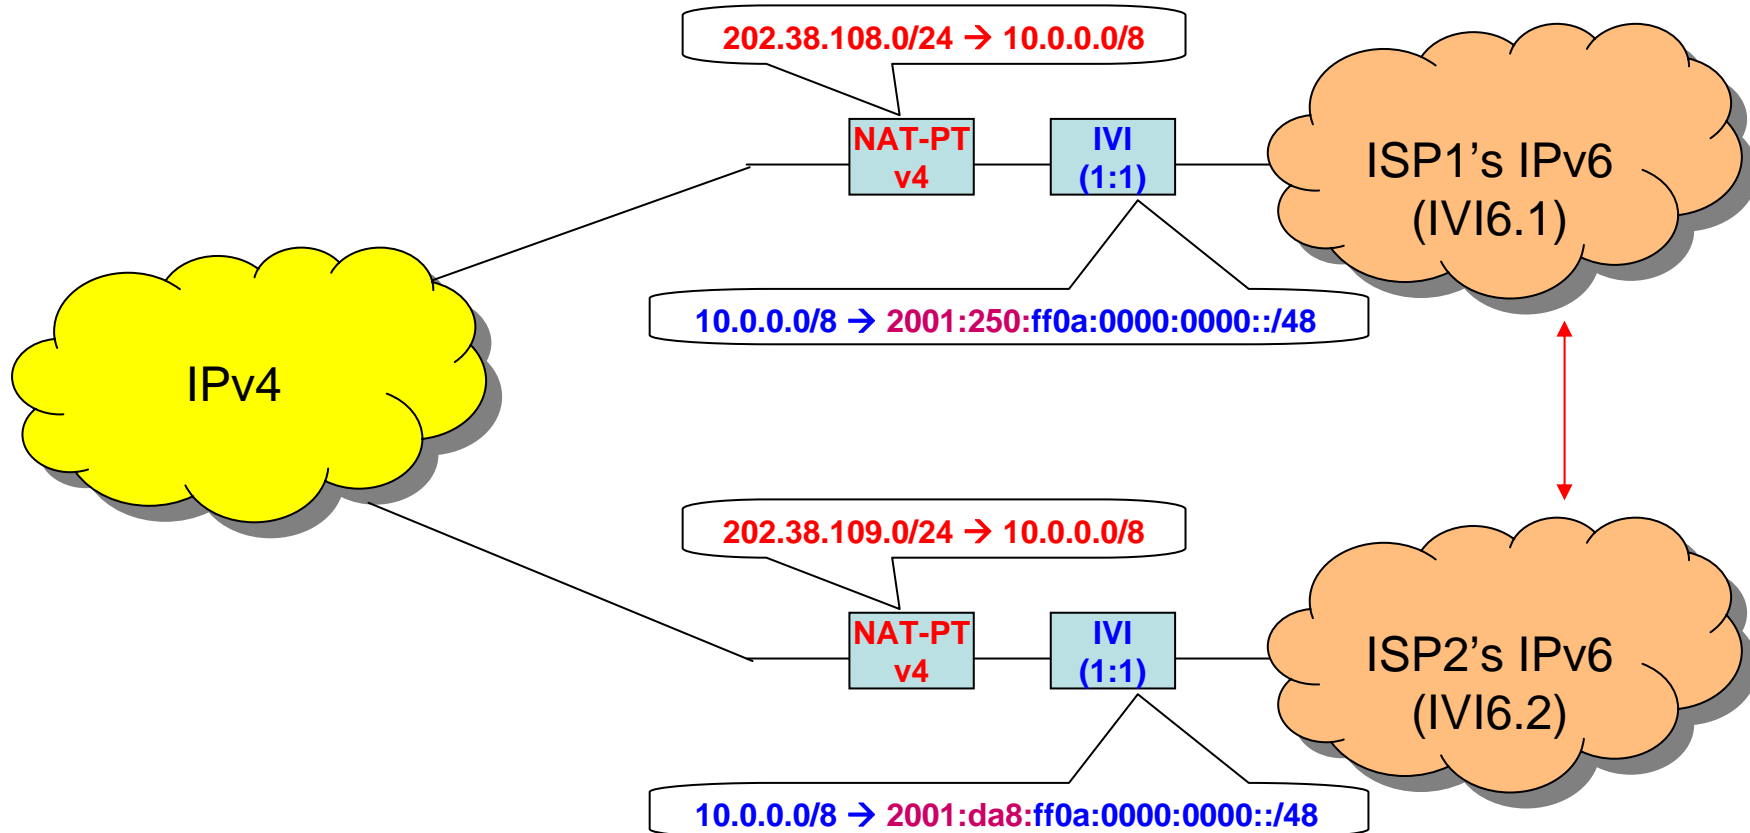
IVI Deployment Scenarios (1)



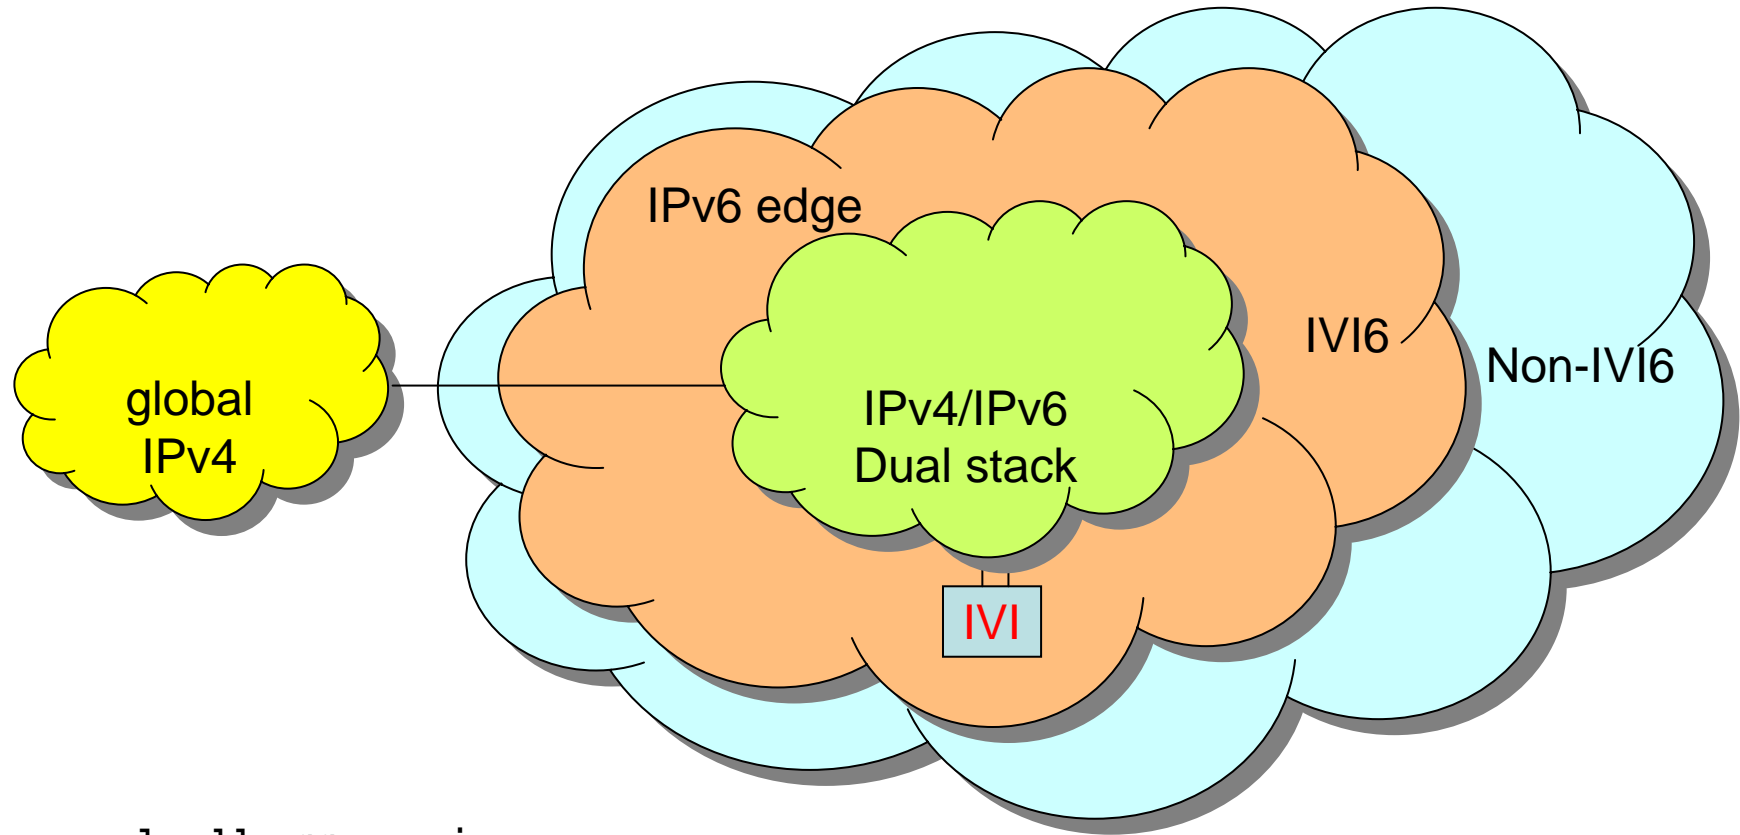
IVI Deployment Scenarios (2)



IVI Deployment Scenarios (3)



IVI Deployment Scenarios (4)



IVI general address mapping

2001:DB8:FF00::/40

2001:DB8:FFFF::/48,

2001:DB8:ABCD:FF00::/56

2001:DB8:ABCD:FFFF::/64

2001:DB8:XXXX:XXXX:XXXX:XXXX::/96

backbone scope (implemented)

site scope

sub-site scope

subnet scope

IPv4 mapped alike scope

General IVI address mapping

Prefix len of IVIG46	Prefix len of IVI4	Prefix len of IVI6
/40	/24	/64
	/32	/72
/48	/24	/72
	/32	/80
/56	/24	/80
	/32	/88
/64	/24	/88
	/32	/96
/96	/24	/120
	/32	/128

Multiple IVI can be configured

IVIG64

2001:DB8:ABCD:FFF0::/64

2001:DB8:ABCD:FFF1::/64

2001:DB8:ABCD:FFFF::/64

IVI4

202.38.108.0/32

202.38.108.1/32

202.38.108.15/32

IVI6

2001:DB8:ABCD:FFFF:ca26:6000::/96


2001:DB8:ABCD:FFFF:ca26:6001::/96

2001:DB8:ABCD:FFFF:ca26:600f::/96

Some URLs

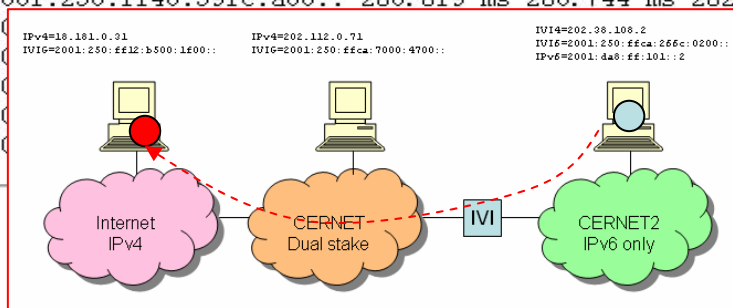
- The IVI (1:1) has been running at CNGI-CERNET2 for more than two years.
- IVI6 server for global IPv4
 - <http://202.38.114.1/>
- IVI6 server for global IPv6
 - [http://\[2001:250:ffca:2672:0100::0\]/](http://[2001:250:ffca:2672:0100::0]/)
- IVI server for stub IPv4
 - <http://202.38.114.129/>
- IVI open source for Linux (1:1)
 - <http://202.38.114.1/impl/>

From IVI6 host traceroute6 IIVI646




Host N in in pure IPv6 environment

```
traceroute to 2001:250:ff12:b500:1f00:: (2001:250:ff12:b500:1f00::), 30 hops max, 40 byte packets to not_ivi
1 2001:250:ffca:266c:100:: 0.902 ms 0.884 ms 0.849 ms 2001:250:ffca:266c:100:: 202.38.108.1
2 2001:250:c000:63::1 1.210 ms 1.302 ms 1.378 ms 2001:250:c000:63::1 not_ivi
3 2001:250:c000:20::1 1.668 ms 1.766 ms 1.917 ms 2001:250:c000:20::1 not_ivi
4 2001:250:c000:2::2 2.915 ms 3.042 ms 3.095 ms 2001:250:c000:2::2 not_ivi
5 2001:250:ff0a:0:100:: 4.302 ms 4.283 ms 4.284 ms 2001:250:ff0a:0:100:: 10.0.0.1
6 2001:250:ffca:703d:4100:: 6.878 ms 7.676 ms 7.658 ms 2001:250:ffca:703d:4100:: 202.112.61.65
7 * 2001:250:ffca:703d:f100:: 5.879 ms * * 202.112.61.65
8 2001:250:ffca:7035:b100:: 11.638 ms 11.434 ms 11.356 ms 2001:250:ffca:7035:b100:: 202.112.53.177
9 2001:250:ffca:703d:9e00:: 5.074 ms 5.532 ms 5.399 ms 2001:250:ffca:703d:9e00:: 202.112.61.158
10 2001:250:ffca:7035:1200:: 5.325 ms 4.358 ms 5.162 ms 2001:250:ffca:7035:1200:: 202.112.53.18
11 2001:250:ffcb:b5c2:7d00:: 92.976 ms 91.484 ms 91.458 ms 2001:250:ffcb:b5c2:7d00:: 203.181.194.125
12 2001:250:ffc0:cb74:9100:: 209.784 ms 208.310 ms 224.348 ms 2001:250:ffc0:cb74:9100:: 192.203.116.145
13 2001:250:ffcf:e7f0:8300:: 206.548 ms 206.539 ms 206.649 ms 2001:250:ffcf:e7f0:8300:: 207.231.240.131
14 2001:250:ff40:391c:2d00:: 240.147 ms 239.321 ms 238.206 ms 2001:250:ff40:391c:2d00:: 64.57.28.45
15 2001:250:ff40:391c:2a00:: 263.962 ms 263.894 ms 261.707 ms 2001:250:ff40:391c:2a00:: 64.57.28.42
16 2001:250:ff40:391c:700:: 276.193 ms 276.179 ms 275.508 ms 2001:250:ff40:391c:700:: 64.57.28.7
17 2001:250:ff40:391c:a00:: 280.819 ms 280.744 ms 282.437 ms 2001:250:ff40:391c:a00:: 64.57.28.10
18 2001:250:ffc0:559:dd00:: 280.070 ms 2001:250:ffc0:559:dd00:: 192.5.89.221
19 2001:250:ffc0:559:ed00:: 280.742 ms 2001:250:ffc0:559:ed00:: 192.5.89.237
20 2001:250:ffcf:d28f:6e00:: 280.081 ms 2001:250:ffcf:d28f:6e00:: 207.210.143.110
21 2001:250:ff12:a800:1900:: 280.6591 ms 2001:250:ff12:a800:1900:: 18.168.0.25
22 2001:250:ff12:b500:1f00:: 280.5856 ms 2001:250:ff12:b500:1f00:: 18.181.0.31
```



完成



From IPv4 host traceroute IPv4

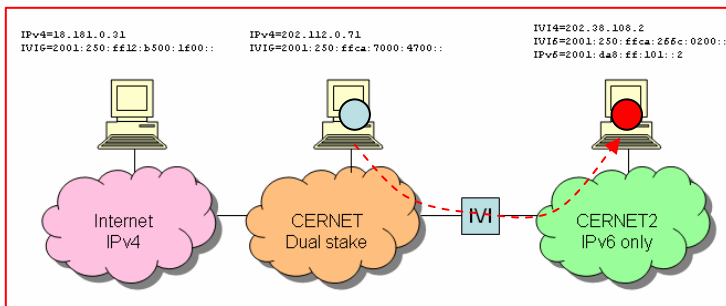
Mozilla Firefox

文件(F) 编辑(E) 查看(V) 历史(S) 书签(B) 工具(T) 帮助(H)

http://202.112.0.71:8033/cgi-bin/nph-ivi4

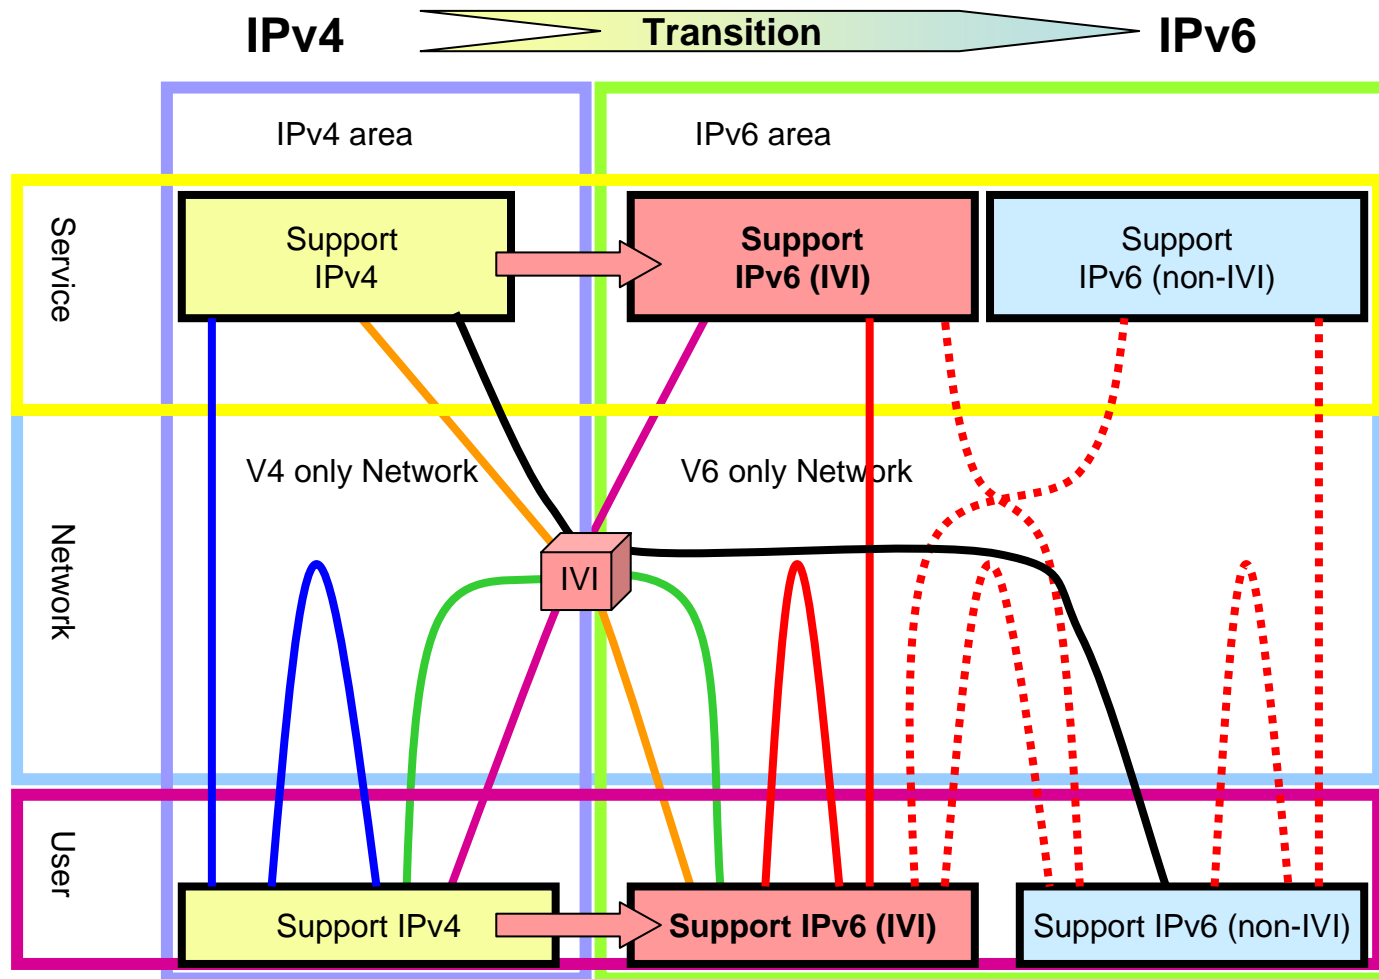
Host C in the IPv4 environment

```
1 * 202.112.0.65 1 ms *      0.0.0.0      0.0      AS0
2 * 202.112.53.73 1 ms 1 ms 0.0.0.0      0.0      AS0
3 202.112.53.178 0 ms 0 ms 0 ms 202.112.53.178 202.112 AS13746
4 202.112.61.242 9 ms 1 ms 1 ms 202.112.61.242 202.112 AS15858
5 202.38.17.186 1 ms 1 ms 1 ms 202.38.17.186 202.38 AS4538
6 202.38.17.186 1 ms 1 ms 1 ms 202.38.17.186 202.38 AS4538
7 202.38.17.186 2 ms 2 ms 2 ms 202.38.17.186 202.38 AS4538
8 202.38.17.186 2 ms 2 ms 2 ms 202.38.17.186 202.38 AS4538
9 202.38.17.186 2 ms 2 ms 2 ms 202.38.17.186 202.38 AS4538
10 202.38.108.2 2 ms 3 ms 4 ms 202.38.108.2 202.38 AS27650
```



完成

Evolution of the IVI Addresses and Services



Remarks

- The IVI is a prefix-specific and explicit address mapping scheme.
- Both IPv6 initiated and IPv4 initiated communications can be supported.
- No affect to both IPv4 and IPv6 routing. It is scalable and reliable.
- The deployment can be done incrementally and independently.
- Depending on the mapping rule, the gateway can be in any part inside the ISP's network.
- The IVI comes the closest to the end-to-end address transparency model.
- The IVI scheme encourages the transition.