


IVI - IPv4/IPv6 Coexistence and Transition

Xing Li


2009-03-04

Unexpected situation

What Should Have Happened



What Is Happening?



Why Is This Happening?

No transition plan

Declared victory before the hard part started

No real long term plan

No realistic estimation of costs

No support for the folk on the front lines

Victory will be next month

Why Is This Happening?

No transition plan

Declared victory before the hard part started

No real long term plan

No realistic estimation of costs

No support for the folk on the front lines

Victory will be next month

This Describes:

a - The invasion of Iraq

b - IPv6

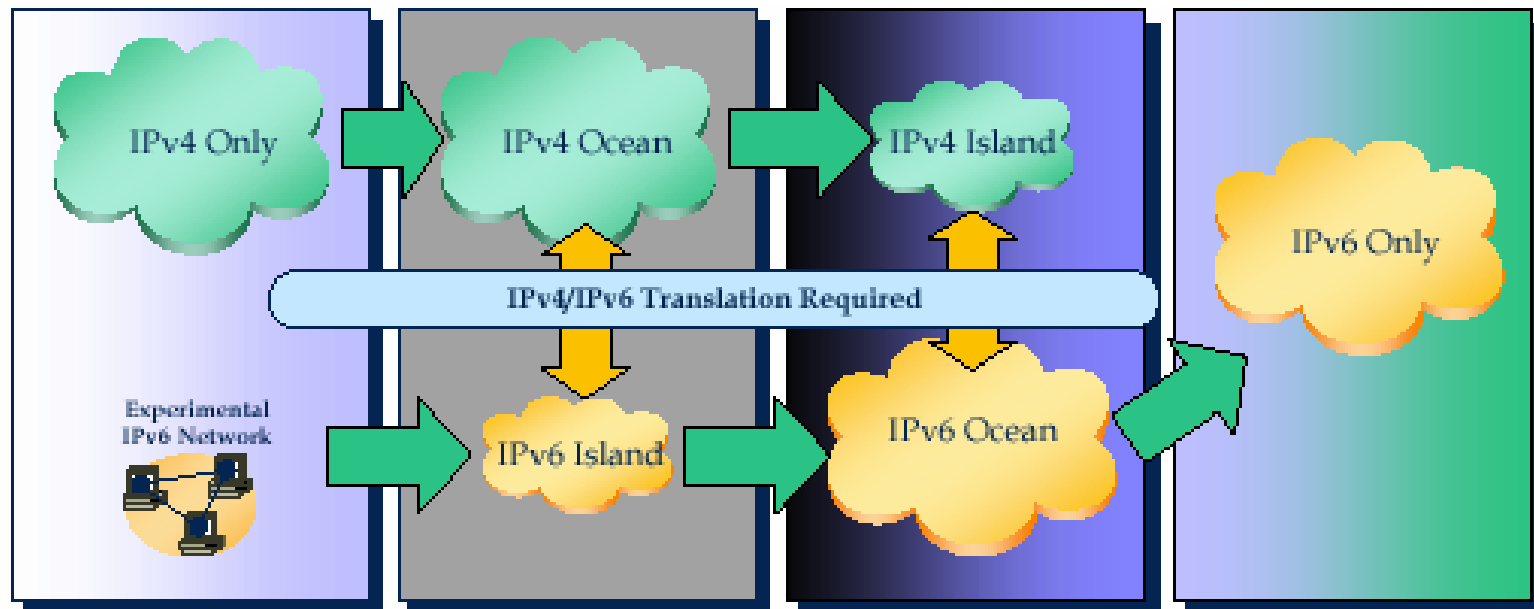
c - DNSSec

d - All of the above

Transition

- IPv4 and IPv6 is NOT compatible
- There is NO flag day
- Transition is REALLY difficult

IPv6 Road Map



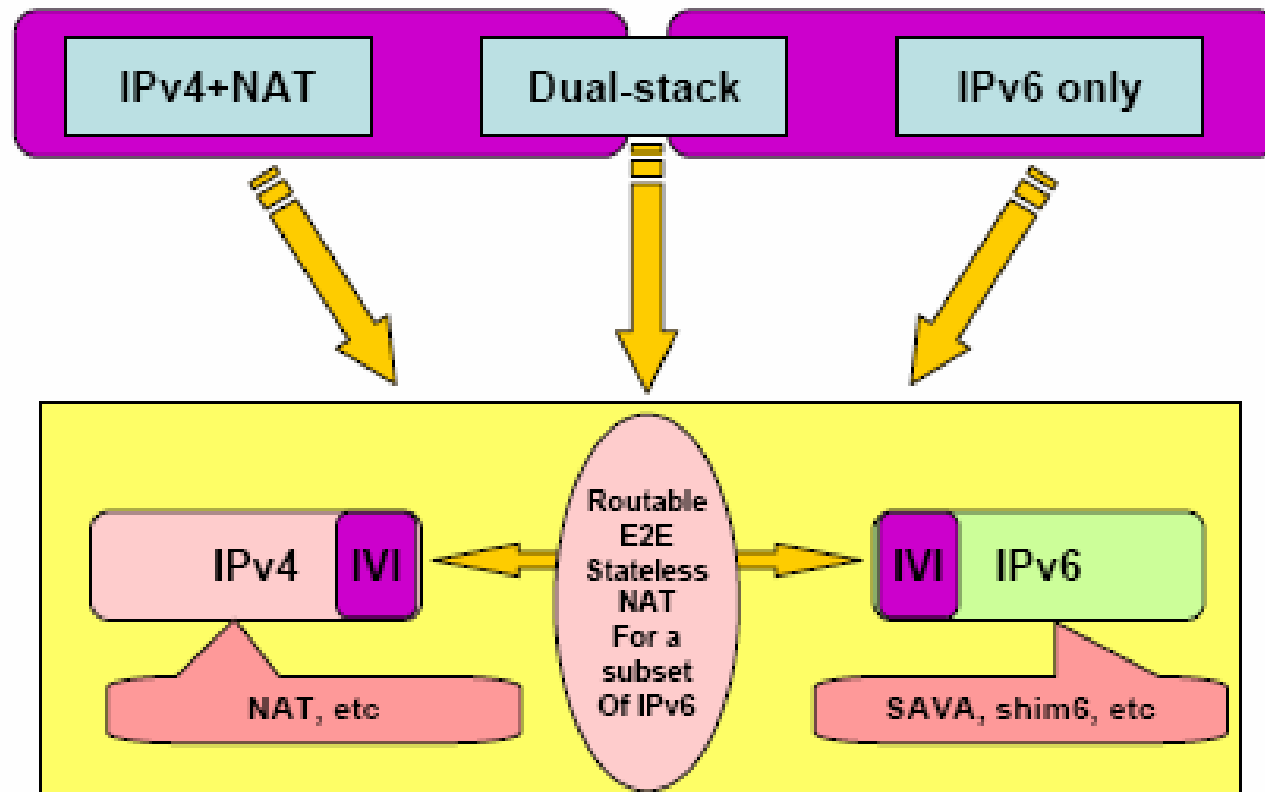
Transition methods

- Dual-stack
- Tunneling
- Translation

The IP infrastructure at crossroad




Crossroad



IPv6 Workshop, Sigcomm 2007

Background



CERNET (IPv4)

2,000 universities
connected
20M users

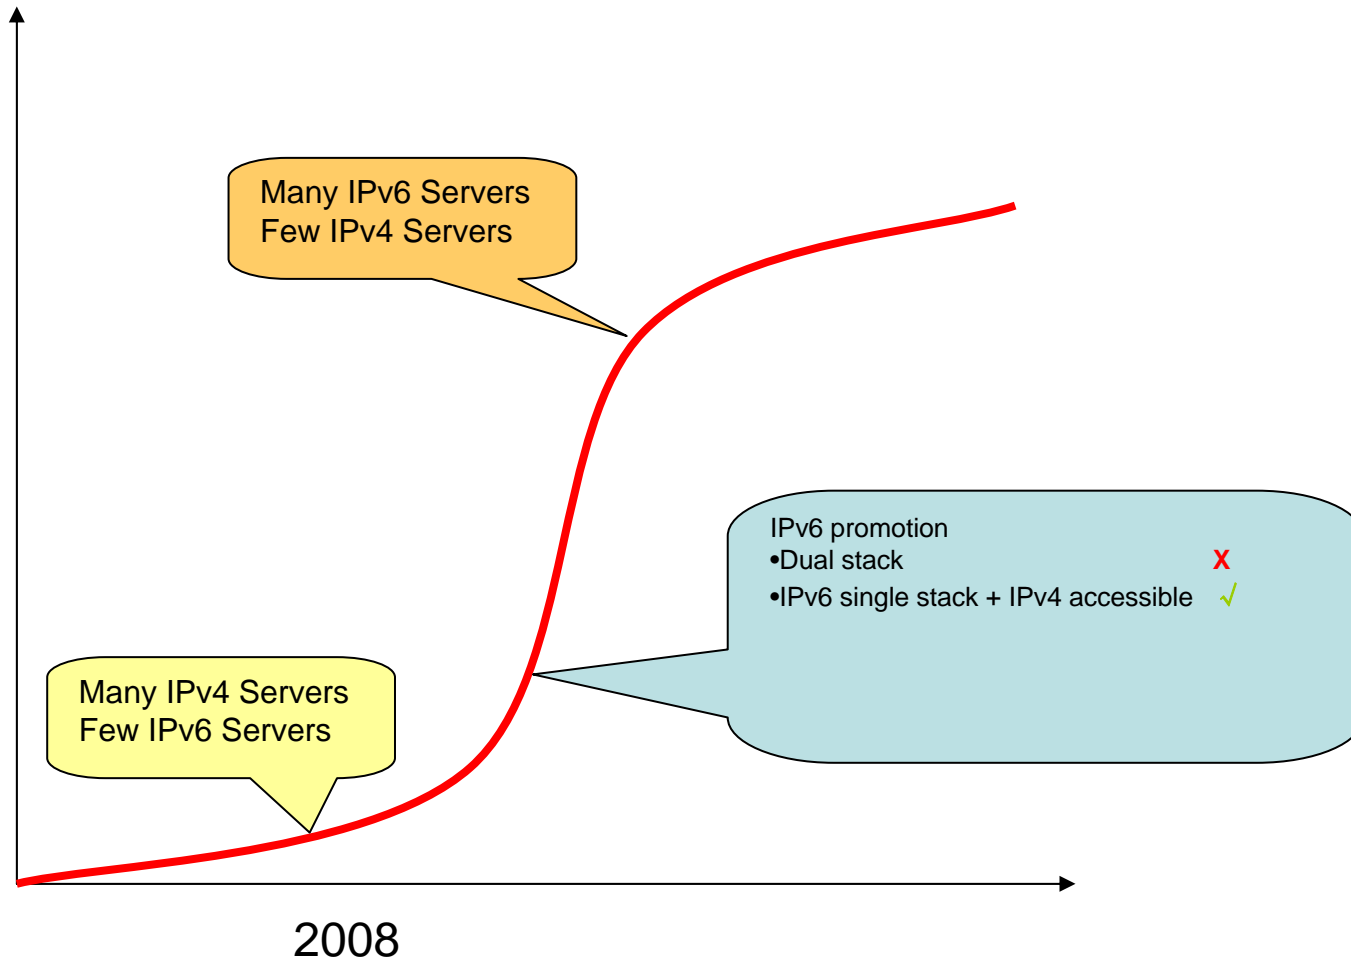
CNGI-CERNET2 (IPv6)

100 universities
connected
400K users


The lessons learned

- The only viable option for future Internet is IPv6
 - The transitions can only start when the part of it is pure IPv6
- The scenarios of building new IPv6 network for the unwired population
 - The cost-effective way for building a new infrastructure
- The natural transition
 - Construction and operation single stack costs less than dual-stack
 - Construction and operation simple (stateless) network costs less than complex (stateful) network
- The resources should be shared via inter-communication
 - The IPv6 servers should be IPv4 accessible
 - The IPv4 servers should be IPv6 accessible


IPv6 S Curve



The IVI model



IVI address format

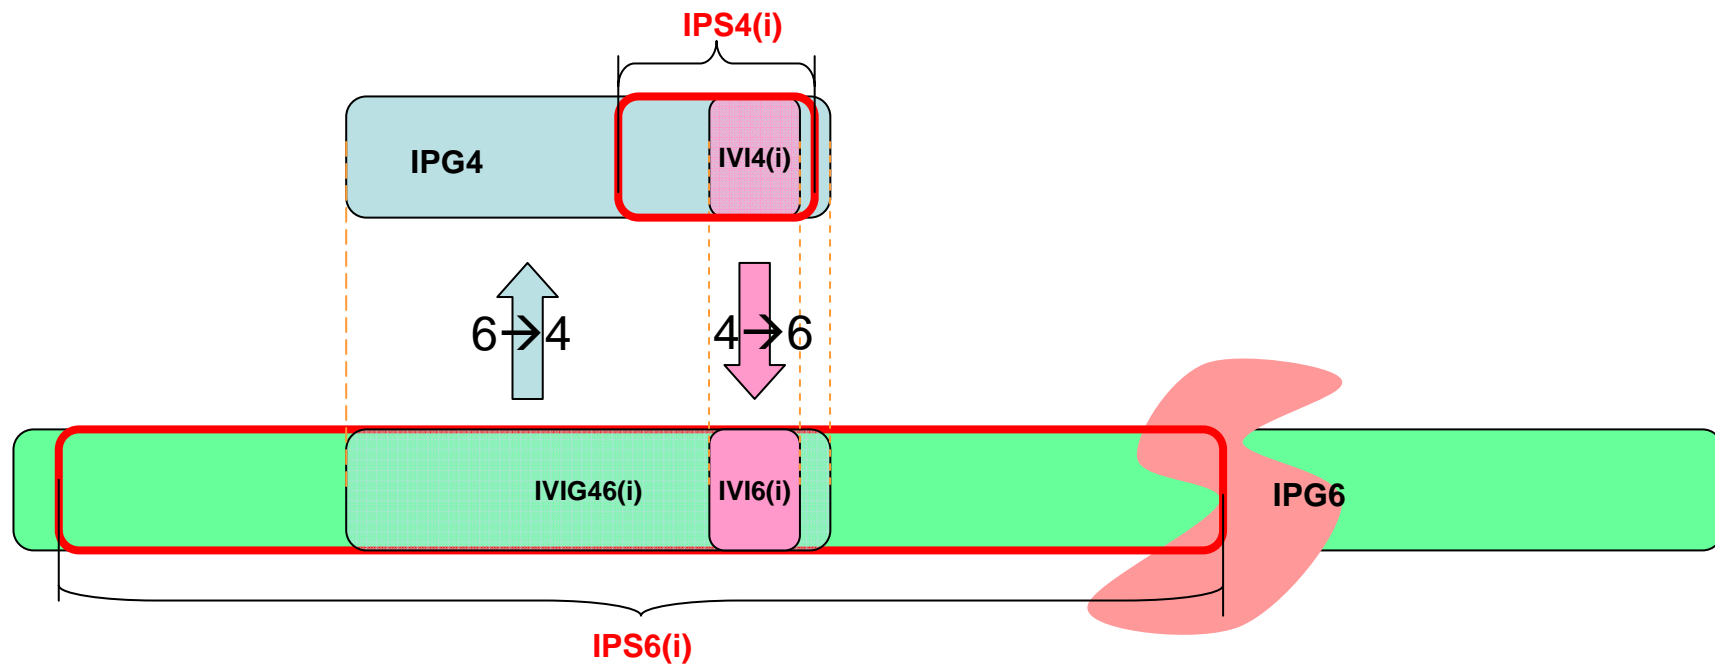


For example

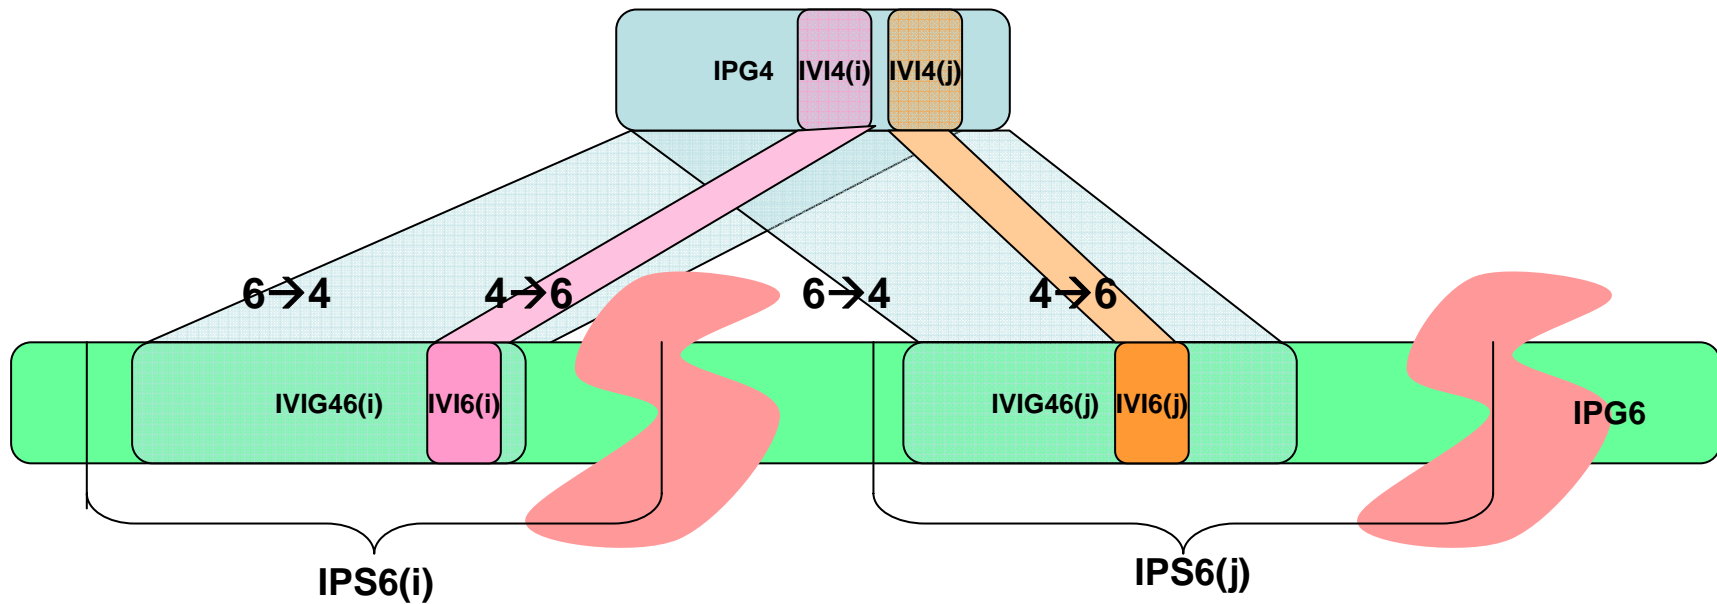
LIR consists of ISP prefix (usually /32) and IVI flag
CERNET/CNGI-CERNET2's selection

- LIR = 2001:da8:ff00::/40
- **ISP's IVI service IPv4 address mapping**
 - 202.38.108.0/24 → 2001:250:ffca:266c:0000::/64
- ISP's non-IVI service IPv4 address mapping
 - 202.38.96.0/20 → 2001:250:ffca:2660:0000::/60
- Other ISP's IPv4 address mapping
 - 0.0.0.0 → 2001:da8:ff00::/40
 - 18.181.0.31/32 → 2001:250:ff12:b500:1f00::/72


Address Mapping (1)



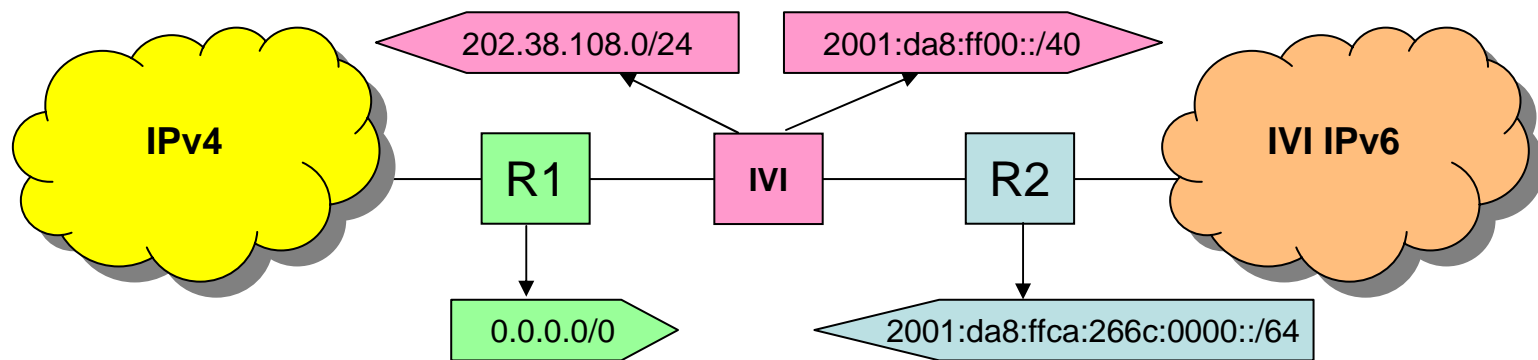
Address Mapping (2)



Conceptual example



IVI routing



Longest prefix match

IVI DNS service

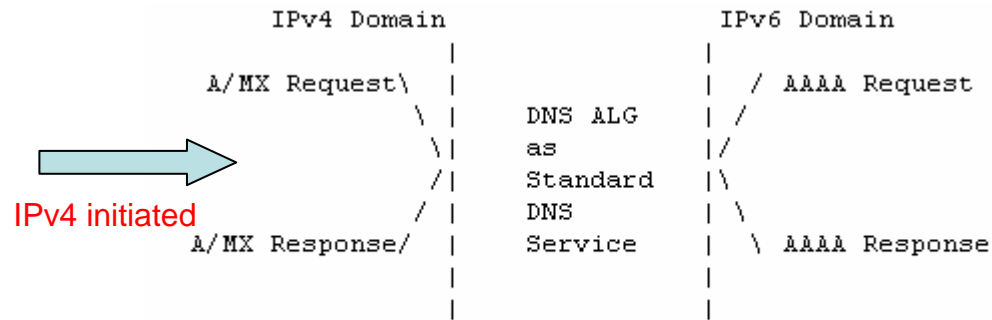


Figure 4: Normal DNS Service

- Normal DNS
 - Algorithm based

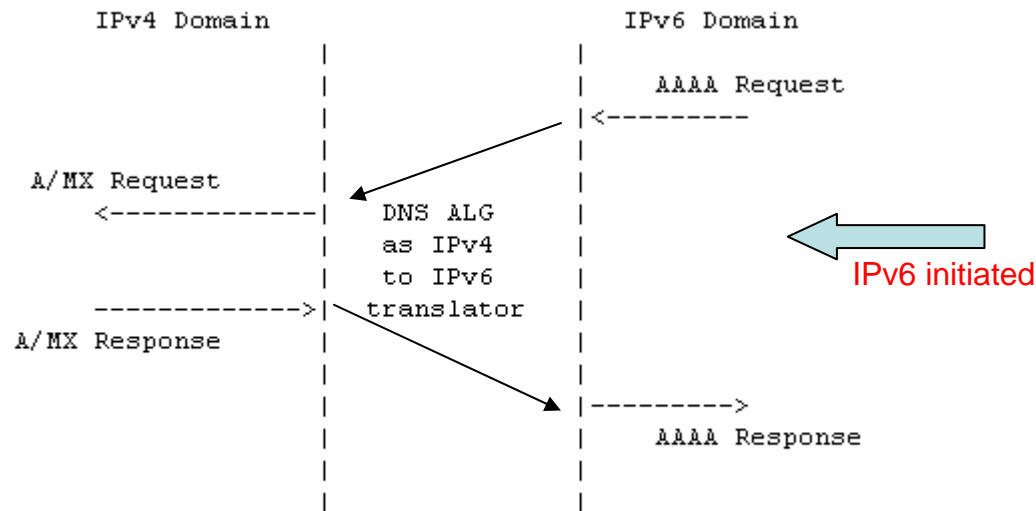


Figure 5: DNS Record Translation Service

- DNS translation
 - Algorithm based

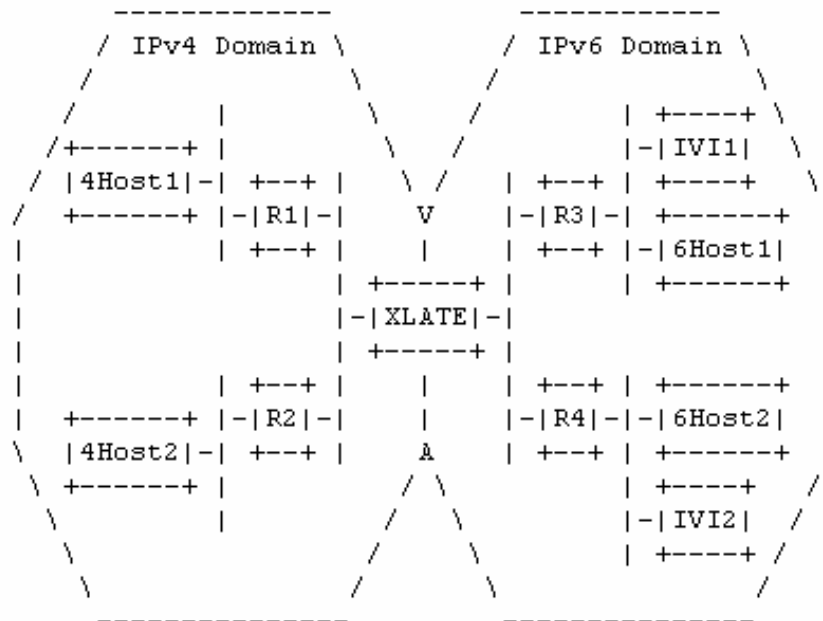
Stateless (1:1) operation

- Stateless
 - SIIT extension
 - Based on ISP's prefix
 - The mapping between IPv4 and IPv6 is based on algorithm
 - Support both IPv6 initiated and IPv4 initiated communications

Stateful (1:N) operation

- **Stateful**
 - NAT-PT (NAPT-PT) improvement
 - IPv4 address multiplexing
 - **Based on ISP's prefix**
 - Support IPv6 initiated communication

IVI Reachability





Route Advertisements:

R1: its IPv4 LAN	R3: its IPv6 LAN
R2: its IPv4 LAN	R3: its IVI /64
XLATE: IPv4 IVI prefix possible IPv4 overlay prefix	R4: its IPv6 LAN R4: its IVI /64 XLATE: IVI /40

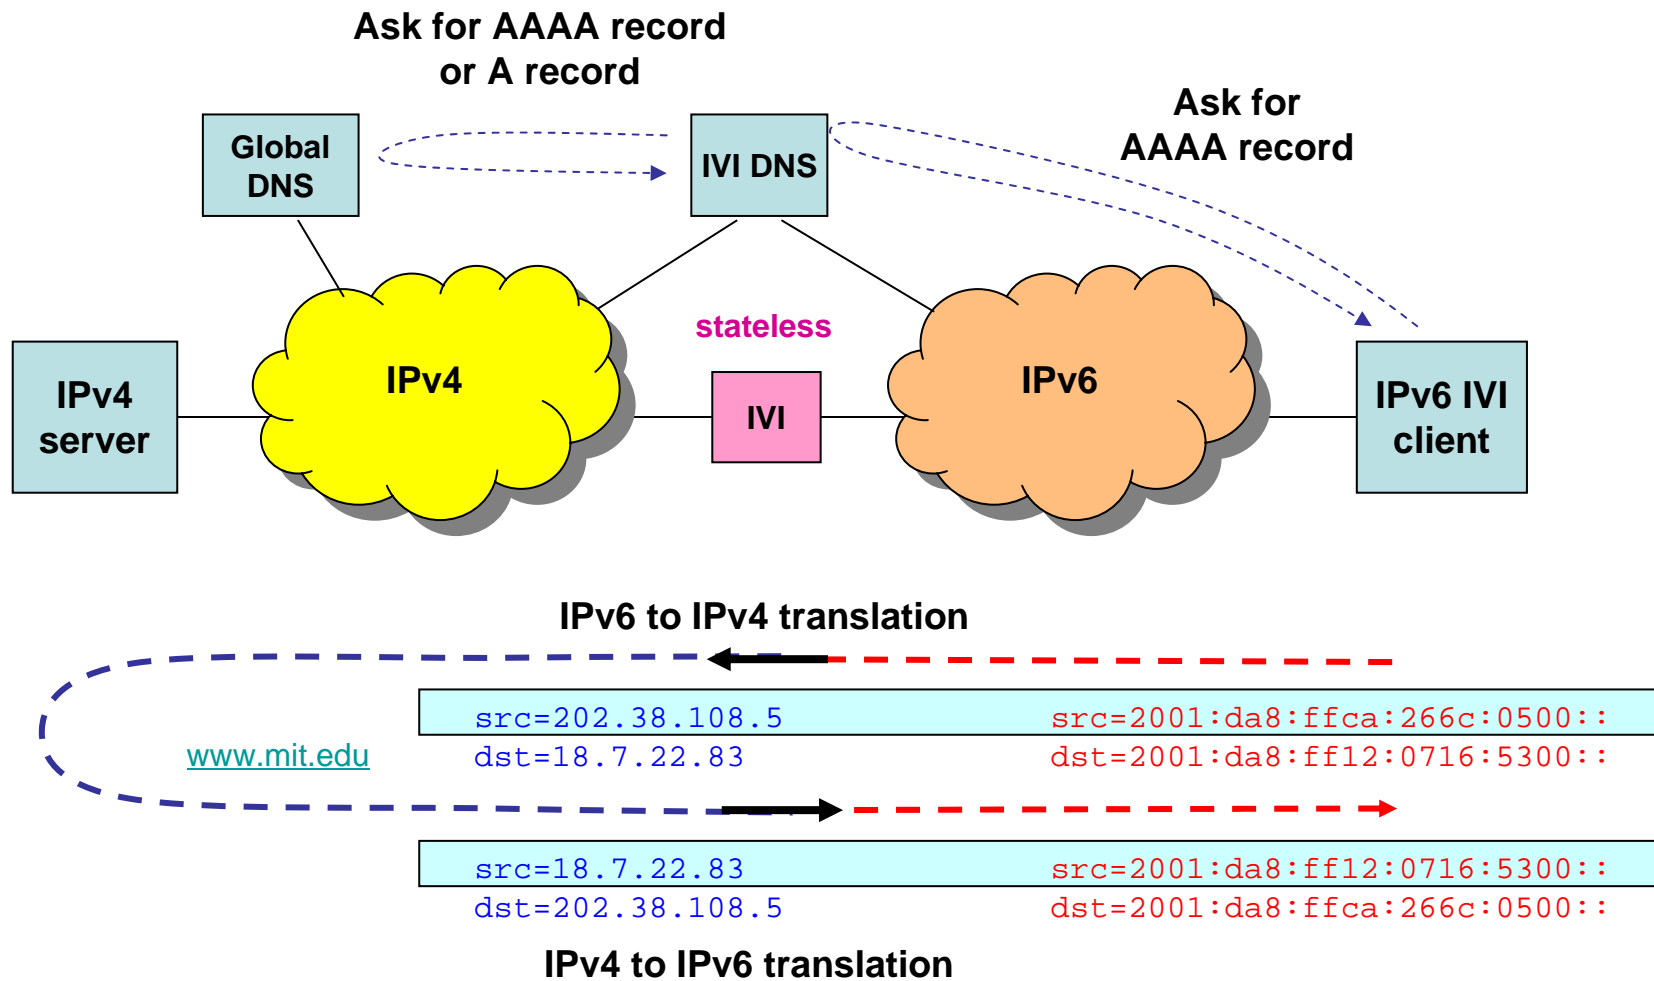
Figure 3: IVI Reachability example

- 4Host1 → IVI1 (stateless)
- IVI1 → 4Host1 (stateless)
- 4Host1 --X 6Host1
- 6Host1 → 4Host1 (stateful)
- 4Host1 → 4Host2
- 4Host2 → 4Host1
- 6Host1 → IVI1
- IVI1 → 6Host1
- 6Host1 → 6Host2
- 6Host2 → 6Host1
- IVI1 → IVI2
- IVI2 → IVI1

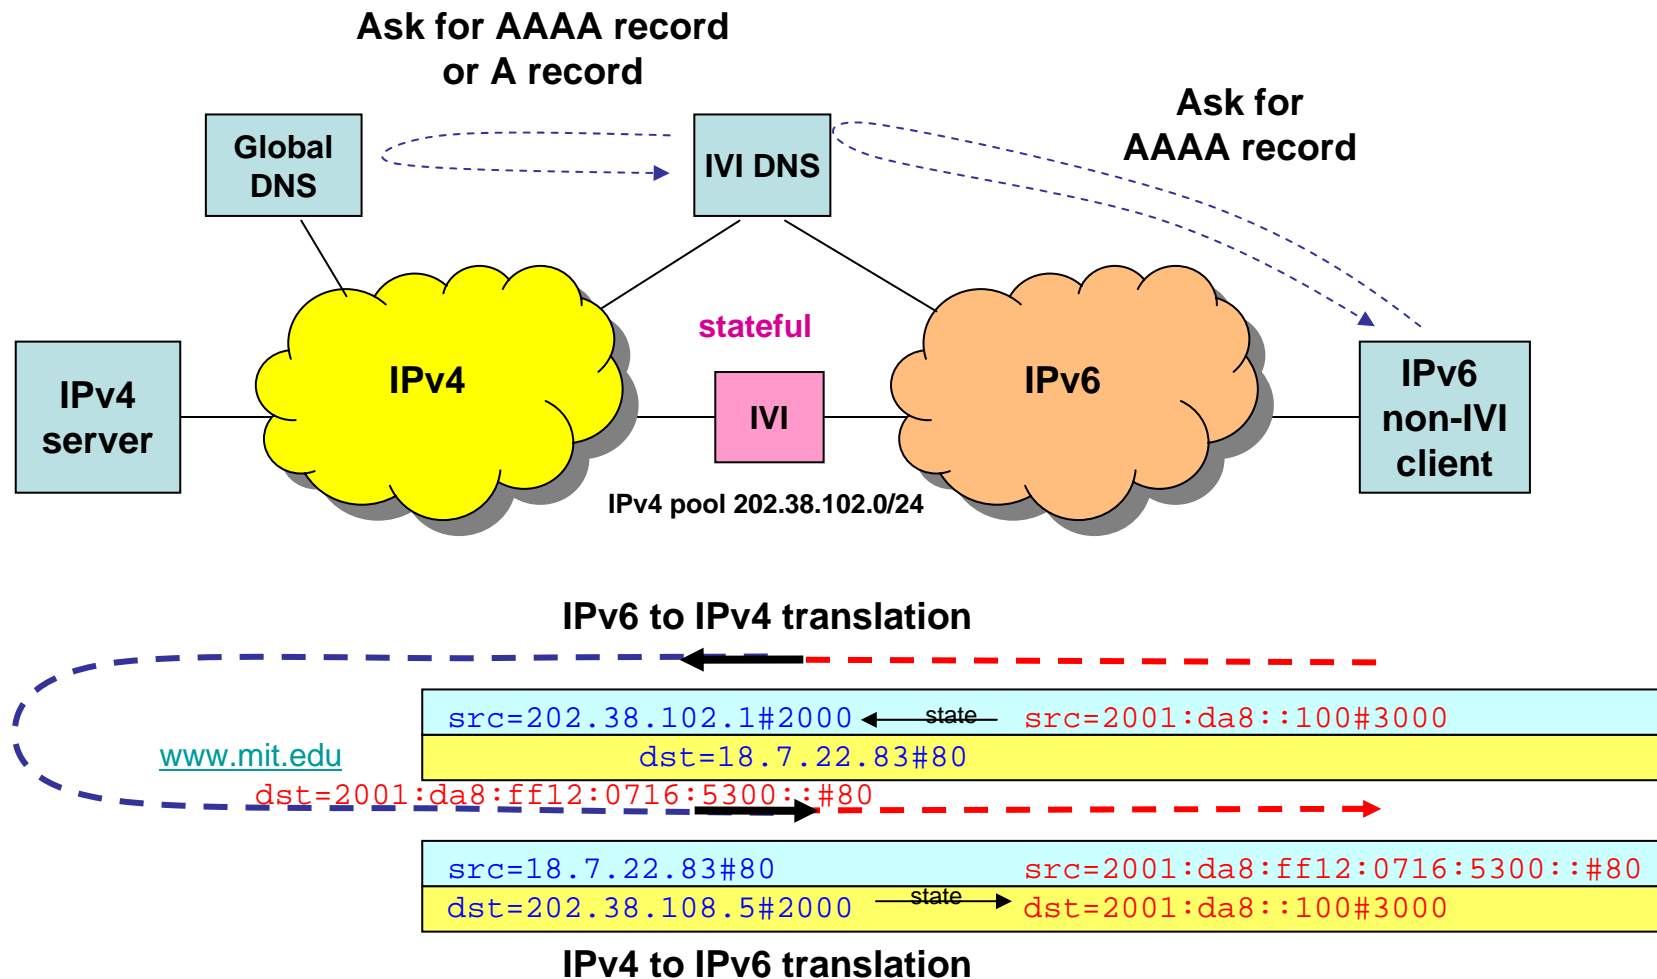
IPv4 initiated communication (1:1)



IPv6 initiated communication (1:1)



IPv6 initiated communication (1:N)



IVI ICMP extension

```
ivitraceroute6 www.mit.edu

src_ivi4=202.38.97.205 src_ivi6=2001:da8:ffca:2661:cd00::
dst_host=www.mit.edu
dst_ip4=18.7.22.83 dst_ivi6=2001:da8:ff12:716:5300::

traceroute to 2001:da8:ff12:716:5300:: (2001:da8:ff12:716:5300::)
30 hops max, 40 byte packets to not_ivi

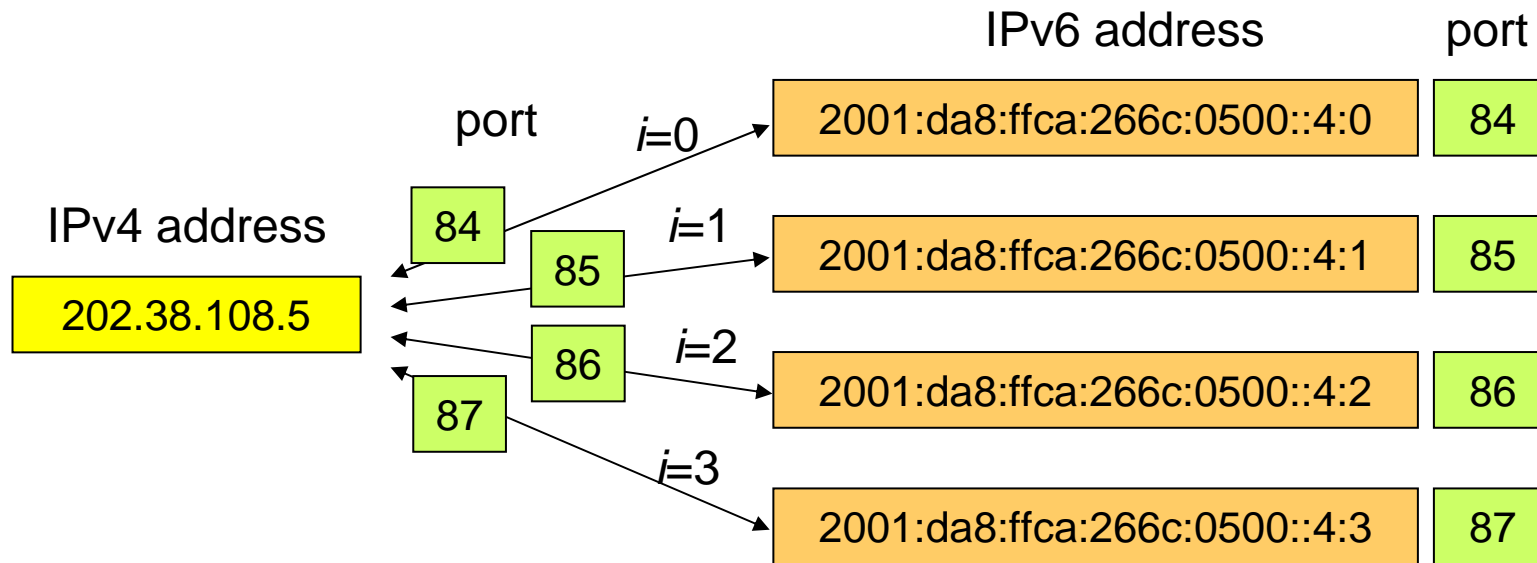
 1  2001:da8:ff0a:0:100::      0.304 ms 0.262 ms 0.190 ms
    10.0.0.1
 2  2001:da8:ffca:7023:fe00::  0.589 ms * *
    202.112.35.254
 3  2001:da8:ffca:7035:4900::  1.660 ms 1.538 ms 1.905 ms
    202.112.53.73
 4  2001:da8:ffca:703d:9e00::  0.371 ms 0.530 ms 0.459 ms
    202.112.61.158
 5  2001:da8:ffca:7035:1200::  0.776 ms 0.704 ms 0.690 ms
    202.112.53.18
 6  2001:da8:ffcb:b5c2:7d00::  89.382 ms 89.076 ms 89.240 ms
    203.181.194.125
 7  2001:da8:ffc0:cb74:9100::  204.623 ms 204.685 ms 204.494 ms
    192.203.116.145
 8  2001:da8:ffcf:e7f0:8300::  249.842 ms 249.945 ms 250.329 ms
    207.231.240.131
 9  2001:da8:ff40:391c:2d00::  249.891 ms 249.936 ms 250.090 ms
    64.57.28.45
10  2001:da8:ff40:391c:2a00::  259.030 ms 259.110 ms 259.086 ms
    64.57.28.42
11  2001:da8:ff40:391c:700::   264.247 ms 264.399 ms 264.364 ms
    64.57.28.7
12  2001:da8:ff40:391c:a00::   271.014 ms 269.572 ms 269.692 ms
    64.57.28.10
13  2001:da8:ffc0:559:dd00::   274.300 ms 274.483 ms 274.316 ms
    192.5.89.221
14  2001:da8:ffc0:559:ed00::   274.534 ms 274.367 ms 274.517 ms
    192.5.89.237
15  * * *
16  2001:da8:ff12:a800:1900::  276.032 ms 275.876 ms 276.090 ms
    18.168.0.25
17  2001:da8:ff12:716:5300::  276.285 ms 276.370 ms 276.214 ms
    18.7.22.83
```

- Operation
 - IPv4 → IPv6
 - IPv6 → IPv4

```
ivitraceroute 202.38.108.2

 1  202.112.0.65 6 ms 2 ms 1 ms
 2  202.112.53.73 4 ms 6 ms 12 ms
 3  202.112.53.178 1 ms 1 ms 1 ms
 4  202.112.61.242 1 ms 1 ms 1 ms
 5  202.38.17.186 1 ms 1 ms 1 ms
    202.38 AS4538
 6  202.38.17.186 1 ms 1 ms 1 ms
    202.38 AS4538
 7  202.38.17.186 2 ms 2 ms 2 ms
    202.38 AS4538
 8  202.38.17.186 2 ms 2 ms 2 ms
    202.38 AS4538
 9  202.38.17.186 4 ms 4 ms 3 ms
    202.38 AS4538
10  202.38.108.2 2 ms 3 ms 3 ms
```

Stateless 1:N operation



Port number keep the same

The multiplexing ratio

- If the multiplexing ratio is 256。
- One IPv4 /8 can support 4.3 billion IPv6 hosts, same as the size of the global IPv4 space.

IVI multicast

- IVI supports PIM SSM
 - Group address mapping
 - RPF → mapped IPv6 address
 - PIM Spare-mode ALG

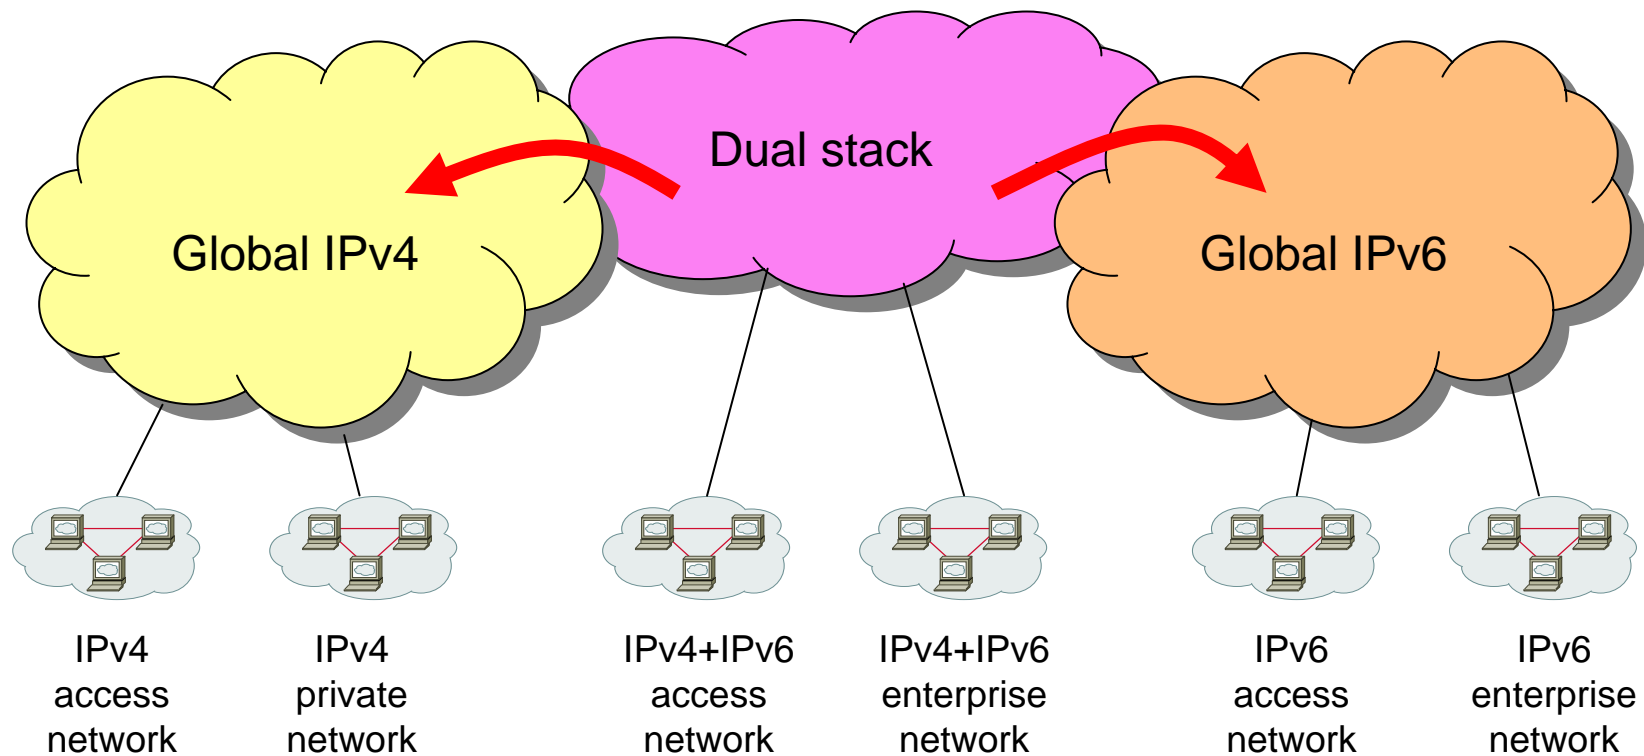
IVI Multicast Group Address Mapping

IPv4 Group Address	IPv6 Group Address
232.0.0.0/8	ff3e:0:0:0:0:0:f000:0000/96
232.255.255.255/8	ff3e:0:0:0:0:0:f0ff:ffff/96

Transition mechanisms

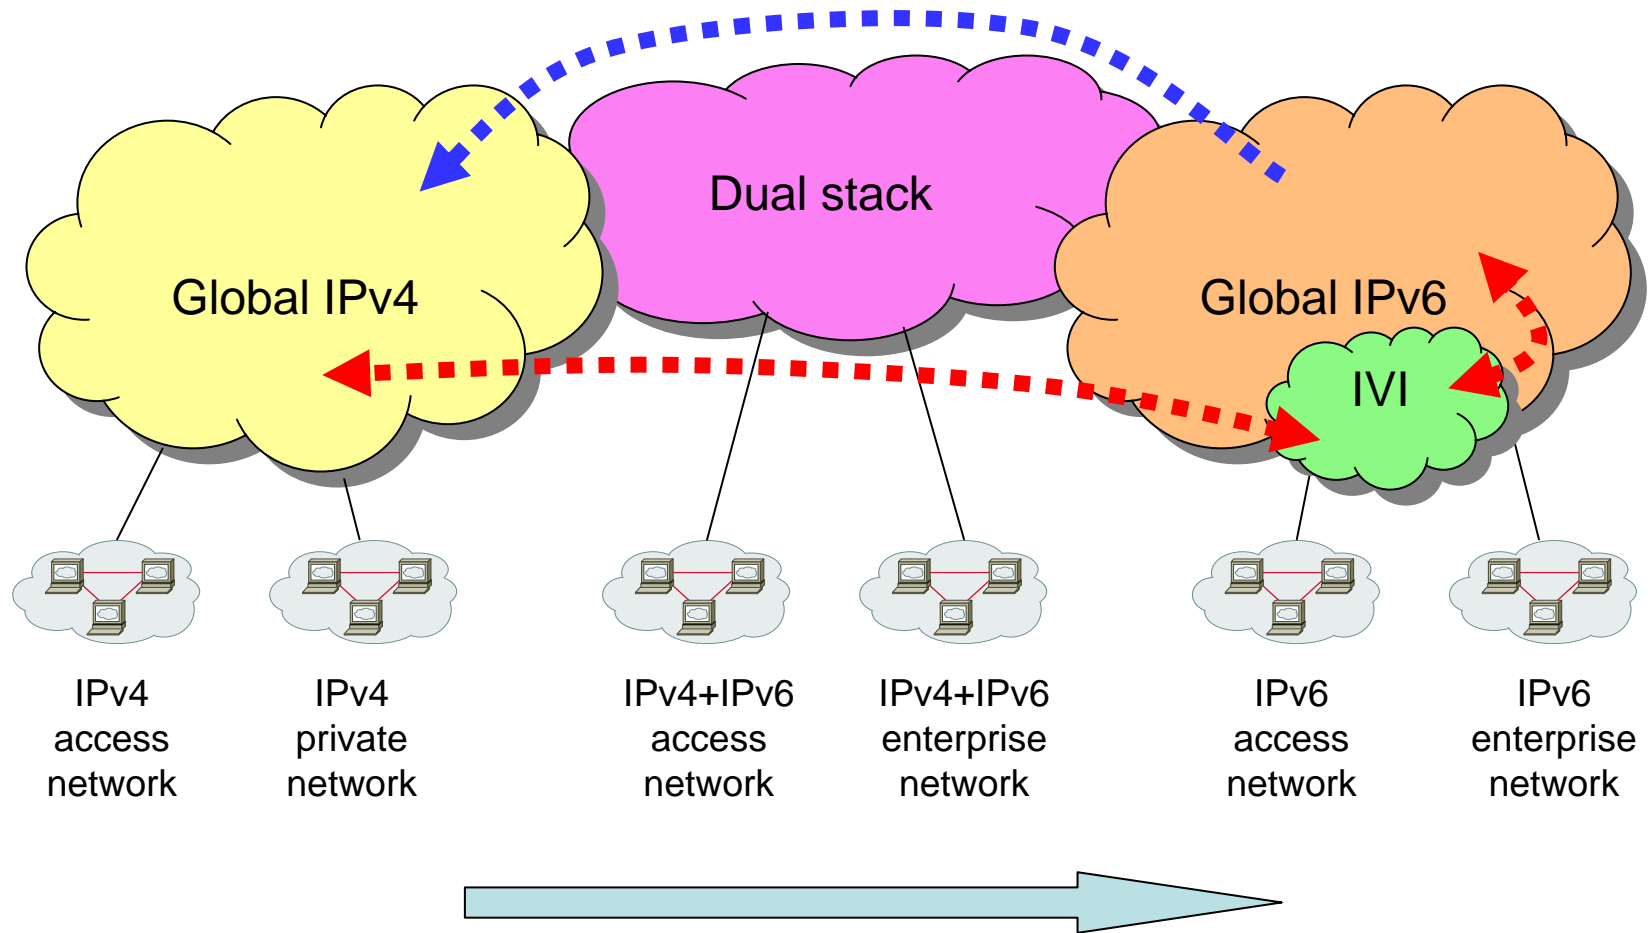
- When IPv4 addresses are running out
 - IPv4 + NAT
 - Short term solution
 - Pure IPv6
 - Cannot reach the global IPv4, unacceptable
 - Dual stack
 - The cost increases, ISPs want others to deploy dual stack
 - IVI IPv6
 - The cost is the same as the single stack, but the IPv6 host can be reached by global IPv4

Dual stack



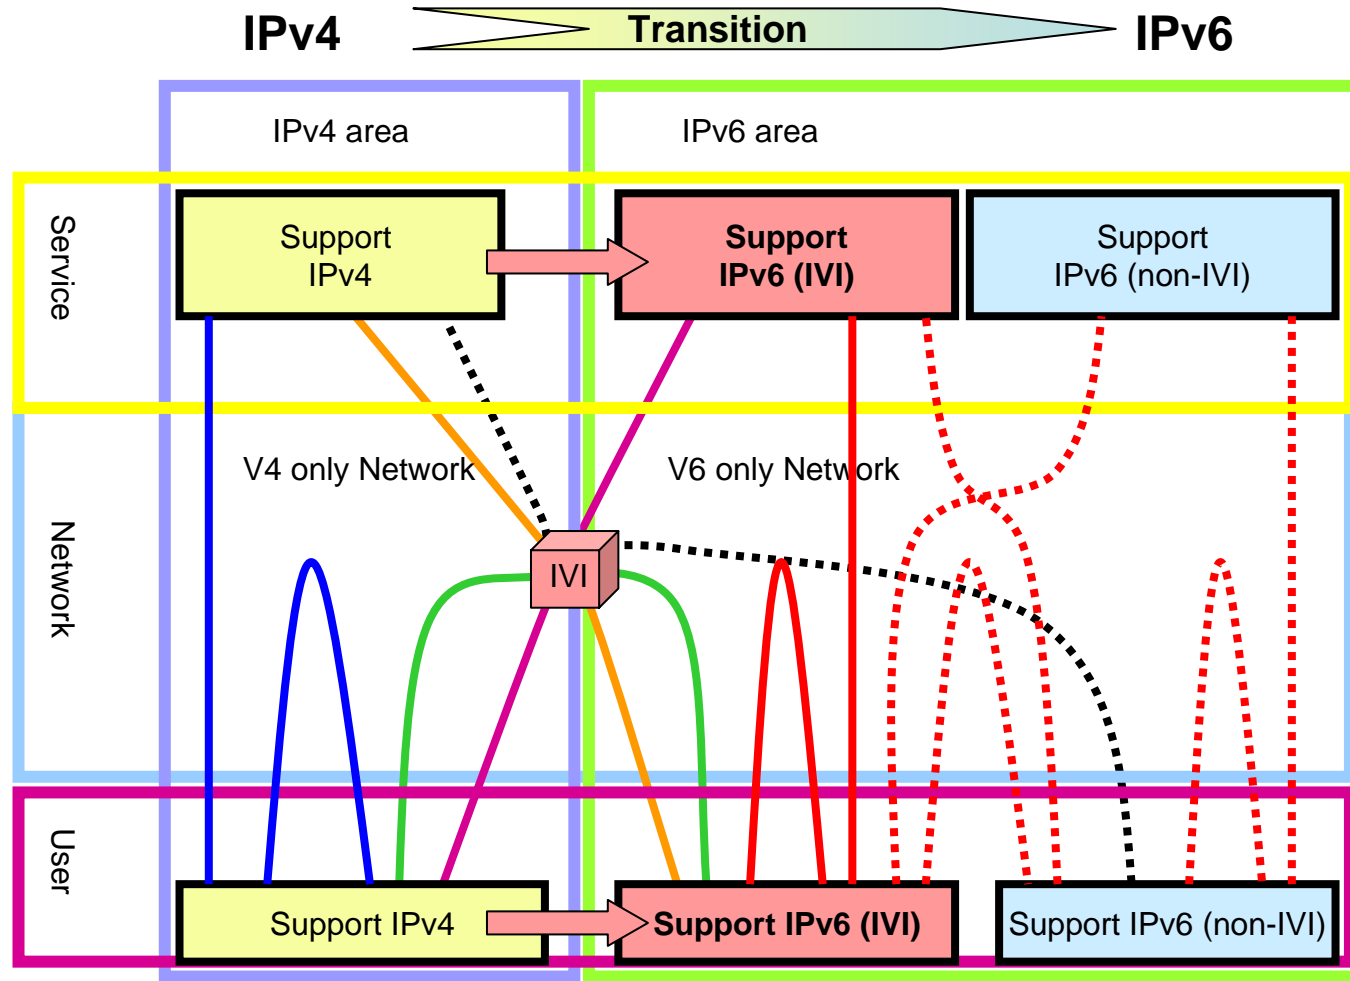
- Not easy

IVI



- Encourage transition

Transition



<http://www.ivi2.org/IVI/>

IVI - Mozilla Firefox

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

http://www.ivi2.org/IVI/

IVI

IVI source code download

The IVI IPv4/IPv6 packet translation implementation as a Linux kernel patch is available below.

- [IVI v0.5 kernel patch](#) for [Linux kernel 2.6.12](#)
- [IVI v0.5 kernel patch](#) for [Linux kernel 2.6.18](#)

The IVI A/AAAA DNS proxy implementation is available below.

- [IVIDNS v0.1 C code](#)

For installing and configuration, please follow the instructions in the source code packages.


IVI test servers

- [Access single-stack IPv6 server \[2001:250:ffca:2672:100::\] = 202.38.114.1 via IPv4](#)
- [Access single-stack IPv6 sever \[2001:250:ffca:2672:100::\] via IPv6](#)
- [Access IPv4 server \(202.38.114.129\) cross single-stack IPv6 network](#)

IVI references

IETF drafts

- [Prefix-specific and Stateless Address Mapping \(IVI\) for IPv4/IPv6 Coexistence and Transition \(01\)](#)
- [IVI Update to SIIT and NAT-PT \(01\)](#)

 copyright CERNET Center 2005-2009

完成

开始 IVI - ... 2 Mic... 收件箱... Multi... 100% 15:03