DECLARATION

I, Alexa Morris, based on my personal knowledge and information, hereby declare as follows:

1. I am Managing Director of the IETF Administration LLC and have held that position since the LLC was formed in August 2018. Prior to that, starting on January 1, 2008, I was the Executive Director of the Internet Engineering Task Force, which was an activity of the Internet Society. Since the business of IETF did not change in any materially relevant manner with the formation of the LLC, I will collectively refer to both the activity and the LLC as IETF.

2. One of my responsibilities with IETF has been to act as the custodian of Internet-Drafts and records relating to Internet-Drafts. I am familiar with the record keeping practices relating to Internet-Drafts, including the creation and maintenance of such records.

3. I hereby declare that all statements made herein are of my own knowledge and information contained in the business records of IETF and are true, and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements may be punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

4. If depositions regarding the information in this declaration are required, the deposition should be taken by phone or videoconference or, if it must be in person, should be in California.

5. Since 1998, it has been the regular practice of the IETF to publish Internet-Drafts and make them available to the public on its website at www.ietf.org (the IETF website). The IETF maintains copies of Internet-Drafts in the ordinary course of its regularly conducted activities.

6. Any Internet-Draft published on the IETF website was reasonably accessible to the public and was disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable diligence could have located it. In particular, the Internet-Drafts were indexed and searchable on the IETF website.

7. Internet-Drafts are posted to an IETF online directory. When an Internet-Draft is published, an announcement of its publication that describes the Internet-Draft is disseminated. Typically, that dated announcement is made within 24 hours of the publication of the Internet-Draft. The announcement is kept in the IETF email archive and the date is affixed automatically.

8. The records of posting the Internet-Drafts in the IETF online repository are kept in the course of the IETF's regularly conducted activity and ordinary course of business. The records are made pursuant to established procedures and are relied upon by the IETF in the performance of its functions.

9. It is the regular practice of the IETF to make and keep the records in the online repository.

10. Exhibit 1 is a true and correct copy of an announcement of the publication of draft-xli-behave-ivi-00.txt, titled "Prefix-specific and Stateless Address Mapping (IVI) for IPv4/IPv6 Coexistence and Transition." I have determined that an announcement of the publication of this Internet-Draft was made on July 6, 2008. Therefore, based on the normal practice of the IETF, that Internet-Draft was reasonably available to the public within 24 hours of that announcement. At that time, the Internet-Draft would have been disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located it.

11. Exhibit 2 is a true and correct copy of the slide presentation titled, "Prefix-specific and Stateless Address Mapping (IVI) for IPv4/IPv6 Coexistence and Transition draft-xli-behave-ivi-00." The IETF records for which I am responsible, and which I rely upon to be accurate, indicate that this presentation was made at an IETF meeting on July 29, 2008. Therefore, based on the normal practice of the IETF, that Internet-Draft was reasonably available to the public within 24 hours of that announcement. At that time, the Internet-Draft would have been disseminated or otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising reasonable diligence, could have located it.

Pursuant to Section 1746 of Title 28 of United States Code, I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct and that the foregoing is based upon personal knowledge and information and is believed to be true.

Date: 10/23/23

By: Alexa Morris

Alexa Morris

4880-9037-7097

EXHIBIT 1

Network Working Group
X. Li
Internet-Draft
M. Chen
Intended status: Standards Track
C. Bao
Expires: January 7, 2009
H. Zhang

J. Wu
CERNET Center/
Tsinghua University

July 6, 2008

Prefix-specific and Stateless Address Mapping (IVI) for IPv4/IPv6 Coexistence and Transition

draft-xli-behave-ivi-00.txt

Status of this Memo

documents at any
time. It is inappropriate to use Internet-Drafts as reference
material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at
http://www.ietf.org/shadow.html.

This Internet-Draft will expire on January 7, 2009.

Abstract

This document presents the concept and practice of the prefix-
specific and stateless address mapping mechanism (IVI) for IPv4/IPv6
coexistence and transition. In this scheme, subsets of the IPv4

addresses are embedded in prefix-specific IPv6 addresses and these IPv6 addresses can therefore communicate with the global IPv6 networks directly and can communicate with the global IPv4 networks via stateless (or almost stateless) gateways. The IVI scheme supports the end-to-end address transparency, incremental deployment and performance optimization in multi-homed environment. This document is a comprehensive report on the IVI design and its deployment in large scale public networks. Based on the IVI scenario, the corresponding address allocation and assignment policies are also proposed.

Table of Contents

1. Introduction

This document presents the concept and practice of the prefix-specific and stateless address mapping mechanism (IVI) for IPv4/IPv6 coexistence and transition.

The experiences for the IPv6 deployment in the past 10 years strongly indicate that for a successful transition, the IPv6 hosts need to communicate with the global IPv4 networks [JJI07]. However, the current transition methods do not fully support this requirement [RFC4213]. For example, dual-stack hosts can communicate with both the IPv4 and IPv6 hosts, but the IPv4 address depletion problem makes the dual-stack approach inapplicable [COUNT]. The tunneled architectures can link the IPv6 islands cross IPv4 networks, but they cannot help the communication between two address families [RFC3056] [RFC5214] [RFC4380]. The translation architectures can relay the communications for the hosts located in IPv4 and IPv6 networks, but the current implementation of this kind of architecture is not scalable and it cannot maintain the end-to-end address transparency [RFC2766] [RFC3142] [RFC4966] [RFC2775].

However, since IPv4 and IPv6 are different protocols with different addressing structure, the translation mechanism is still necessary for the communication between the two address families. There are several ways to implement the translation. One is the stateless IP/ICMP translation algorithm (SIIT), which provides a mechanism for the translation between IPv4 and IPv6 packet headers (including ICMP headers) without requiring any per-connection state. But, SIIT does not specify the address assignment and routing scheme [RFC2766]. For example, when SIIT is used for the IPv4 mapped IPv6 addresses [::FFFF:ipv4-addr/96] and IPv4 compatible IPv6 addresses [::ipv4-address/96]), these addresses violate the aggregation nature of the IPv6 routing [RFC4291]. The other translation mechanism is NAT-PT, which has serious technical and operational difficulties and IETF has reclassified it from proposed standard to historic status. But in the same document, it suggested that a revised, possibly restricted version of NAT-PT can be a suitable solution for the communication between IPv4 and IPv6 hosts [RFC4966]. Recently, several mechanisms are proposed in this direction, for example NAT64 translates the IPv4 server address by adding or removing a /96 prefix, and translates the IPv6 client address by installing mappings in the normal NAT manner [I-D.bagnulo-behave-nat64].

In this document, we follow the basic specification of SIIT, but we define the address assignment and routing scheme (IVI). Our IVI mechanism is related to the NAT-PT and NAT64, but differs from them significantly. First, it is stateless (or almost stateless) in both the IPv4-to-IPv6 mapping direction, as well as in the IPv6-to-IPv4 mapping direction. Secondly, it supports address transparency. Thirdly, it supports both client-server applications and the peer-to-peer applications cross IPv4 and IPv6 address families without using NAT-traversal techniques. Finally, it can satisfy most of the basic and advanced requirements for the IPv4 to IPv4 transition as specified by the Internet Drafts [I-D.v6ops-nat64-pb-statement-req].

2. Terms and Abbreviations

The following terms and abbreviations are used in this document:

IVI: IV means 4 and VI means 6 in Roman representation, so IVI means mapping and translation between IPv4 and IPv6.

ISP(i): A specific Internet service provider "i".

IPG4: An address set containing all IPv4 addresses, the addresses in this set are mainly used by IPv4 hosts at the current stage.

IPS4(i): A subset of IPG4 allocated to ISP(i).

IVI4(i): A subset of IPS4(i), the addresses in this set will be mapped to IPv6 via IVI rule and physically used by IPv6 hosts of ISP(i).

IPG6: An address set containing all IPv6 addresses.

IPS6(i): A subset of IPG6 allocated to ISP(i).

IVIG46(i): A subset of IPS6(i), an image of IPG4 in IPv6 address family via IVI mapping rule.

IVI6(i): A subset of IVIG46(i), an image of IVI4(i) in IPv6 address family via IVI mapping rule.

IVI gateway: The mapping and translation gateway between IPv4 and IPv6 based on IVI scheme.

IVI DNS: Providing IVI Domain Name Service (DNS).

The key words MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [RFC2119].

3. The Overview of the IVI Mechanism

The IVI is a prefix-specific and stateless address mapping scheme which can be carried out by individual ISPs.

IVI mapping and translation mechanism is implemented in an IVI gateway which connects to both IPv4 and IPv6 networks. The SIIT stateless translation is implemented in the IVI gateway [RFC2765].

A unique, prefix-specific and stateless mapping scheme is defined between IPv4 addresses and subsets of IPv6 addresses, so each provider-independent IPv6 address block (usually a /32) will have a small portion of IPv6 addresses, which is the image of the totality of the global IPv4 addresses.

Each provider can borrow a portion of its IPv4 addresses and maps them into IPv6 based on the above mapping rule. These special IPv6 addresses will be physically used by IPv6 hosts. The original IPv4 form of the borrowed addresses is the image of these special IPv6 addresses.

The packets generated from the global IPv4 addresses and sent to the special IPv6 addresses are routed to the IPv4 interface of the IVI gateway via the IPv4 routing protocol and the packets generated from the special IPv6 addresses and sent to the global IPv4 addresses are routed to the IPv6 interface of the IVI gateway via the IPv6 routing protocol. The processes in both directions are symmetric. In addition, the special IPv6 addresses can communicate with the global IPv6 networks.

The IVI scheme related issues, for example the IVI DNS support, the multiplexing of the public IPv4 addresses, the IVI multicast support, etc. can be solved without involving any major change in the current Internet protocol.

3.1. Address Mapping

The IVI address mapping is defined based on individual ISP's prefix as shown in the following figure.

IVI Address Mapping

```
   | 0                   |32 |40                    |72
127|

-----------------------------------------------------------
-------
   |                   |FF |                      |
|

-----------------------------------------------------------
-------
   |<-  IPv6 prefix  ->|   |<-  IPv4 address   ->|<- zero
padding ->|
```

Figure 1

where bit 0 to bit 31 are the prefix of ISP(i)'s /32 (e.g. IPS6=2001:DB8::/32), bit 32 to bit 39 are all one's as the identifier of IVI, bit 40 to bit 71 are embedded global IPv4 space (IPG4) presented in hexadecimal format. (e.g. 2001:DB8:ff00::/40). Because this mapping is 1-to-1 defined by the IVI mapping rule, it is stateless and it has feature of end-to-end address transparency.

(1) The ISP(i) uses a subset of ISP4(i) defined as IVI4(i), and maps it into IPv6 as IVI6(i). The IVI6(i) is physically used by IPv6 hosts inside ISP(i)'s IPv6 network and the IVI4(i) cannot be used by IPv4 hosts. Therefore, IVI6(i) is the special IPv6 address block which can communicate with both address families.

(2) Based on the above mapping rule, the ISP(i) uses a subset of ISP6(i) defined as IVIG46(i), and maps it into IPv4 as IPG4. The IVIG46(i) is virtually used by global IPv4 hosts and it cannot be used by IPv6 hosts, except the portion of IVI6(i).

The mapping of the different address sets and the relations are shown in the following figure.

IVI Address Mapping Relation

```
             |<-------IPG4-------------------->|
             |        |<----IPS4(i)----->|      |
```

```
                |           |<-IVI4(i)->|          |
                |           |           |          |
                |           |    /\     |          |
                |           |    ||     |          |
                |           | mapping   |          |
                |           |    ||     |          |
                |           |    \/     |          |
                |           |           |          |
                |           |<-IVI6(i)->|          |
                |<------IPG46(i)----------------->|
          |<--------IPS6(i)----------------------------->|
      |<-----------
IPG6------------------------------------------->|
```

Figure 2

where IVI4(i) and IVI6(i) are representing the same entities in IPv4 and IPv6 address families, respectively. Similarly, IPG4 and IVIG46(i) are representing the same entities in IPv4 and IPv6 address families, respectively. In addition, IVI4(i) is a subset of IPG4 and IVI6(i) is a subset of IVIG46(i).

3.2. Routing and Forwarding

Based on the IVI address mapping rule, the routing is straightforward, as shown in the following figure.

IVI Routing

```
    /-----
\                                                    /----
-\
   (Global )  ----192.168.1.2  -------------
2001:DB8::2----   (Global )
   (IPv4   )--|R1|-------------|IVI gateway|------------|
R2|---(IPv6   )
   (network)  ----  192.168.1.1------------2001:DB8::1
```

```
----   (network)
    \-----/
\-----/
```

Figure 3

where

(1) Router R1 has IPv4 route of IVI4(i)/k (k is the prefix length of

IVI4(i)) with next-hop equals to 192.168.1.1 and this route is distributed to global IPv4 networks with proper aggregation.

(2) Router R2 has IPv6 route of IVIG46(i)/40 with next-hop equals to 2001:DB8::1 and this route is distributed to global IPv6 networks with proper aggregation.

(3) IVI gateway has IPv6 route of IVI6(i)/(40+k) with next hop equals to 2001:DB8::2. IVI gateway also has IPv4 default route 0.0.0.0/0 with next hop equals to 192.168.1.2 .

Note that the routes described above can be learned/ inserted by dynamic routing protocols in the IVI gateway neighboring (IGP) or peering (BGP) with R1 and R2.

The address reachability matrix of the IPv4, IVI and IPv6 is shown in the following figure.

IVI reachability Matrix

	IPG4	IVI	IPG6
IPG4	OK	OK	NO
IVI	OK	OK	OK
IPG6	NO	OK	OK

Figure 4

Since both IVI4(i) and IVI6(i) are aggregated to IPS4(i) and IPS6(i) in ISP(i)'s border routers respectively, there will be no affect to the global IPv4 and IPv6 routing tables [RFC4632].

If IVI4(i) and IVI6(i) has 1-to-1 mapping relationship, then IVI is stateless and it can support multi-homing.

Since IVI can be implemented independently in each ISP's network, it can be incrementally deployed.

3.3. IVI Communication Scenarios

Scenario 1:

Assume that there are IPv4 address A and ISP(1) IVI-mapped IPv6 address A', an arbitrary IPv4 address B and ISP(1) IVI-mapped IPv6 address B', as well as an arbitrary IPv6 address C'. If ISP(1) deploys IVI, then A' is a physical IPv6 host and B is a physical IPv4 host. A' can communicate with B via the IVI gateway. Note that in this scenario A' is actually communicating with B', an image of B, and B is actually communicating with A, an image of A'. Since A' is an IPv6 address inside ISP6(1), it can also communicate with arbitrary IPv6 host C'. This can form an early stage of IPv4/IPv6 coexistence and transition.

IVI Communication Scenario 1

```
             ------------------------
           /          IPv4            \
          |                             |
          |     A<------->B             |
           \              /           /
             ------------/------------
                 |      /
              ------   /
             |IVI GW| /
              ------ /
                 |  /
             ------/----------------------------------
           /      /                                     \
          |     A'<-->B'                                 |
          |                                              |
           \           IPv6                C'           /
             -----------------------------------------
```

Figure 5

Scenario 2:

Assume that there are IPv4 address A, ISP(1) IVI-mapped IPv6 address A' and ISP(2) IVI-mapped IPv6 address A''. Similarly, assume that there are IPv4 address B, ISP(1) IVI-mapped IPv6 address B' and ISP(2) IVI-mapped IPv6 address B''. If both ISP(1) and ISP(2) deploy IVI, then A' and B'' are physical IPv6 hosts. In addition, if ISP(1) and ISP(2) do not know the IVI deployment on the other end, then A' can still communicate with B'' through A and B via two IVI gateways.
Note that in this scenario A' is actually communicating with B', an ISP(1)'s version image of B, and B'' is actually communicating with A'', an ISP(2)'s version image of A. Since there are two IVI gateways involved, the routing is not optimal.

IVI Communication Scenario 2

```
           /          IPv4            \
          |                            |
          |        A<--------->B       |
           \         \       /        /
             ---------\-----/--------
                 |     \   /     |
              ------    \ /   ------
             |IVI GW|    .   |IVI GW|
              ------    / \   ------
                 |     /   \     |
             ---------/-----\-------------------------
           /         /       \                         \
          |        A'<->B'    \                         |
          |                    \                        |
          |               A''<->B''                     |
           \           IPv6                            /
             -----------------------------------------
```

Figure 6

Scenario 3:

Assume that there are IPv4 address A and ISP(1) IVI-mapped IPv6 address A'. Similarly, assume that there are IPv4 address B and ISP(2) IVI-mapped IPv6 address B''. If both ISP(1) and ISP(2) deploy IVI, then A' and B'' are physical IPv6 hosts. In addition, if ISP(1) and ISP(2) by contrast know the IVI deployment on the other end, then A' can communicate with B'' directly. Since it is the communication in IPv6, the routing is optimal. This can form a later stage of the transition.

IVI Communication Scenario 3

```
             ------------------------
           /           IPv4           \
          |                            |
          |        A           B       |
           \                          /
             ------------------------
                 |               |
               ------          ------
             |IVI GW|        |IVI GW|
               ------          ------
                 |               |
             -----------------------------------------
           /                                           \
          |        A'<---------->B''                    |
          |                                             |
           \            IPv6                           /
             -----------------------------------------
```

Figure 7

4. Design Considerations

The components of the IVI scheme include: address mapping, network-layer header translation, transport-layer header translation, fragmentation/MTU handling, ICMP handling, application layer gateway, IPv6 source address selection and IPv4 over IPv6 support.

4.1. Address Mapping

The address mapping rule is defined in Section 3.1.

In addition, depending on the implementation scope of the IVI gateway, IVIG46(i) block can also be defined as 2001:DB8:FFFF::/48, 2001:DB8:ABCD:FF00::/56 or 2001:DB8:ABCD:FFFF::/64, etc. A special case is to define IVIG46(i)=2001:DB8:XXXX:XXXX:XXXX:XXXX::/96, then the mapping rule is similar to the method of translating the IPv4 server address proposed in [I-D.bagnulo-behave-nat64].

4.2. Network-layer Header Translation

IPv4 [RFC791] [RFC791] and IPv6 [RFC2460] are different protocols with different network layer header format, the translation of the IPv4 and IPv6 headers must be performed [MVB98] [RFC2765] as shown in the following figures.

IPv4 to IPv6 Header translation based on IVI scheme

IPv4 Field	Translated to IPv6
Version (0x4)	Version (0x6)
IHL	discarded
Type of Service	discarded
Total Length	Payload Length = Total Length -IHL * 4
Identification	discarded
Flags	discarded
Offset	discarded

```
   Time to Live          Hop Limit
   Protocol              Next Header
   Header Checksum       discarded
   Source Address        IVI address mapping
   Destination Address   IVI address mapping
   Options               discarded

-----------------------------------------------------------
--
```

Figure 8

IPv6 to IPv4 Header translation based on IVI scheme

```
-----------------------------------------------------------
--
   IPv6 Field            Translated to IPv4 Header

-----------------------------------------------------------
--
   Version (0x6)         Version (0x4)
   Traffic Class         discarded
   Flow Label            discarded
   Payload Length        Total Length = Payload Length +
20
   Next Header           Protocol
   Hop Limit             TTL
   Source Address        IVI address mapping
   Destination Address   IVI address mapping
   -                     IHL = 5
   -                     Header Checksum recalculated
```

Figure 9

4.3. Transport-layer Header Translation

Since the TCP and UDP headers [RFC793] [RFC768] consist of check sums which include the IP header, the recalculation and updating of the transport-layer headers must be performed [RFC2765].

4.4. Fragmentation and MTU Handling

When the packet is translated by the IVI gateway, due to the different sizes of the IPv4 and IPv6 headers, the IVI6 packets will be at least 20 bytes larger than the IVI4 packets, which may exceed the MTU of the next link in the IPv6 network. Therefore, the MTU handling and translation between IPv6 fragmentation headers and fragmentation field in the IPv4 headers are necessary, which is performed in the IVI gateway according to SIIT [RFC2765].

4.5. ICMP Handling

For ICMP message translation between IPv4 and IPv6, IVI follows the ICMP/ICMPv6 message correspondence as defined in SIIT [RFC2765]. Note that the ICMP message may be generated by an intermediate router whose IPv6 address does not belong to IVIG46(i). Since ICMP translation is important to the path MTU discovery, the inverse mapping for unmapped addresses is defined in this document. In the current prototype, a pseudo IPv4 address is generated in such a way that the first 16 bits are the IPv4 address of the IVI gateway, and the last 16 bits are the AS number of the current domain. This prevents translated ICMP messages from being discarded due to unknown

or private IP source. A small IPv4 address block should be reserved to identify the non-IVI mapped IPv6 addresses.

4.6. Application Layer Gateway

Due to the features of 1-to-1 address mapping and stateless, IVI can support most of the existing applications, such as HTTP, SSH, Telnet and Microsoft Remote Desktop Protocol. However, some applications are designed such that IP addresses are used to identify application-layer entities (e.g. FTP). In these cases, application layer gateway (ALG) is unavoidable, but it can be integrated into the IVI gateway. A list of applications which support the IVI scheme will be given in a later version of this document.

4.7. IPv6 Source Address Selection

Since each IPv6 host may have multiple addresses, it is important for the host to use an IVI6(i) address to reach the global IPv4 networks. The short-term work around is to use IVI6(i) as the default IPv6 address of the host. The long-term solution requires that the application be able to select the source addresses for different services.

4.8. IPv4 over IPv6 Support

The IVI scheme can support the IPv4 over IPv6 service (NAT646), i.e. a stub IPv4 network can be connected to an IVI gateway to reach the IPv6 network and via another IVI gateway to reach the global IPv4 network [RFC4925]

A more interesting scenario is to integrate the functions of the first IVI gateway into the end-system. In this case, the application softwares are IPv4-based and there is no need to have ALG support in the IVI gateway when it is communicating with IPv4 hosts.

5. DNS Configuration and Mapping

The DNS [RFC1035] service is important for the IVI scheme.

5.1. DNS Configuration for the IVI6(i) Addresses

For providing authoritative DNS service for IVI4(i) and IVI6(i), each host name will both have an A record and an AAAA record pointing to IVI4(i) and IVI6(i), respectively. Note that the same name always points to a unique host, which is an IVI6(i) host and it has IVI4(i) representation via the IVI gateway.

5.2. DNS Mapping for the IVIG46(i) Addresses

For resolving the IVI IPv6-mapped global IPv4 space (IVIG46(i)), each ISP must provide customized IVI DNS service for the IVI6(i) hosts. The IVI DNS server is in dual stack environment. When the IVI6(i) host queries an AAAA record for an IPv4 only domain name, the IVI DNS server will query the A record and map it to IVIG46(i) with ISP's IPv6 prefix and return an AAAA record to the IVI6(i) host.

6. Multiplexing of the Global IPv4 Addresses

Since public-IPv4 address is a scarce resource, the effective use of the IPv4 address is important for the IVI scheme. The multiplexing techniques are temporal multiplexing, port multiplexing,

spatial
multiplexing and multiplexing using IPv4 NAT-PT techniques.

6.1. Temporal Multiplexing

The IVI6 can be temporally multiplexed inside the ISP(i)'s /32. This is to say that the ISP can dynamically assign IVI6(i) to an end system when it requests the IPv4 communication service and release the IVI6(i) when the communication is finished. For temporal multiplexing, the features of stateless and end-to-end address transparency are maintained.

6.2. Port Multiplexing

To further increase the utilization ratio of the public IPv4 addresses, the port multiplexing inside the ISP(i)'s /32 can be deployed [RFC2766] [RFC4966]. This is to say that a single IPv4 address (IVI4(i)) can be used for multiple IVI6(i) addresses. The mapping scheme is to use the least significant bits in the IVI6(i) to define the multiple mapping and combine the transport-layer port number to perform uniquely the mapping from IVI4(i) to IVI6(i).

IVI Address Mapping for Port Multiplexing

Ratio	IVI6(i) range
1-to-1	2001:DB8:ffxx:xxxx:xx00:: - 2001:DB8:ffxx:xxxx:xx00::

```
   1-to-2^1 2001:DB8:ffxx:xxxx:xx00:: -
2001:DB8:ffxx:xxxx:xx00::1
   ......

   1-to-2^4 2001:DB8:ffxx:xxxx:xx00:: -
2001:DB8:ffxx:xxxx:xx00::15

-----------------------------------------------------------
-----
```

Figure 10

Based on this method, the mapping gain can be adjusted incrementally depending on the requirements. For example, zero bit means 1-to-1 mapping, and it is stateless. One bit means 1-to-1 mapping, and it has two states. Four bits means 1-to-16 mapping, etc. In the case of one-to-many mapping, when two IVI6(i)s have the same port number, the IVI gateway will map one of the port number to an unused port number and maintain the mapping table (IVI4(i) plus port number). Since the one-to-many mapping loses the feature of being stateless

and may loses the end-to-end address transparency, the proper use of the one-to-many mapping is the balancing of tradeoffs [RFC4966].

The tradeoffs are: (1) the number of the port number (2^16); (2) the gain of the IPv4 address utilization; (3) half association (3-tuple: source IP address, source port, transport protocol) or full association (5-tuple: source IP address, source port, destination IP address, destination port, transport protocol); (4) the number of states in the IVI gateway; (5) the average concurrent port used in an IPv6 host and; (6) the collision ratio of the port number.

6.3. Spatial Multiplexing

The spatial multiplexing means that for different operation modes of server and client, the different port multiplexing ratios can be applied. There are basically three cases.

(1) Server: we suggest having 1-to-1 mapping between IVI4(i) and IVI6(i), because it has the advantages of end-to-end address transparency, being stateless, having multi-homing support and providing services via well-known ports.

(2) Client with self-initiated connection: we suggest having 1-to-2^N mapping between IVI4(i) and IVI6(i) (N is a positive integer greater than 1), i.e. one IVI4(i) can support several IVI6(i) users to access the IPv4 network. By adjusting N, The number of states can be controlled. In this case, the port number is randomly generated by the client operating system. The IVI gateway maintains the port mapping table to avoid collision. There is no need to

modify the
client operating system and/or client application.

(3) Client with peer initiated connection: we suggest having 1-to-2^M mapping between IVI4(i) and IVI6(i) (M is a positive integer greater than 1 and may be smaller than N), i.e. one IVI4(i) can support several IVI6(i) users as the peer-to-peer hosts for the IPv4 network. By adjusting M, The number of states can be controlled. In this case, we can define "pseudo-well-known port number", which is unique for IVI4(i) and known to the peers. However, modification of the client operating system and/or client application may be necessary. By combining address and pseudo-well-known port number, the feature of end-to-end address transparency can still be maintained.

6.4. Multiplexing using IPv4 NAT-PT

If the private IPv4 address (e.g. 10.0.0.0/8) is used as the IPv4 address under the IVI scheme, combining conventional NAT-PT and NAT-traversing techniques, the public IPv4 addresses can also be multiplexed. The advantage of this method is that IPv4 NAT-PT

equipments are widely available and can be deployed immediately.
Moreover, the mapped prefix-specific IPv6 addresses (IVI6(i)) are no longer behind the NAT box in IPv6 and can be accessed by any IPv6 hosts.

7. IVI Multicast Support

The IVI scheme can support IPv4/IPv6 communication of the protocol-independent specific-source sparse-mode multicast (PIM SSM) [RFC3171] [RFC3569] [RFC4607].

(1) The IVI group address mapping rule: There will be 2^24 group addresses for IPv4 SSM. The corresponding IPv6 SSM group addresses can be defined as shown in the following figure.

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

Figure 11

(2) The IVI multicast source address selection: The source address in IPv6 has to be IVI6(i) in order to perform reverse path forwarding (RPF) as required by PIM-SM.

(3) The multicast protocol: The inter operation of PIM-SM for address families IPv4 and IPv6 can either be implemented via the application layer gateway or via the static join based on IGMPv3 and MLDv2 in IPv4 and IPv6, respectively.

The Any Source Multicast (ASM) cannot be supported in the cross address-family environment, since IPv6 does not support the MSDP [RFC4611], and IPv4 does not support the embedded RP [RFC3956].

8. IVI Implementation and Preliminary Testing Results

The IVI scheme presented in this document is implemented in the Linux OS and the source code can be downloaded [LINUX]. The example of the configuration is shown in Appendix A.

The IVI gateway based on the Linux implementation has been deployed between CERNET (IPv4 and partially dual-stack) [CERNET] and CNGI-CERNET2 (pure IPv6) [CERNET2] since March 2006. The pure IPv6 web servers using IPv6 addresses (IVI) behand IVI gateway can be accessed by the IPv4 hosts [IVI4], and also by the global IPv6 hosts [IVI6].

In addition, two traceroute results are presented in Appendix B to show the address mapping of the IVI scheme.

9. Features of IVI

The basic features of the IVI scheme are:

(1) Special IPv6 addresses can communicate with the global IPv6 network directly and can communicate with the global IPv4 network via IVI gateways.

(2) When the mapping is 1-to-1, the IVI gateway is stateless and can support multi-homing. When mapping is 1-to-2^N (N!=0), the IVI gateway is stateful, but the number of state can be controlled.

(3) When the mapping is 1-to-1, the IVI scheme has the advantages of

end-to-end address transparency. When mapping is 1-to-2^M, by introducing pseudo-well-known ports, the feature of end-to-end address transparency can also be maintained.

(4) The IVI addresses are globally routable.

(5) The IVI scheme is incrementally deployable.

(6) Based on the multiplexing techniques, the global IPv4 addresses can be effectively used.

The IVI scheme can satisfy most of the basic and advanced requirements for the IPv4 to IPv4 transition as specified by the Internet Drafts [I-D.v6ops-nat64-pb-statement-req].

For the basic requirements (MUST):

(1) No need to change the end system (IPv4 and IPv6).

(2) Support v4-initiated and v6-initiated short-lived local handle.

(3) Support interaction with dual-stack hosts.

(4) The standard IPv4 NAT can easily be integrated into the system.

(5) Do not violate standard DNS semantics.

(6) No affect to IPv6 routing.

(7) Support TCP, UDP, ICMP.

(8) Can handle fragmentation.

For the advanced requirements (SHOULD):

(1) Support multicast (SSM).

(2) Support operational flexibility.

(3) Support central Management.

Other requirements specified by the IETF RFC or the IETF drafts will be studied in a later version of this document.

10. Address Policy and IVI Address Evolution

Based on the IVI scheme, we propose to modify IPv4 address-allocation and IPv6 address-assignment policies [RFC1744] [RFC2008] for IPv4/IPv6 coexistence and transition as follows.

10.1. IPv6 Address Assignment Policy

(1) Reserve 2001:DB8:ff00::/40 for each 2001:DB8::/32 (2001:DB8::/32 is the documentation address, which represents all /32s [RFC4291]).

(2) Encourage ISPs to deploy their IPv6 networks and to install their IVI gateways.

(3) Encourage ISPs to use a subset (i.e. IVI4(i)) of their own IPv4 address blocks and map it into IPv6 via the IVI scheme (i.e. IVI6(i)) for their initial deployment of IPv6. For severs using the 1-to-1 mapping, and for clients using the 1-to-2^N mapping. In this way, the scarce IPv4 addresses can be effectively used. This special IPv6 block can communicate with the global IPv6 networks directly and communicate with the global IPv4 networks via IVI gateways.

(4) Encourage ISPs to increase the size of IVI4(i). When IVI4(i)=IPS4(i), the IPv4 to IPv6 transition for ISP(i) will be accomplished.

10.2. IPv4 Address Allocation Policy

(1) The remaining IPv4 address should be dedicated for the IVI transition use, i.e. using these blocks for the IVI6(i) deployment. The users using IVI6(i) can access the IPv6 networks directly and the IPv4 networks via the IVI gateways.

(2) Based on multiplexing techniques, the global IPv4 addresses can be used effectively. For example, with a reasonable port multiplexing ratio (say 16), one /8 can support 268M hosts. If 10 /8s can be allocated for the IVI use, it will be 2.6 billion addresses, possibly enough even for the unwired population in the world. The 43.0.0.0/8 could be a good candidate for the initial trial [APNIC].

10.3. Evolution of the IVI Addresses and Services

The IVI scheme is an effective method for transparent IPv4/IPv6 coexistence and smooth IPv4/IPv6 transition. Unlike the existing transition techniques which treat the IPv6 addresses equally [JSG2008], the IVI scheme suggests dividing the current IPv6 addresses into IVI6 addresses and non-IVI6 addresses. The IVI6

addresses, due to their nature as images of IVI4, can communicate with the global IPv4 networks via IVI gateways and they can also communicate with the global IPv6 networks directly. Therefore, the ISPs should use the IVI6 addresses for the initial deployment of their IPv6 infrastructure and this should be the IPv4/IPv6 coexistence stage. When IVI4(i)=IPS4(i) for most of ISP(i), the rest of the IPv6 addresses (non IVI6(i)) can be used for the further development of the global Internet, as shown in the following figure.

IPv4/IPv6 Address Coexistence and Evolution

```
-------------------------------------------------
              IPv4 area    |            IPv6 area
           ---------- ---  ----------------
--------------
   Service |    IPv4    |  |  IPv6 IVI    |   |IPv6 non-
IVI|
           --------------  ----------------
--------------
            |       \           /       | \     /        |
            |        \         /        |  \   /         |
            |         \  ---  /         |   \ /          |
            |          \|   |/          |   /\           |
   Network  |       - --|IVI|----       |  /  \          |
            |      |    |GW |    |      | |    |         |
            |      |   /|   |\   |      | |    |         |
            |      |  /  ---  \  |      | |    |         |
            |  /\  | |         | |  /\  | | /\ |    /\   |
            | /  \ | |         | | /  \ |/ /  \ \  /  \  |
          ---------------   ----------------
--------------
   User   |    IPv4     |   |  IPv6 IVI    |  |IPv6 non-
IVI|
          ---------------   ----------------
--------------
                  Coexistence    =>     Transition
```

Figure 12

11. Security Considerations

This document presents the prefix-specific and stateless address mapping scheme (IVI) for the IPv4/IPv6 coexistence and transition. The IPv4 security and IPv6 security issues should be addressed by related documents of each address family and are not included in this document.

However, the specific security issues for the IVI gateway implementation should be studied and addressed during the development of the IVI mechanisms.

12. IANA Considerations

The address allocation and assignment policies discussed in this document may have impact to IANA operation.

July 2008

13. Principal Authors

Xing Li

Maoke Chen

Congxiao Bao

Hong Zhang

Jianping Wu

14. Contributors

The authors would like to acknowledge the following contributors in the different phases of the IVI development: Ang Li, Yuncheng Zhu, Junxiu Lu and Yu Zhai.

The authors would like to acknowledge the following contributors who provided helpful inputs concerning the IVI concept: Bill Manning, David Ward, Lixia Zhang, Jun Murai and Fred Baker.

15. Acknowledgments

The authors thank to the funding supports of the CERNET, CNGI-CERNET2, CNGI Research and Development, China "863" and China "973" projects.

16. Appendix A. The IVI gateway configuration example

IVI Configuration Example

```
   #!/bin/bash
   # open forwarding
   echo 1 > /proc/sys/net/ipv6/conf/all/forwarding
   echo 1 > /proc/sys/net/ipv4/conf/all/forwarding

   # config route for IVI6 = 2001:da8:ffca:2661:cc00::/70,
   #                  IVI4 = 202.38.97.204/30

   # configure IPv6 route
   route add -A inet6 2001:da8:ffca:2661:cc00::/70 \
   gw 2001:da8:aaae::206 dev eth0

   # config mapping for      source-PF = 2001:da8::/32
   # config mapping for destination-PF = 2001:da8::/32

   # for each mapping, a unique pseudo-address (10.0.0.x/8)
   # should be configured.
   # ip addr add 10.0.0.1/8 dev eth0

   # IPv4-to-IPv6 mapping, multiple mappings can be done
via multiple
   # commands.
   # mroute IVI4-network IVI4-mask pseudo-address interface
\
   # source-PF destination-PF
   /root/mroute 202.38.97.204 255.255.255.252 10.0.0.1 \
   eth0 2001:da8:: 2001:da8::

   # IPv6-to-IPv4 mapping
```

```
# mroute6 destination-PF destination-PF-pref-len
/root/mroute6 2001:da8:ff00:: 40
```

Figure 13

17. Appendix B. The traceroute results

ivitraceroute

```
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
```

Figure 14

Note that the non-IVI IPv6 addresses are mapped to 202.38.17.186, which is defined in this document (the first two sections are the IPv4 prefix of /16 of the IVI gateway interface and the last two sections are the autonomous system number 4538).

ivitraceroute6

```
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_ivig=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
```

Figure 15

Note that all of the IPv4 addresses can be mapped to prefix-specific IPv6 addresses (for example 18.7.22.83 is mapped to 2001:da8:ff12:716:5300::).

18. References

18.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", RFC 1035, November 1987.

[RFC2008] Rekhter, Y. and T. Li, "Implications of Various Address Allocation Policies for Internet Routing", RFC 2008, October 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2765] Nordmark, E., "Stateless IP/ICMP Translation Algorithm (SIIT)", RFC 2765, February 2000.

[RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, Feb. 2000.

[RFC3056] Carpenter, B. and K. Moore, "Connection of ipv6 domains via ipv4 clouds", RFC 3056, February 2001.

[RFC3171] Albanna, Z., Almeroth, K., Meyer, D., and M. Schipper, "IANA Guidelines for IPv4 Multicast Address Assignments", RFC 3171, August 2001.

[RFC3956] Savola , P. and B. Haberman , "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, November 2004.

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, Feb. 2006.

[RFC4380] Huitema , C., "Teredo: Tunneling ipv6 over udp through network address translations (nats)", RFC 4380, Feb. 2006.

[RFC4607] Holbrook , H. and B. Cain , "Source-Specific Multicast for IP", RFC 4607, August 2006.

[RFC4611] McBride, M., Meylor, J., and D. Meyer, "Multicast Source

Discovery Protocol (MSDP) Deployment Scenarios", RFC 4611, August 2006.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing

(CIDR): The Internet Address Assignment and Aggregation
Plan", RFC 4632, August 2006.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site
Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214,
March 2008.

[RFC768] Postel , J., "User Datagram Protocol", RFC 768,
August 1981.

[RFC791] Postel, J., "Internet Protocol", RFC 791, September 1981.

[RFC793] Postel, J., "Transmission Control Protocol", RFC 793,
Semptember 1981.

18.2. Informative References

[APNIC] Ito, K., "Large IPv4 address space Usage trial for Future
IPv6 Deployment", http://www.apnic.net/meetings/25/
program/policy/ito-large-ipv4-trial.pdf .

[CERNET] "CERNET Homepage:
http://www.edu.cn/english_1369/index.shtml".

[CERNET2] "CNGI-CERNET2 Homepage:
http://www.cernet2.edu.cn/index_en.htm".

[COUNT] "IPv4 address count down: http://penrose.uk6x.com/".

[I-D.bagnulo-behave-nat64]
Bagnulo, M., Matthews, P., and I. van Beijnum, "NAT64/
DNS64: Network Address and Protocol Translation from IPv6
Clients to IPv4 Servers",

draft-bagnulo-behave-nat64-00 (work in progress),
June 2008.

[I-D.v6ops-nat64-pb-statement-req]
Bagnulo, M., Baker, F., and I. van Beijnum, "IPv4/IPv6
Coexistence and Transition: Requirements for solutions",
draft-ietf-v6ops-nat64-pb-statement-req-00 (work in
progress), May 2008.

[IVI4] "Test homepage for the IVI4(i): http://202.38.114.1/".

[IVI6] "Test homepage for the IVI6(i):
http://[2001:250:ffca:2672:0100::0]/".

[JJI07] Joseph, D., Chuang, J., and I. Stocia, "Modeling the
Adoption of new Network Architectures", EECS Department,
University of California, Berkeley Tech. Rep. UCB/
EECS-2007-41, April 2007.

[JSG2008] "A Report of Japaness Study Group on Internet's Smooth
Transition to IPv6:
http://www.soumu.go.jp/joho_tsusin/eng/pdf/080617_1.pdf",
June 2008.

[LINUX] "Source Code of the IVI implementation for Linux:
http://linux.ivi2.org/impl/".

[MVB98] Fiuczynski, M., Lam, V., and B. Bershad , "The design and
implementation of an ipv6/ipv4 network address and
protocol translator", Proceedings of the USENIX Annual
Technical Conference (NO 98), June 1998.

[RFC1744] Huston, G., "Observations on the Management of the
Internet Address Space", RFC 1744, December 1994.

[RFC2775] Carpenter, B., "Internet Transparency", RFC 2775,
Feb. 2000.

[RFC3142] Hagino, J. and K. Yamamoto, "An IPv6-to-IPv4 Transport
Relay Translator", RFC 3142, June 2001.

[RFC3569] Bhattacharyya, S., "An Overview of Source-Specific
Multicast (SSM)", RFC 3569, July 2003.

[RFC4925] Li, X., Dawkins, S., Ward, D., and A. Durand, "Softwire
Problem Statement", RFC 4925, July 2007.

[RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network
Address Translator - Protocol Translator (NAT-PT) to
Historic Status", RFC 4966, July 2007.

Authors' Addresses

Xing Li
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 62785983
Email: xing@cernet.edu.cn

Maoke Chen
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 62785983
Email: mk@cernet.edu.cn

Congxiao Bao
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 62785983

Email: congxiao@cernet.edu.cn

Hong Zhang
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 62785983
Email: neilzh@gmail.com

Jianping Wu
CERNET Center/Tsinghua University
Room 225, Main Building, Tsinghua University
Beijing 100084
CN

Phone: +86 62785983
Email: jianping@cernet.edu.cn

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

EXHIBIT 2

Prefix-specific and Stateless Address Mapping (IVI) for IPv4/IPv6 Coexistence and Transition
draft-xli-behave-ivi-00

Xing Li, Maoke Chen, Congxiao Bao,
Hong Zhang and Jianping Wu

IETF-72, Dublin, behave, 29 July 2008

Outline

- Introduction
- IVI scheme
- Design considerations
- Testing result
- Transition
- Address Policy
- Conclusions

Introduction

- The experiences for the IPv6 deployment in the past 10 years strongly indicate that the IPv6 hosts need to communicate with the global IPv4 networks.
- In this document, we follow the basic specification of SIIT, but we define the address assignment and routing scheme (IVI).
  - It is stateless (or almost stateless) in both the IPv4-to-IPv6 mapping direction, as well as in the IPv6-to-IPv4 mapping direction
  - It supports address transparency.
  - It supports both IPv6 initiated communication and the IPv4 initiated communication without using NAT-traversal techniques.

IVI Scheme

- The IVI is a prefix-specific and explicit bidirectional 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 multiplexing techniques can be used.
- Ref:
  - http://www.ietf.org/internet-drafts/draft-xli-behave-ivi-00.txt

Terms and Abbreviations of IVI

- **General**
  - **IVI.**
  - **ISP(i)**
- **IPv4**
  - **IPG4:** An address set containing all IPv4 addresses, the addresses in this set are mainly used by IPv4 hosts at the current stage.
  - **IPS4(i):** A subset of IPG4 allocated to ISP(i).
  - **IVI4(i):** A subset of IPS4(i), the addresses in this set will be mapped to IPv6 via IVI rule and physically used by IPv6 hosts of ISP(i).
- **IPv6**
  - **IPG6:** An address set containing all IPv6 addresses.
  - **IPS6(i):** A subset of IPG6 allocated to ISP(i).
  - **IVIG46(i):** A subset of IPS6(i), an image of IPG4 in IPv6 address family via IVI mapping rule.
  - **IVI6(i):** A subset of IVIG46(i), an image of IVI4(i) in IPv6 address family via IVI mapping rule.
- **Components**
  - **IVI gateway**
  - **IVI DNS**

Address Mapping (1)

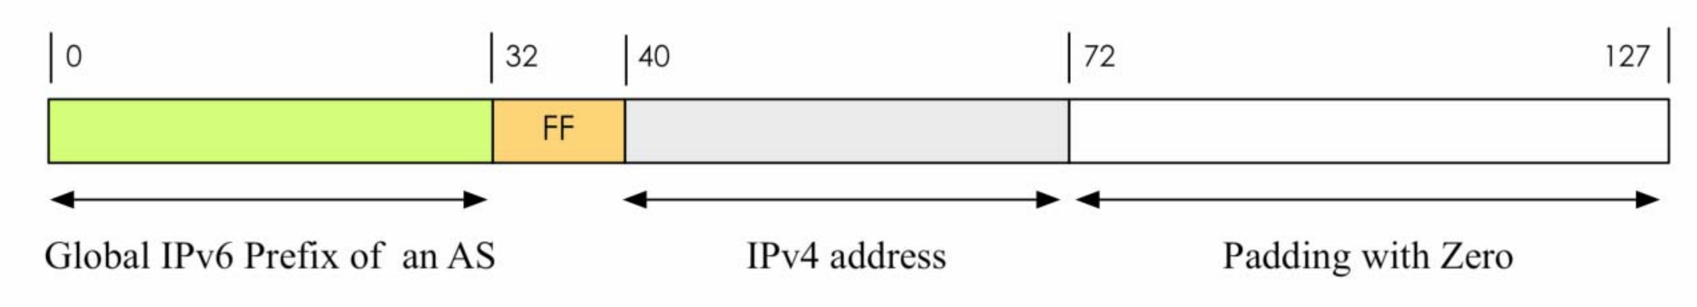

Mapping Rule: IPv4 addresses are embedded from bit 40 to bit 72 of the IPv6 addresses of a specific /32.

Example:

ISP's IPv6 /32 (ISP6)	2001:250::/32
image of global IPv4 (IVIG46):	2001:250:ff00::/40
borrowed IPv4 address (IVI4):	202.38.108.0/24
mapped IVI IPv6 address (IVI6):	2001:250:ffca:266c::/64

Address Mapping (2)

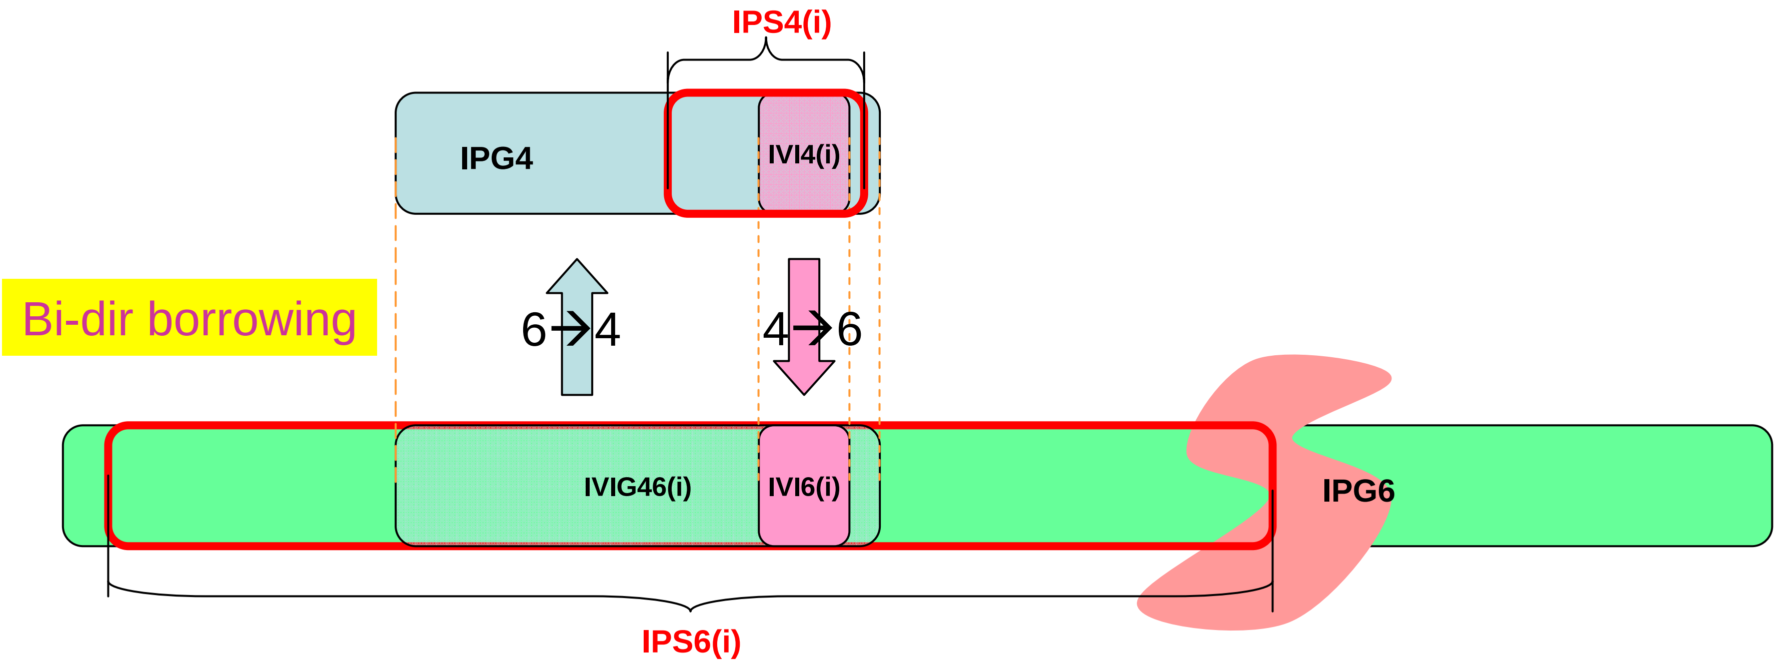

Address Mapping (3)

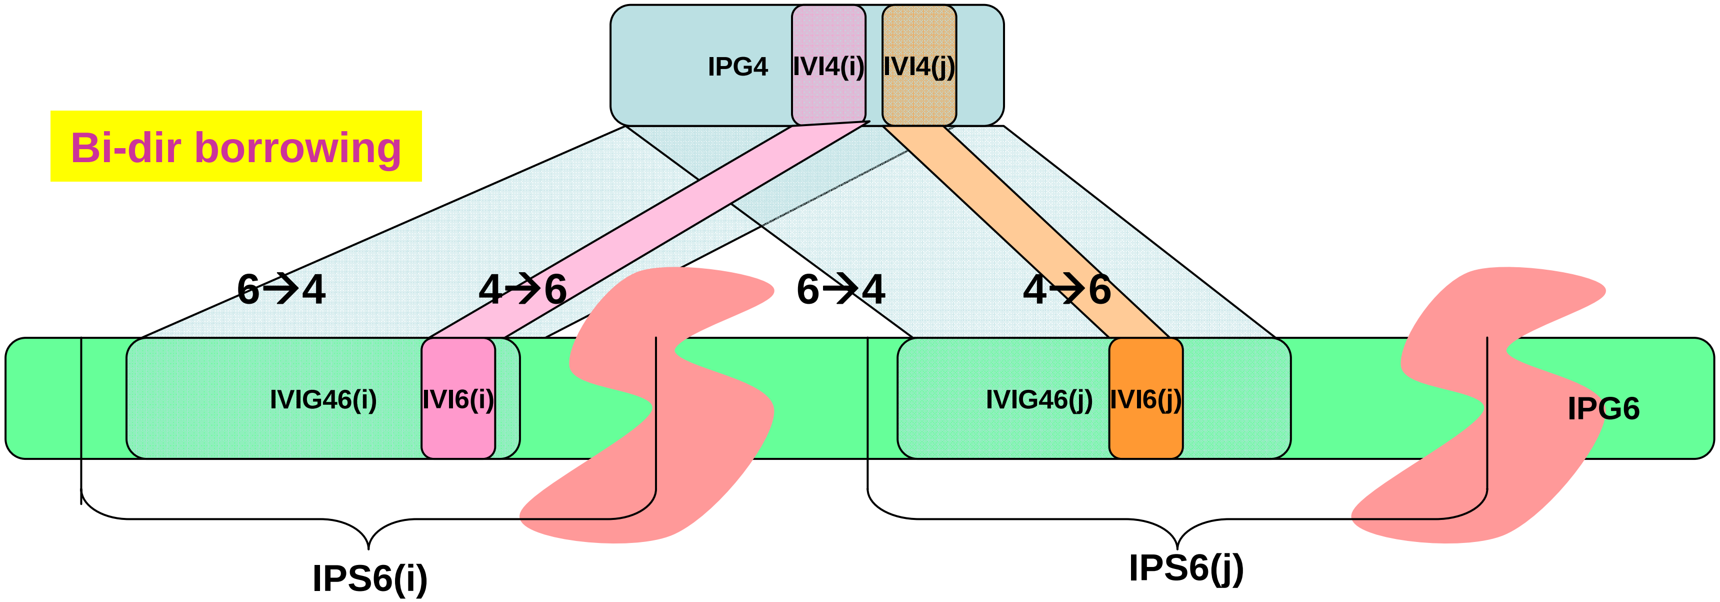

Routing and Forwarding

Routing and mapping configuration example

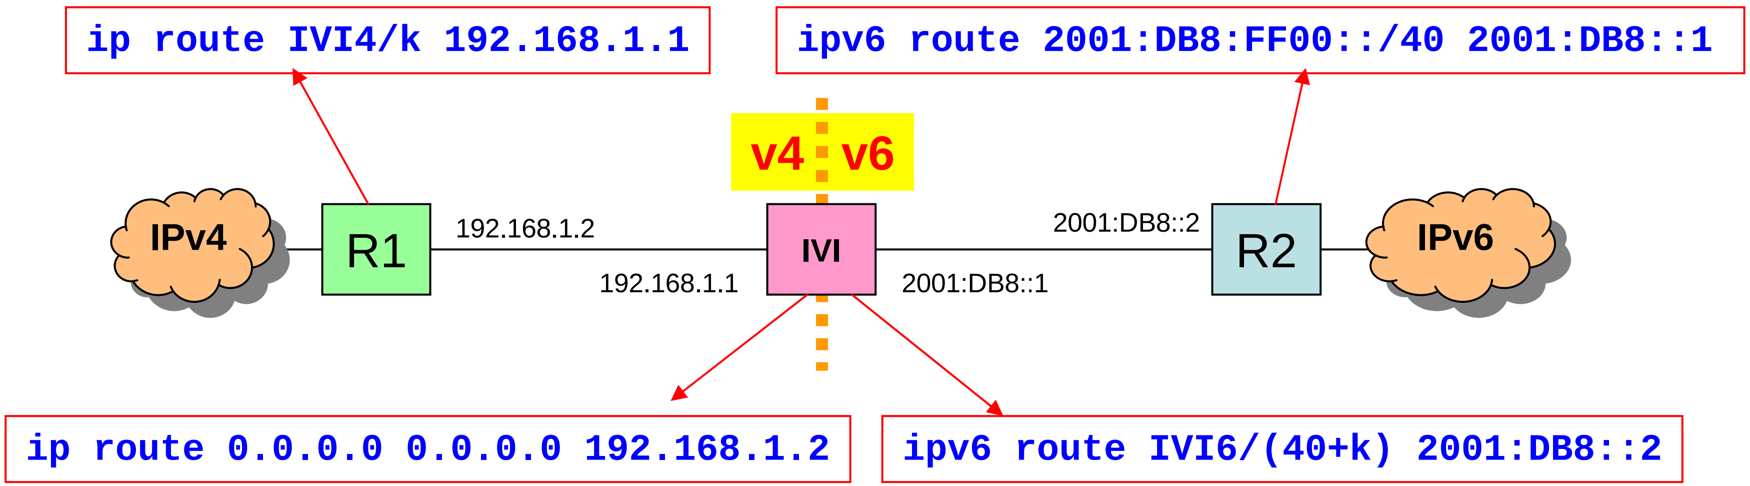


```
mroute IVI4-network IVI4-mask pseudo-address interface source-PF destination-PF
mroute6 destination-PF destination-PF-pref-len
```

IVI Reachability Matrix

	IPG4	IVI	IPG6
IPG4	**OK**	**OK**	**NO**
IVI	**OK**	**OK**	**OK**
IPG6	**NO**	**OK**	**OK**

IVI Communication Scenarios (1)

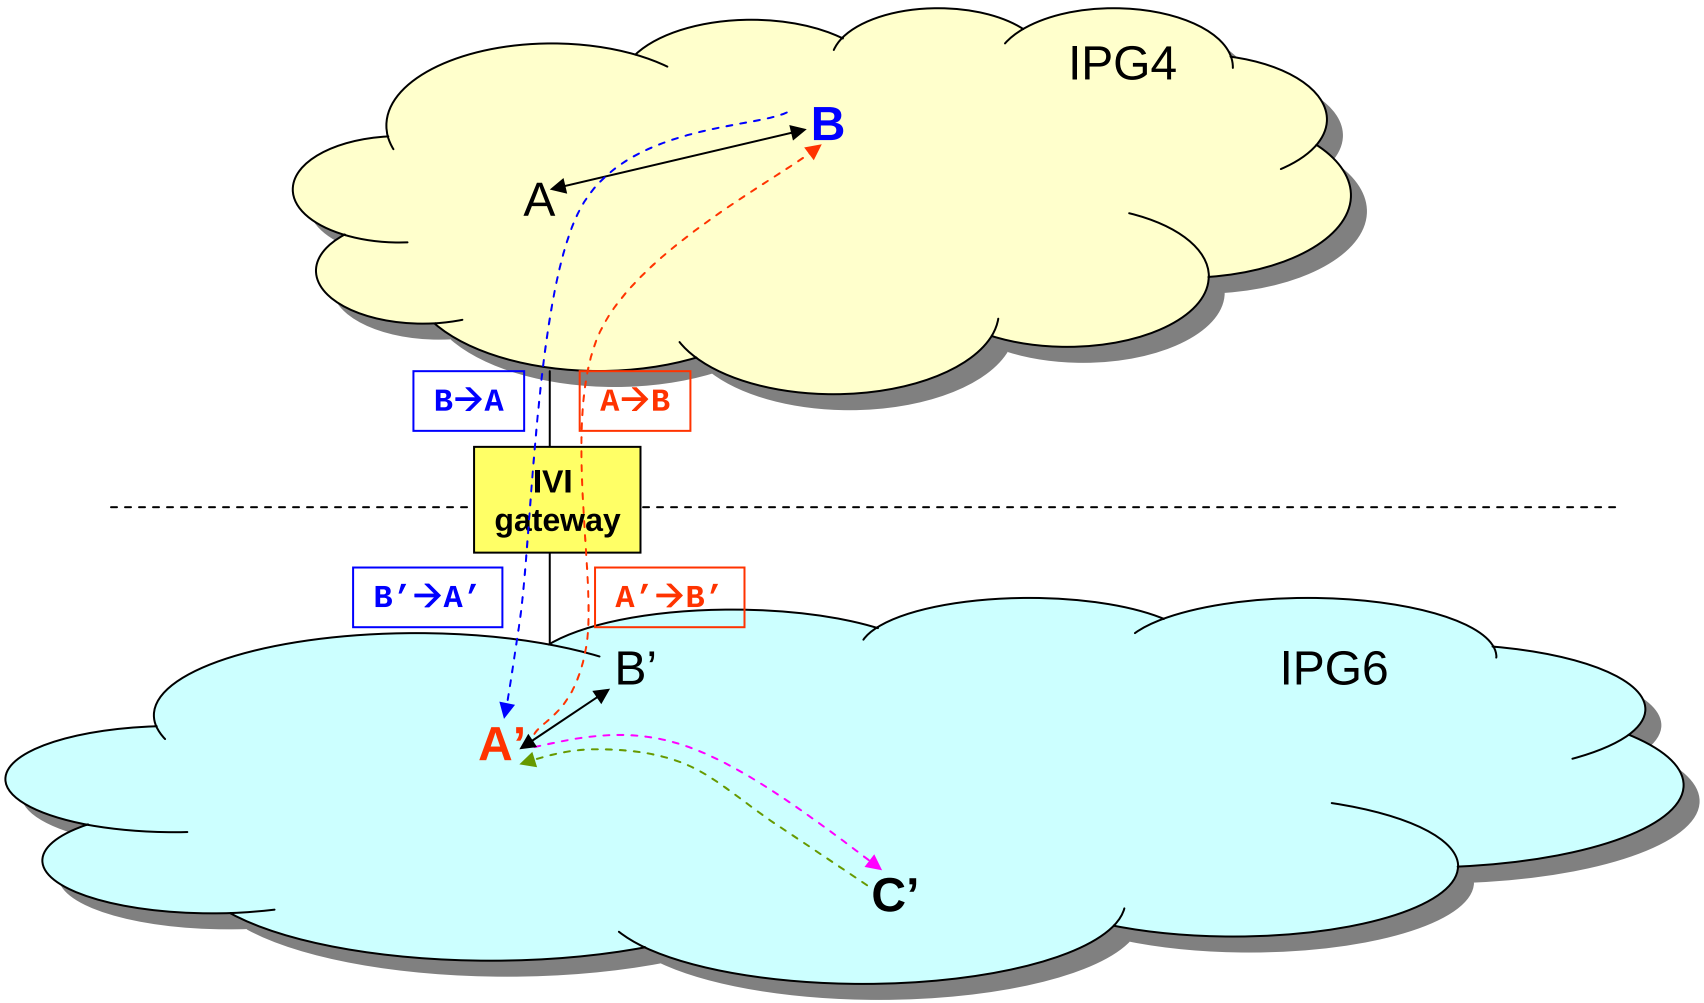

- A' ←→ B
- A' ←→ C'

IVI Communication Scenarios (2)

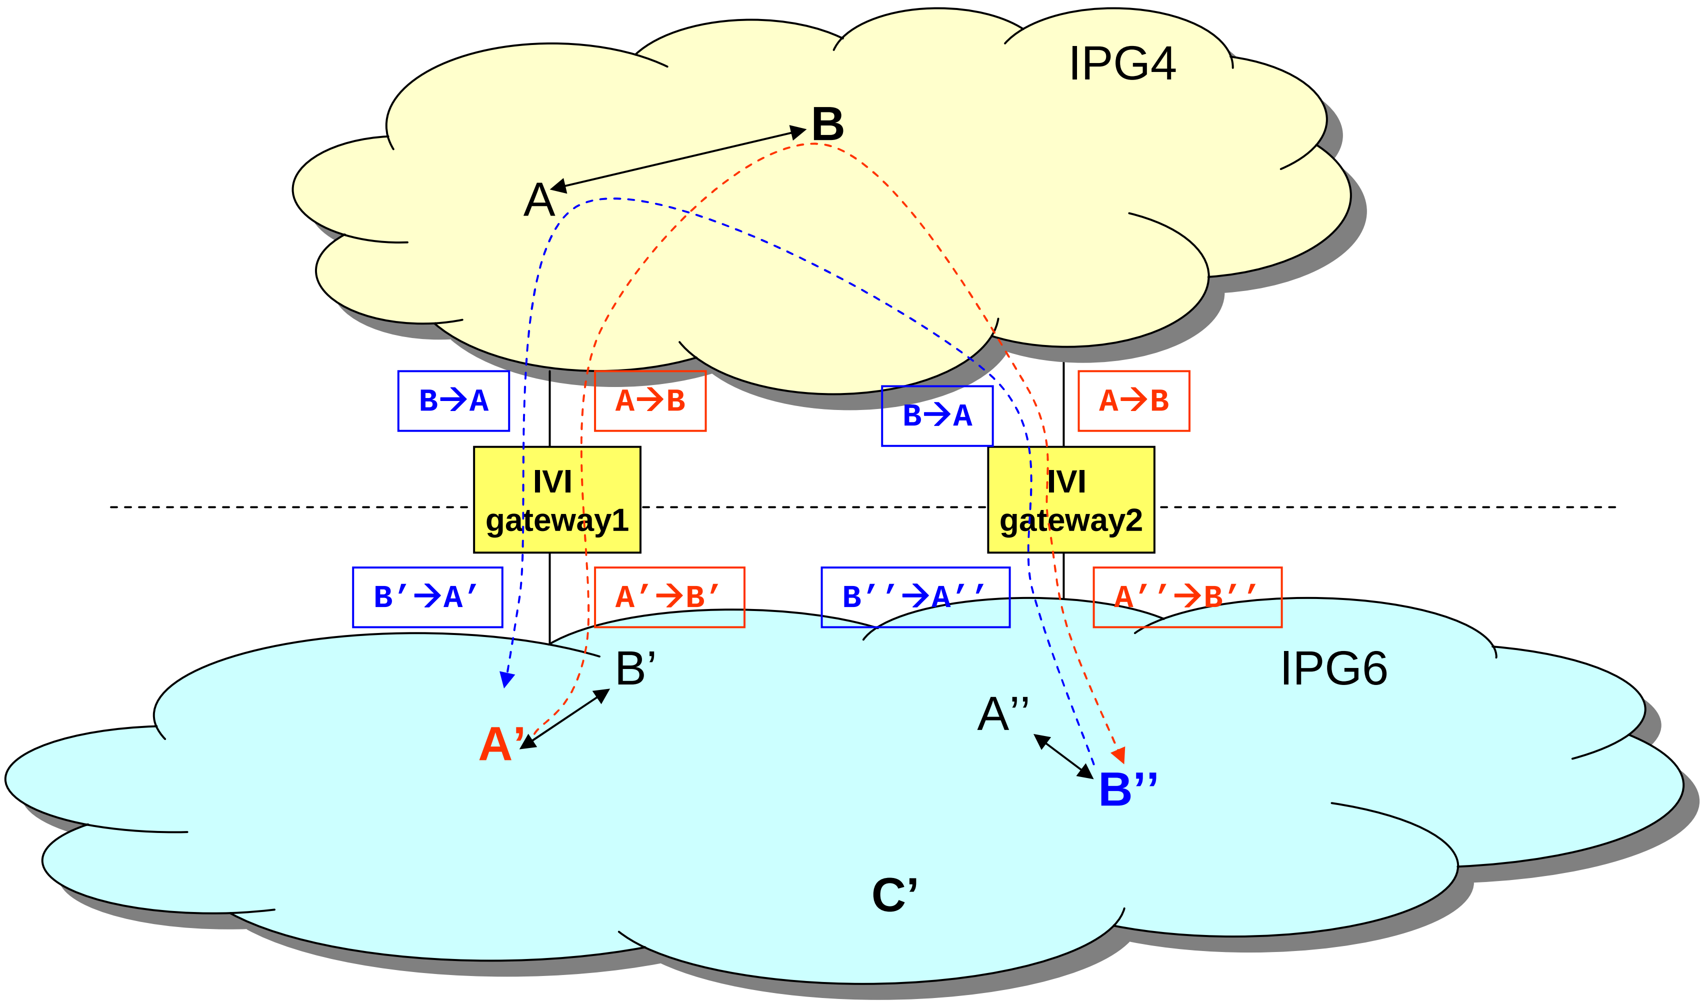

- A' ←(B&A)→ B''

IVI Communication Scenarios (3)

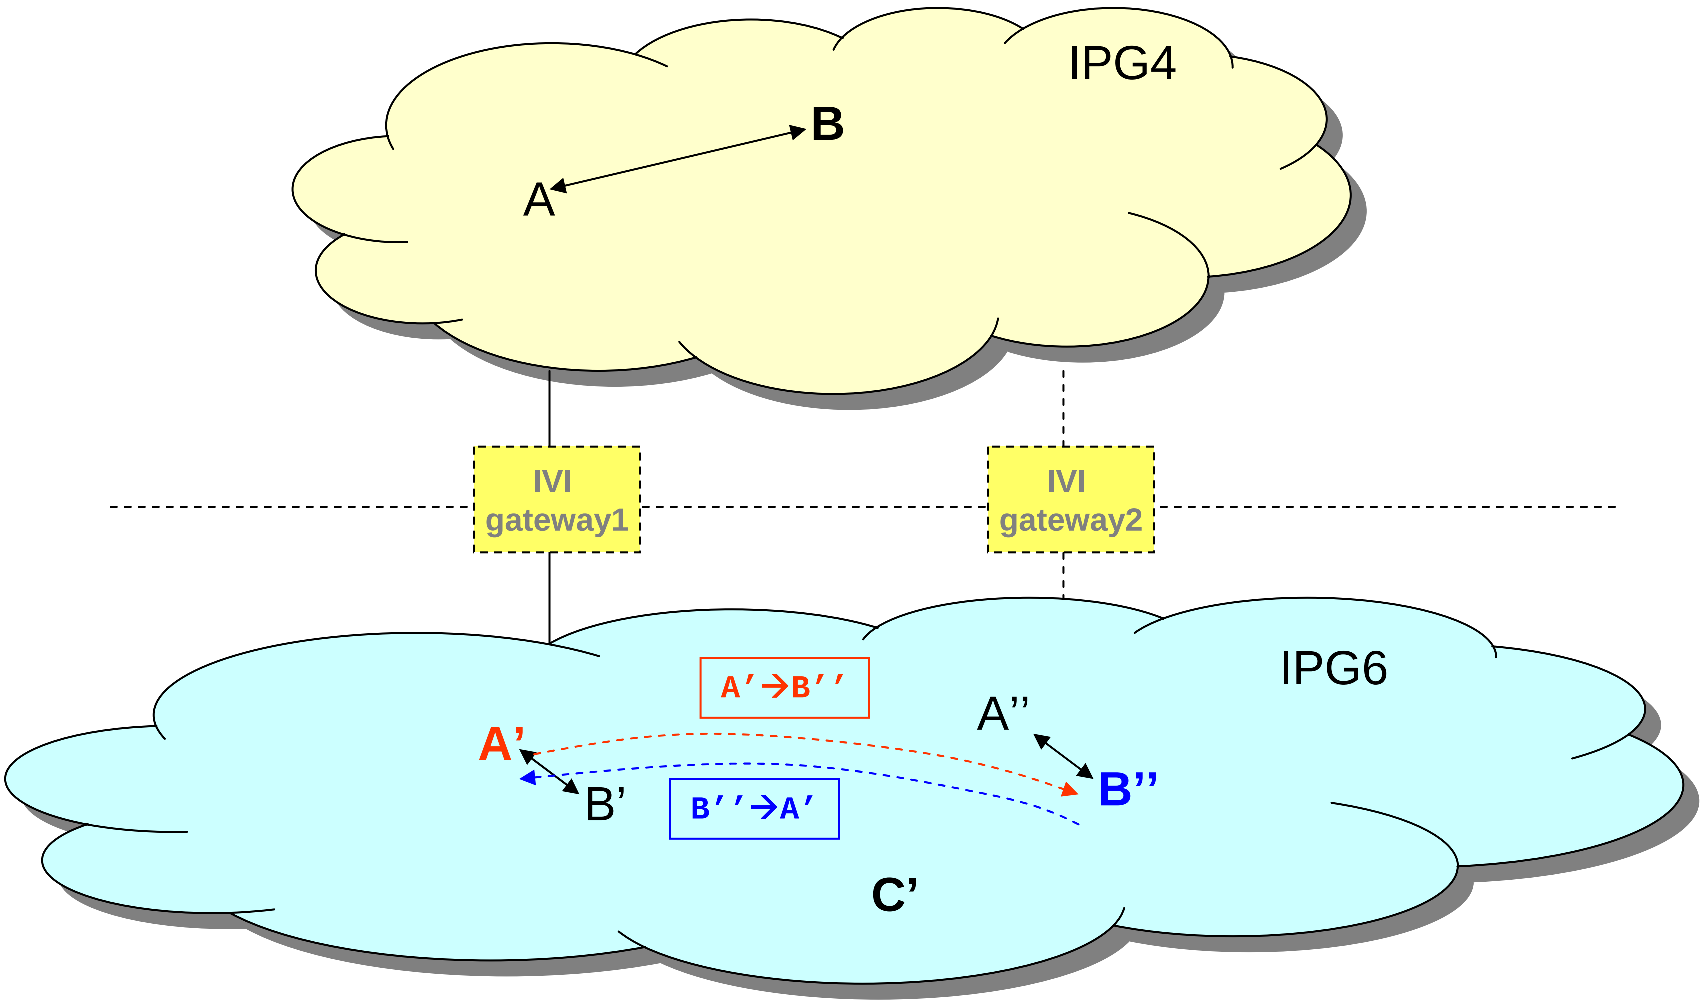

- A' ←→ B''

Design Considerations

- Address Mapping (general)
- Network-layer Header Translation (SIIT)
- Transport-layer Header Translation (SIIT)
- Fragmentation and MTU Handling (SIIT)
- ICMP Handling (SIIT + extension)
- Application Layer Gateway (SIIT)
- IPv6 Source Address Selection
- IPv4 over IPv6 Support
- IVI DNS
- Multiplexing of the Global IPv4 Addresses
- Multicast support

Address Mapping (general)

- 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**

ICMP + Extension

- The ICMP message may be generated by an intermediate router whose IPv6 address does not belong to IVIG46(i). Since ICMP translation is important to the path MTU discovery, the inverse mapping for unmapped addresses is defined in this document.
- In the current prototype, a pseudo IPv4 address is generated
  - First 16 bits are the IPv4 address of the IVI gateway
  - The last 16 bits are the AS number of the current domain. This prevents translated ICMP messages from being discarded due to unknown or private IP source.
- A small IPv4 address block should be reserved to identify the non-IVI mapped IPv6 addresses.
  - Similar to 4-byte AS AS23456

IPv6 Source Address Selection

- Since each IPv6 host may have multiple addresses, it is important for the host to use an IVI6(i) address to reach the global IPv4 networks.
  - The short-term work around is to use IVI6(i) as the default IPv6 address of the host.
  - The long-term solution requires that the application be able to select the source addresses for different services.
- IVI6 address configuration
  - DHCPv6 is required

IPv4 over IPv6 Support

- The IVI scheme can support the IPv4 over IPv6 service (NAT6-4-6), i.e. a stub IPv4 network can be connected to an IVI gateway to reach the IPv6 network and via another IVI gateway to reach the global IPv4 network
- A more interesting scenario is to integrate the functions of the first IVI gateway into the end-system. In this case, the application software are IPv4-based and there is no need to have ALG support in the IVI gateway when it is communicating with IPv4 hosts.

DNS Configuration and Mapping

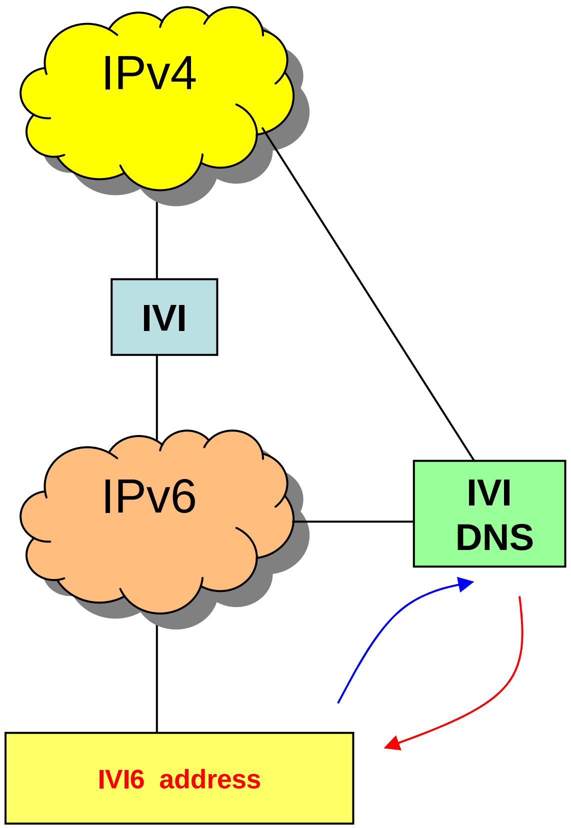

- For providing primary DNS service for IVI4(i) and IVI6(i), each host will have both A and AAAA records
- Authoritative DNS server
  - Example
  - www.ivi2.org A 202.38.108.2
  - www.ivi2.org AAAA 2001:250:ffca:266c:200::
- For resolving IVIG46(i) for IVI6(i), use IVI DNS to do the dynamic mapping based 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::
- Implementation scope
  - Host
  - DNS server provided via DHCPv6
  - ISP

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
  - Home server 1:M mapping (via IPv4 initiated communication)
  - Client 1:N mapping (via IPv6 initiated communication)
- Multiplexing using IPv4 NAT-PT
  - Cascade IPv4 NAT-PT and IVI (1:1 mapping)

Port multiplexing – IPv6 initiated

- Example:
  - **202.38.108.5#100 ←→ 2001:250:ffca:266c:0500::81#100**
  - **202.38.108.5#101 ←→ 2001:250:ffca:266c:0500::82#100**
  - **202.38.108.5#102 ←→ 2001:250:ffca:266c:0500::83#100**
  - **202.38.108.5#103 ←→ 2001:250:ffca:266c:0500::84#100**
- In the case of port collision, map to an unused port.

Port multiplexing – IPv4 initiated

- The remote IPv4 host can reach different IVI6s via different port number (pseudo-well-known port number)
  - **202.38.108.2#81 --> IVI61=2001:250:ffca:266c:0200::81#81**
  - **202.38.108.2#82 --> IVI61=2001:250:ffca:266c:0200::82#82**
  - **202.38.108.2#83 --> IVI61=2001:250:ffca:266c:0200::83#83**
  - **202.38.108.2#84 --> IVI61=2001:250:ffca:266c:0200::84#84**
- This can be provided via SRV DNS record.

Multicast support

- SSM is supported for the IVI
  - no MSDP in IPv6
  - no embedded RP in IPv4
  - It is also possible to build a gateway for ASM
- Group address mapping rule (there will be 2^{24} group ID available)
  - **`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`**
- For the cross address family SSM
  - the source address in IPv6 has to be IVI6 for the RPF scheme
- The inter operation of PIM-SM in IPv4 and IPv6
  - Application layer gateway
  - Static join using IGMPv3 and MLDv2

IVI Deployment Scenarios (1)

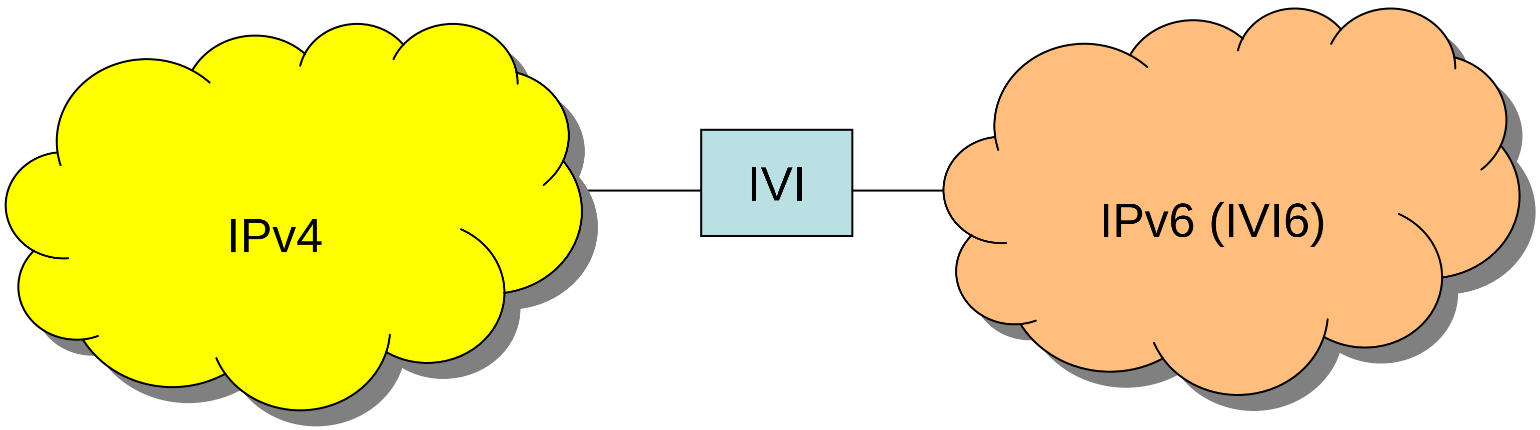

IVI Deployment Scenarios (2)

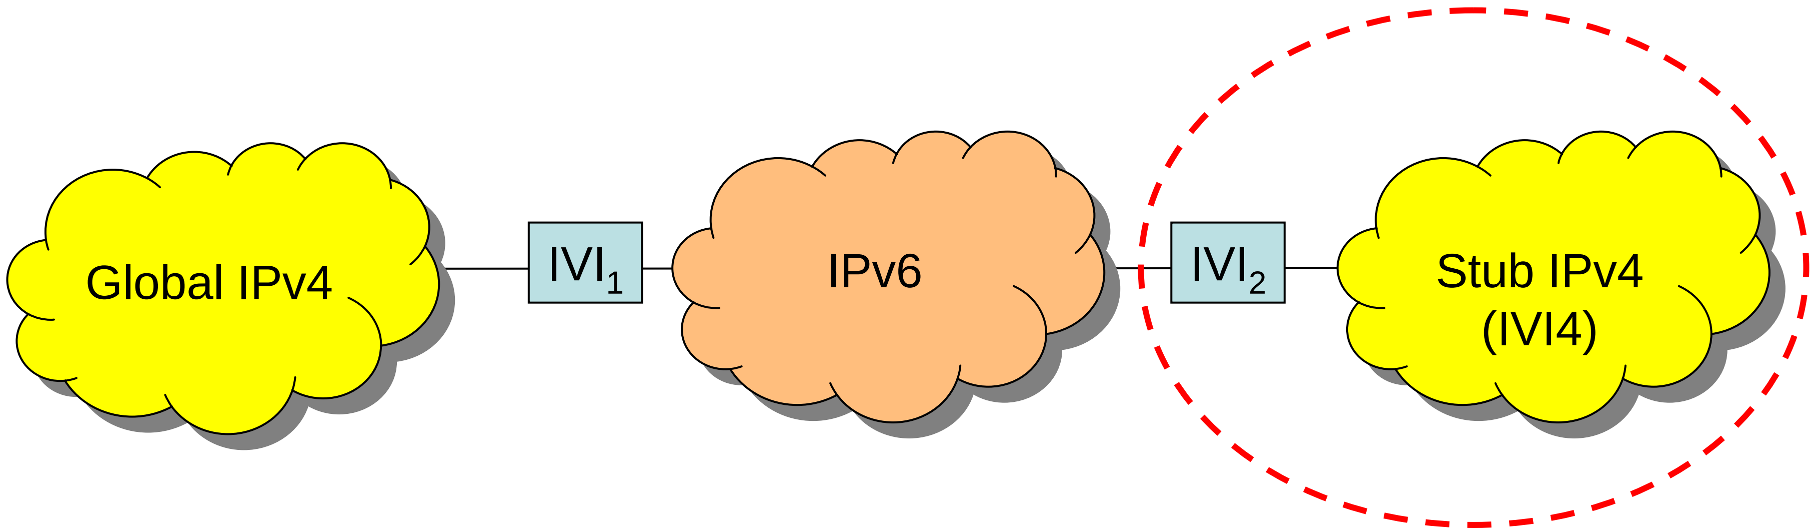

IVI Deployment Scenarios (3)

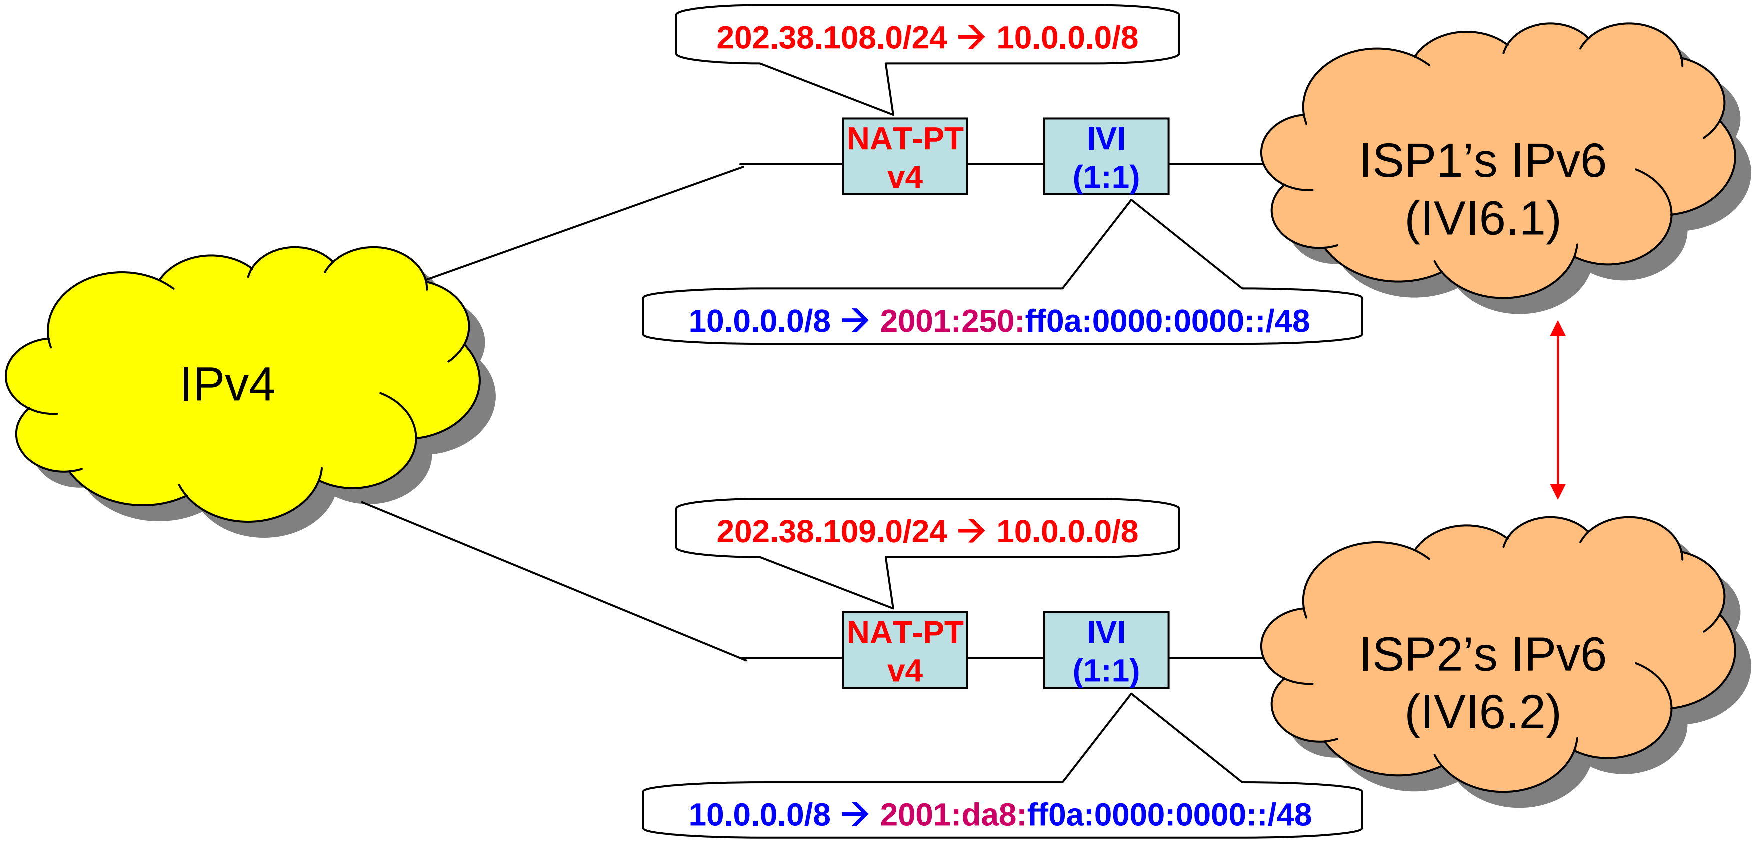

IVI Deployment Scenarios (4)

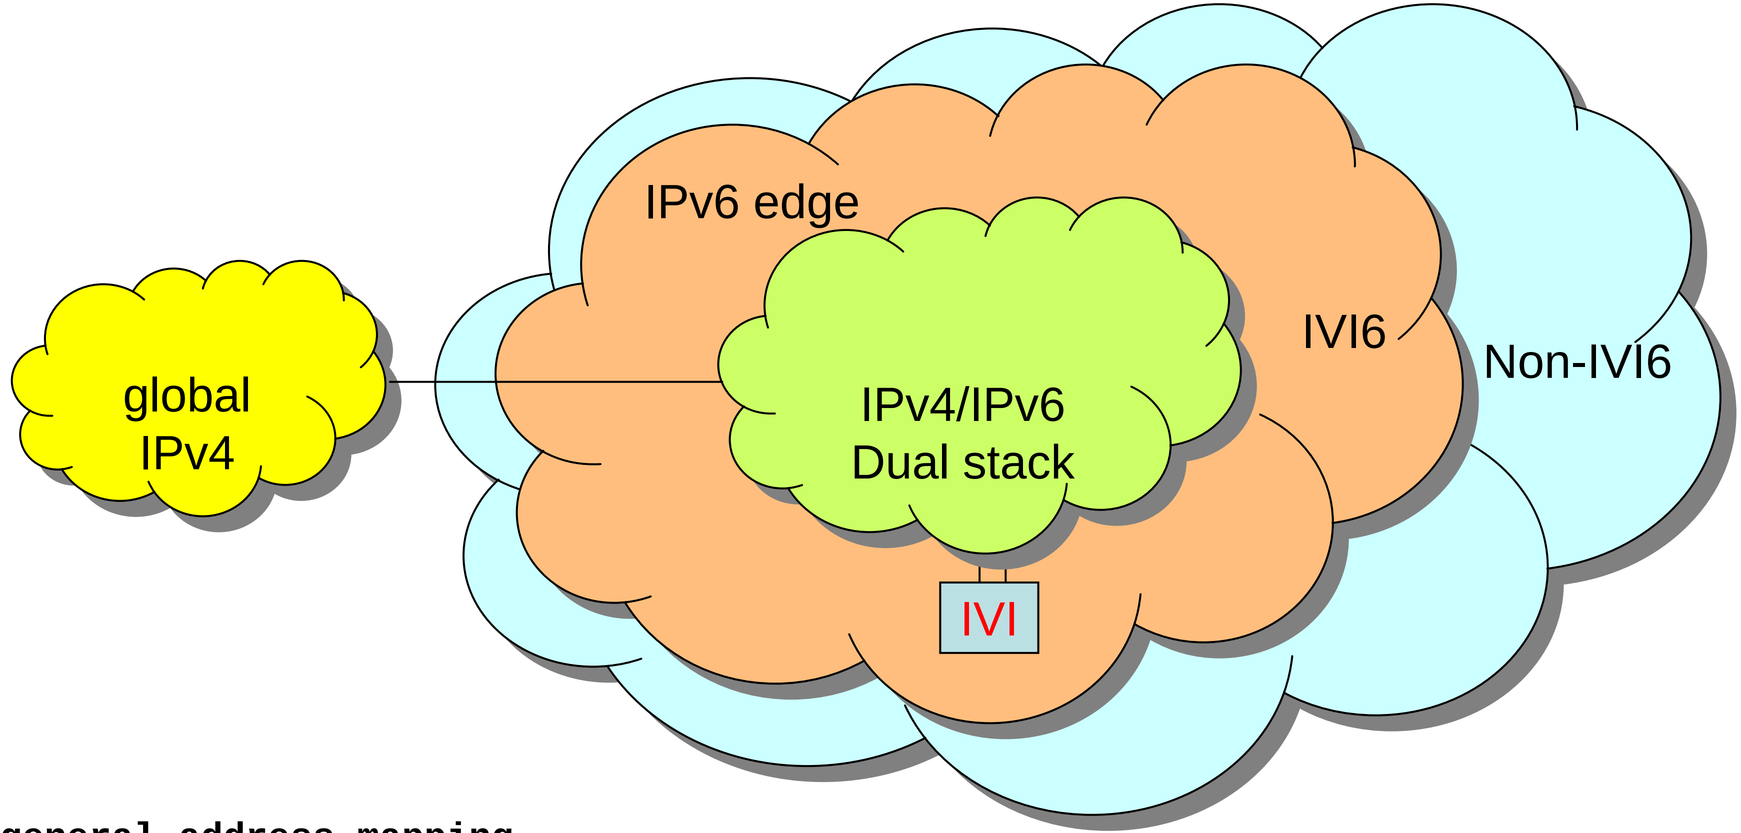

IVI general address mapping

2001:DB8:FF00::/40	**backbone scope (implemented)**
2001:DB8:FFFF::/48,	**site scope**
2001:DB8:ABCD:FF00::/56	**sub-site scope**
2001:DB8:ABCD:FFFF::/64	**subnet scope**
2001:DB8:XXXX:XXXX:XXXX:XXXX::/96	**IPv4 mapped alike scope**

Implementation and Preliminary Testing Results

- The IVI scheme presented in this document is implemented in the Linux OS
  - The source code can be downloaded [http://202.38.114.1/impl/].
- CERNET (IPv4 and partially dual-stack) and CNGI-CERNET2 (pure IPv6) since March 2006 (basic implementation).
  - IVI6 server for global IPv4
    - http://202.38.114.1/
  - IVI6 server for global IPv6
    - http://[2001:250:ffca:2672:0100::0]/
  - IVI server for stub IPv4
    - http://202.38.114.129/

IVI Hosts Installation in CNGI-CERNET2

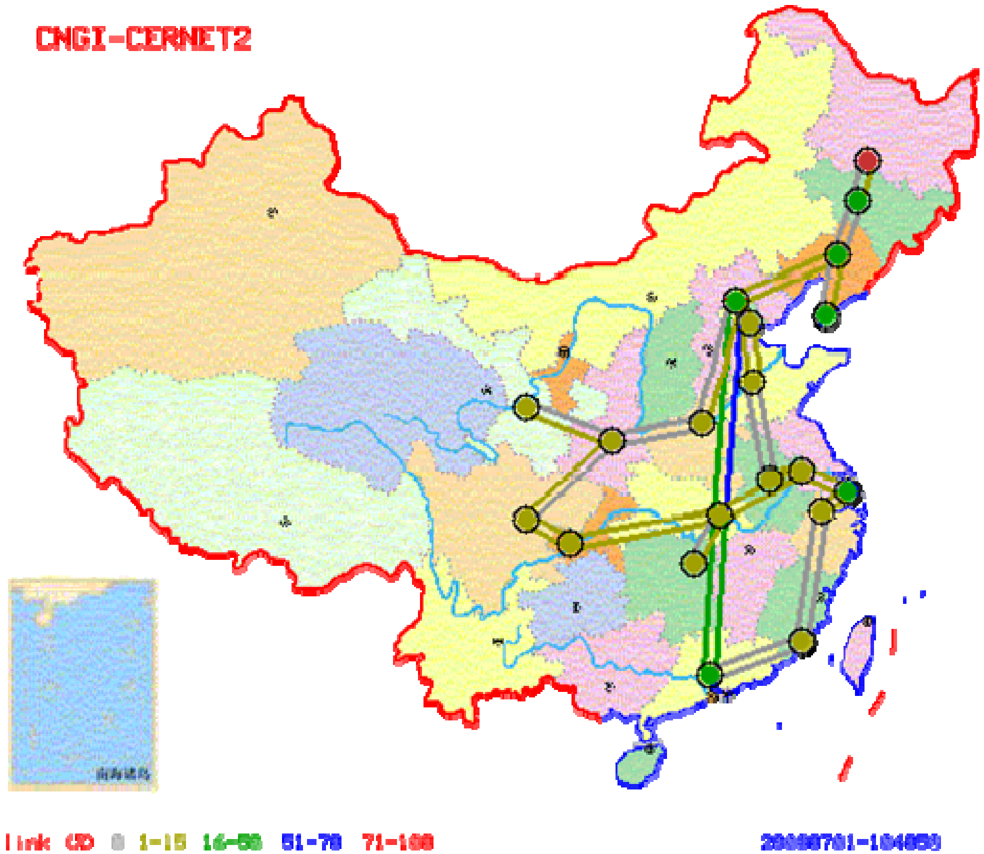

From IVI6 host traceroute6 IVIG46

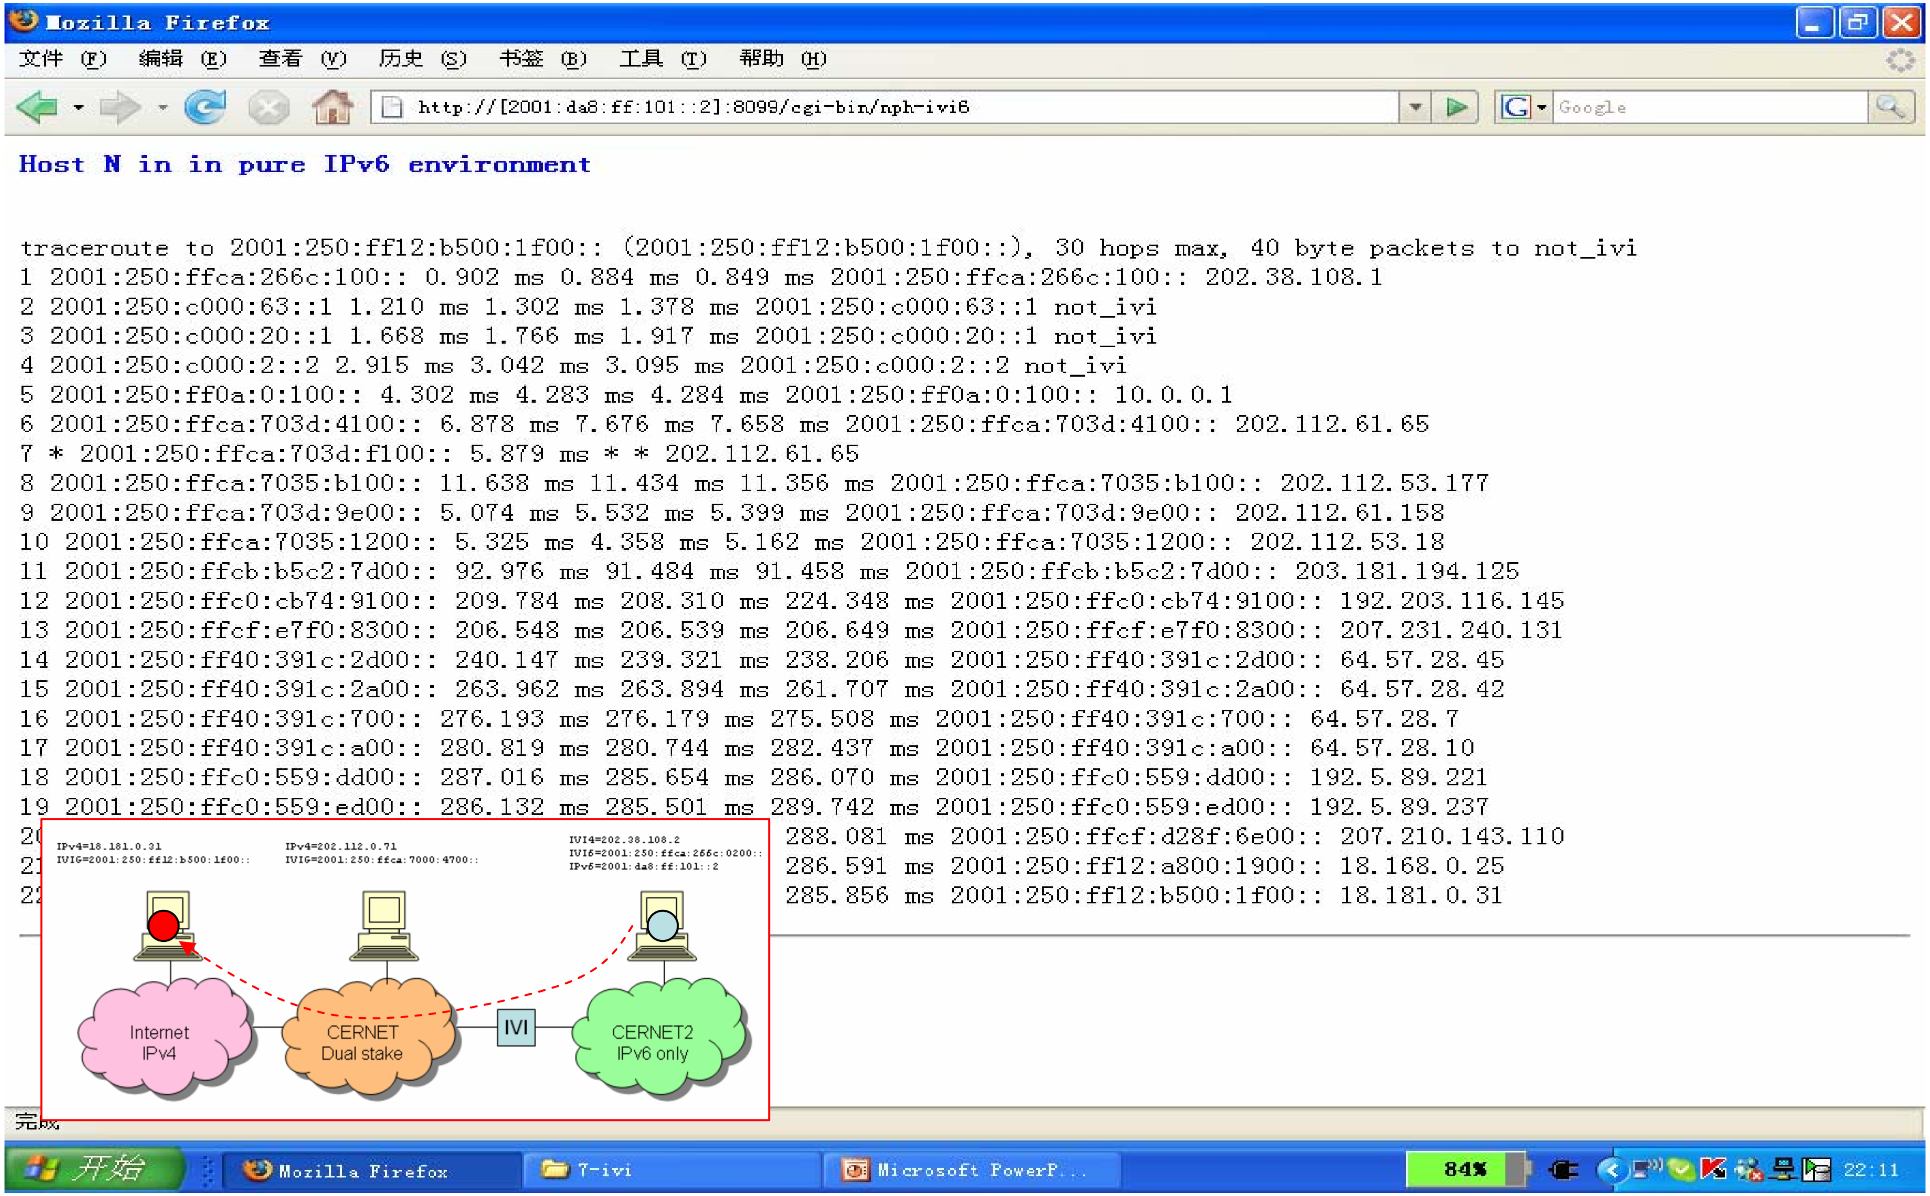

From IPv4 host traceroute IVI4

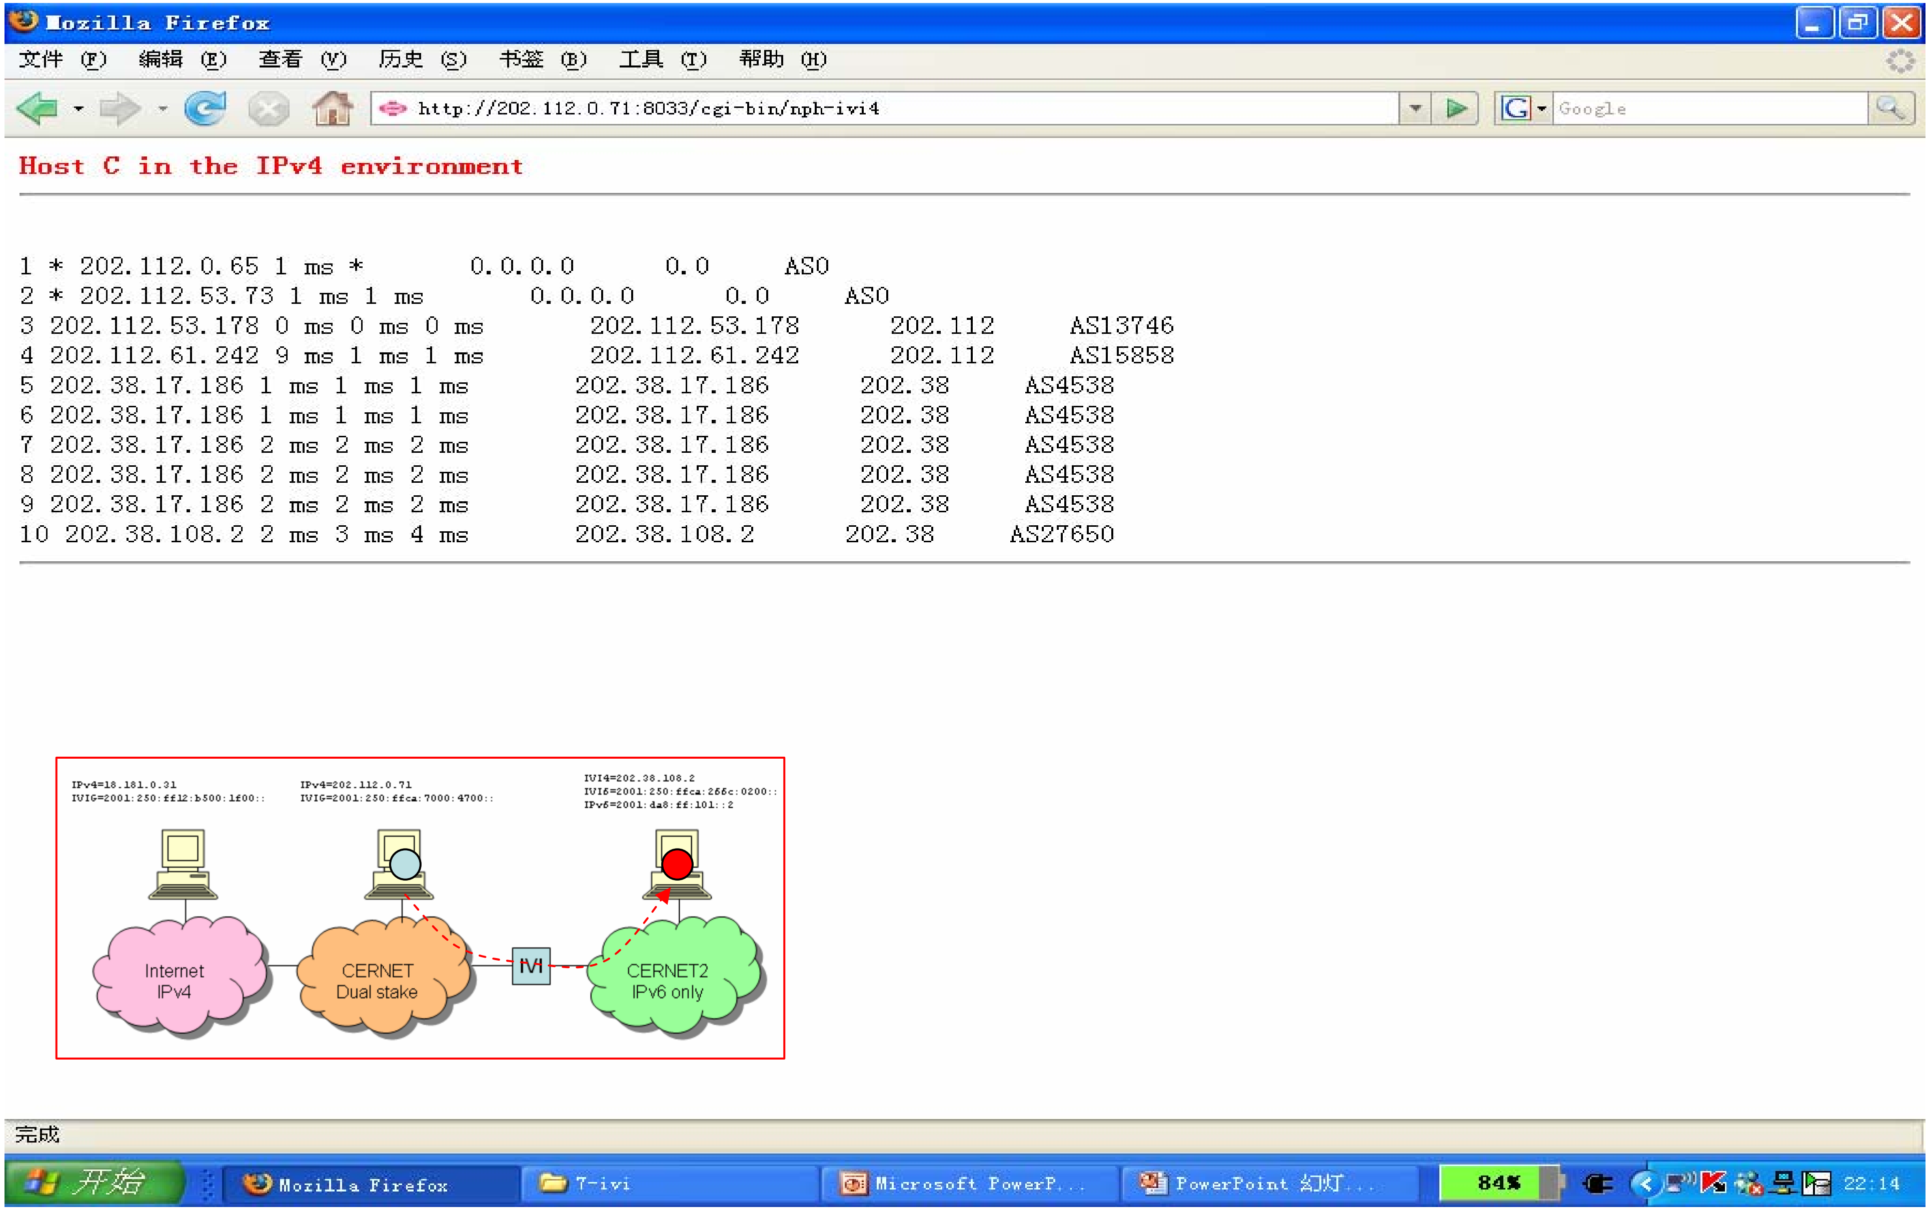

Features of IVI

1. No need to change the end system (IPv4 and IPv6).
2. Support v4-initiated and v6-initiated communications.
3. Support interaction with dual-stack hosts.
4. The standard IPv4 NAT can easily be integrated into the system.
5. Do not violate standard DNS semantics.
6. No affect to both IPv4 and IPv6 routing.
7. Support TCP, UDP, ICMP
8. Can handle fragmentation.
9. Support incremental deployment
10. Support multicast (SSM)

Address Policy and IVI Address Evolution

- IPv6 Address Assignment Policy
- IPv4 Address Allocation Policy
- Evolution of the IVI Addresses and Services

IPv6 Address Assignment Policy

- Encourage ISPs to deploy their IPv6 networks and to install their IVI gateways.
  - Reserve 2001:DB8:ff00::/40 for each 2001:DB8::/32
  - Encourage ISPs to use a subset (i.e. IVI4(i)) of their own IPv4 address blocks and map it into IPv6 via the IVI scheme (i.e. IVI6(i)) for their initial deployment of IPv6.
    - For severs using the 1-to-1 mapping, and for clients using the 1-to-2^N mapping.
    - In this way, the scarce IPv4 addresses can be effectively used.
    - This IVI6 can communicate with the global IPv6 networks directly and communicate with the global IPv4 networks via IVI gateways.
- Encourage ISPs to increase the size of IVI4(i). When IVI4(i)=IPS4(i), the IPv4 to IPv6 transition for ISP(i) will be accomplished.

IPv4 Address Allocation Policy

- The remaining IPv4 address should be dedicated for the IVI transition use, i.e. using these blocks for the IVI6(i) deployment.
  - The users using IVI6(i) can access the IPv6 networks directly and the IPv4 networks via the IVI gateways.
- Based on multiplexing techniques, the global IPv4 addresses can be used effectively.
  - For example, with a reasonable port multiplexing ratio (say 16), one /8 can support 268M hosts. If 10 /8s can be allocated for the IVI use, it will be 2.6 billion addresses, possibly enough even for the unwired population in the world.
- The 43.0.0.0/8 could be a good candidate for the initial trial

Measures for address space exhaustion

1. For continuous development of the Internet since 2011, the combination of **the transition to a new address system (IPv6)** and **sharing of one address by multiple users (using NAT/NAPT)** must be performed from three viewpoints of *feasibility within a time limit, continuity of service on the Internet, and continuance of effect,*
2. It is appropriate to **introduce the measures in three stages**: before exhaustion, early, and middle stages of exhaustion.

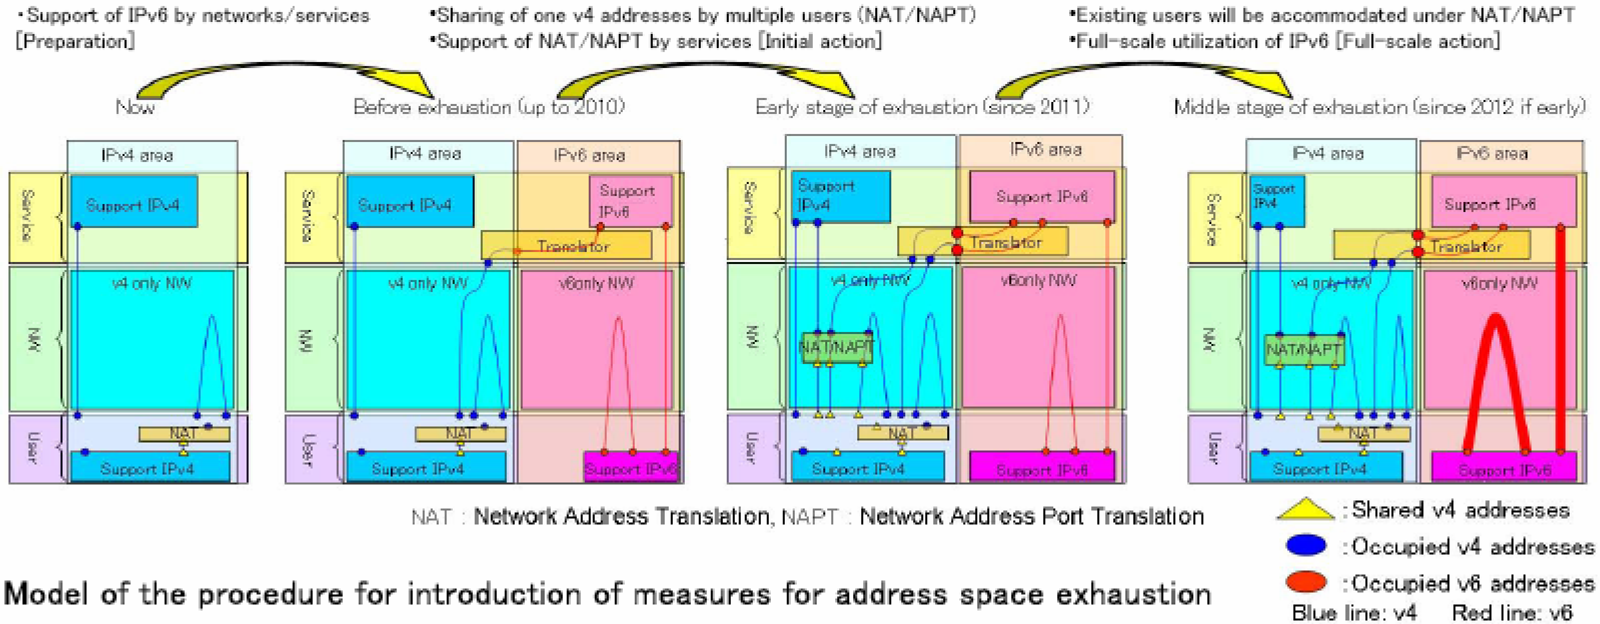

Model of the procedure for introduction of measures for address space exhaustion

6

From the June 2008 Report of the Japanese Study Group on Internet's Smooth Transition to IPv6

Evolution of the IVI Addresses and Services

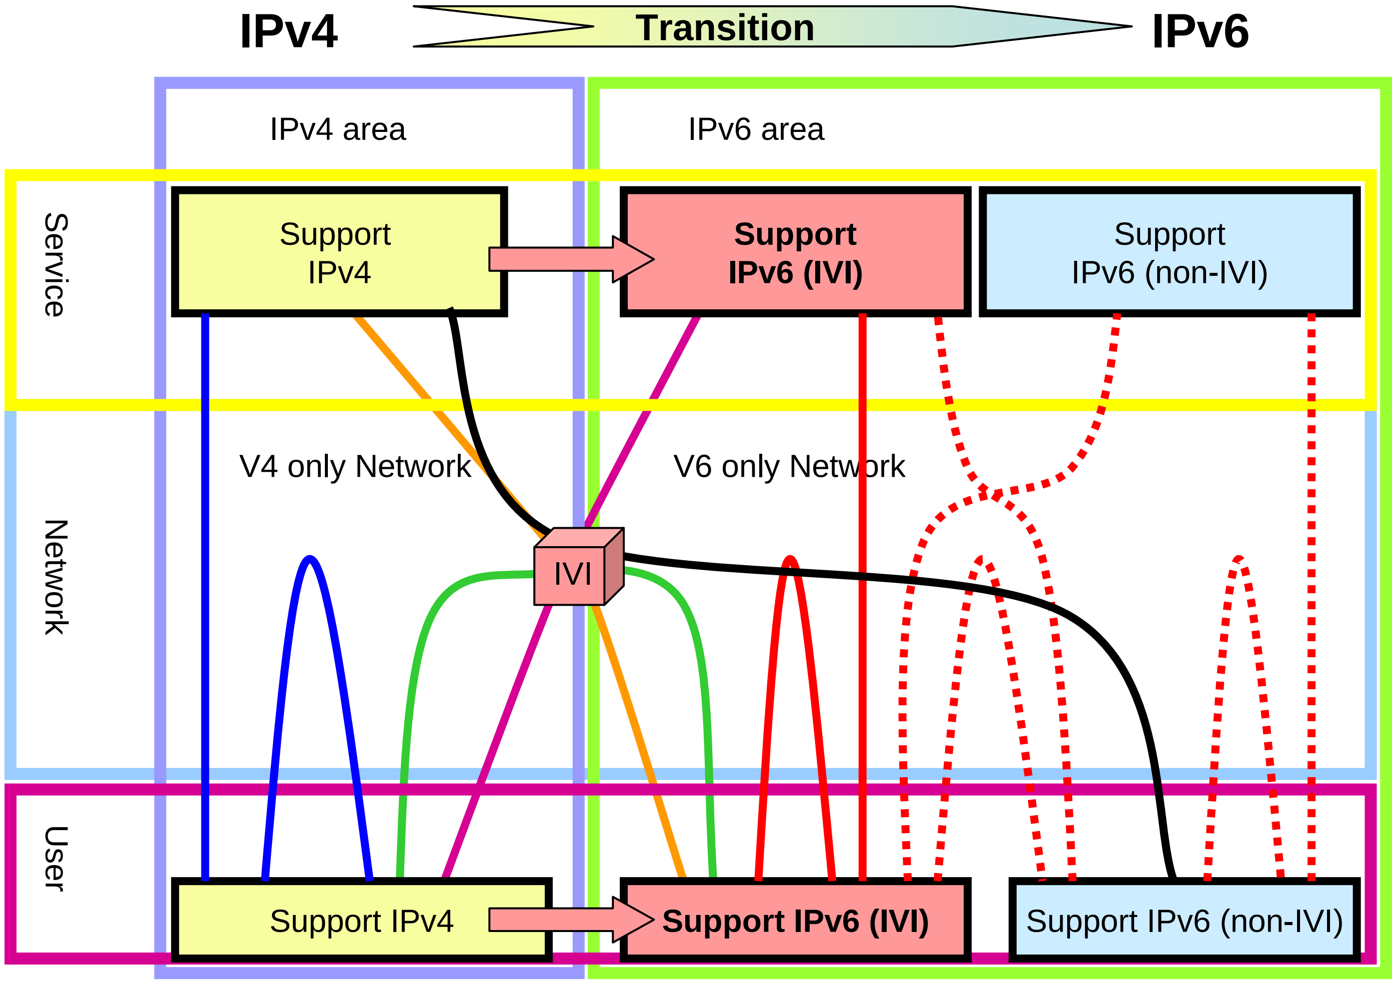

Remarks for the transition (1)

- The existing IPv4 users may not have motivation to transit to IPv6.
- Provide IVI6(i) for new Internet users, so they can have IPv4 connectivity and new IPv6 services. Then the existing IPv4 users may want to use IVI6(i). Therefore, more and more IPv4 addresses are borrowed by IPv6 networks as IVI6(i).
- When the number of services and users which support IPv6 (via IVI) reaches a critical mass, non-IVI IPv6 addresses can be used.

Remarks for the transition (2)

	Utilization of NAT/NAPT (Sharing of IPv4 addresses)	Reallocation of the assigned IPv4 addresses	Transition to IPv6	**IVI**
Feasibility within a time limit	✓	Doubtful	Extremely difficult	√
Service continuity	Limited	✓	✓	√
Permanent effect	Doubtful	NG	✓	√

Modified based on the June 2008 Report of the Japanese Study Group on Internet's Smooth Transition to IPv6

Discussion

- Why select a subset of the IPv6 addresses, rather than allow the whole IPv6 addresses to access the IPv4
- Mathematics of mapping
  - Because of the different size of the two address families, there must exist constrains.
- A subset is enough for the initial deployment
  - The IVI6 subnet is much, much larger than the global IPv4 whenIPv4 multiplexing techniques are used), even only a small portion of the public IPv4 addresses are borrowed by IVI.
  - Every IPv6 host can communicate with the global IPv4, not every IPv6 address (IPv4 class E address cannot communicate with class A, B, C).
- The standard NAT-PT methods also require the reservation of a similar size of the public IPv4 addresses in the pool.
  - These methods are maintaining a pool of public IPv4 addresses in NAT-PT box
- This subset supports the v6 and v4 initiated communications.
  - P2P
  - Pseudo-well-know-port, DNS SRV record

Comparisons (1:1 mapping example)

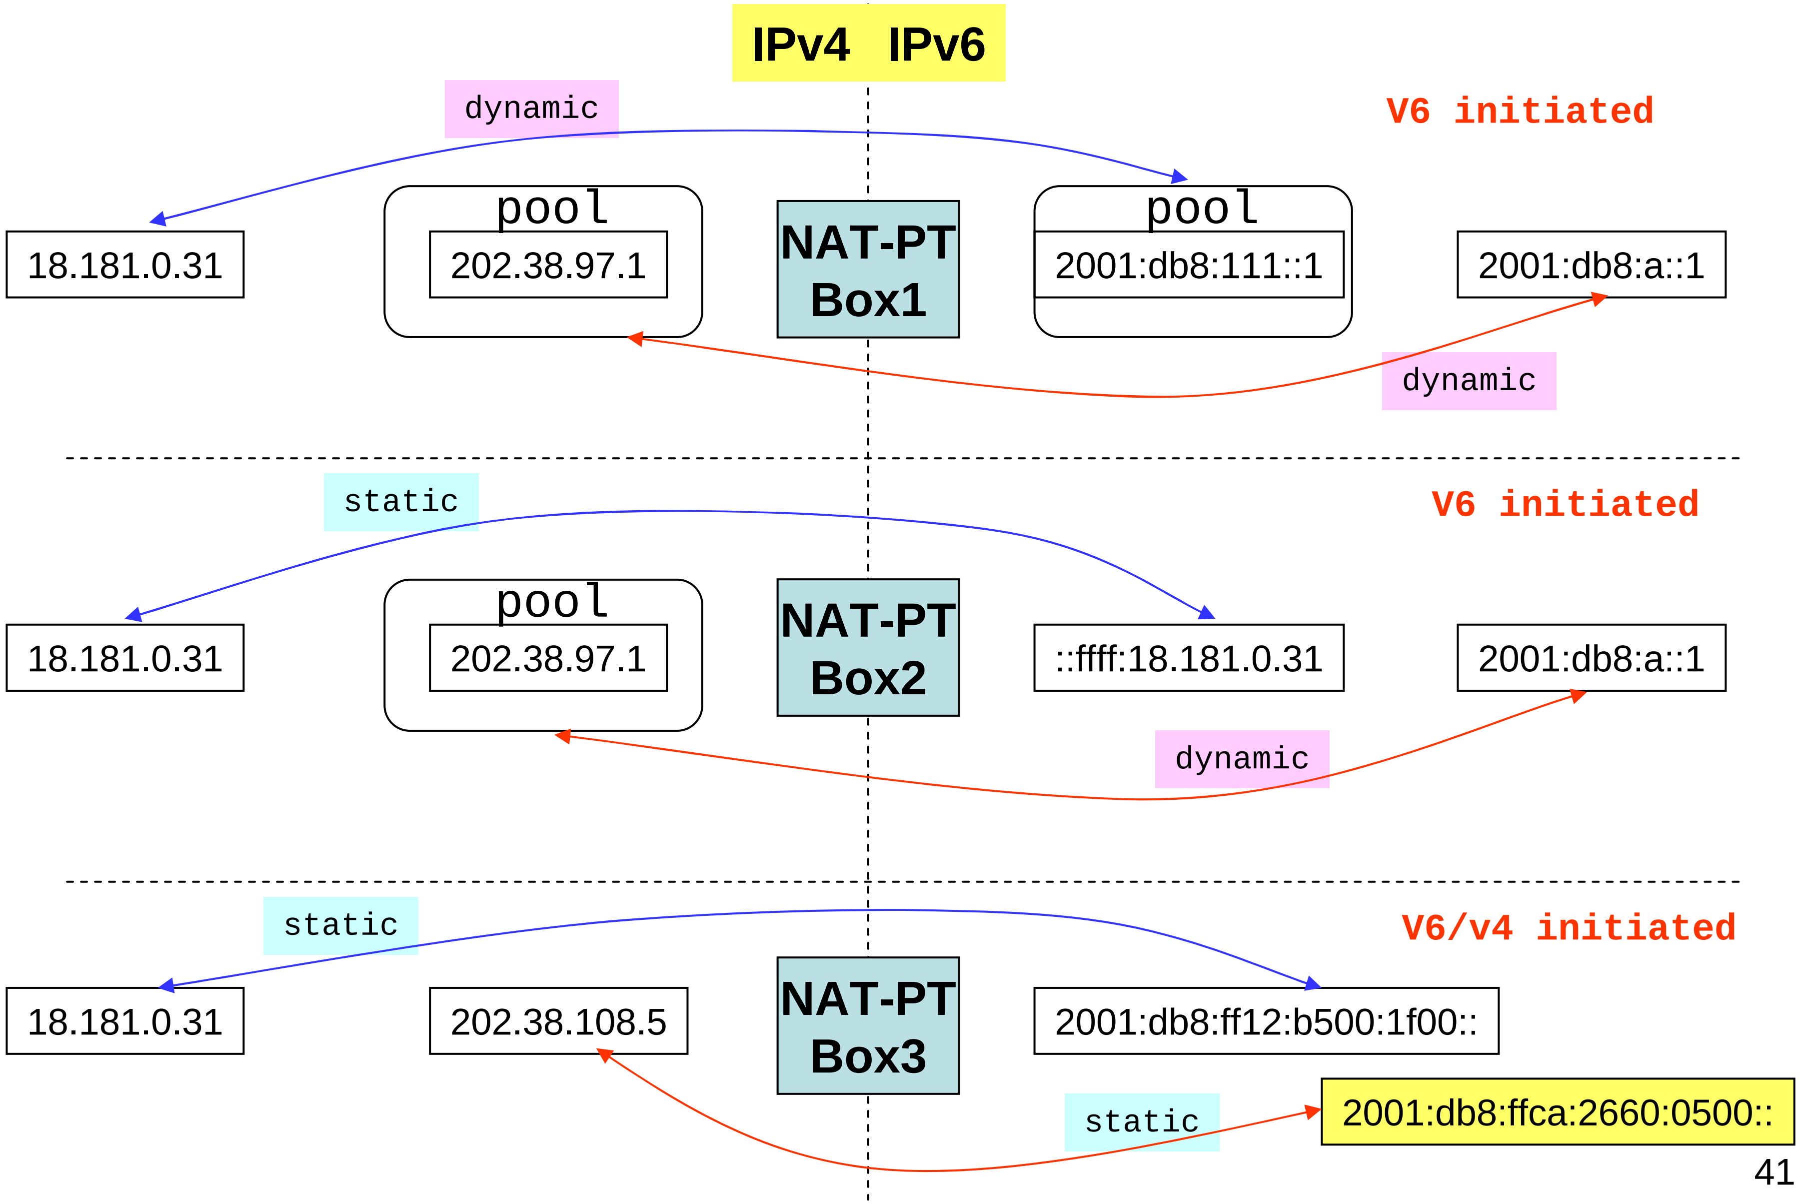

Conclusions

- The IVI is a prefix-specific and explicit bidirectional 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.