Stateless 1:N IVI and Stateless 1:N double IVI

X. Li, C. Bao, H. Zhang

November 12, 2009

Outline

- Introduction
- Stateless 1:N IVI
  - Design concept
  - Address mapping algorithm
  - Extended address format
- Stateless 1:N double IVI
  - Port mapping algorithm
- Comparisons
- Summary

Introduction

- Stateless 1:1 IVI
  - Cannot use IPv4 address effectively
  - Support both IPv6 initiated and IPv4 initiated communications
- Stateless 1:N IVI
  - Can share the IPv4 addresses among IPv6 hosts
  - Support both IPv6 initiated and IPv4 initiated communications
- Stateless 1:N dIVI
  - Can share the IPv4 addresses among IPv6 hosts
  - Support both IPv6 initiated and IPv4 initiated communications
  - Do not need ALG

Stateless 1:1 IVI

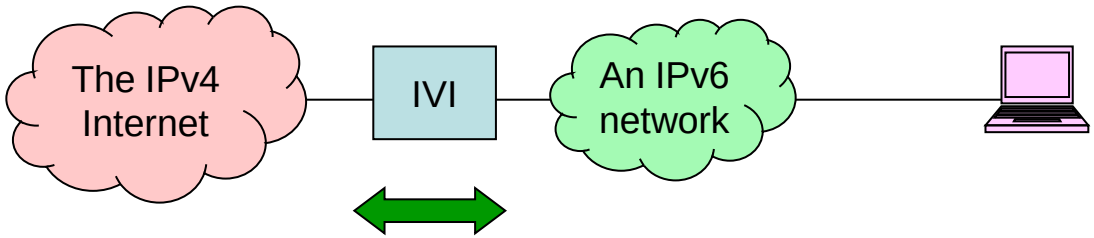

Algorithm

202.38.114.1 ←→ 2001:250:ffca:2672:0100::0

Scenario 1 “an IPv6 network to the IPv4 Internet”
Scenario 2 “the IPv4 Internet to an IPv6 network”
Scenario 5 “an IPv6 network to an IPv4network”
Scenario 6 “an IPv4 network to an IPv6 network”

Stateless 1:N IVI

- The port number multiplexing technique is used, similar to the ones used in NAT44
- For example,
  - If the port multiplexing ratio is 256, each IPv6 host with IPv4-translatable address can use 256 concurrent port numbers when communicating with IPv4 Internet. Therefore, one /24 is equivalent to one /16.
  - There is no port number restriction (65,536) when these IPv6 hosts communicate with the IPv6 Internet.

Design concept of stateless 1:N IVI

- Keep the basic features of the stateless translation (IVI)
- An IPv4 public address is shared by N IPv6 hosts (in this example N=4).
- The port number range for each IPv6 host is predefined.
- The port number range is encoded in the IPv6 address (IPv4-translatable address) and no signaling scheme is required.

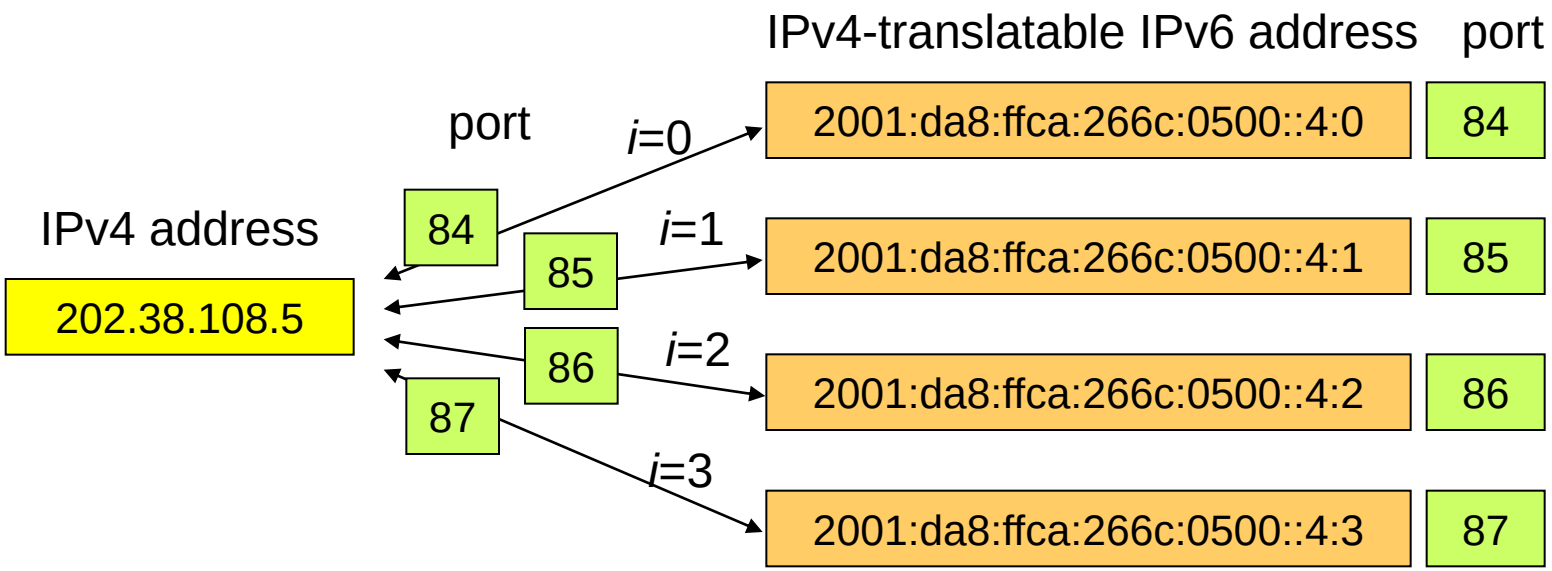

Address-mapping algorithm

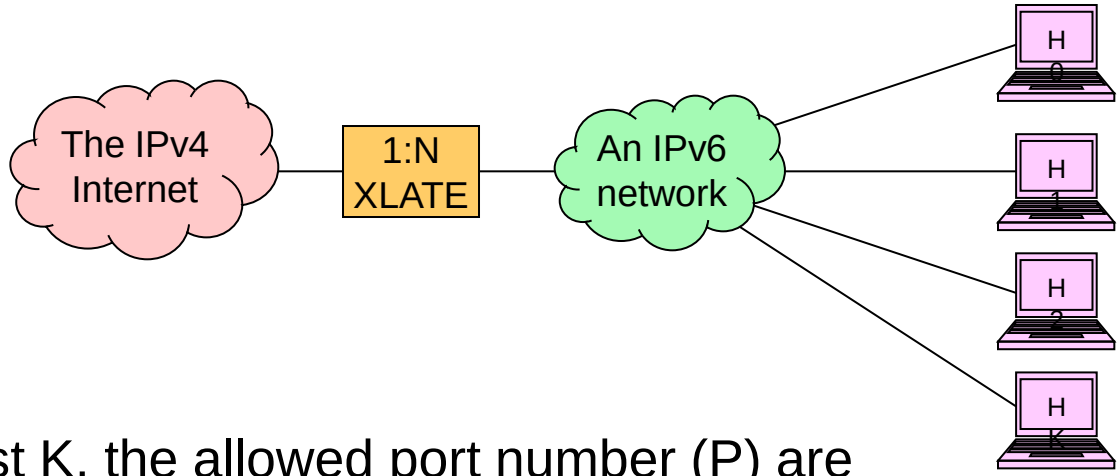

- For host K, the allowed port number (P) are
  - **P=j*N + K (j=0, 1,..., N-1).**
- For the destination port number (P), the packets will be sent to host
  - **K=(P%N) (% is the Modulus Operator).**

- For example: If N=256, then
  - Host K=5 is only allowed to use port numbers 5, 261, 517, 773, ..., 65,285 as the source port
  - The packets with these port numbers as the destination port number will be send to host K=5.

Extended address format (1)

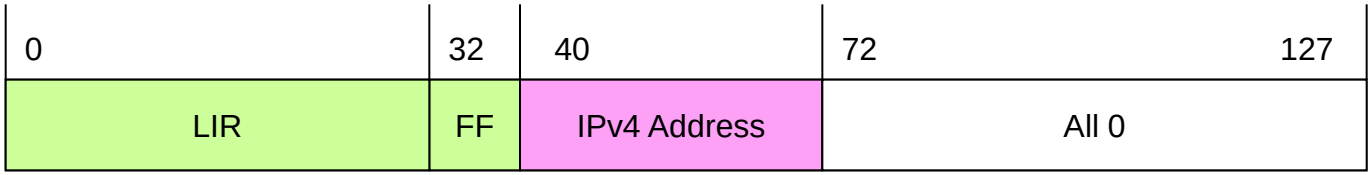

IPv4-converted address format

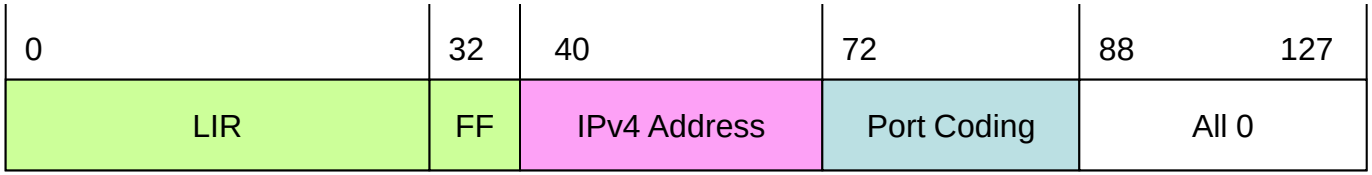

Extended IPv4-translatable address format

- Use 16 bits to encode the port number range (Port Coding)

Extended address format (2)

(4 bits)	Index Range(12 bits)	Multx ratio	# of Ports
0	000-000	1	65,536
1	000-001	2	32,768
2	000-003	4	16,384
3	000-007	8	8,192
4	000-00f	16	4,096
5	000-01f	32	2,048
6	000-03f	64	1,024
7	000-07f	128	512
8	000-0ff	256	256
9	000-1ff	512	128
A	000-3ff	1,024	64
B	000-7ff	2,048	32
C	000-fff	4,096	16

Figure 5: Transport layer port number coding

- Use 16 bits to encode the port number range
  - 4 bits: multiplexing ratio
  - 12 bits: the host index

The translator behavior

- From the IPv4 Internet to an IPv6 network
  - Src.a → map to IPv4-converted → Src.a6
  - Src.p → unchanged → src.p
  - **Dst.a → map to extended IPv4-translatale → Dst.a6**
  - Dst.p → unchanged → Dst.p
- From an IPv6 network to the IPv4 Internet
  - **Src.a6 → (check, map to IPv4 or drop) → Src.a**
  - **Src.p → (check, unchanged or drop) → Src.p**
  - Dst.a6 → (map to IPv4 or drop) → Dst.a
  - Dst.p → (unchanged or drop) → Dst.p

The IPv6 end system requirements

- The IPv6 end system requirements
  - Client: the end system must generate the source port numbers in the range defined by the extended IPv4-translatable address format
    - Option 1: Modify the IPv6 end system
    - Option 2: Use gateway to map the port number range to the ones defined by extended IPv4-translatable address
  - Server: different well-known ports will be served by different IPv6 hosts
    - The http redirection techniques can be used without providing the port information.

Remarks for stateless 1:N IVI

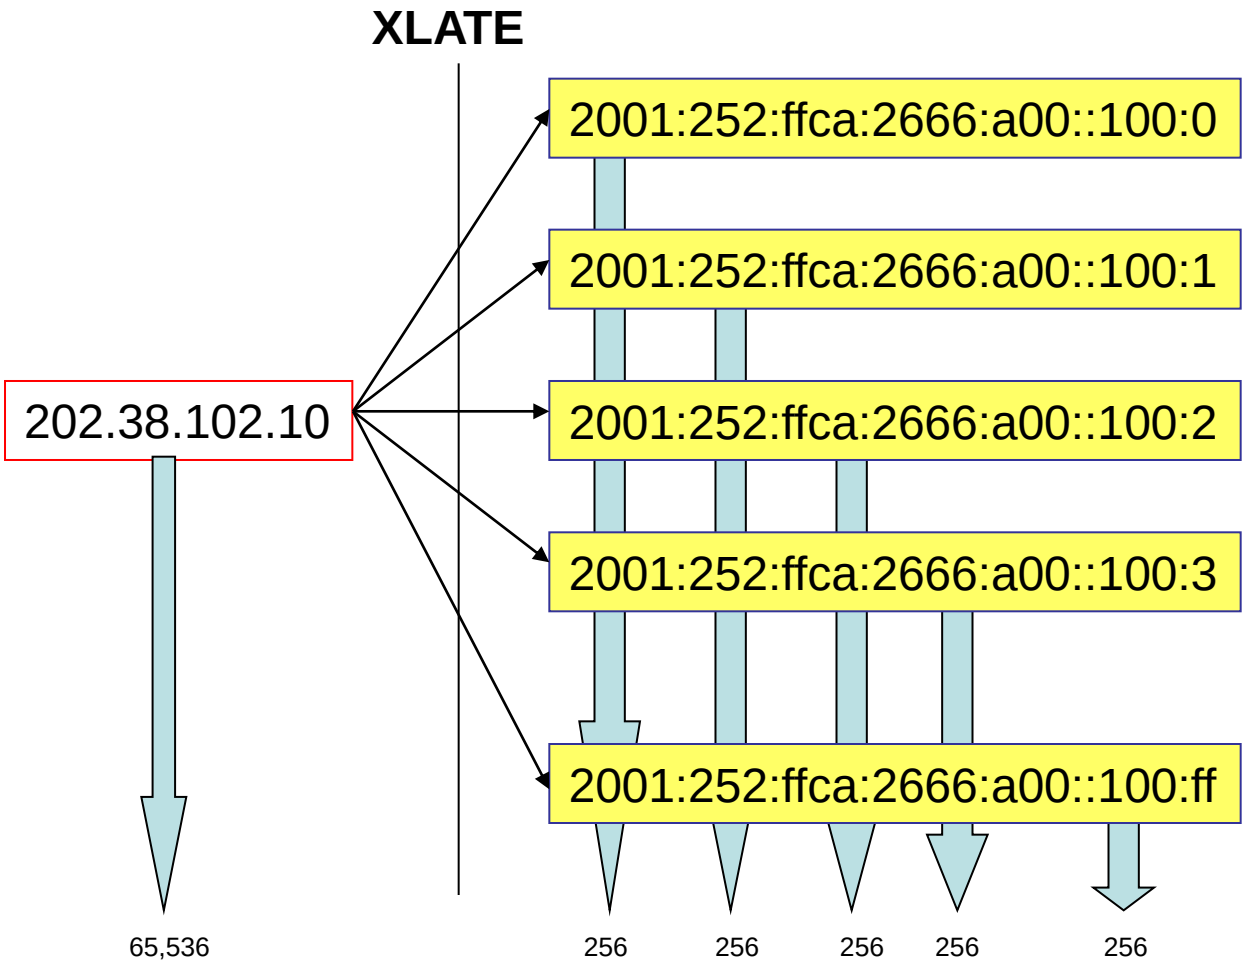

Stateless 1:N double IVI

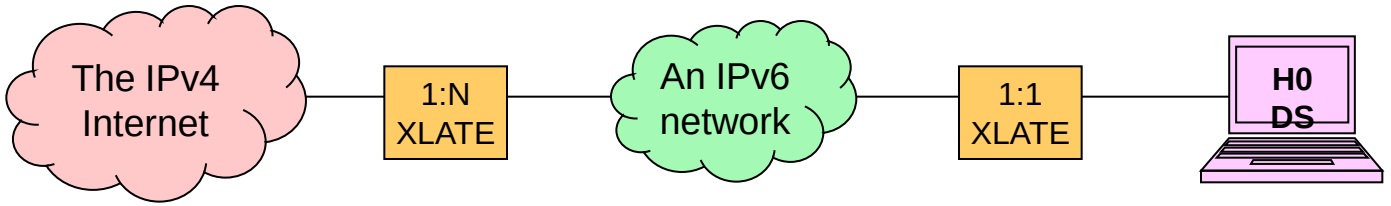

- If we use a gateway to map the port number, then we can use the same gateway to translate the IPv6 packets to IPv4.

- 1:N XLATE
  - Same as stateless 1:N IVI
- 1:1 XLATE
  - For IPv4-translatable and IPv4-converted addresses
    - 1:1 address translation (stateless)
    - Port number mapping (maintain states close to the end system)
  - Otherwise
    - IPv6 routing/IPv6 cut-through

Port number mapping algorithm

- The port number mapping device maintains a database of allowed port numbers defined by the extended IPv4-translatable address format.
- The end system randomly generated source port can be mapped to the allowed port based on the extended IPv4-translatable address format and the database.
- Different mapping rules may apply to
  - The Well Known Ports (0-1,023).
  - The Registered Ports (1,024-49,151).
  - The Dynamic and/or Private Ports (49,152-65,535).

Home gateway implementation

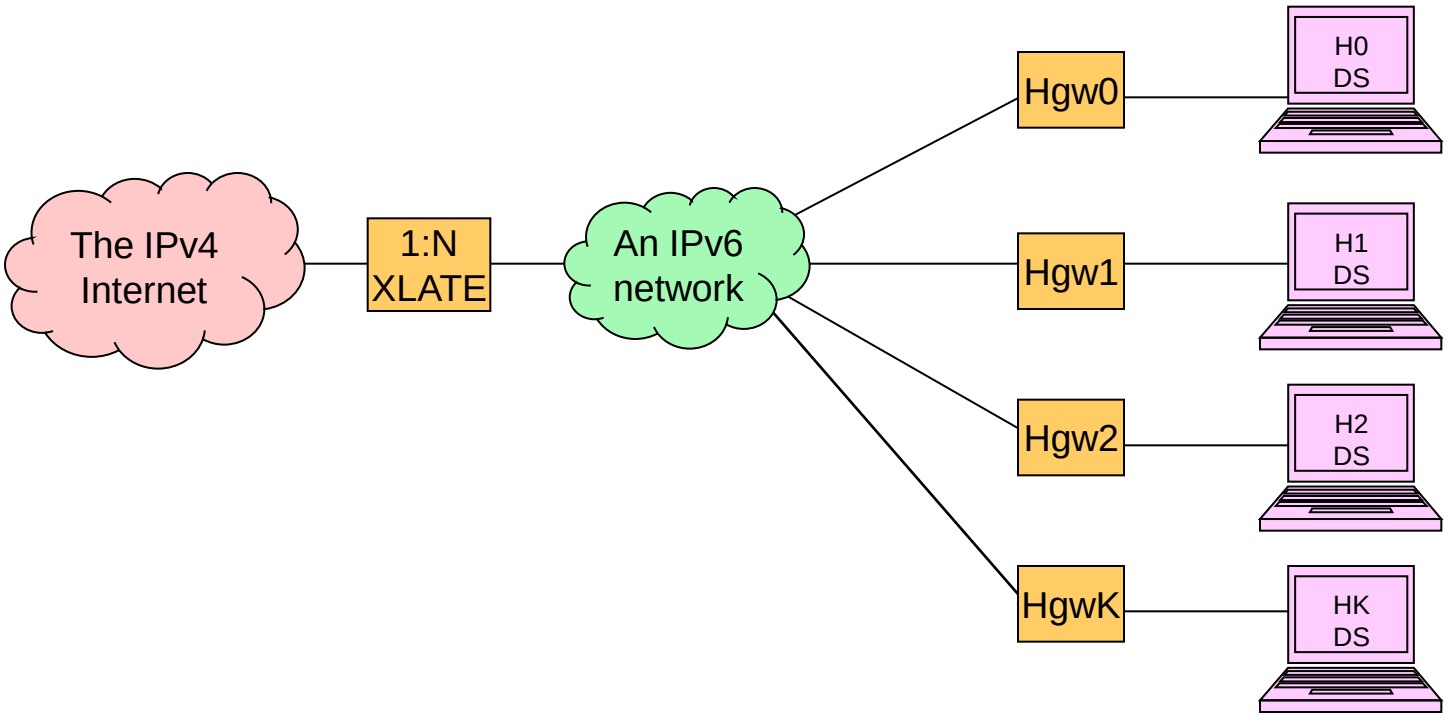

End system implementation

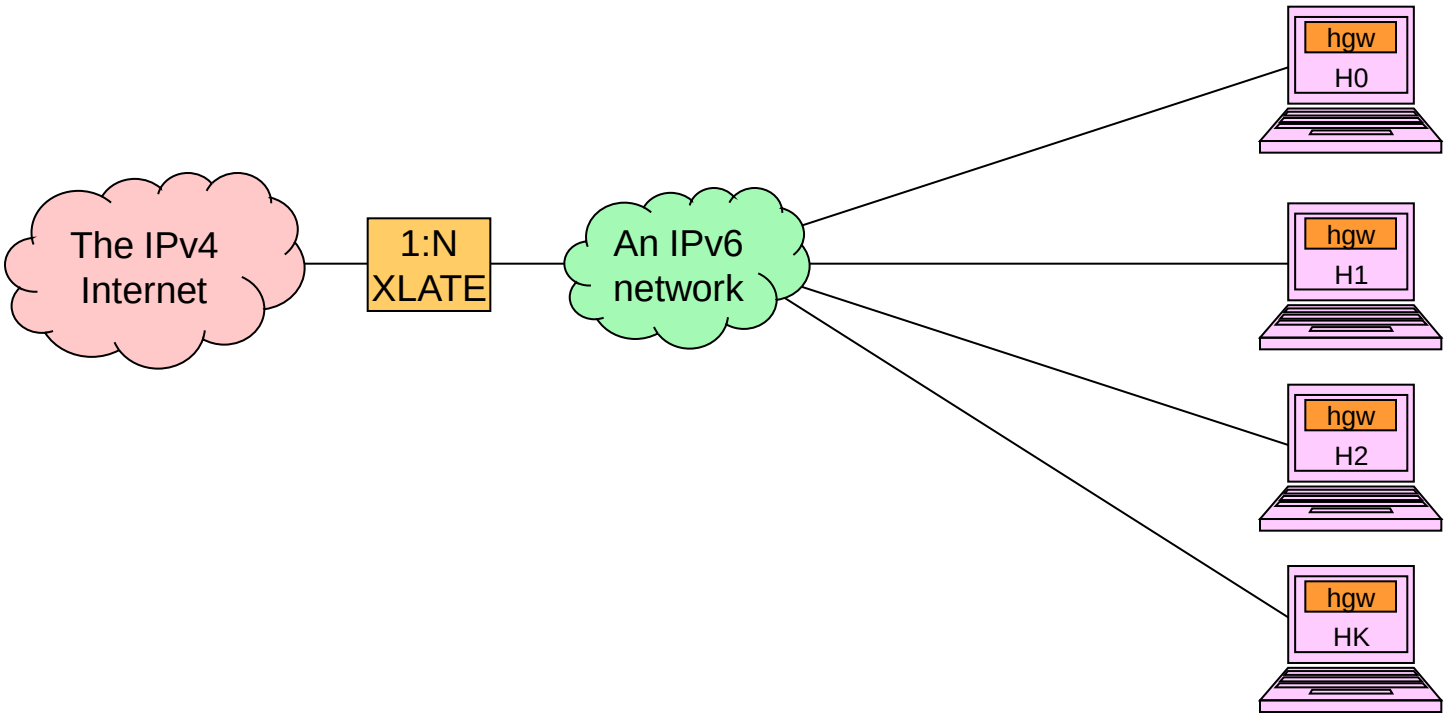

Remarks for stateless 1:N dIVI

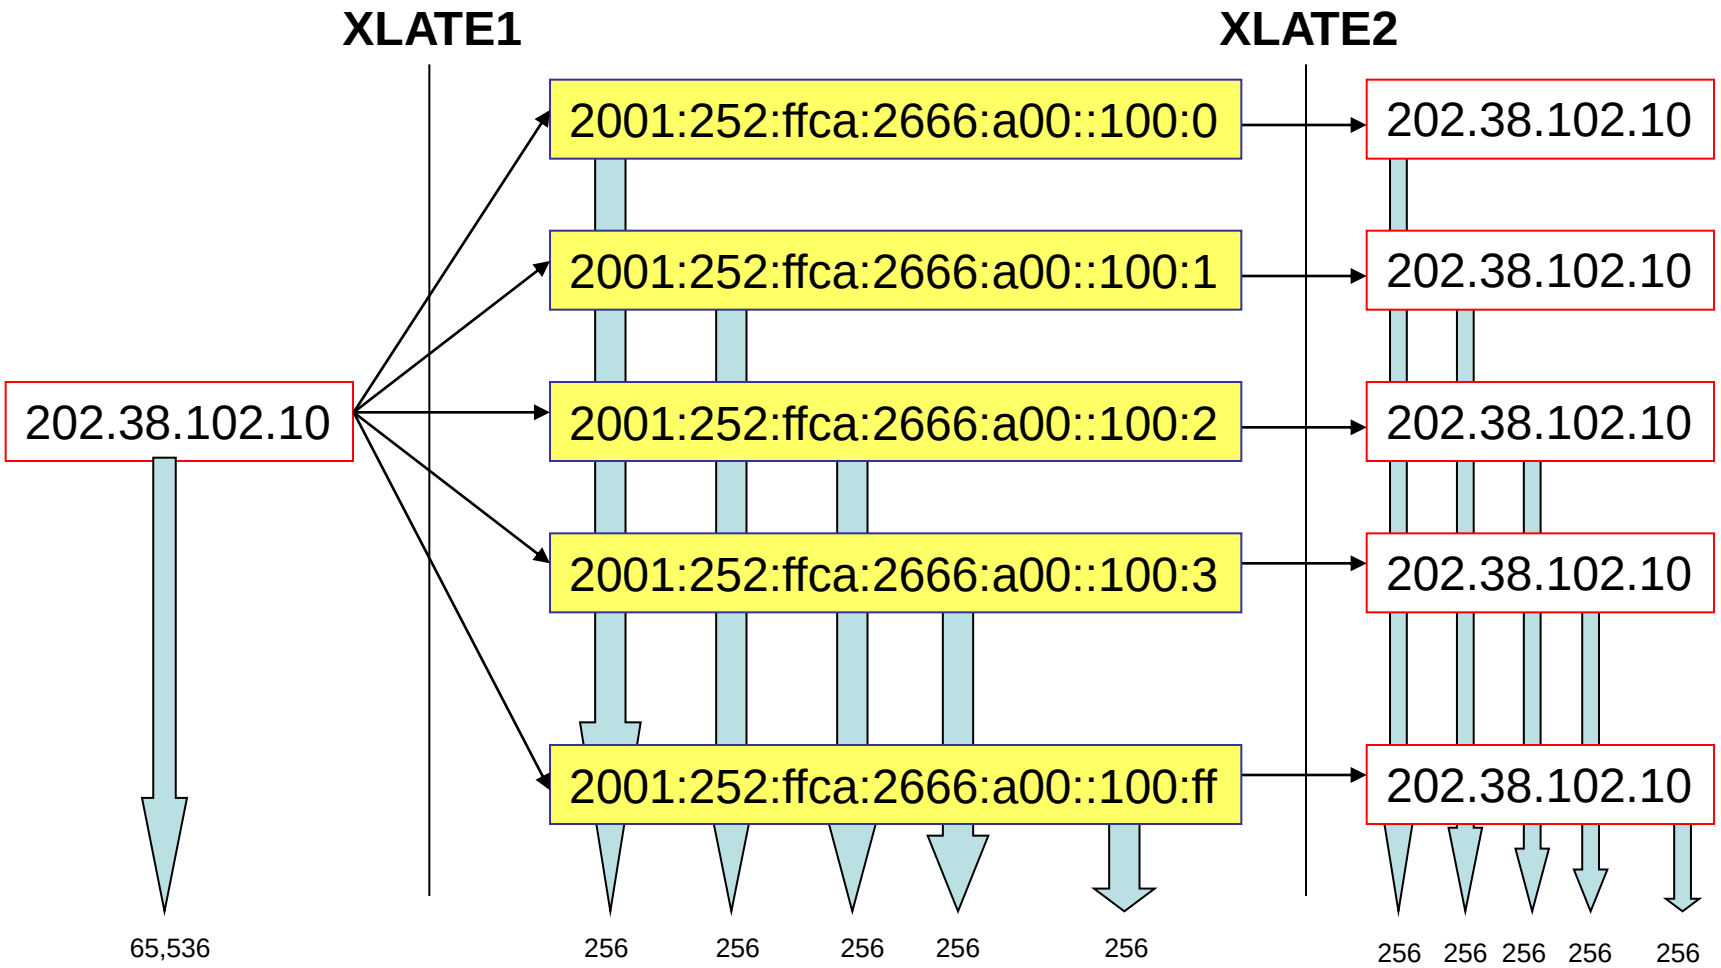

Prototype and testing environment

- The prototypes are running using Linux

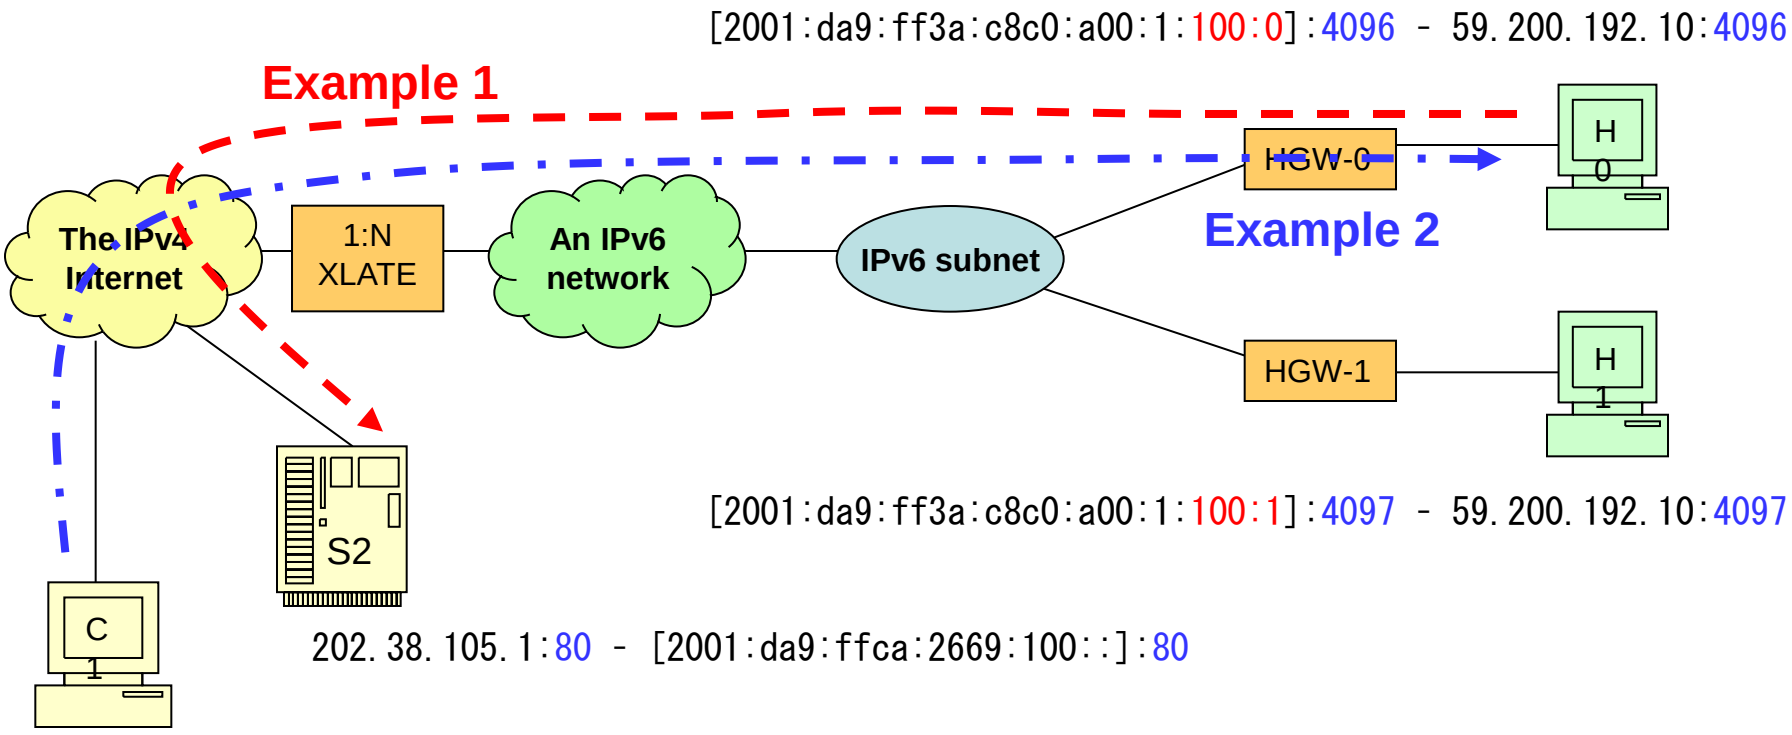

Example 1 (client)

- An address-sharing end system Host0 (58.200.192.10) in an IPv6 network behind home gateway initiates communication with Host S2 (http://202.38.105.1:80) in the IPv4 Internet
  - On the end system Host0
  - Src#p= 58.200.192.10:1881 (random port)
  - Dst#p= 202.38.105.1:80 (server port)

  - On an IPv6 network
  - Src#p= [2001:DA9:FF**3A:C8C0:A**00:0:**100:0**]:8192 (home gateway mapped port)
  - Src#p= [2001:252:ffca:2669:100::]:80 (server port)

  - On the IPv4 Internet
  - Src#p= 58.200.192.10:8192 (home gateway mapped port)
  - Dst#p= 202.38.105.1:80 (server port)

Example 2 (server)

- Host C1 (125.34.46.137) in the IPv4 Internet initiates communication with address-sharing end system Host0 (http://58.200.192.10:4096) in an IPv6 network behind home gateway.
  - On the IPv4 Internet
  - Src#p= 125.34.46.137:1856 (random port)
  - Dst#p= 58.200.192.10:4096 (server port)

  - On an IPv6 network
  - Src#p= [2001:DA9:ff7d:222e:8900::]:1856 (random port)
  - Dst#p= [2001:DA9:FF**3A:C8C0:A**00:0:**100:0**]:4096 (server port)

  - On the address-sharing end system Host0
  - Src#p= 125.34.46.137:1856 (random port)
  - Dst#p= 58.200.192.10:4096 (server port)

www.ivi2.org/demo.html

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/demo.html

Back 访问最多 Windows Media Windows 免费 Hotmail 自定义链接

dIVI demo

H0 as a client http://58.200.192.10:4096/cgi-bin/nph-ivi-client
H0 as a server http://58.200.192.10:4096/cgi-bin/nph-ivi-server

H1 as a client http://58.200.192.10:4097/cgi-bin/nph-ivi-client
H1 as a server http://58.200.192.10:4097/cgi-bin/nph-ivi-server

Address-sharing dIVI Demo (client) H0

• As an IPv4 client to connect to an IPv4 test server

server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8192 (256).0

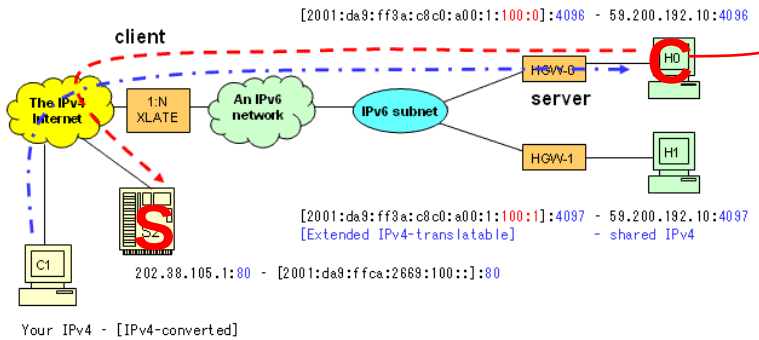

Address-sharing dIVI Demo (server) H0

• As an IPv4 server to report your host information

server 58.200.192.10:4096 client 2001:da9:ff7d:2235:4800:: 125.34.53.72 7561

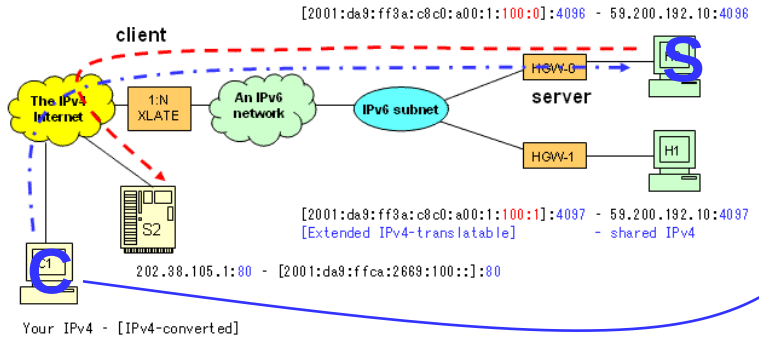

Reload homepage for H0 and H1

- **H0**

- **As an IPv4 client to connect to an IPv4 test server**
  - server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8448 (256).0
  - server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8192 (256).0
  - server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8704 (256).0
- **As an IPv4 server to report your host information**
  - server 58.200.192.10:4096 client 2001:da9:ff7d:2235:4800:: 125.34.53.72 7608
  - server 58.200.192.10:4096 client 2001:da9:ff7d:2235:4800:: 125.34.53.72 7612

- **H1**

- **As an IPv4 client to connect to an IPv4 test server**
  - server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8193 (256).1
  - server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8449 (256).1
  - server 202.38.105.1:80 client 2001:da9:ff3a:c8c0:a00:: 58.200.192.10 8705 (256).1
- **As an IPv4 server to report your host information**
  - server 58.200.192.10:4097 client 2001:da9:ff7d:2235:4800:: 125.34.53.72 7618
  - server 58.200.192.10:4097 client 2001:da9:ff7d:2235:4800:: 125.34.53.72 7622

Comparison (1)

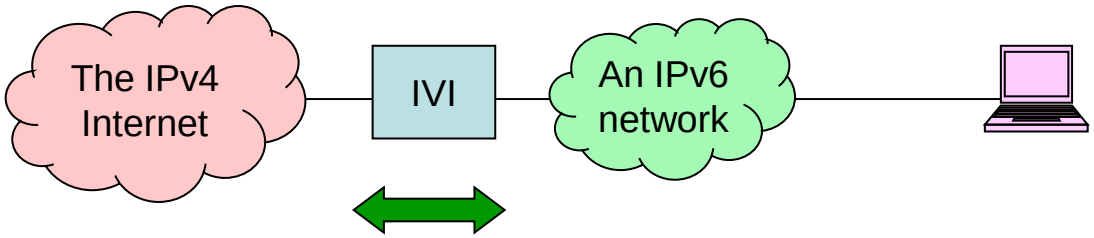

Algorithm

202.38.114.1 ↔ 2001:250:ffca:2672:0100::0

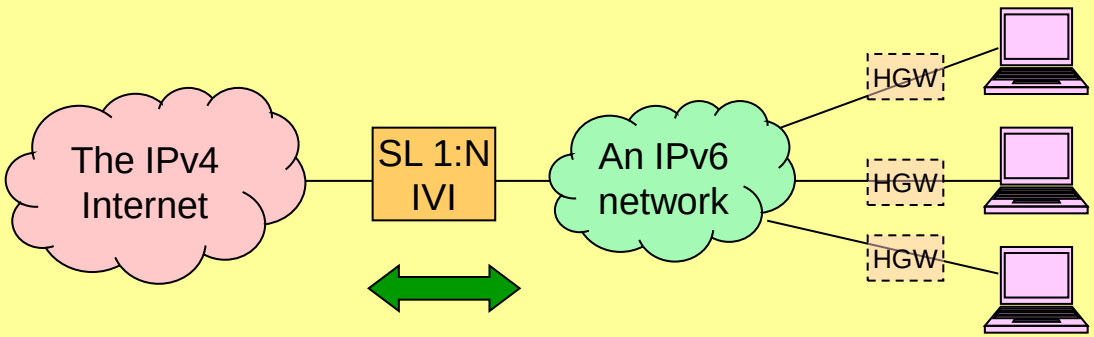

Algorithm

58.200.192.10#P0 ↔ 2001:da9:ff3a:c8c0:a00:0000:100:0#P0

58.200.192.10#P1 ↔ 2001:da9:ff3a:c8c0:a00:0000:100:1#P1

58.200.192.10#P2 ↔ 2001:da9:ff3a:c8c0:a00:0000:100:2#P2

Comparison (2)

- Stateless 1:1 IVI
  - server/client full-function (bi-direction initiated)
  - Restricted IPv6 addresses (few, 1:1 mapping)
  - Stateless and scalable
  - Need DNS64 and ALG
- Stateless 1:N IVI (dIVI)
  - Server/client limited function (bi-direction initiated)
  - Restricted IPv6 addresses (many, 1:N mapping)
  - Stateless and scalable
  - Do not need DNS64 and ALG (dIVI)

Summary (1)

- There is no change to the IP model
  - For the gateway implementation, there is no change to the end system and the applications.
  - It has less restriction compared with NAT44 and the NAT transversal techniques can be used
    - NAT44: both address and port number are changed
    - 1:N dIVI: only the port number is changed
  - The address and port routing is via IPv6 addresses and the operation complexity is less than what is done via RFC1918 addresses.
  - The core translator (1:N XLATE) is stateless
  - The home gateway maintains states for port number mapping only

Summary (2)

- This technique can help for the IPv4/IPv6 coexistence and encourage the transition
  - An IPv6 network can communicate with the IPv4 Internet with two directional initiation.
    - The IPv4 addresses can be used by multiple IPv6 hosts, which partially solve the IPv4 address depletion problem.
    - The total of N (N<65,536) concurrent sessions per IPv6 host can be used when communicating with the IPv4 Internet.
  - An IPv6 network can communicate with the IPv6 Internet with two directional initiation.
    - The total of 65,536 concurrent sessions per IPv6 host can be used when communicating with the IPv6 Internet.