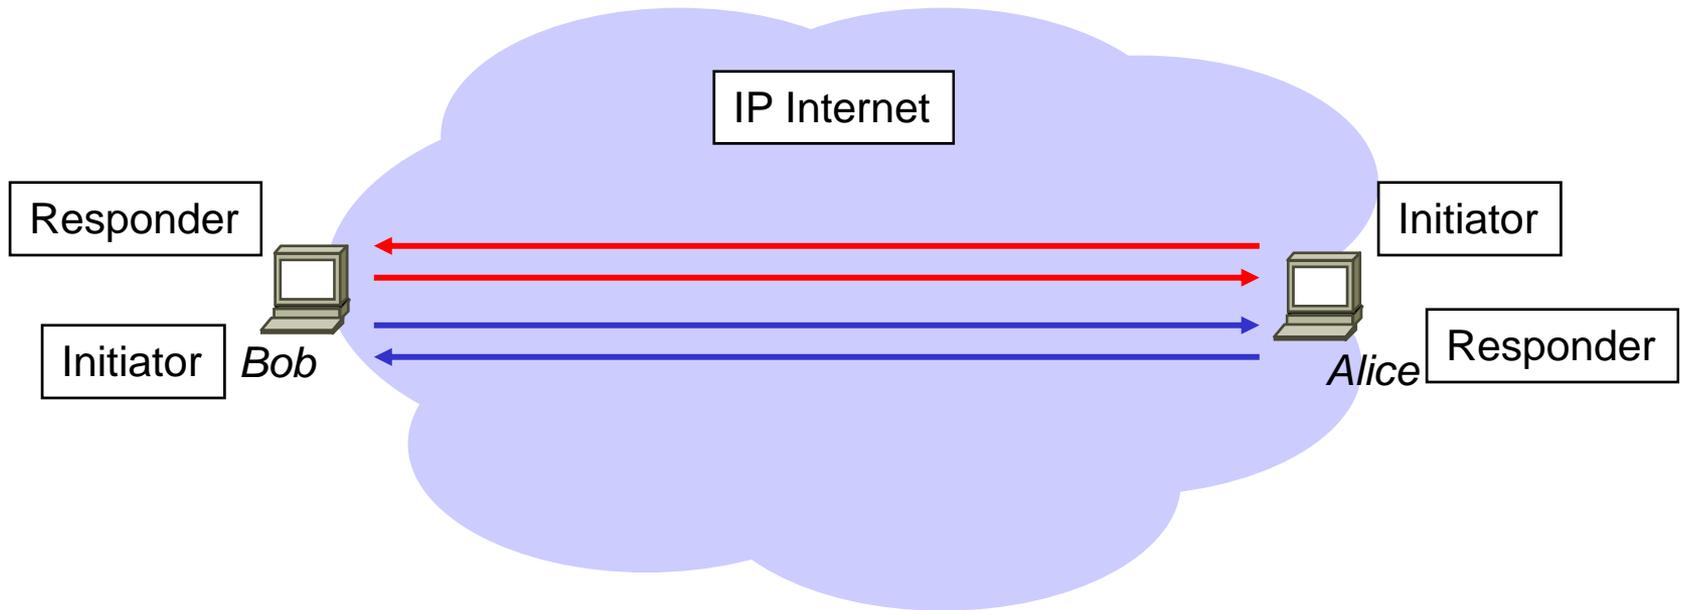


Data Networks

Lecture 26: IPv6 and NAT

Ideal Internet



- Every host has a globally unique IP address
- Bi-directional connectivity is a fundamental property

32-Bit IP Address Space Is Too Small

- Upper bound: 31% of IP address space is covered by aggregated routing table
- Poor utilization
- Increasing demand
 - Always-on access (e.g., DSL, cable modem)
 - Internet enabled devices (e.g., mobile phones, PDAs)
- Fear of exhaustion leads to aggressive conservation
- IP addresses are increasingly difficult to obtain

Trend #1

- Deploy networks using IPv6 addresses
 - enormous 128-bit address space
 - globally unique
 - hosts and routers run IPv6
- Not yet very common, but picking up steam
 - E.g. US Dept of Defense has mandated IPv6 by 2008
 - i.e. They won't buy anything that is not IPv6 ready by 2008

IPv6 Addressing

- 128-Bit Addresses
 - Standard unicast addresses 128 bits long
 - 3-bit prefix: 001
 - 8×10^{18} nodes / cm² of the Earth
- Aggregation Levels (RFC 2374)
 - Use address prefix like CIDR to aggregate
 - 13 bits top-level (major providers)
 - 24 bits next-level (intermediate level)
 - 16 bits site-level (like current subnetting)
 - 64 bits interface (like current MAC addresses)

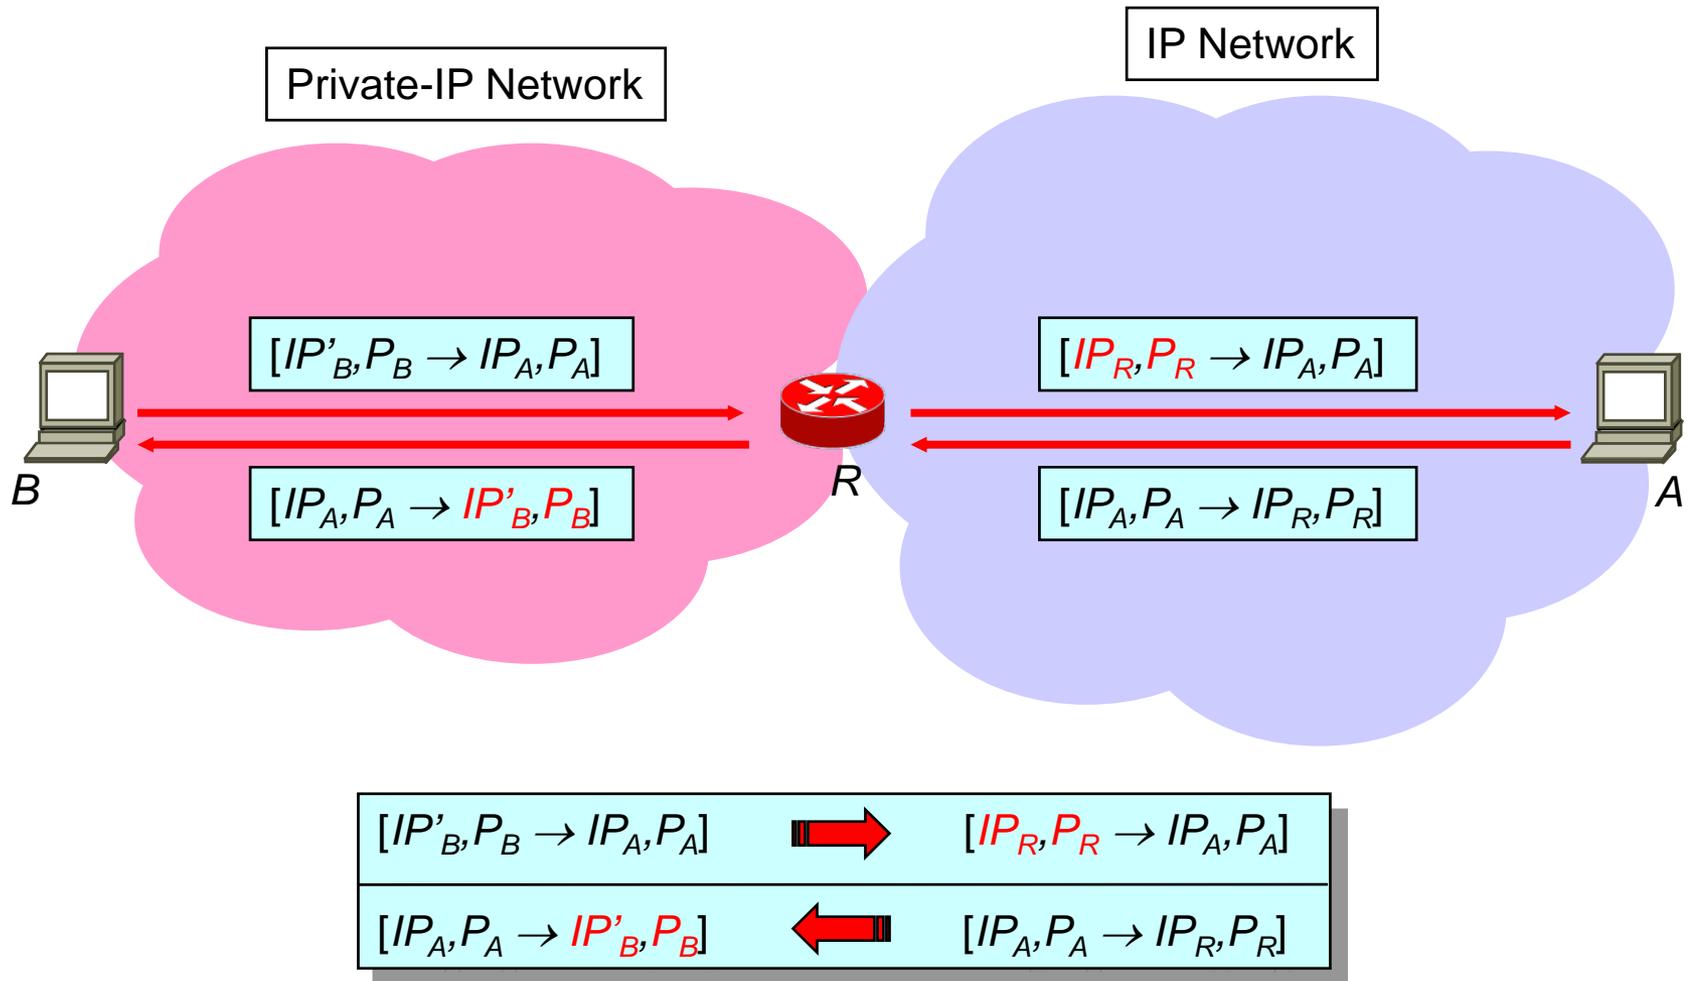
Special IPv6 Addresses

- IPv4-compatible IPv6 address for dual stack hosts
 - Dual stack: Run both IPv4 and IPv6 on a host for backward compatibility
 - Zero extend (prefix) a 32bit IPv4 address to 128bit
- IPv4-mapped IPv6 address for IPv4 only hosts
 - IPv4 only hosts can still be identified by IPv6 hosts
 - Prefix 32bit IPv4 address by 0xFFFF then zero extend to 128bit
 - Indicates this host does not actually understand IPv6

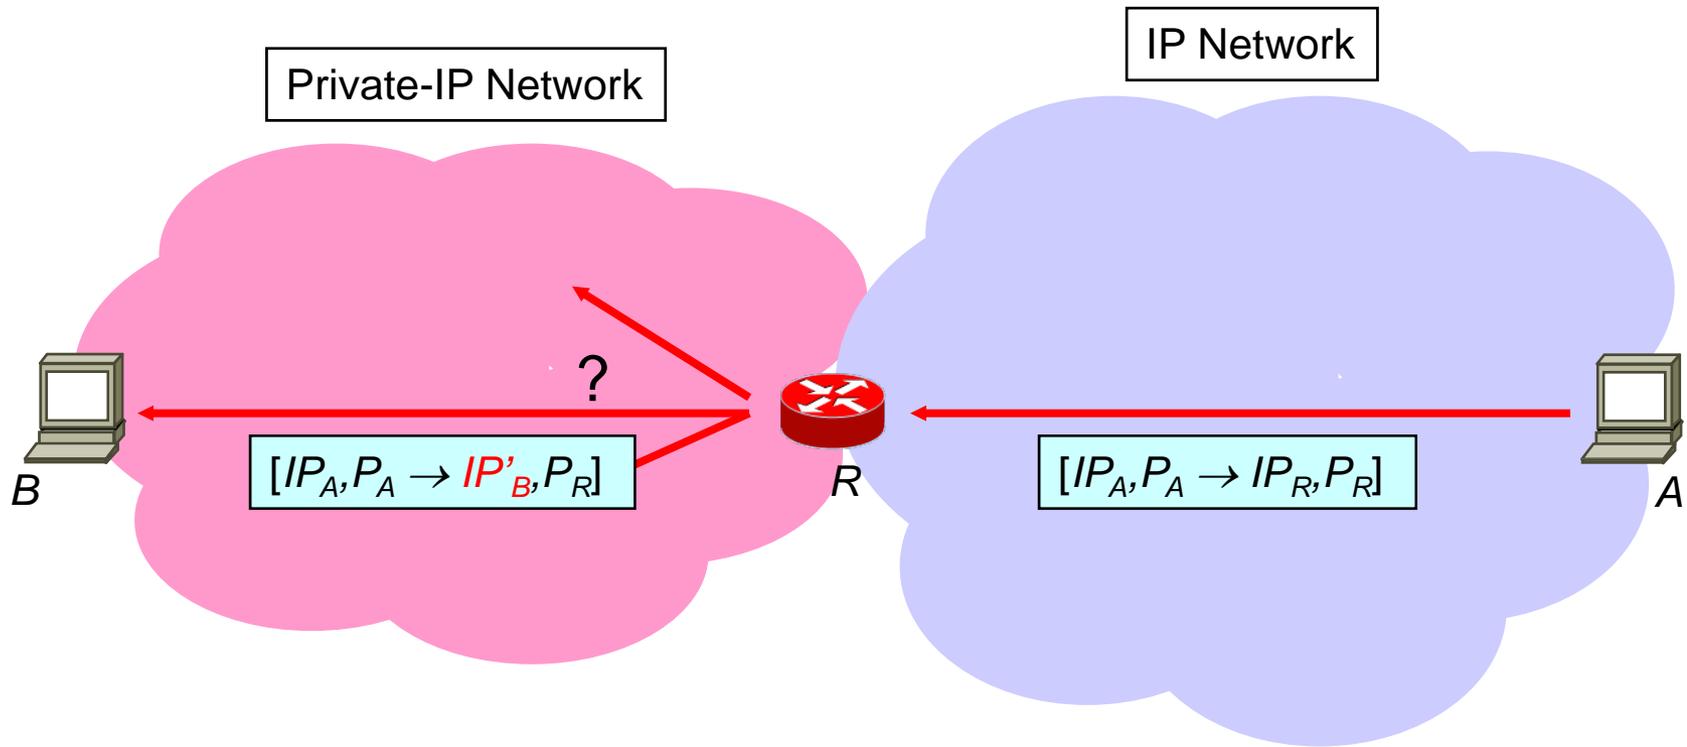
Trend #2

- Deploy networks using reusable-IP addresses
 - a.k.a. private-IP addresses
 - IP network prefixes 10/8, 172.16/12, 192.168/16
 - not globally unique, not routable
 - hosts and routers remain running IP
- Need Network Address Translation (NAT) Device
 - This is basically your Linksys cable router, or your Windows Internet connection sharing

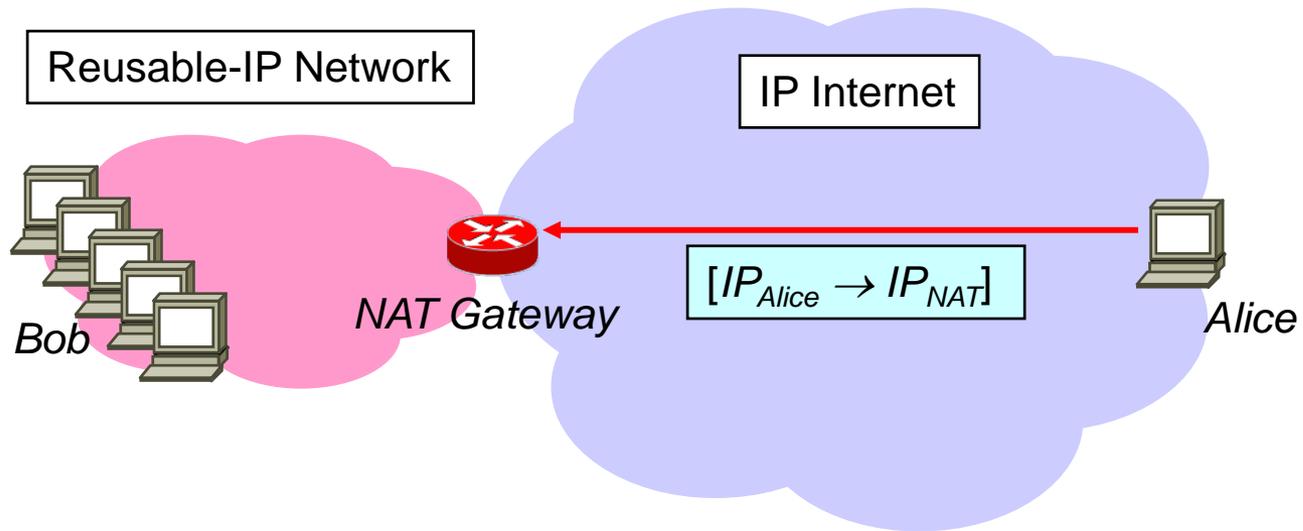
Translator: Out-Bound



Translator: In-Bound



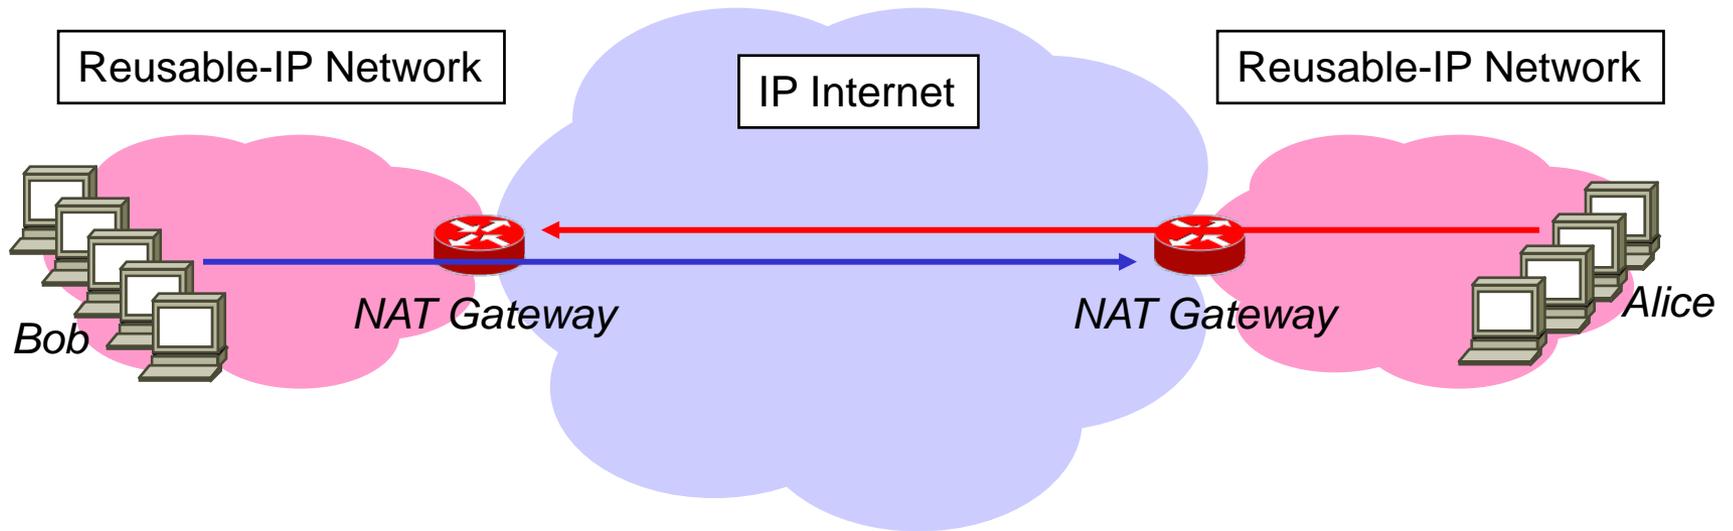
The Problem with NAT



- Bob has no globally unique IP address and so Alice cannot directly address Bob

NAT cannot provide bi-directional connectivity

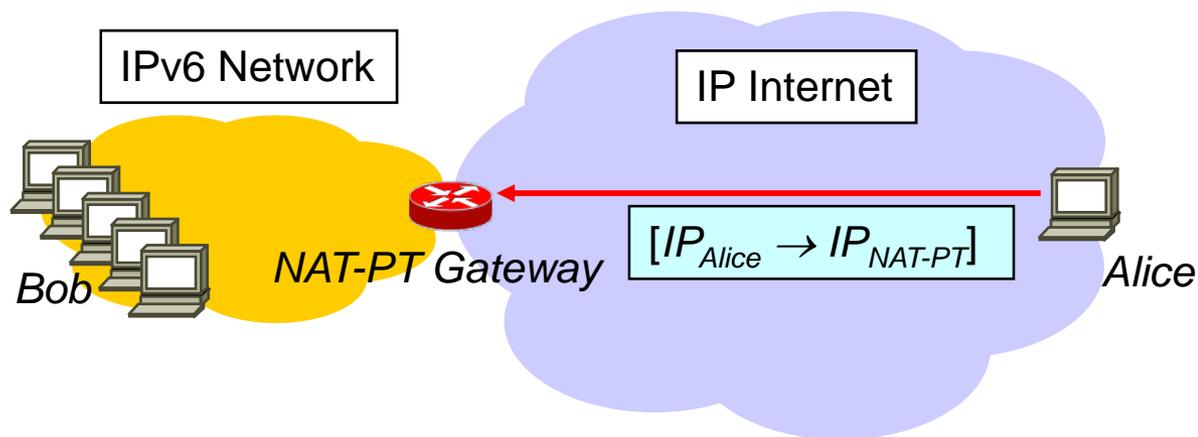
Even Worse



- No connectivity between Alice and Bob at all!
- Think about Peer-to-Peer applications

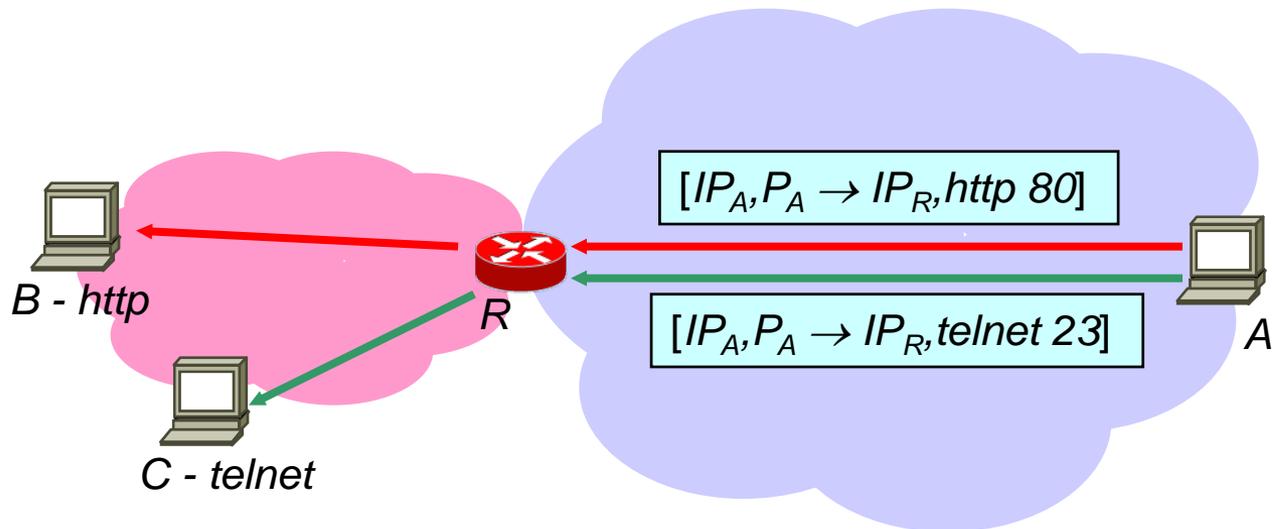
Using IPv6 Addresses

- IPv6 can be fully compatible with IP
- Key: every IPv6 host must consume a globally unique IPv4 address!
 - i.e. IPv4 compatible IPv6 address
- In reality, many IPv6 networks will be IPv6-only
 - connect to IP Internet via NAT-Protocol Translation (PT) gateway
 - NAT-PT has the same problem as NAT



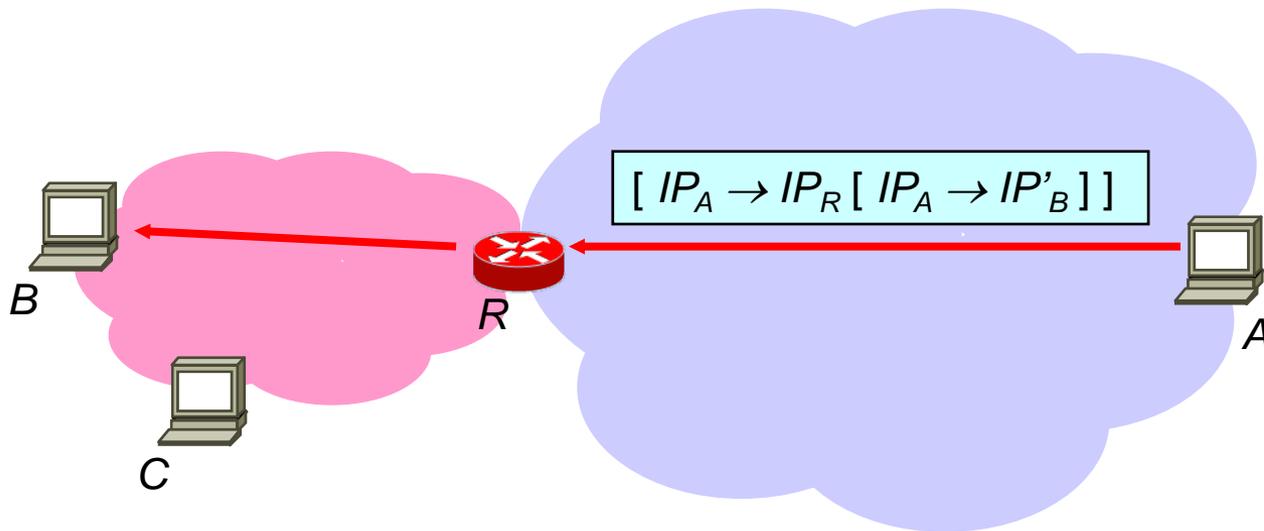
What Can be Done?

- Some network entity needs to know that the destination host is private-IP
- 1. Translator gateway knows...
 - Can statically bind port numbers to private-IP hosts
 - Poor connectivity and flexibility



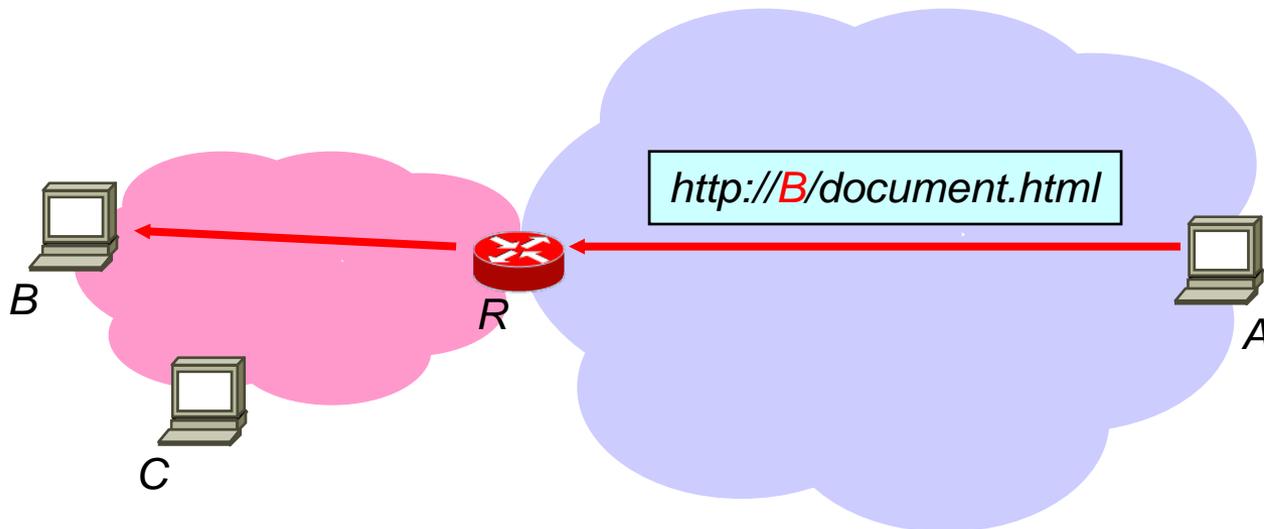
What Can be Done?

- **2.** If the initiating IP host knows...
 - Can encapsulate packets to destination directly
 - Requires DNS and existing hosts to be heavily modified
 - Difficult to deploy



What Can be Done?

- **3.** If the application knows...
 - The real destination can be encoded in higher level protocols
 - Too application specific
- More possibilities, but all experimental



Reality: Heterogeneous Internet Address Spaces

- NAT and NAT-PT cannot provide bi-directional connectivity
- Key problems: **IP to reusable-IP** and **IP to IPv6** connectivity
- NAT is just bad?
 - Local incentive to deploy, but does global harm
 - Worked fine for client-server apps, but not for emerging p2p, not forward looking
 - Reality: Necessary evil

