Reti di Calcolatori

Soluzioni per la carenza di indirizzi IP

Universtità degli studi di Verona
Facoltà di Scienze MM.FF.NN.
A.A. 2010/2011
Laurea in Informatica

Acknowledgement

- *Credits*
  - *Part of the material is based on slides provided by the following authors*
    - *Douglas Comer, "Computer Networks and Internets," 5th edition, Prentice Hall*
    - *Behrouz A. Forouzan, Sophia Chung Fegan, "TCP/IP Protocol Suite," McGraw-Hill*, January 2005

Topics covered

- NAT
- IPv6

Indirizzamento privato: NAT

Indirizzi privati

- IETF ha definito alcuni range di indirizzi all'interno dello spazio di indirizzamento IP da utilizzare solamente in ambito privato
  - *private addresses* o *non-routable addresses*
  - ogni volta che un router pubblico riceve un pacchetto destinato ad un indirizzo IP privato, viene segnalato un errore

Prefisso	Indirizzo iniziale	Indirizzo finale
10.0.0.0/8	10.0.0.0	10.255.255.255
172.16.0.0/12	172.16.0.0	172.31.255.255
192.168.0.0/16	192.168.0.0	192.168.255.255
169.254.0.0/16	169.254.0.0	169.254.255.255

Indirizzi privati: ambito di impiego

- La carenza di indirizzi IP ed il costo degli archi di indirizzamento sono alla base dell'utilizzo degli indirizzi privati
  - le reti con una solo punto di connessione alla Big Internet possono utilizzare l'indirizzamento privato

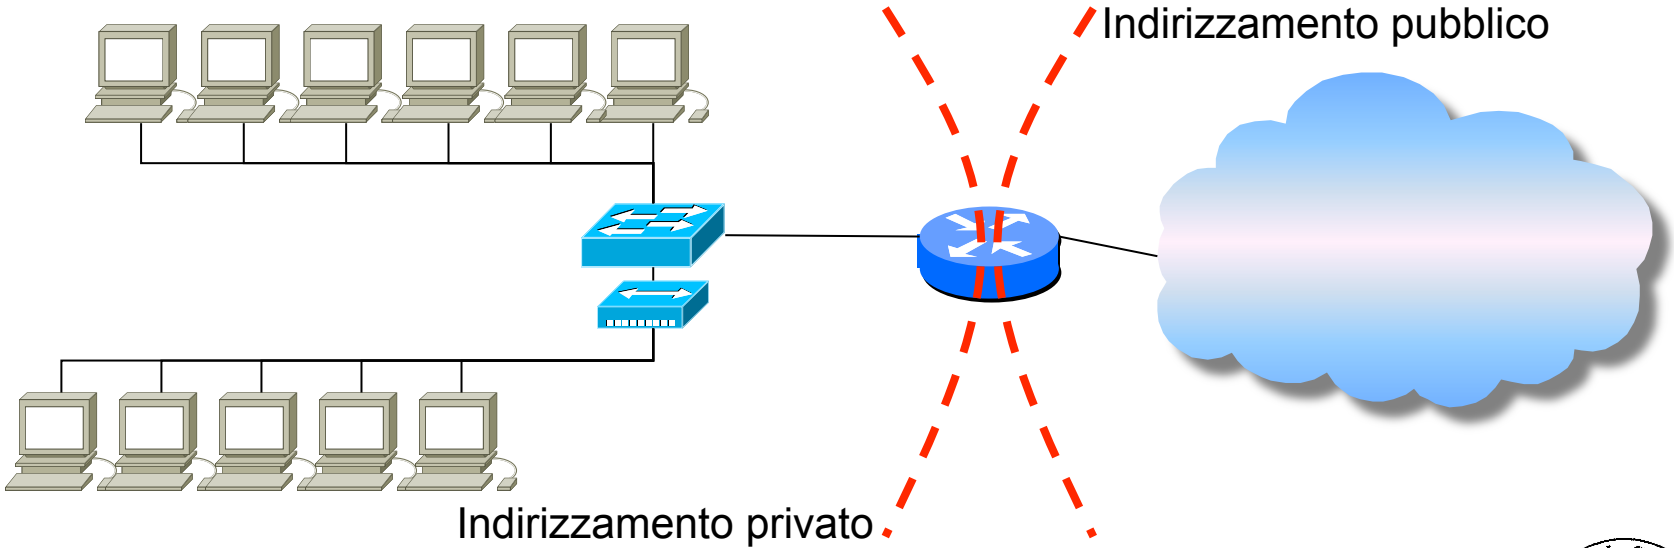

Indirizzi privati: instradamento (1)

- E' necessario introdurre un'ulteriore funzionalità sul bordo tra privato/pubblico per permettere di ricevere i pacchetti all'interno della rete privata

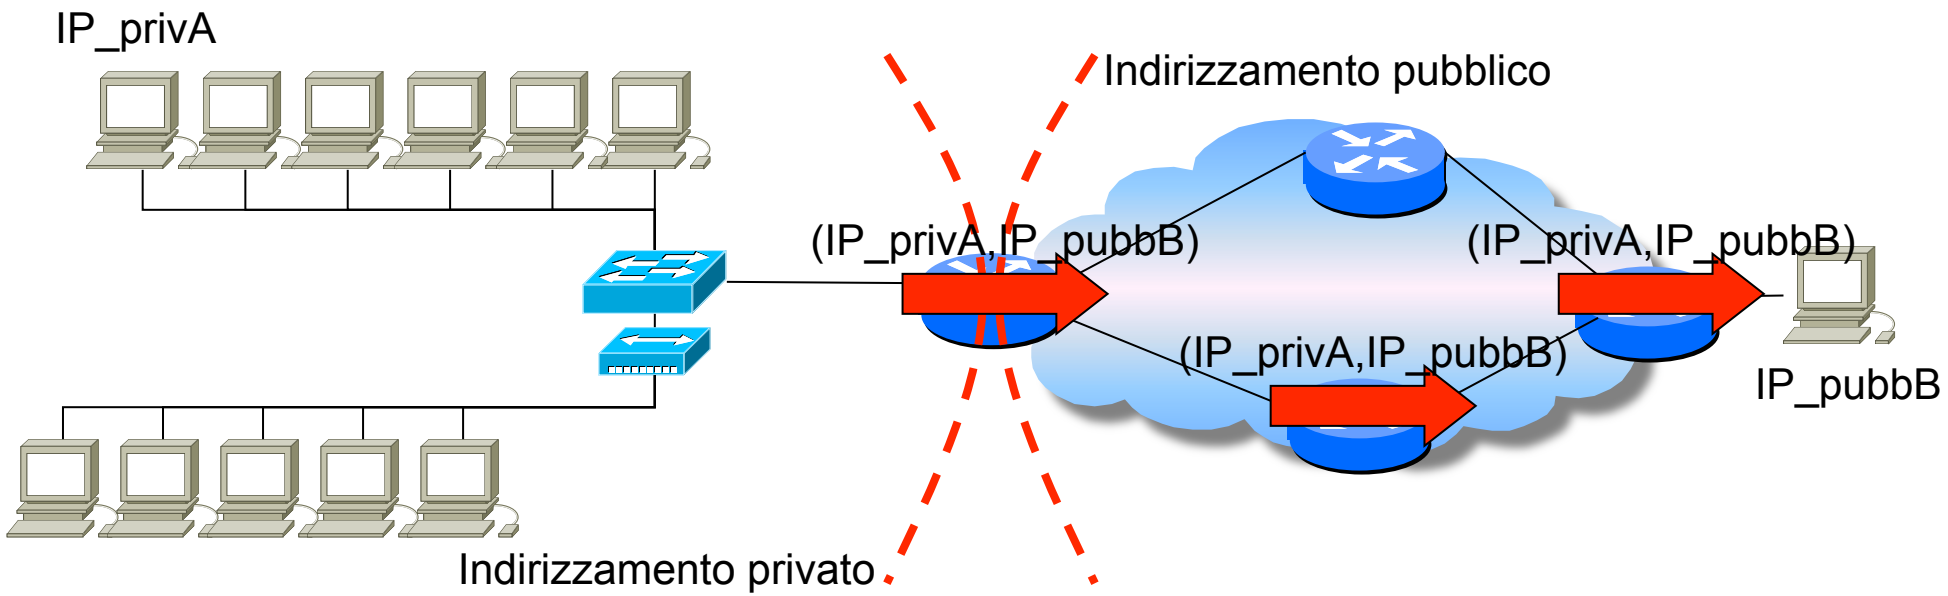

Indirizzi privati: instradamento (2)

- E' necessario introdurre un'ulteriore funzionalità sul bordo tra privato/pubblico per permettere di ricevere i pacchetti all'interno della rete privata

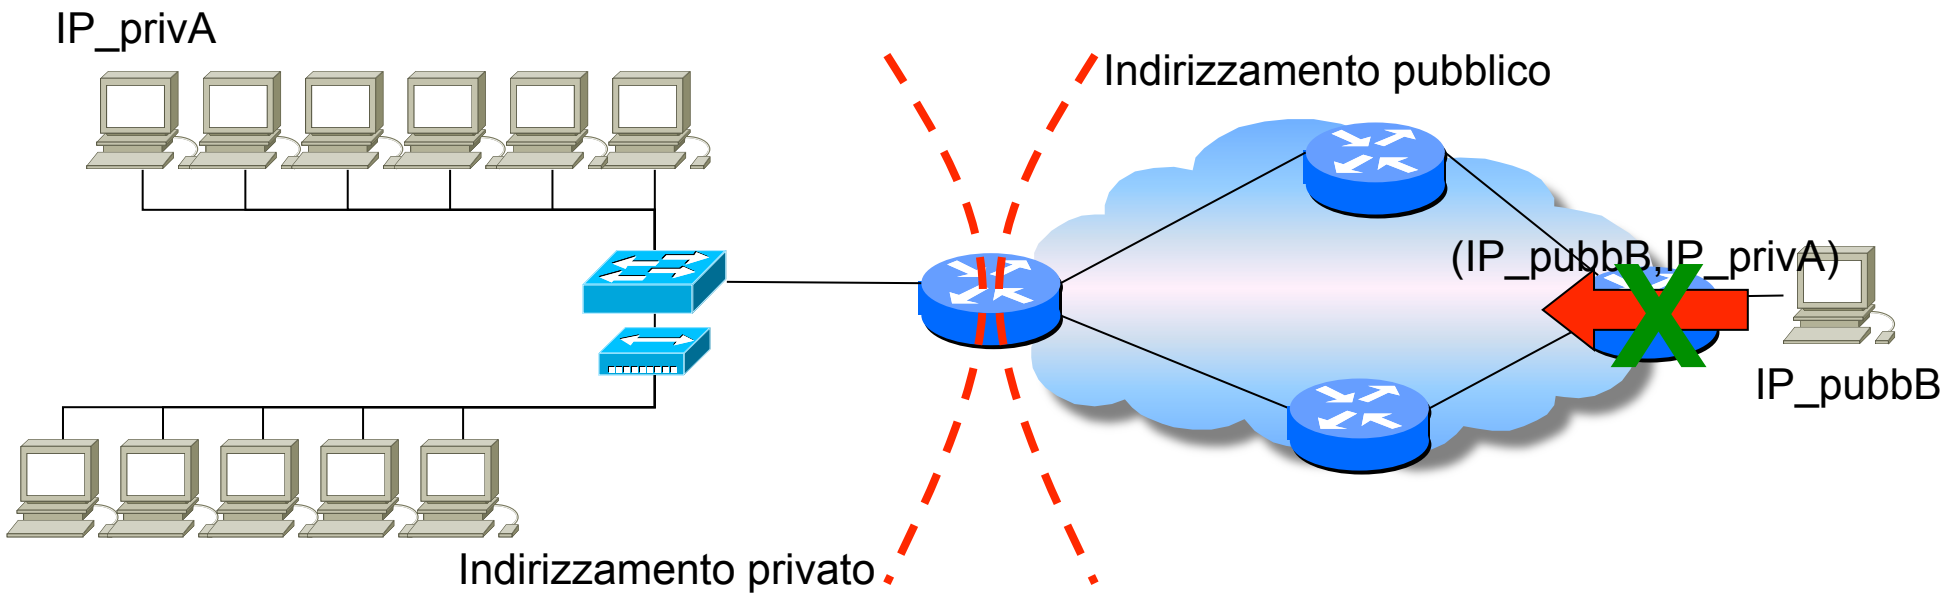

Network Address Translation (1)

- Network Address Translation: funzionalità introdotta per risolvere i problemi di instradamento tra una rete ad indirizzamento privato ed una rete ad indirizzamento pubblico
- Al router di confine tra privato e pubblico viene assegnato un indirizzo pubblico sull' interfaccia verso la rete esterna

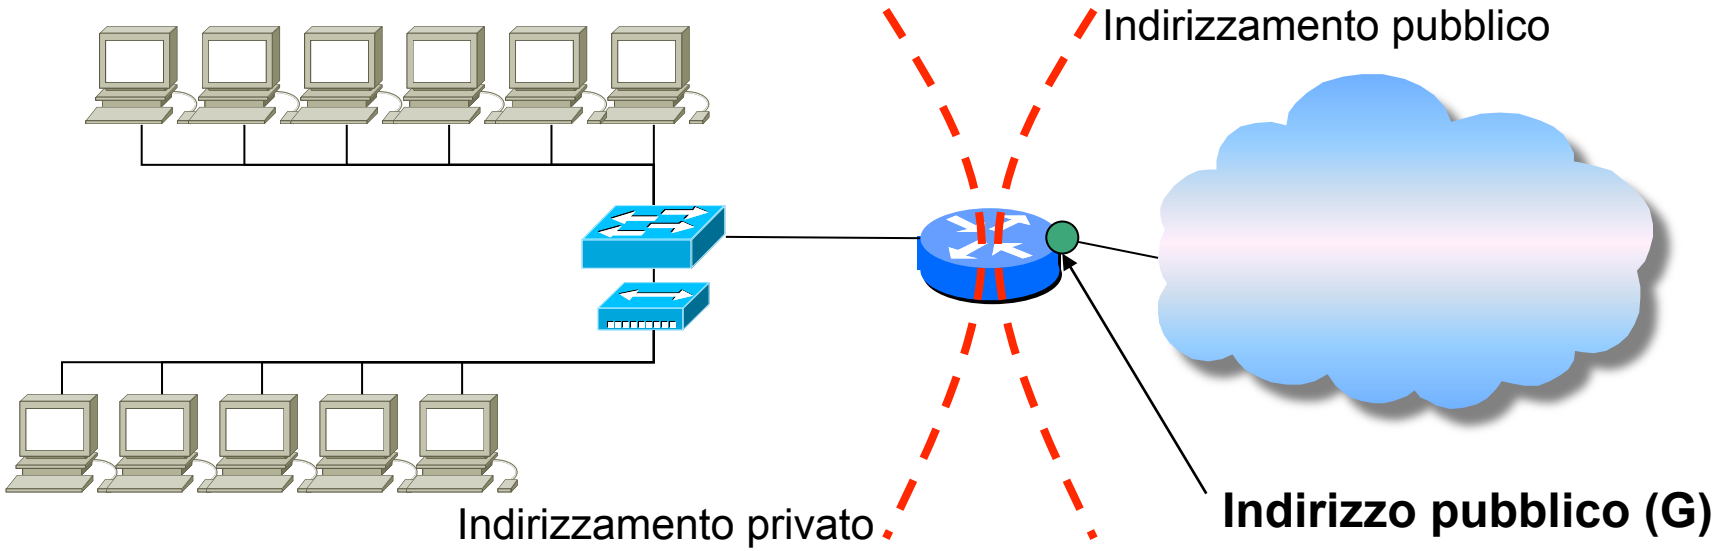

Network Address Translation (2)

- Al router di bordo (privato/pubblico) viene asegnata la funzionalità di **Network Address Translation**
  - NAT traduce l'indirizzo IP dei datagrammi uscenti ed entranti sostituendo
    - l'indirizzo sorgente di ogni pacchetto uscente con il proprio indirizzo pubblico
    - l'indirzzo destinazione di ogni pacchetto entrante con l'indirizzo privato dell'host corretto

Network Address Translation (4)

- NAT traduce l'indirizzo IP dei datagrammi uscenti ed entranti sostituendo
  - l'indirizzo sorgente di ogni pacchetto uscente con il proprio indirizzo pubblico

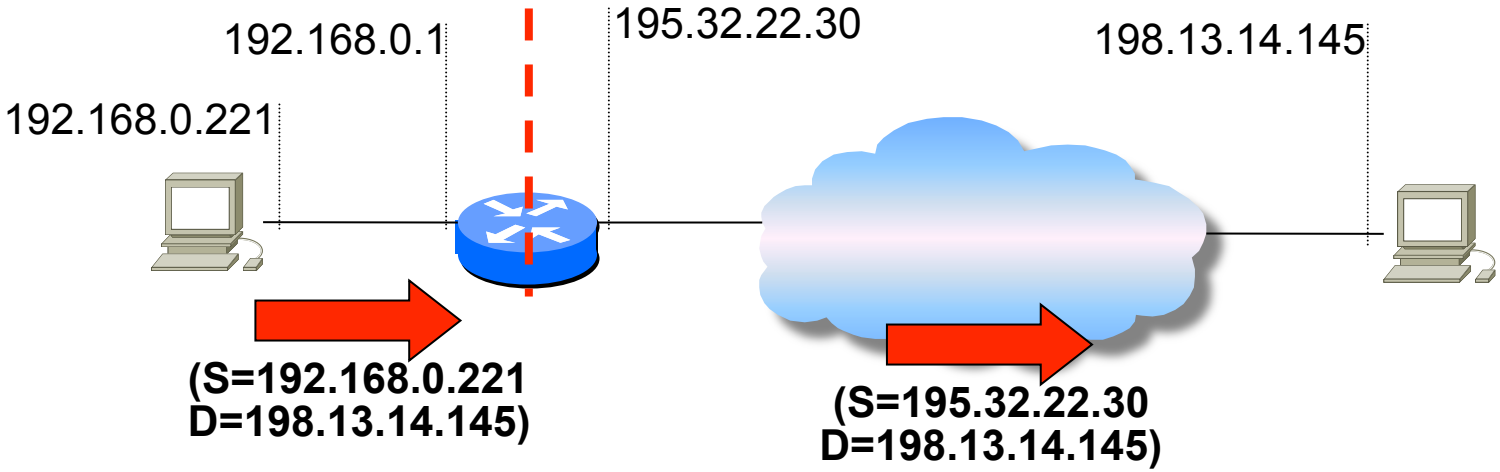

Network Address Translation (5)

- NAT traduce l'indirizzo IP dei datagrammi uscenti ed entranti sostituendo
  - l'indirzzo destinazione di ogni pacchetto entrante con l'indirizzo privato dell'host corretto

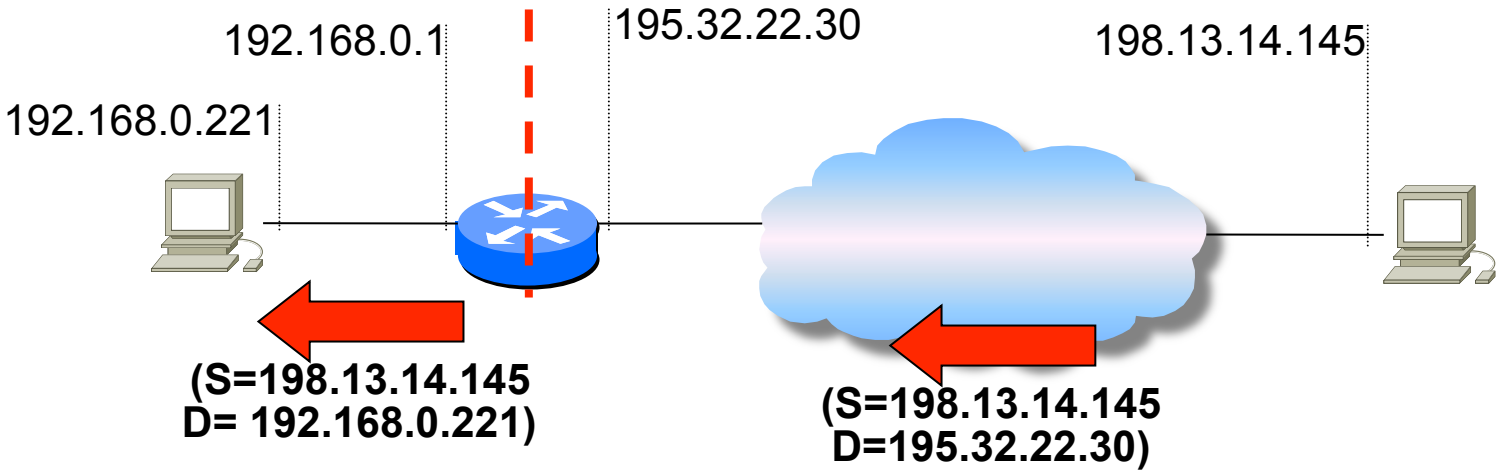

Network Address Translation Table (1)

- Il router NAT mantiene al suo interno una tabella di record con il mapping tra indirizzo privato sorgente della comunicazione ed indirizzo pubblico destinazione della comunicazione

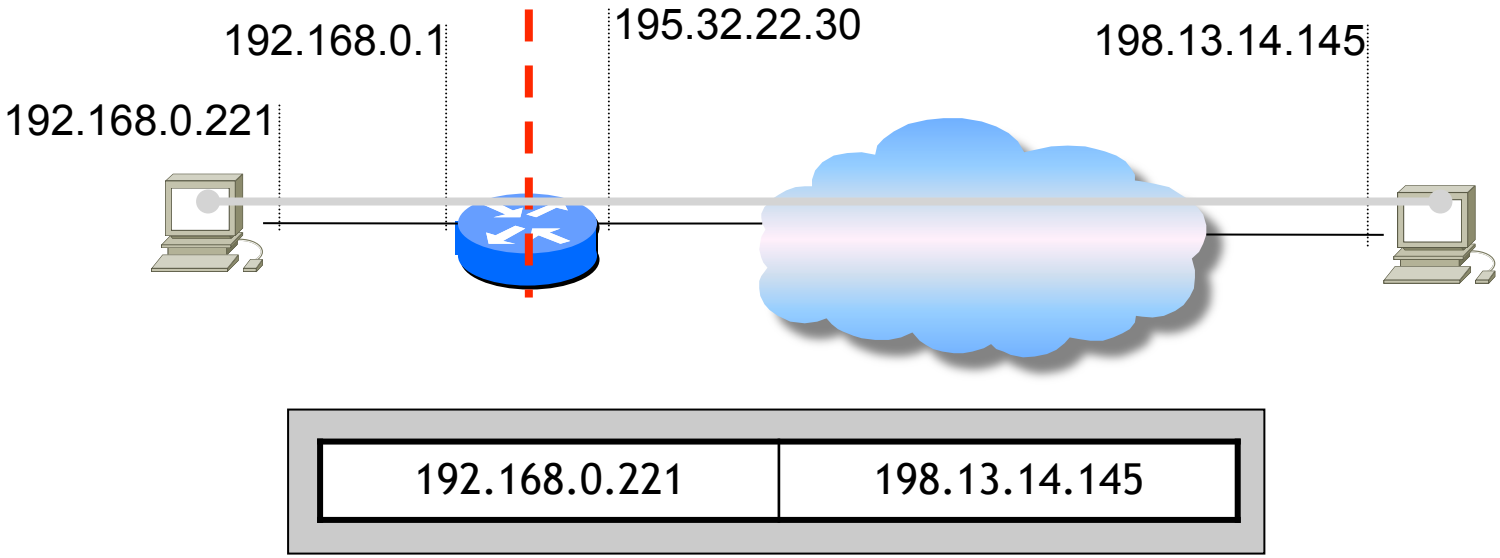

Network Address Translation Table (2)

- Metodi di aggiornamento della NAT Table:
  - Configurazione manuale
    - il gestore della rete configura in modo statico i record della NAT Table
  - Datagrammi uscenti
    - i record vengono creati in modo dinamico ogni volta che un pacchetto verso una data destinazione attraversa il NAT
    - cancellati con meccanismo di timeout

Network Address Translation Table (3)

Configurazione manuale	
Vantaggi	*Svantaggi*
Possibilità permanente di pacchetti in ingresso ed in uscita	Record statici

Datagrammi uscenti	
Vantaggi	*Svantaggi*
Record dinamici	Non permettono l'attivazione di una comunicazione dall'esterno

Limitazioni

- Il NAT basato unicamente sull'indirizzo non permette a differenti host privati di connettersi contemporaneamente allo stesso host pubblico

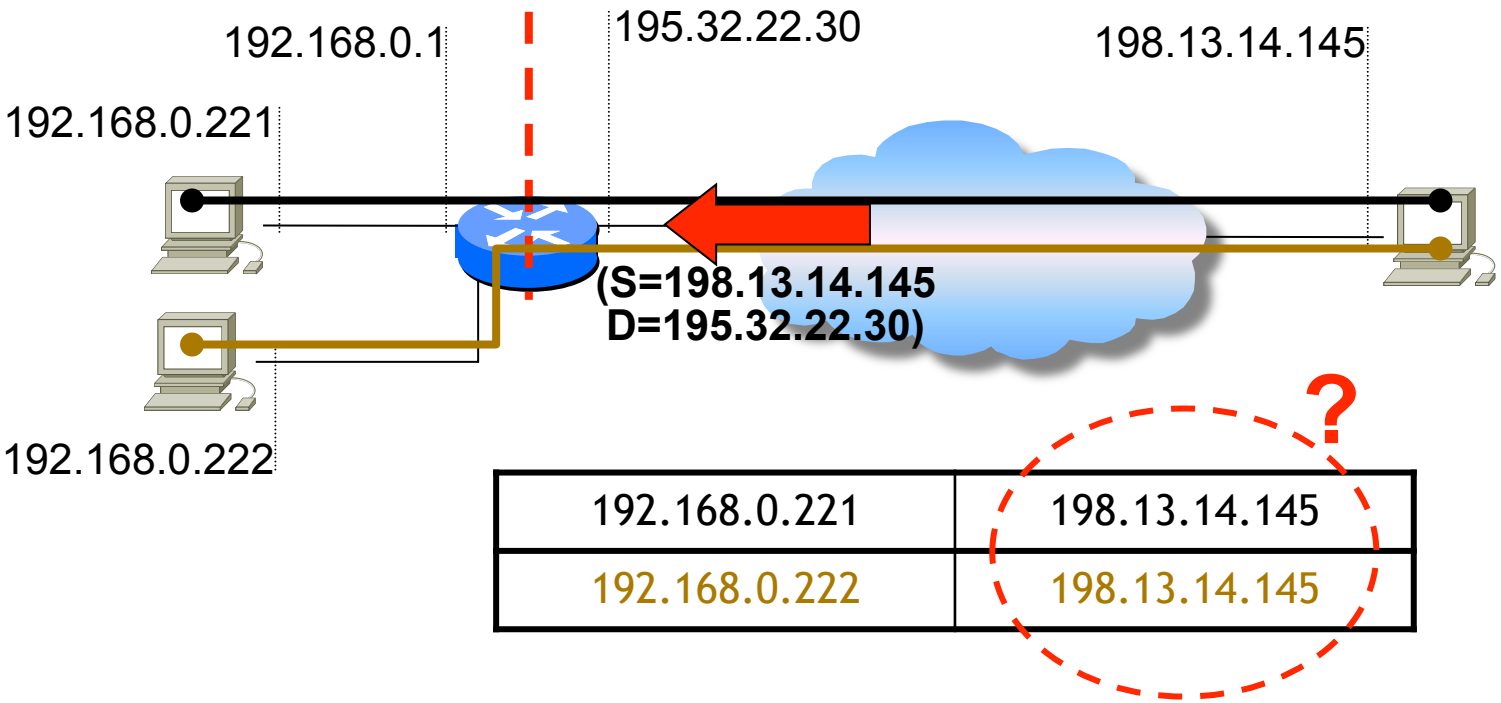

Port mapped NAT (1)

- Il router NAT agisce da gateway di livello 4
  - traduzione sia dell'indirizzo IP che della porta (TCP/UDP)

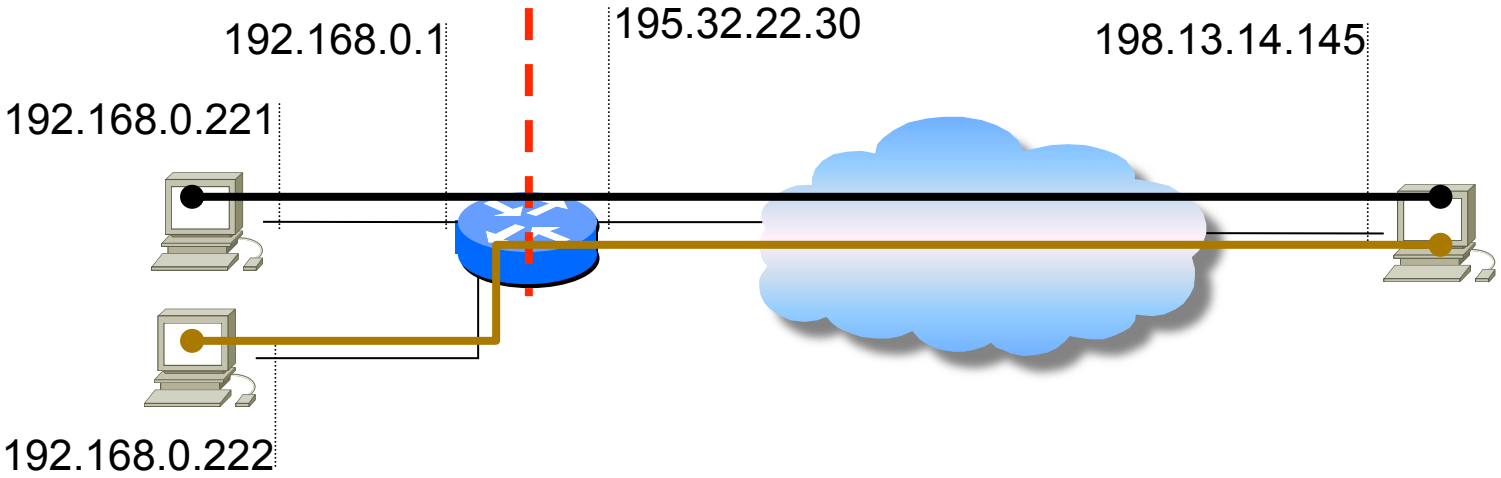

Priv. Addr	Priv. Port	Ext. Addr	Ext. Port	NATport	Prot. 4
192.168.0.221	21023	198.13.14.145	80	14001	TCP
192.168.0.222	13444	198.13.14.145	80	14002	TCP

Port mapped NAT (3)

- Datagrammi uscenti

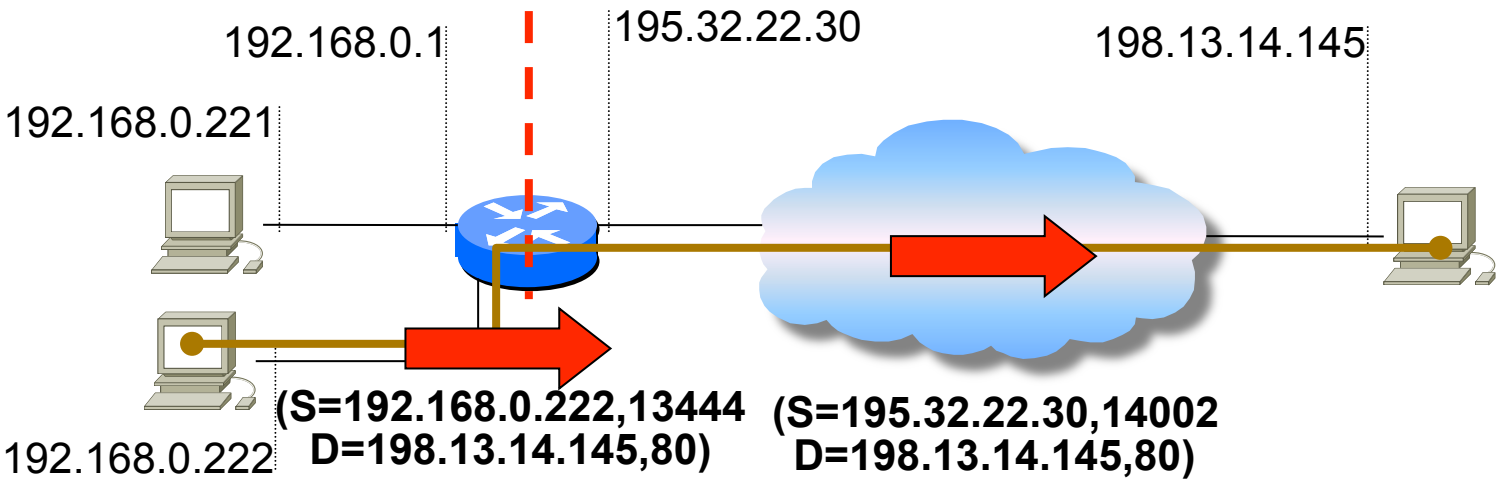

Priv. Addr	Priv. Port	Ext. Addr	Ext. Port	NATport	Prot. 4
192.168.0.222	13444	198.13.14.145	80	14002	TCP

Port mapped NAT (4)

- Datagrammi entranti

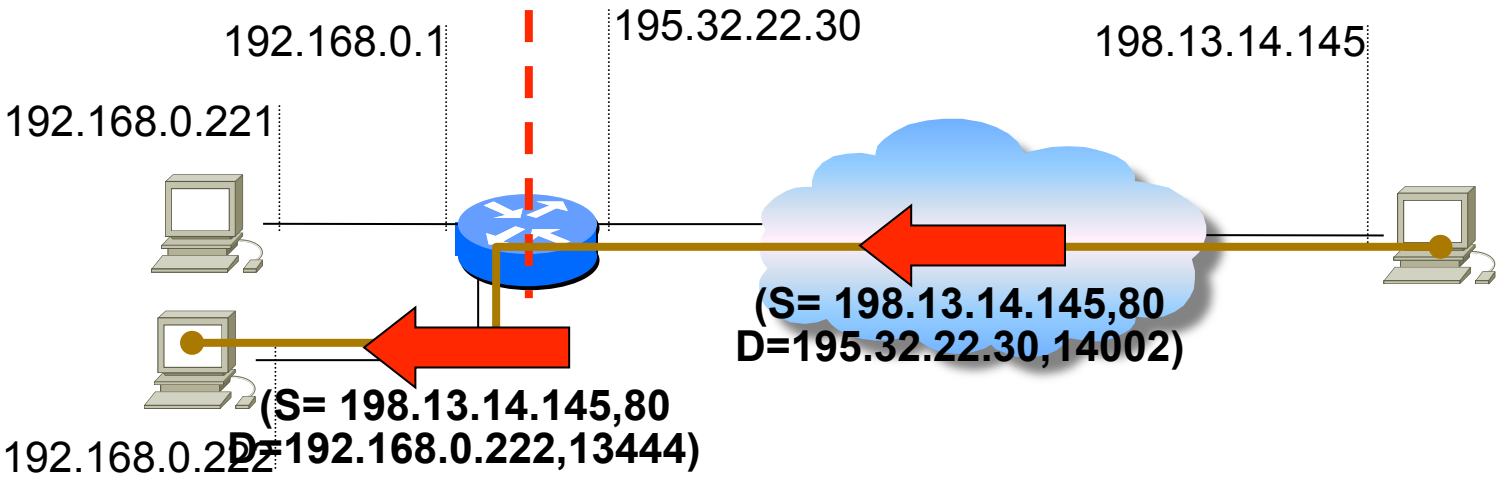

Priv. Addr	Priv. Port	Ext. Addr	Ext. Port	NATport	Prot. 4
192.168.0.222	13444	198.13.14.145	80	14002	TCP

IPv6

The Motivation for Change

- When IP was defined, the 32 bits IP address were selected
  - doing so allowed the Internet could include over a million networks
- The global Internet is growing exponentially
  - Its size is doubling in less than a year
- If the current growth rate maintained
  - each of the possible network prefixes will eventually be assigned
  - and no further growth will be possible

The Motivation for Change (cont'd)

- Motivation for defining a new version of IP?
  - the address space limitation
    - larger addresses are necessary to accommodate continued growth
  - special facilities are needed for some applications
- Consequently, when IP is replaced
  - the new version should have more features
    - For example, is has been argued that a new version of IP should provide a mechanism for carrying real-time traffic to avoid route changes

The Motivation for Change (cont'd)

- A new version of IP should accommodate more complex addressing and routing capabilities
- For example, Google maintains many data centers
  - When a user enters google.com into a browser, it would be efficient if IP passed datagrams to the nearest Google data center
- Many current applications allow a set of users to collaborate
  - To make collaboration efficient
    - Internet needs a mechanism that allows groups to be created or changed
    - It needs a way to send a copy of a packet to each participant in a given group

The Hourglass Model and Difficulty of Change

- Scarcity of available addresses was considered crucial when work began on a new version of IP in 1993
  - no emergency occurred
  - and IP has not been changed
- Think of the importance of IP and the cost to change!
  - IP lies at the center of Internet communication
- Networking professionals argue that Internet communication follows an hourglass model
  - and that IP lies at the position where the hourglass is thin

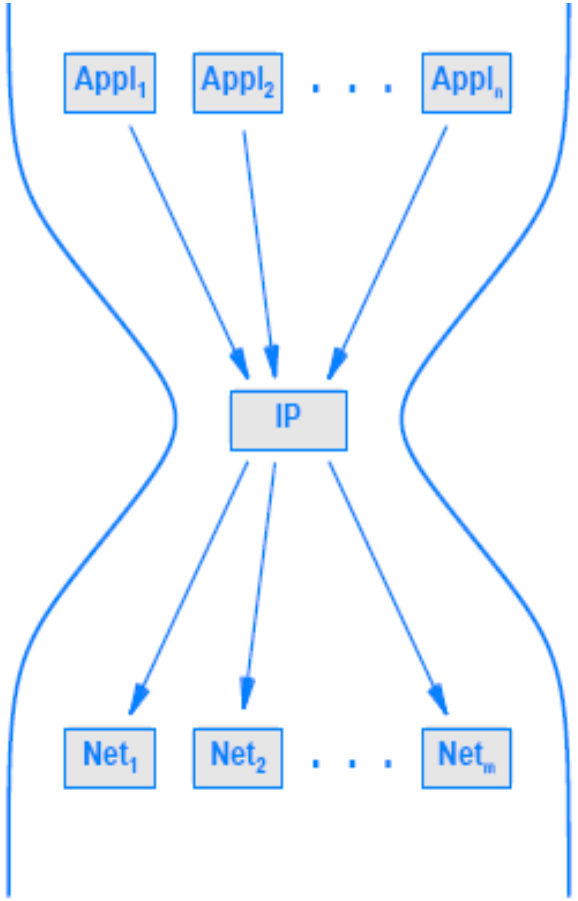

24

A Name and a Version Number

- Researchers selected IP The Next Generation
  - and early reports referred to the new protocol as IPng
  - many competing proposals were made for Ipng
- New IP version number that was selected as a surprise
  - Because the current IP version number is 4 (IPv4)
    - the networking community expected the next official version to be 5
    - version 5 was assigned to an experimental protocol known as ST
  - The new version of IP received 6 as its official version number (IPv6)

IPv6 Features

- IPv6 retains many of the successful features of IPv4 design, such as
  - Like IPv4, IPv6 is connectionless
  - Like IPv4, the header in a datagram contains a maximum number of hops the datagram can take before being discarded
- Despite retaining the basic concepts from the current version, IPv6 changes all the details
- Features of IPv6 can be grouped into a number of broad categories (see next slides)

IPv6 Features

- Address Size
  - Instead of 32 bits, each IPv6 address contains 128 bits.
  - The resulting address space is large enough to accommodate continued growth of the world-wide Internet for many decades
- Header Format
  - The header is completely different from the IPv4 header
  - Almost every field in the header has been changed (some were replaced)
- Extension Headers
  - IPv6 encodes information into separate headers
    - A datagram consists of the base IPv6 header followed by zero or more extension headers, followed by data

IPv6 Features

- Support for Real-Time Traffic
  - a mechanism exists that allows a sender and receiver to establish a high-quality path and to associate datagrams with that path
  - the mechanism is intended for use with audio and video applications
  - the mechanism can also be used to associate datagrams with low-cost paths
- Extensible Protocol
  - IPv6 allows a sender to add additional information to a datagram
  - The extension scheme makes IPv6 more flexible than IPv4
    - and means that new features can be added to the design as needed

IPv6 Datagram Format

- An IPv6 datagram contains a series of headers
  - each datagram begins with a base header
  - followed by zero or more extension headers
  - followed by the payload

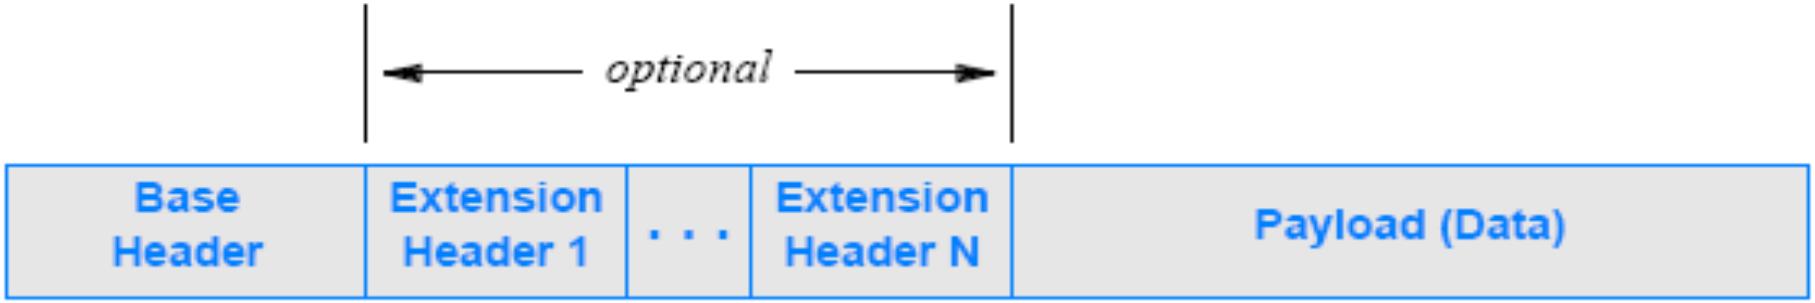

- Fields are not drawn to scale
  - some extension headers are larger than the base header
  - In many datagrams, the size of the payload is much larger than the size of the header

IPv6 Base Header Format

- Although it is twice as large as an IPv4 header, the IPv6 base header contains less fields

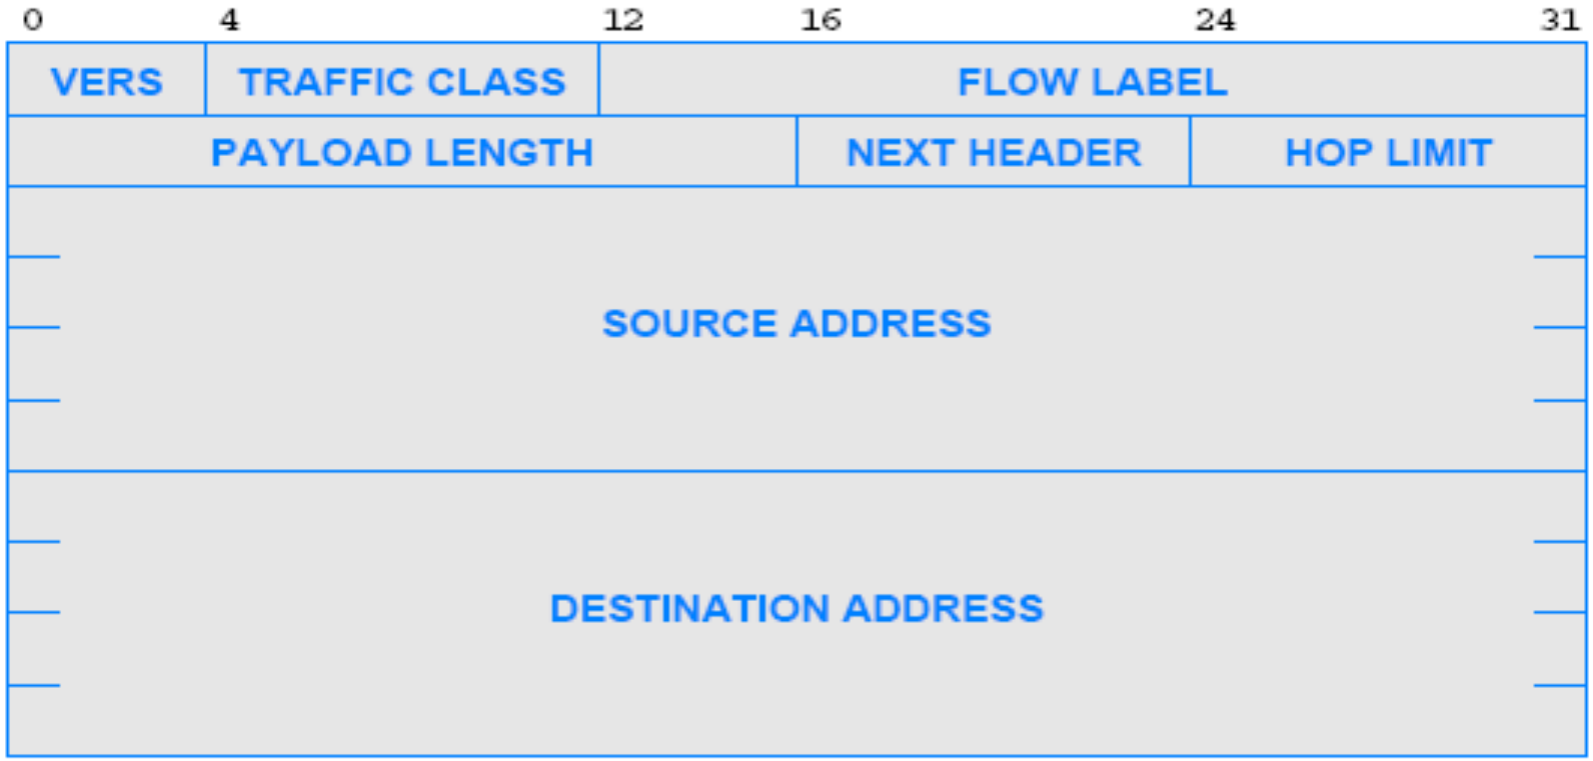

IPv6 Base Header Format

- VERS (Version 6)
- TRAFFIC CLASS
  - specifies the traffic class using a definition of traffic types
  - It is known as differentiated services to specify general characteristics that the datagram needs
  - For example,
    - To send interactive traffic (e.g., keystrokes/mouse) → one might specify a class that has low latency
    - To send real-time audio across the Internet → a sender might request a path with low jitter
- PAYLOAD LENGTH
  - corresponds to IPv4's datagram length field
  - it specifies only the size of the data being carried

IPv6 Base Header Format

- HOP LIMIT
  - corresponds to the IPv4 TIME-TO-LIVE field
- Field FLOW LABEL
  - intended to associate a datagram with a particular path
- NEXT HEADER
  - is used to specify the type of information that follows the current header
  - If the datagram includes an extension header
    - NEXT HEADER field specifies the type of the extension header
  - If no extension header exists
    - NEXT HEADER field specifies the type of data being carried in the payload

Next Header: example

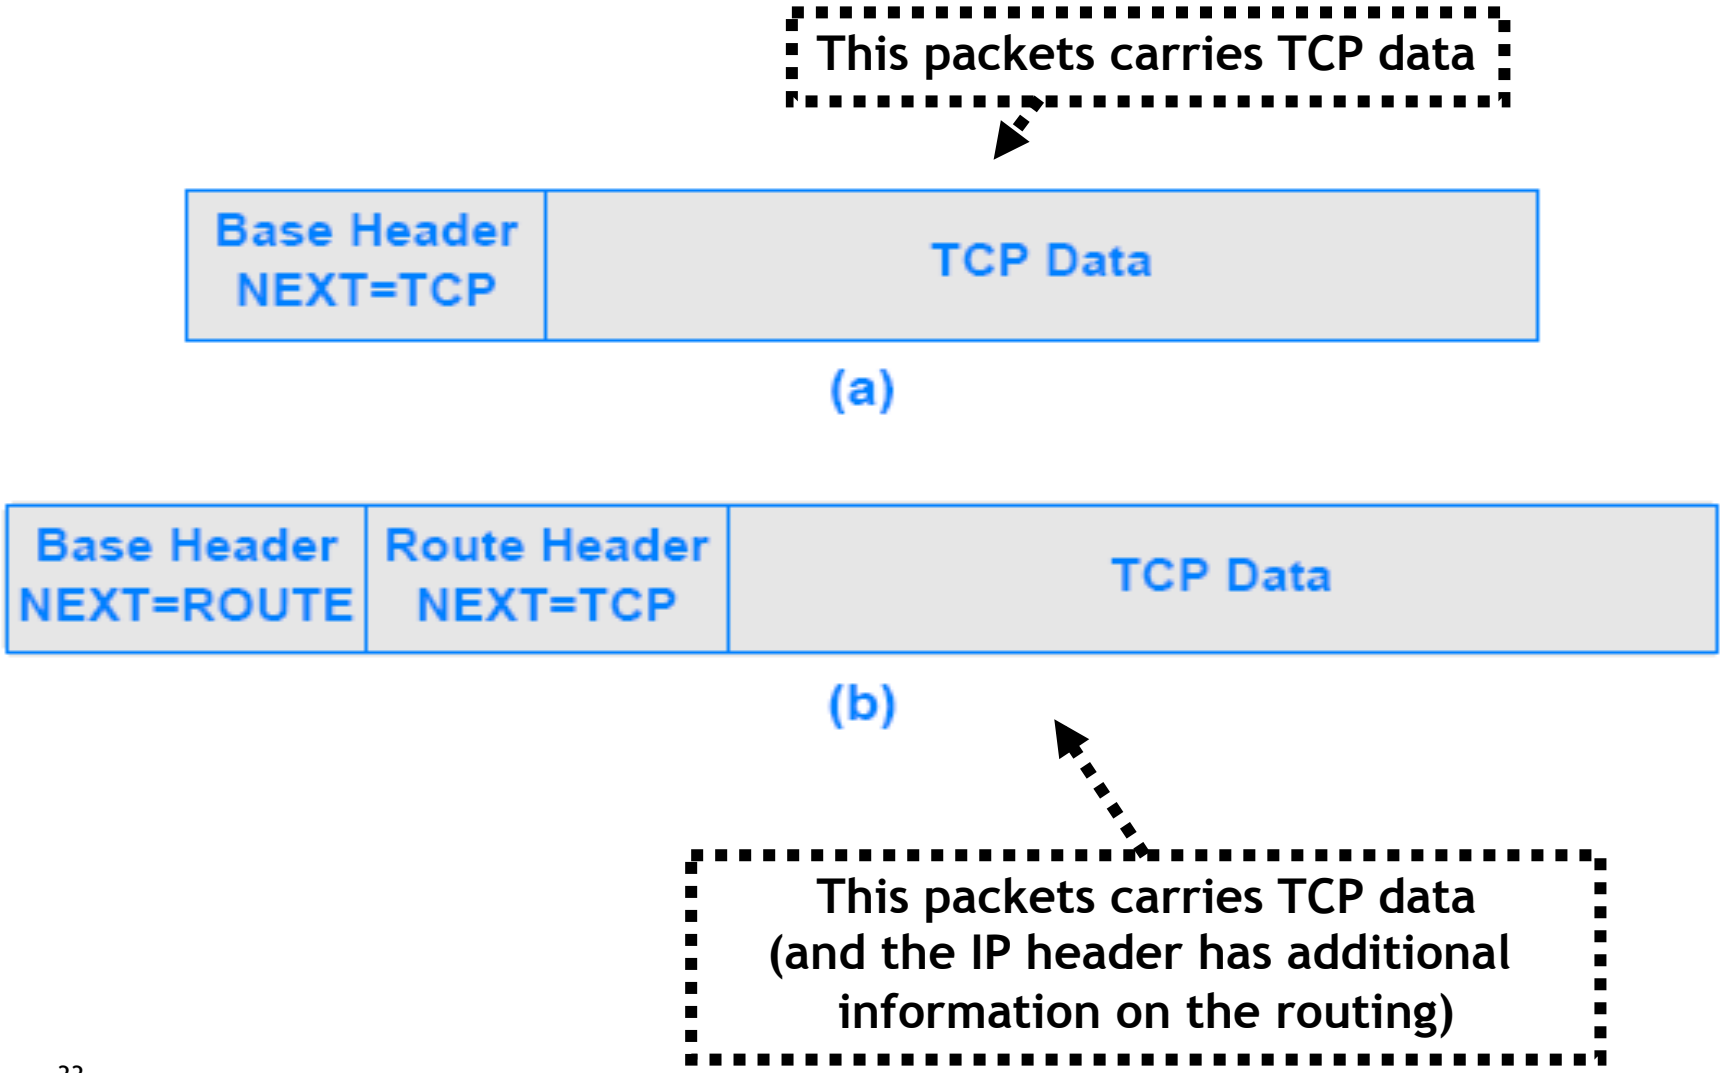

Implicit and Explicit Header Size

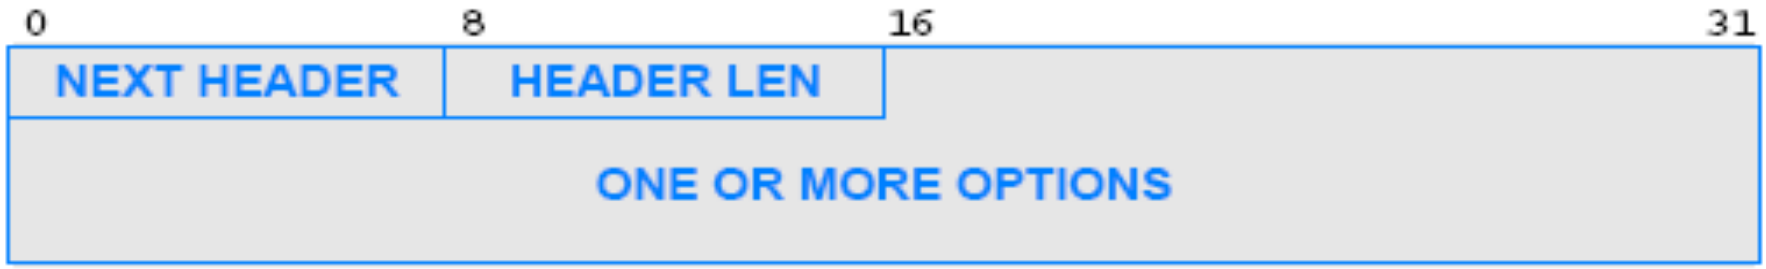

- No ambiguity about the interpretation of the NEXT HEADER
  - the standard specifies a unique value for each possible header
- A receiver processes headers sequentially
  - NEXT HEADER field in each header to determine what follows
- Some header types have a fixed size
  - For example, a base header has a fixed size of exactly 40 octets
- Some extension headers do not have a fixed size
  - the header must contain sufficient information to allow IPv6 to determine where the header ends

Fragmentation, Reassembly, and Path MTU

- IPv6 fragmentation resembles IPv4 fragmentation
- There are some differences between them
- Like IPv4
  - a prefix of the original datagram is copied into each fragment
  - and the payload length is modified to be the length of the fragment
- Unlike IPv4, however
  - It does not include fields for fragmentation in the base header
  - It places the fragment information in a separate fragment extension header
    - the presence of the header identifies the datagram as a fragment

Fragmentation, Reassembly, and Path MTU

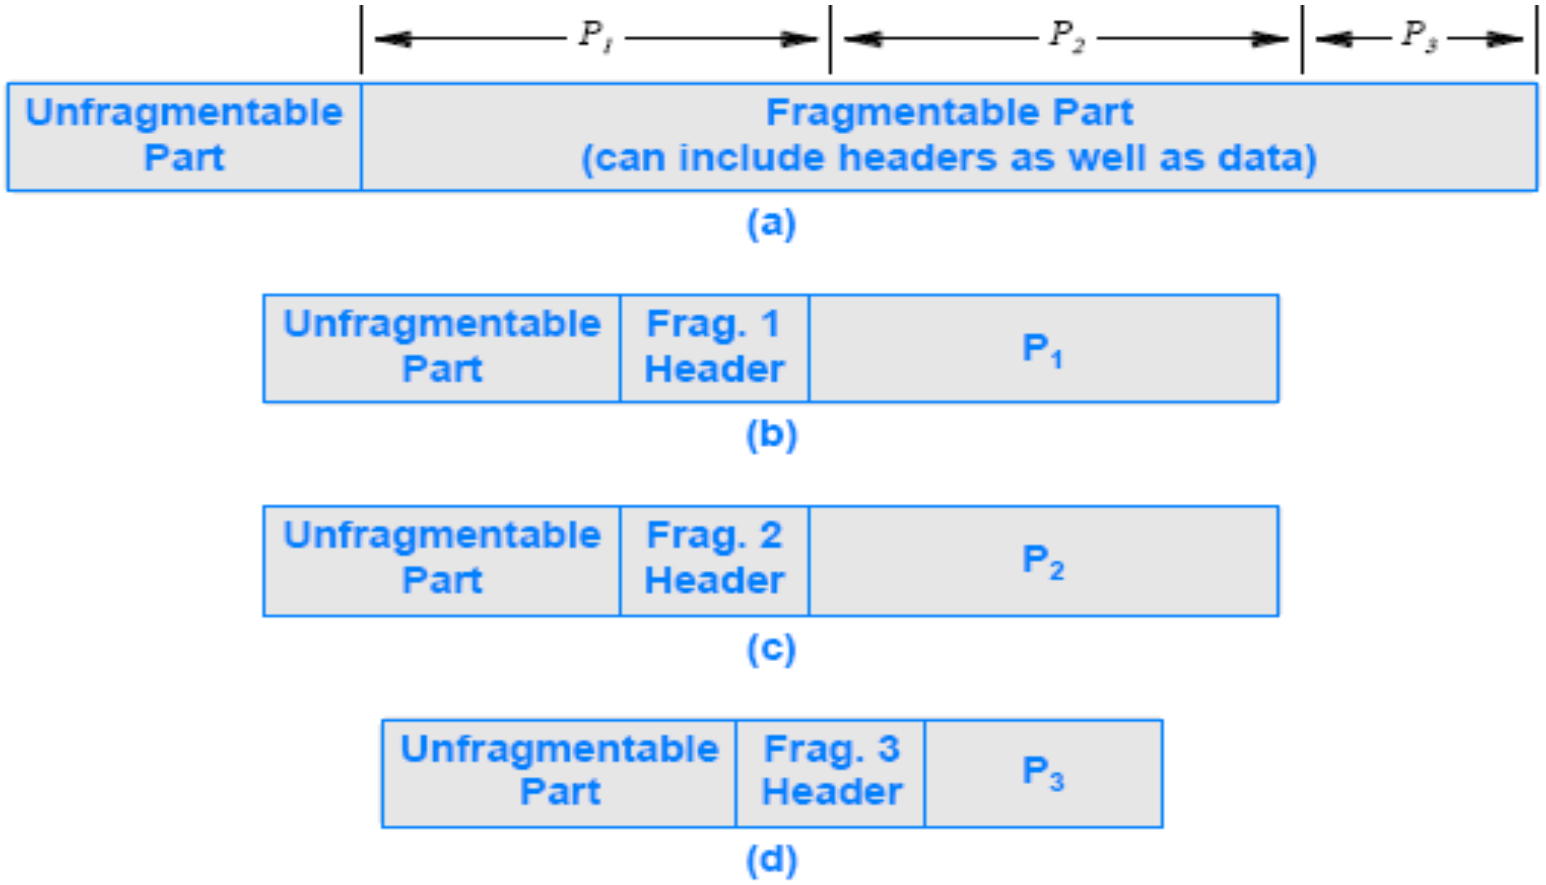

Fragmentation, Reassembly, and Path MTU

- The Unfragmentable Part denotes the base header plus headers that control routing
- To insure that all fragments are routed identically
  - the unfragmentable part is replicated in every fragment
- Fragment size is chosen to be the Maximum Transmission Unit (MTU) of the underlying network
  - the final fragment may be smaller than the others

Fragmentation, Reassembly, and Path MTU

- There is a key point where IPv6 differs dramatically from fragmentation in IPv4
- In IPv4, a router performs fragmentation
  - when the router receives a datagram too large for the network over which the datagram must be sent
- In IPv6, a sending host is responsible for fragmentation
- A router along the path that receives a datagram that is larger than the network MTU
  - will send an error message and discard the datagram

Fragmentation, Reassembly, and Path MTU

- How can a host choose a datagram size that will not result in fragmentation?
  - The host must learn the MTU of each network along the path
    - and must choose a datagram size to fit the smallest
  - The minimum MTU along a path from a source to a destination is known as the path MTU
  - The process of learning the path MTU is known as path MTU discovery

Fragmentation, Reassembly, and Path MTU

- In general, path MTU discovery is an iterative procedure
- A host sends a sequence of various-size datagrams to the destination
  - to see if they arrive without error
- If fragmentation is required
  - the sending host will receive an ICMP error message
  - IPv6 includes a new version of ICMP
- Once a datagram is small enough to pass through without fragmentation
  - the host chooses a datagram size equal to the path MTU

The Purpose of Multiple Headers

- Why does IPv6 use separate extension headers?
- There are two reasons:
  - Economy
  - Extensibility
- Economy is easiest to understand:
  - because it saves space
  - designers expect a given datagram to use only a small subset
  - it is possible to define a large set of features
    - without requiring each datagram header to have at least one field for each
- To understand extensibility
  - consider adding a new feature to a protocol
  - The IPv4 requires a complete change to accommodate new feature
  - In IPv6, however, existing protocol headers can remain unchanged
    - A new NEXT HEADER type is defined as well as a new header format

IPv6 Addressing

- Address details are completely different
  - Like CIDR addresses, the division between prefix and suffix can occur on an arbitrary boundary
  - Unlike IPv4, IPv6 includes addresses with a multi-level hierarchy
  - Although the address assignments are not fixed, one can assume that
    - the highest level corresponds to an ISP
    - the next level corresponds to an organization (e.g., a company)
    - the next to a site, and so on
- IPv6 defines a set of special addresses
  - that differ from IPv4 special addresses
  - Each IPv6 address is one of the three basic types listed in the next slide

IPv6 Addressing

Type	Purpose
unicast	The address corresponds to a single computer. A datagram sent to the address is routed along a shortest path to the computer.
multicast	The address corresponds to a set of computers, and membership in the set can change at any time. IPv6 delivers one copy of the datagram to each member of the set.
anycast	The address corresponds to a set of computers that share a common prefix. A datagram sent to the address is delivered to exactly one of the computers (e.g., the computer closest to the sender).

IPv6 Colon Hexadecimal Notation

- IPv6 address occupies 128 bits
  - writing such numbers can be unwieldy
- Consider a 128-bit number in the dotted decimal notation:
  - 105.220.136.100.255.255.255.255.0.0.18.128.140.10.255.255
- To reduce the number of characters used to write addresses
  - the designers of IPv6 chose a more compact syntactic form known as colon hexadecimal notation, usually abbreviated colon hex
  - each group of 16 bits is written in hex with a colon separating groups
- When the above number is written in colon hex:
  - 69DC : 8864 : FFFF : FFFF : 0 : 1280 : 8C0A : FFFF

IPv6 Colon Hexadecimal Notation

- An additional optimization known as zero compression further reduces the size
  - Zero compression replaces sequences of zeroes with two (2) colons
  - For example, the address:
  - FF0C:0:0:0:0:0:0:B1 → FF0C : : B1
- The large IPv6 address spaces make zero compression especially important
  - the designers expect many IPv6 addresses to contain strings of zeroes
- To help ease the transition to the new protocol
  - The designers mapped existing IPv4 addresses into the IPv6 address space
  - Any IPv6 address that begins with 96-zero bits contains an IPv4 address in the low-order 32-bits