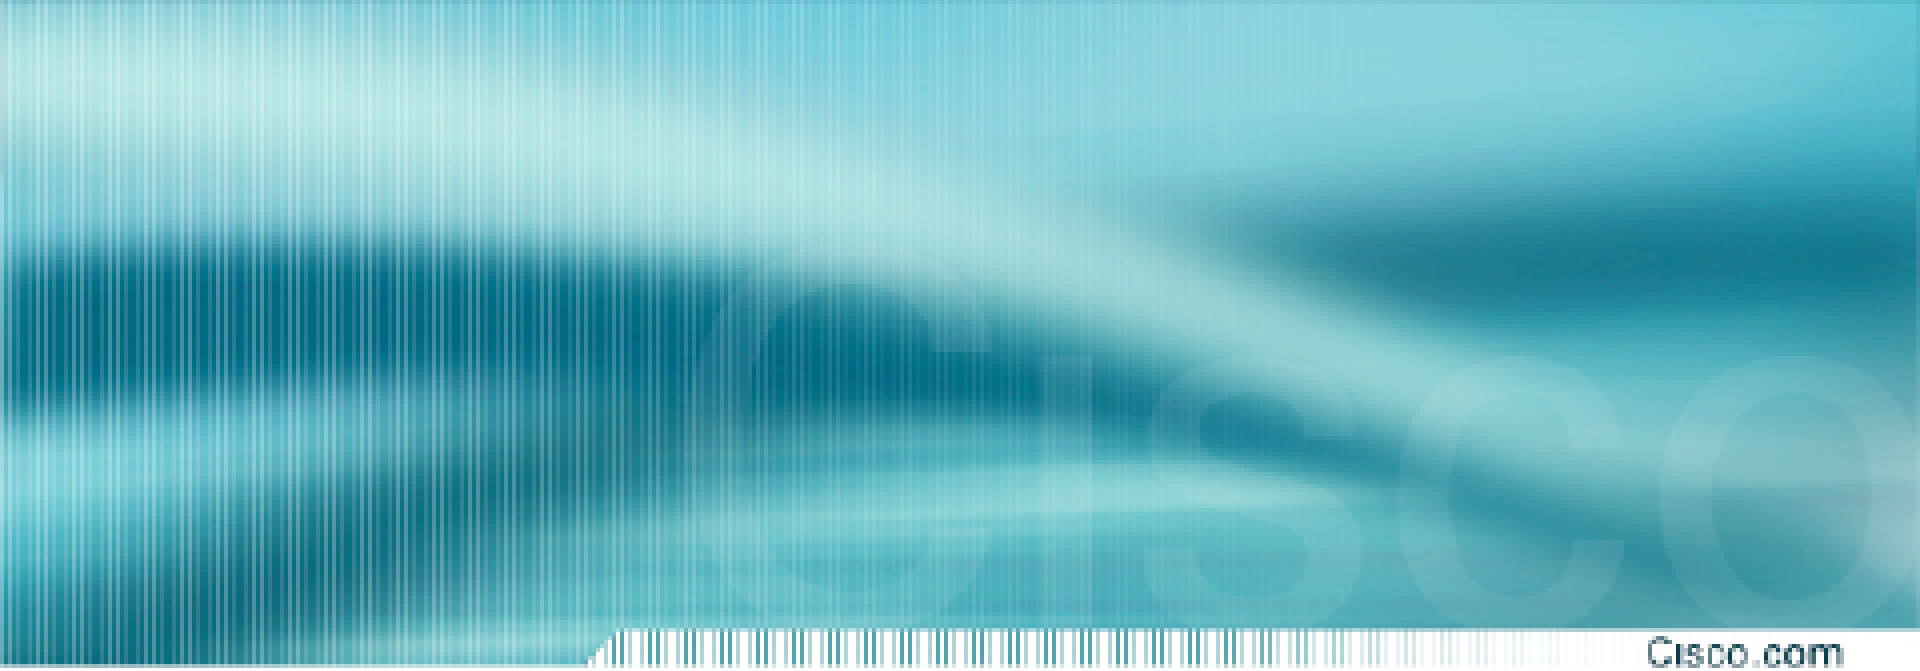

IPv6 Routing Workshop Apricot 2004

Phil Harris

Cisco Systems

Senior Consulting Engineer

pharris@cisco.com

Agenda - Day 1

Cisco.com

- Workshop Orientation
  - Workshop Objectives
  - Attendee Assumptions
  - Classroom Setup
- IPv6 Overview
  - IPv6 - The rationale for a new protocol
  - IPv6 - The v6 Protocol Development (v4 Comparison)
- Initial Network Setup
- Break for the day

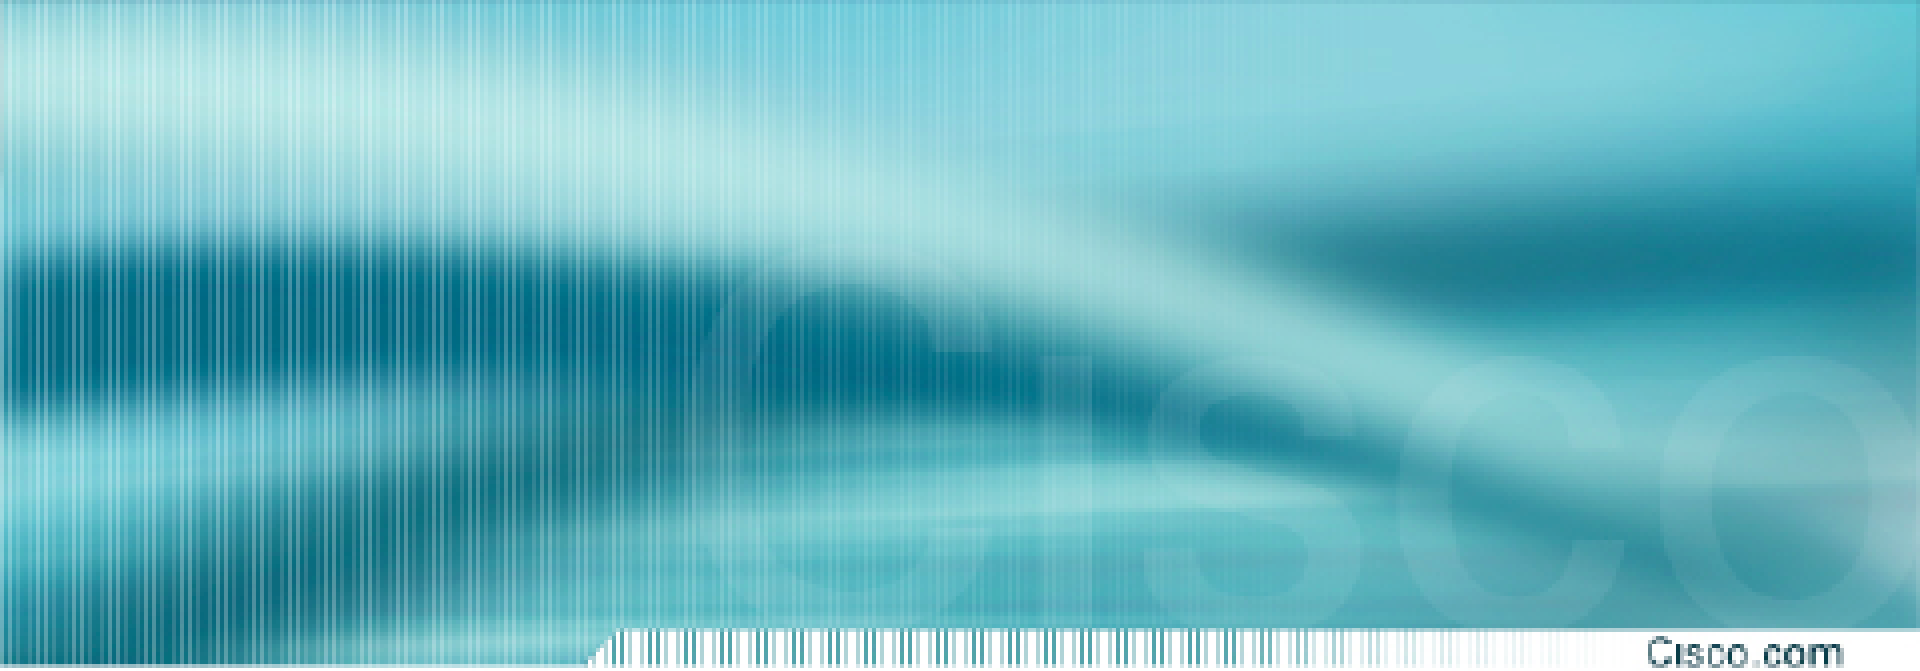

Workshop Orientation

Workshop Objectives

Cisco.com

- Describe the IPv6 Protocol from a requirements and technology perspective
- Give attendees a working knowledge of the protocol and its use and implementation
- Provide “Best Common Practices” (BCP’s) for v6 deployment
- Cover migration/transition strategies
- Discuss troubleshooting techniques

Attendee Assumptions (mine of you)

Cisco.com

- Attendees need to bring (with them)
  - An working knowledge of IPv4 protocols
  - An understanding of IP Routing Protocols (IGP's/EGP's)
  - A Laptop (desktops are heavy) with either wireless (802.11b) or ethernet port(s)
  - Reasonable familiarity with the IOS CLI
  - Willingness to participate :-)

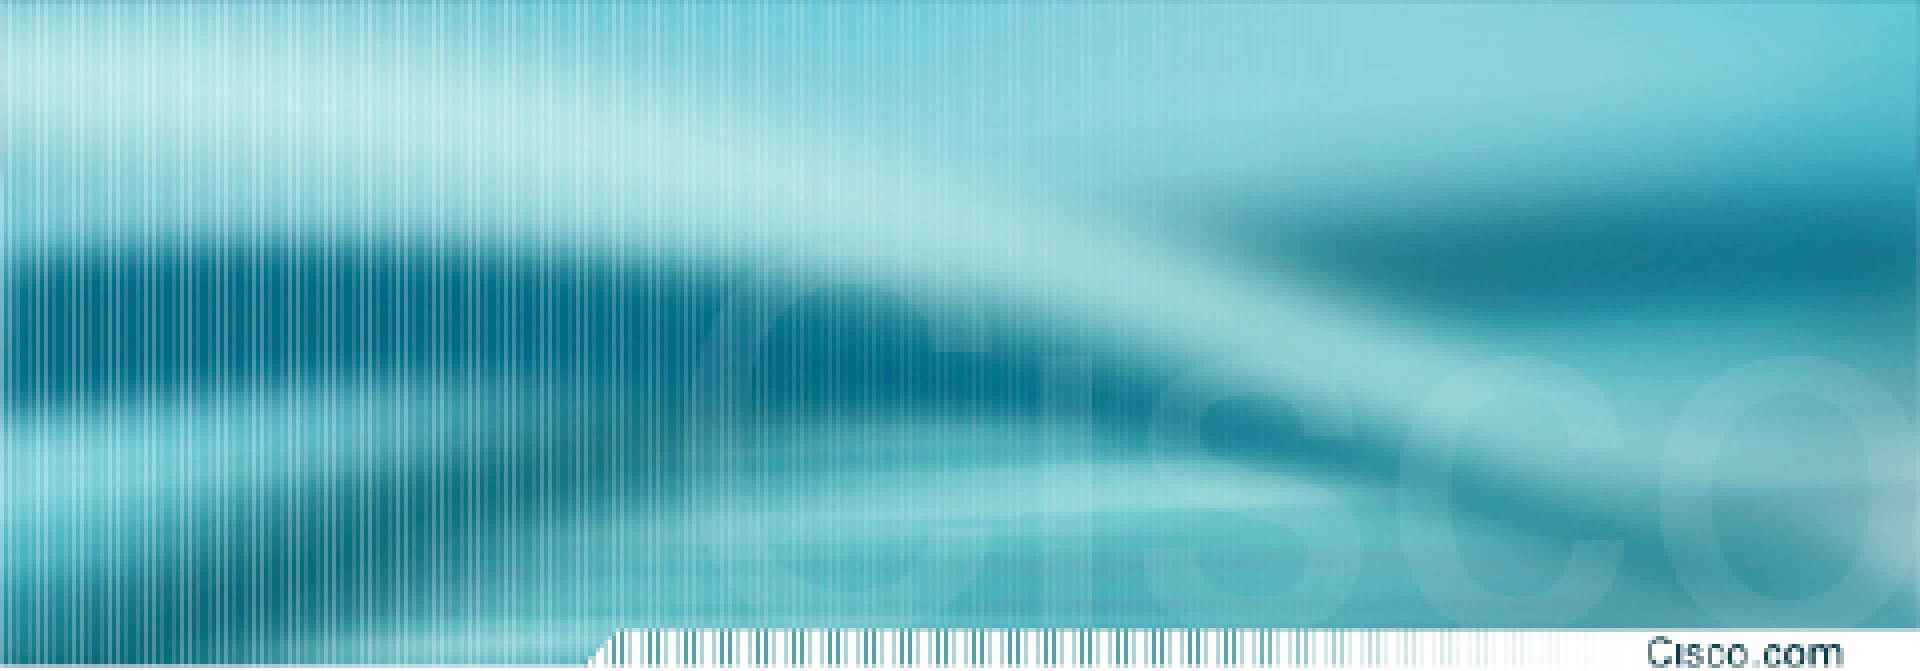

Classroom Setup

Initial Connectivity Check

Cisco.com

- Please boot your computers and attempt the receive an IPv4 address from the DHCP server.
- Everything is on 1 subnet (no need for a Default Gateway/Router)
- We will be using IP addresses not FQDN's
- make sure you can ping 10.1.1.10
- Documentation is via:-

http://10.1.1.9/~pharris/Apricot/v6.htm

Classroom Setup

Cisco.com

- The “Lab” is virtual in nature
- Please configure laptops/computers OS for DHCP
- Wireless devices should be configured with the following parameters:-

  SSID = APRICOT_V6

  WEP = ccaabbaacc (40 bit)

- Each attendee/team will be responsible for one or more routers
- Router Access will be via Telnet (from your Laptop)

Classroom Setup

Cisco.com

- Each router is accessed in the form:-

  telnet 10.0.0.10 {port no.}

- The port number = Router no. + 1300

  e.g. Router 13 = port 1313

- Please remember to “Write memory” after each configuration change

- Please don’t sit there stuck, the instructor and lab technician are here to help

Network Topology

Cisco.com

- The physical topology will remain constant throughout the workshop
- Logical topology will change as we progress through the labs
- Each Router has the following ports assigned:-

  Ethernet 0/0, 1/0, 2/0, 3/0

  Serial 4/0, 5/0

  (We will be using ethernet ports for the workshop)

Network (physical) Topology

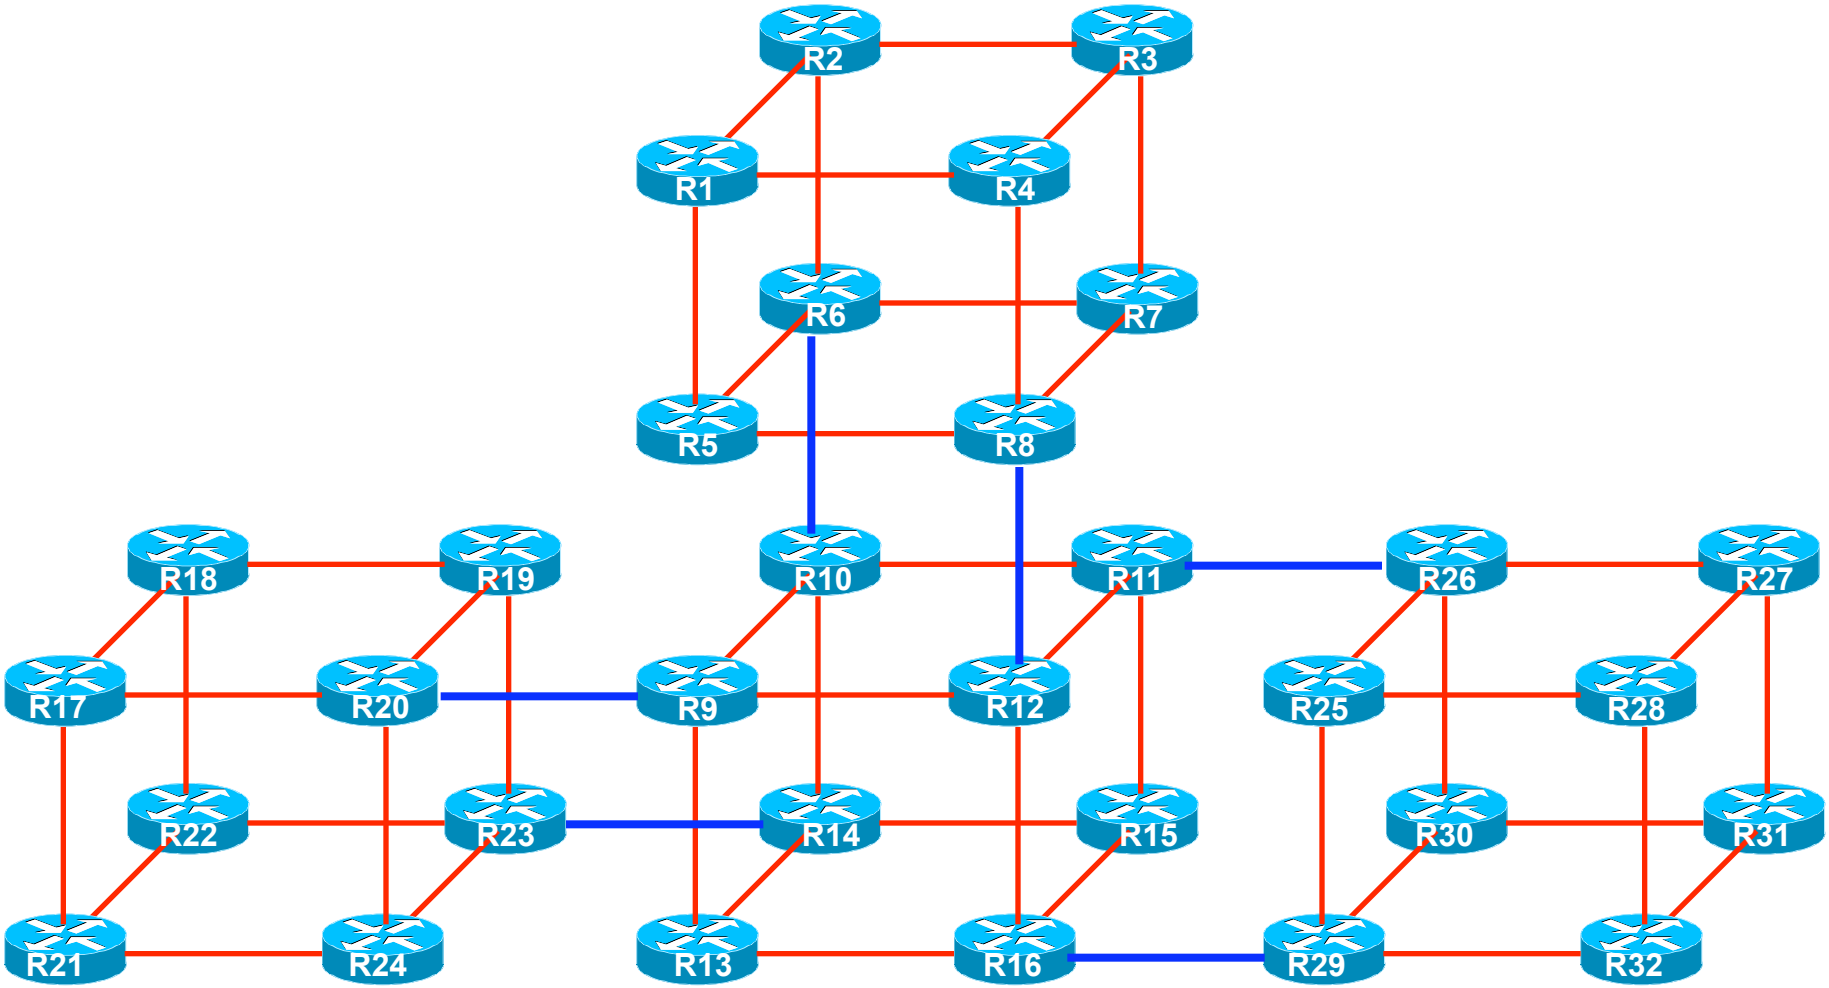

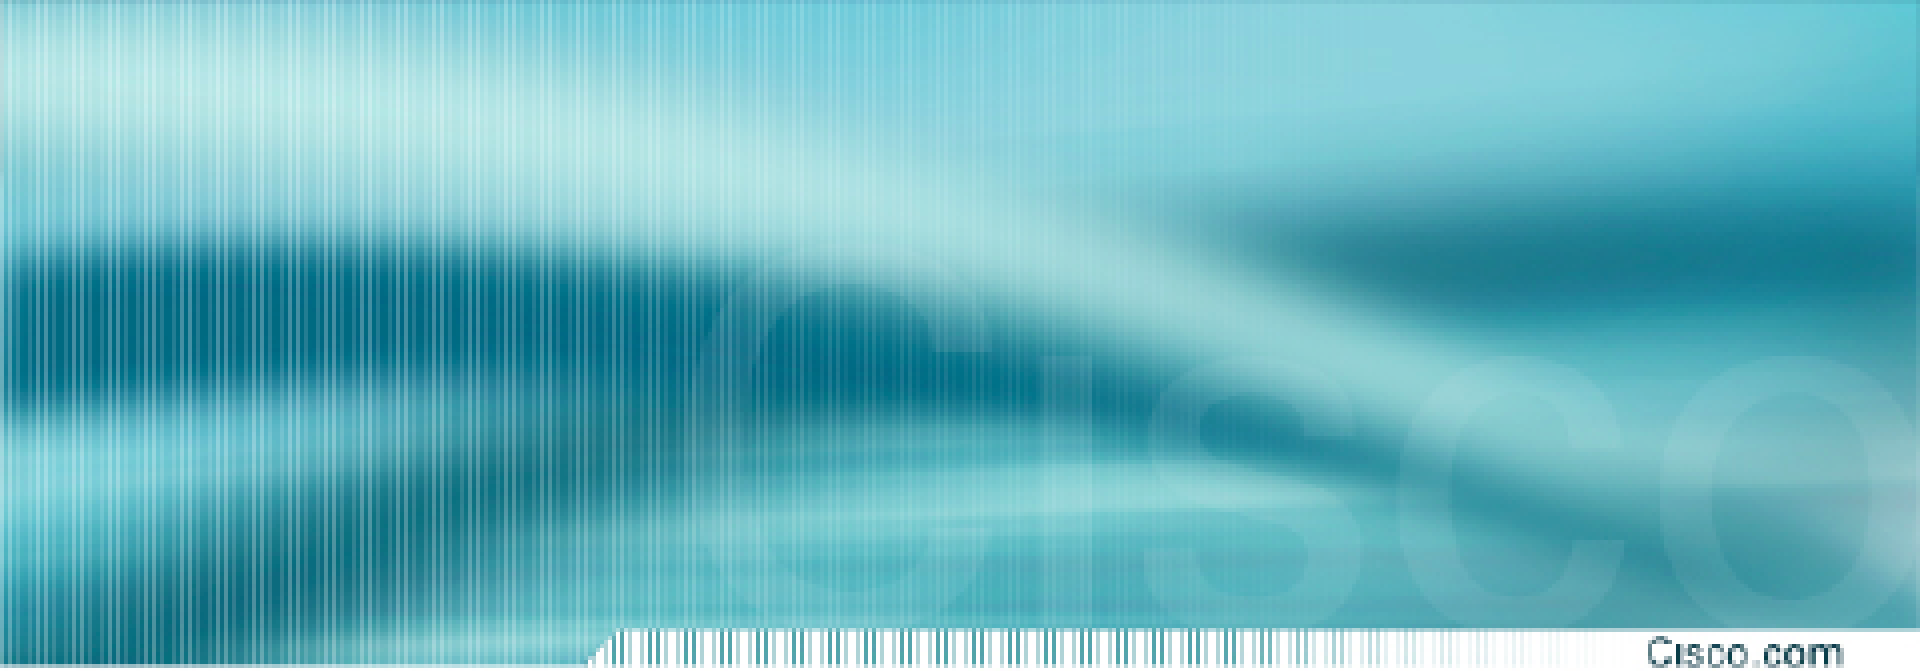

IPv6 Overview

The rationale for a new protocol

A need for IPv6?

- IETF IPv6 WG began in early 90s, to solve addressing growth issues, but

  CIDR, NAT,… were developed

- IPv4 32 bit address = 4 billion hosts

  ~40% of the IPv4 address space is still unused which is different from unallocated

  BUT

- IP is everywhere

  Data, Voice, Audio and Video integration is a Reality

  Regional Registries apply a strict allocation control

- So, the main compelling reason: ***more IP addresses!***

IP Address Allocation History

Cisco.com

1981 - IPv4 protocol published

1985 ~ 1/16 of total space

1990 ~ 1/8 of total space

1995 ~ 1/4 of total space

2000 ~ 1/2 of total space

- This despite increasingly intense conservation efforts

  PPP / DHCP address sharing
  CIDR (classless inter-domain routing)
  NAT (network address translation)
  plus some address reclamation

- Theoretical limit of 32-bit space: ~4 billion devices
  Practical limit of 32-bit space: ~250 million devices

  (see RFC 3194)

The effects of CIDR are diminishing

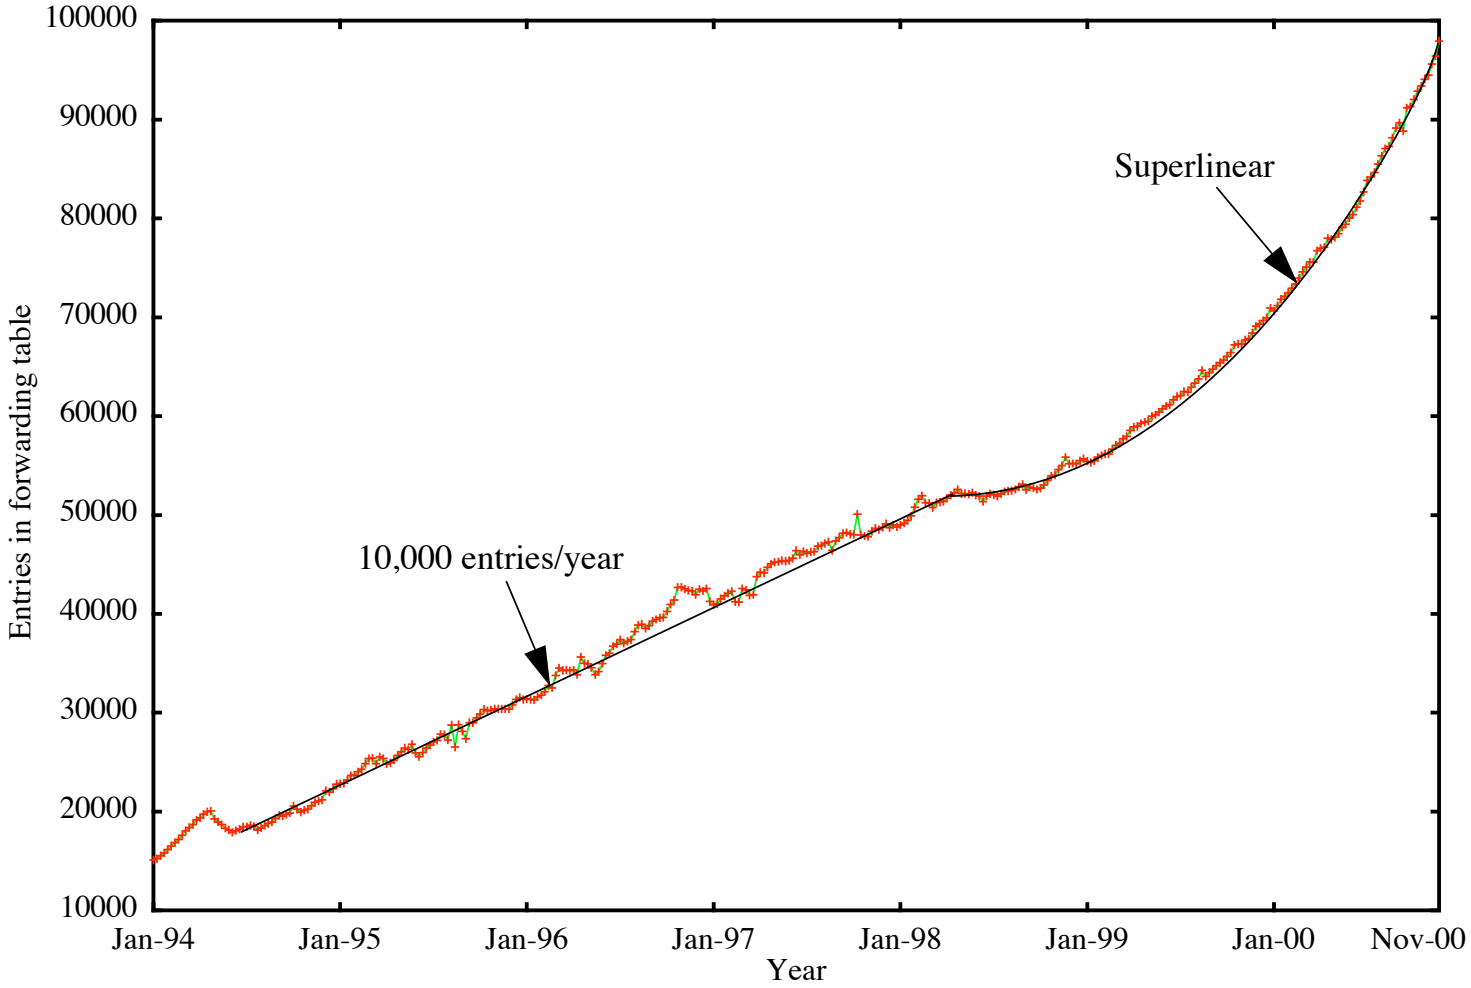

v4 Addressing issue cont.

Active BGP entries

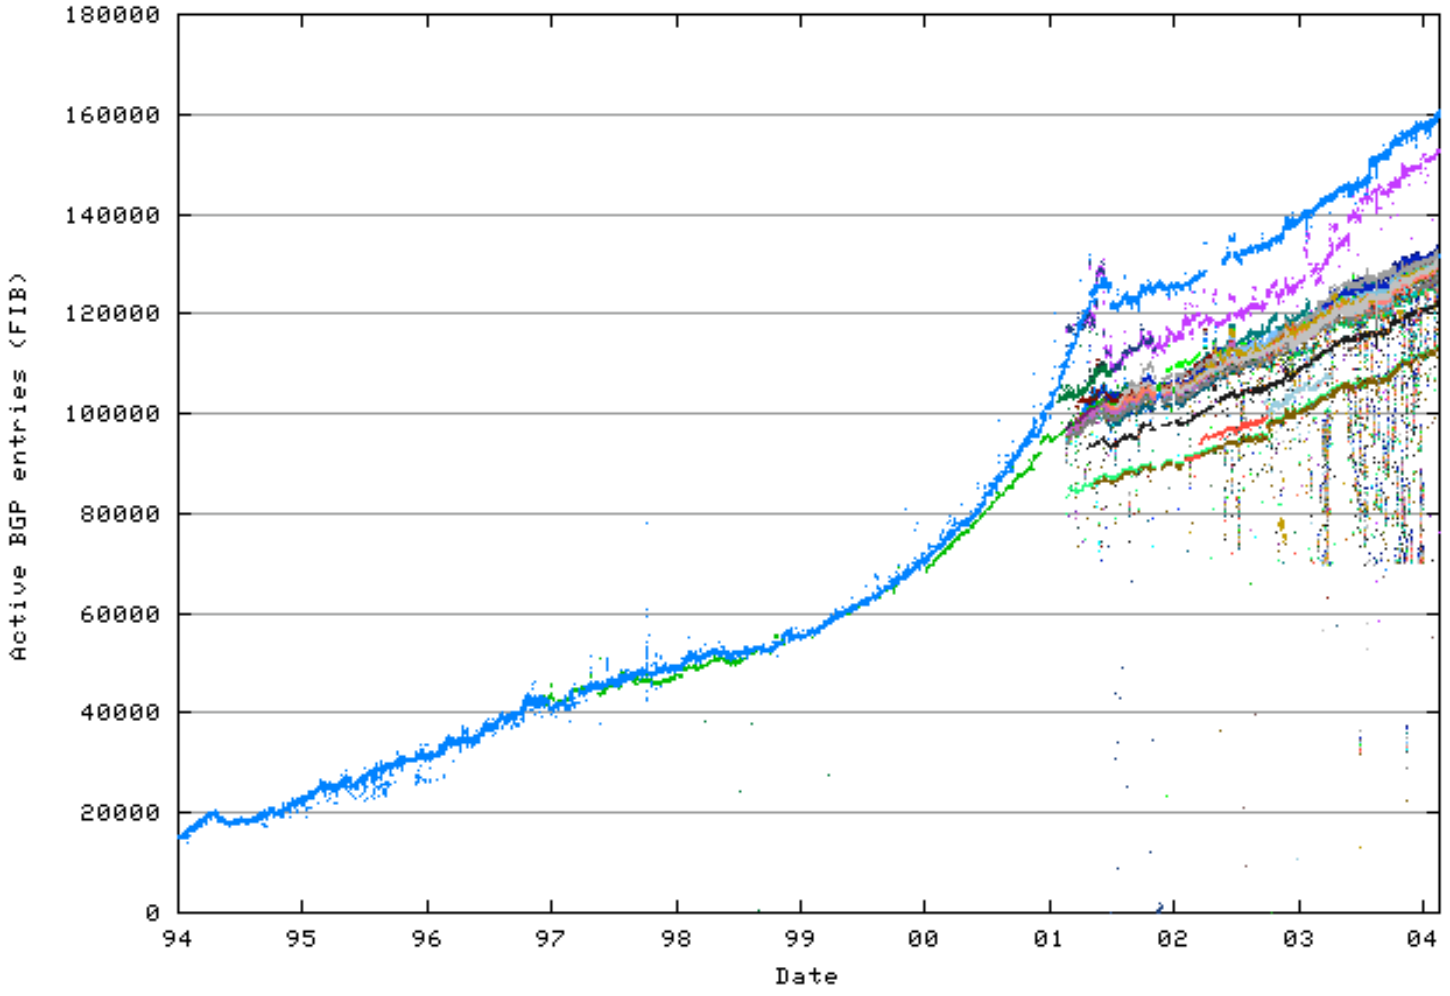

Do We Really Need a Larger Address Space?

Cisco.com

- Internet Users or PC
  - ~530 million users in Q2 CY2002, ~945 million by 2004
  - (Source: Computer Industry Almanac)
  - Emerging population/geopolitical and Address space
- PDA, Pen-Tablet, Notepad,…
  - ~20 million in 2004
- Mobile phones
  - Already 1 billion mobile phones delivered by the industry
- Transportation
  - 1 billion automobiles forecast for 2008
  - Internet access in Planes
- Consumer devices
  - Billions of Home and Industrial Appliances

IPv6 O.S. & Applications support

Cisco.com

- All Operating Systems have an IPv6 stack at some stage of completeness

  All Unix flavours (Sun Solaris, HP Unix, Compaq True64, SGI, IBM AIX, BSD (kame), Linux,…

  Microsoft Windows flavours, MacOS X, Compaq OpenVMS,…

- Focus is now on the Applications

  Ie: Microsoft .NET server, BSD Kame project

- But still need additional vendors

  Ie: Oracle & SAP

- See playground.sun.com/ipv6 and www.hs247.com for latest update

How to get an IPv6 Address?

Cisco.com

How to get address space?

Real IPv6 address space now allocated by APNIC, ARIN and RIPE NCC to ISP

APNIC	**2001:0200::/23 & 2001:0C00::/23**
ARIN	**2001:0400::/23**
RIPE NCC	**2001:0600::/23 & 2001:0800::/23**

6Bone 3FFE::/16

6to4 tunnels 2002::/16

Enterprises will get their IPv6 address space from their ISP.

Further information on www.cisco.com/ipv6

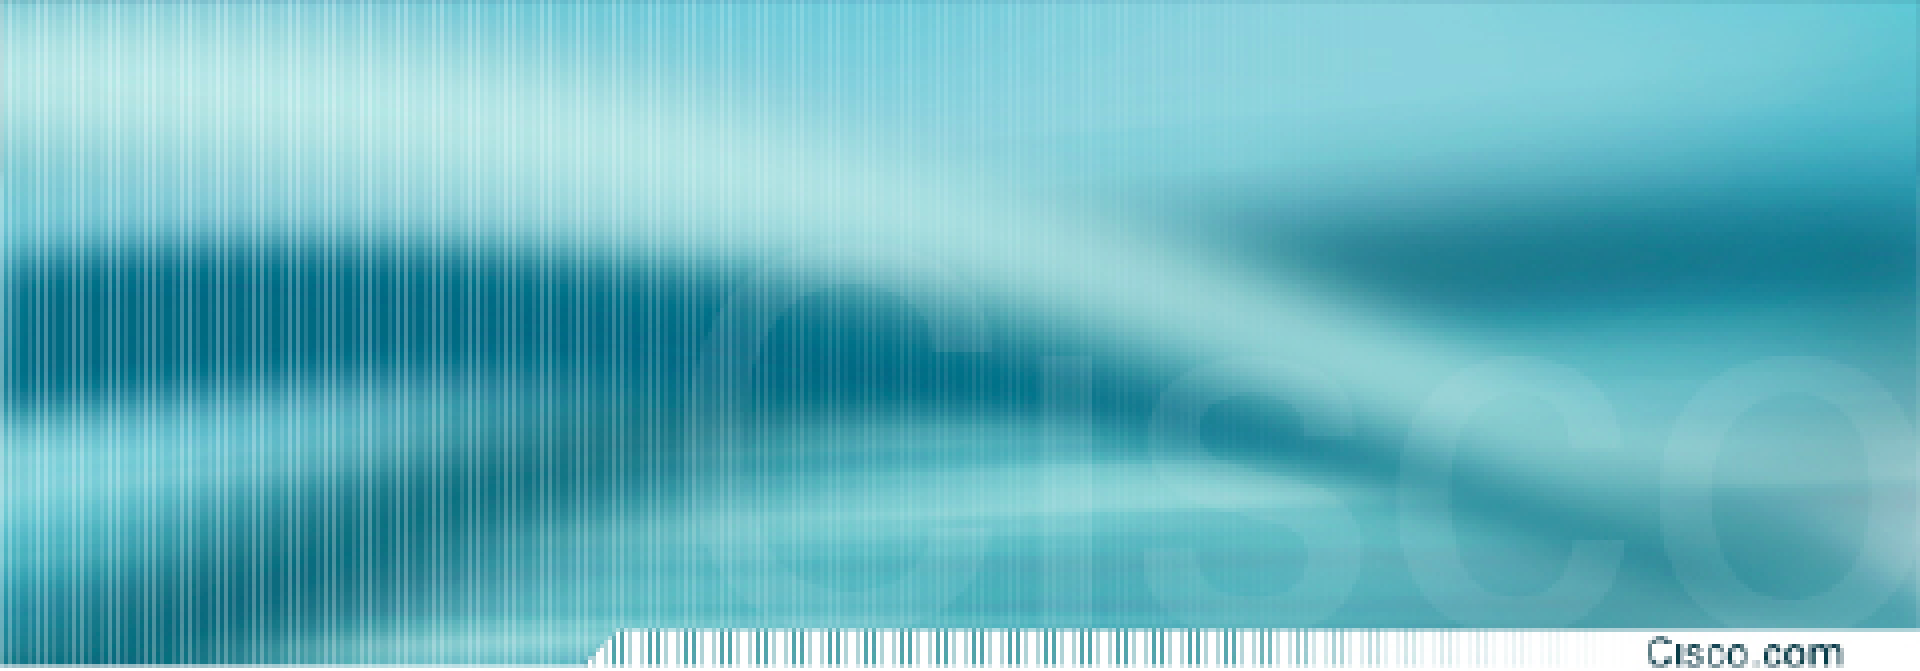

IPv6 Overview

The v6 protocol

Protocol Changes From v4

IPv6 - Major Changes form IPv4

- Larger Address Space

  Address length increased to 128 bits

- Simplified Header Format

  Fixed length, optional headers can be “daisy-chained” together

  v6 header is twice the size of the v4 header (assuming no options set)

- No Layer 3 checksum

- No “hop by hop” segmentation

  Path MTU Discovery

- No Broadcasts

What does it do for:

- **Security**
  - Nothing IP4 doesn’t do - IPSec runs in both
  - but IPv6 mandates IPSec
- **QoS**
  - Nothing IP4 doesn’t do -
    - Differentiated and Integrated Services run in both
    - So far, Flow label has no real use

IPv6 Technology Scope

IP Service	IPv4 Solution	IPv6 Solution
Addressing Range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	Serverless, Reconfiguration, DHCP
Security	IPSec	IPSec Mandated, works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality-of-Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
IP Multicast	IGMP/PIM/Multicast BGP	MLD/PIM/Multicast BGP,Scope Identifier

IPv6 Header

IPv4 & IPv6 Header Comparison

IPv4 Header

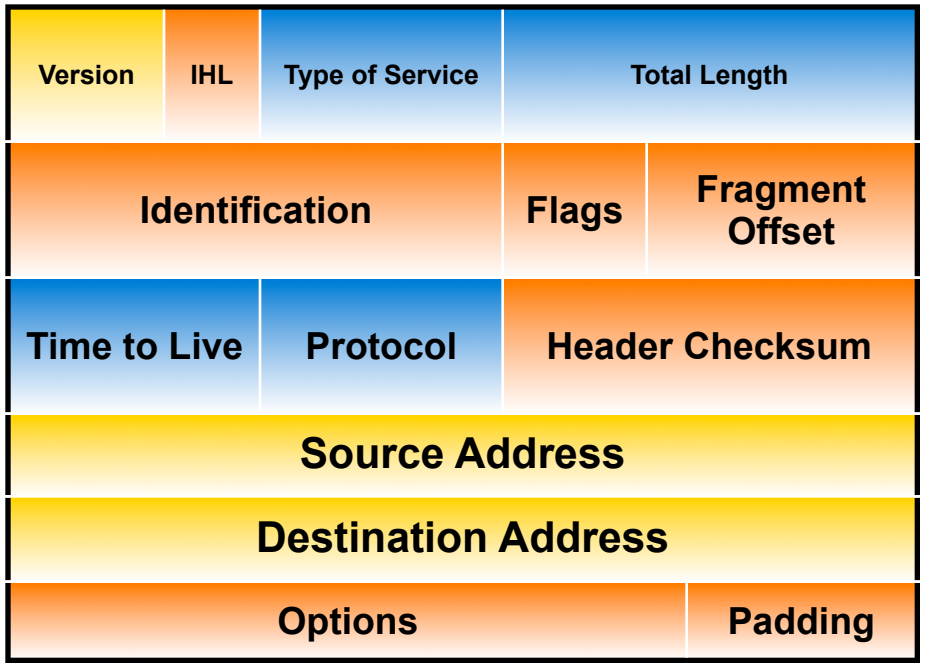

IPv6 Header

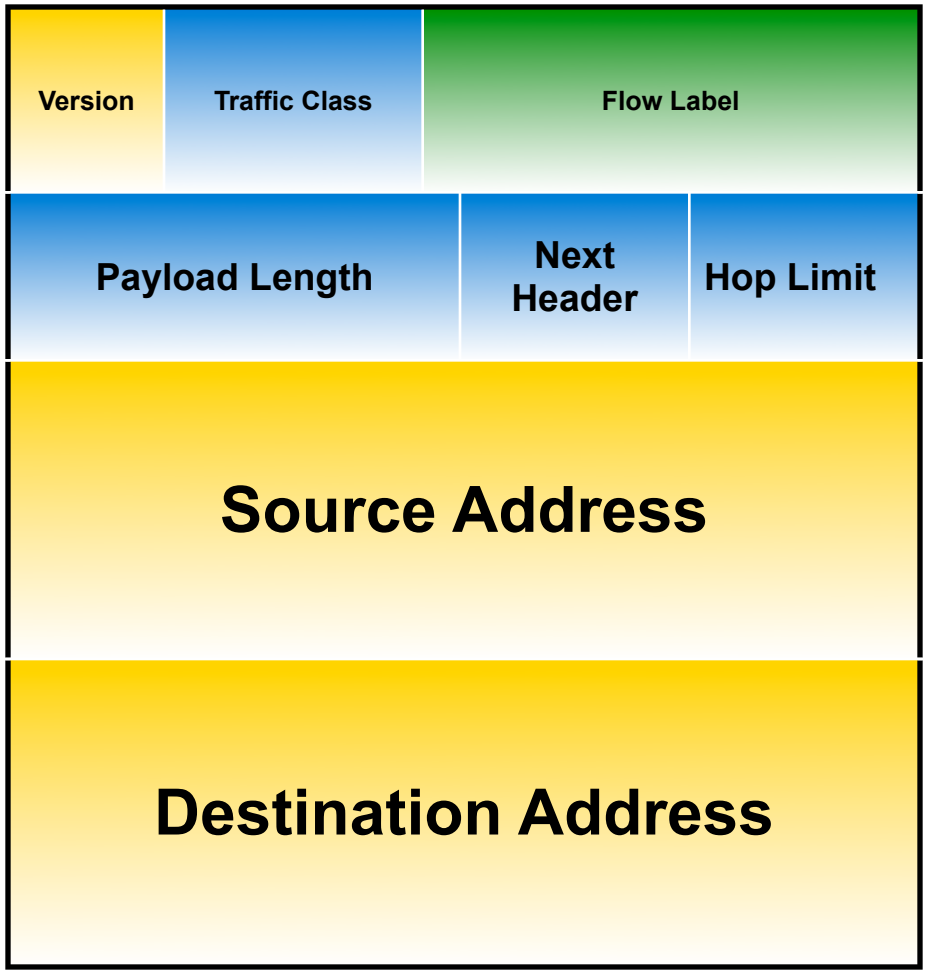

- field's name kept from IPv4 to IPv6
- fields not kept in IPv6
- Name & position changed in IPv6
- New field in IPv6

IPv6 Header Options (RFC 2460)

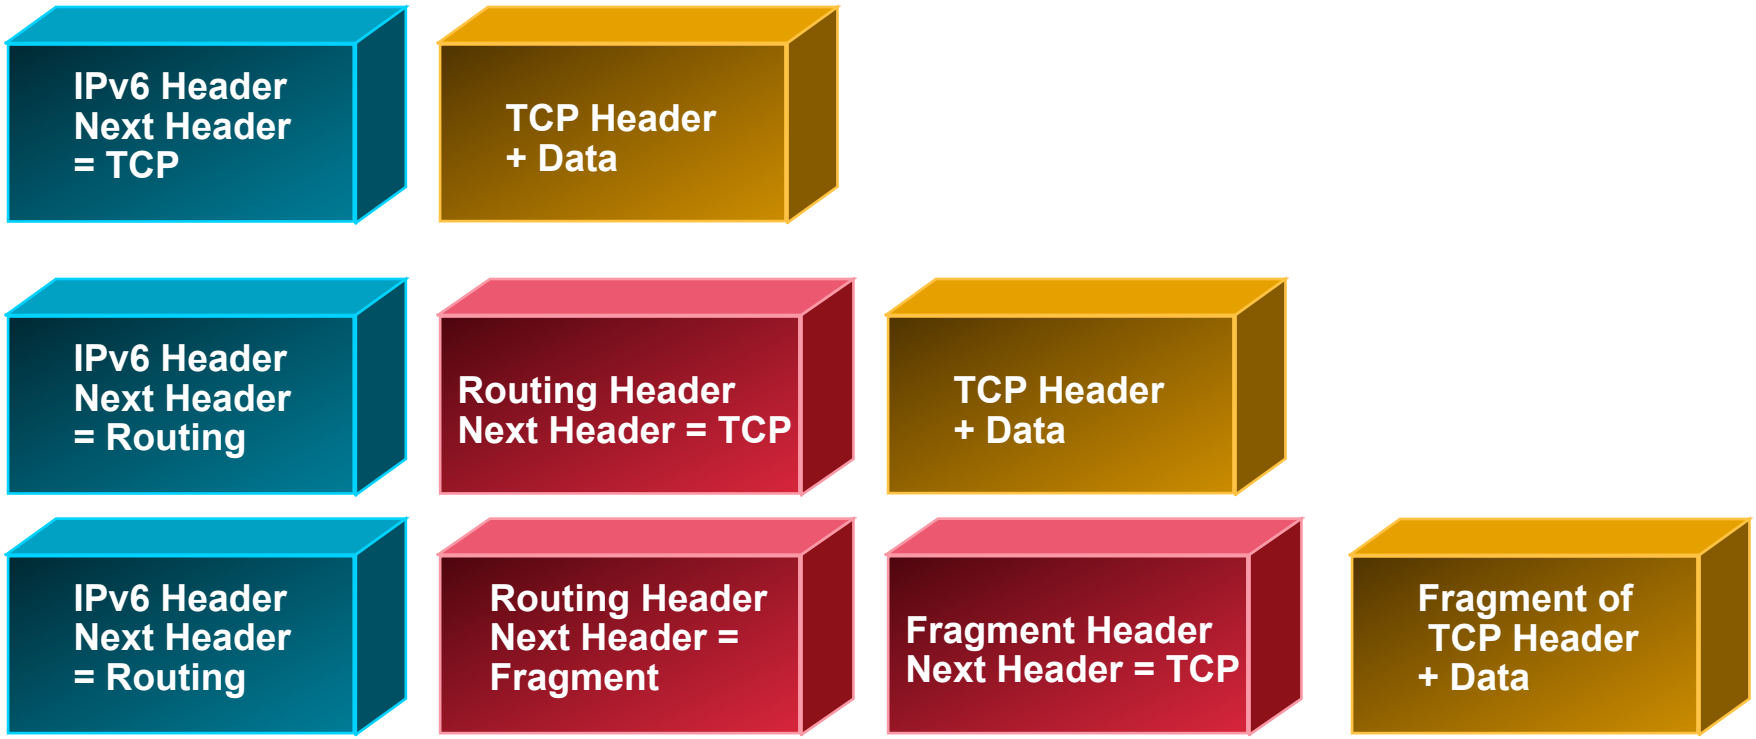

- Processed only by node identified in IPv6 Destination Address field => much lower overhead than IPv4 options
  exception: Hop-by-Hop Options header
- Eliminated IPv4's 40-octet limit on options
  in IPv6, limit is total packet size, or Path MTU in some cases

IPv6 Header Options (RFC2460)

- **Currently defined Headers should appear in the following order**
  - **IPv6 header**
  - **Hop-by-Hop Options header**
  - **Destination Options header**
  - **Routing header**
  - **Fragment header**
  - **Authentication header (RFC 1826)**
  - **Encapsulating Security Payload header (RFC 1827)**
  - **Destination Options header**
  - **upper-layer header (TCP/UDP etc..)**

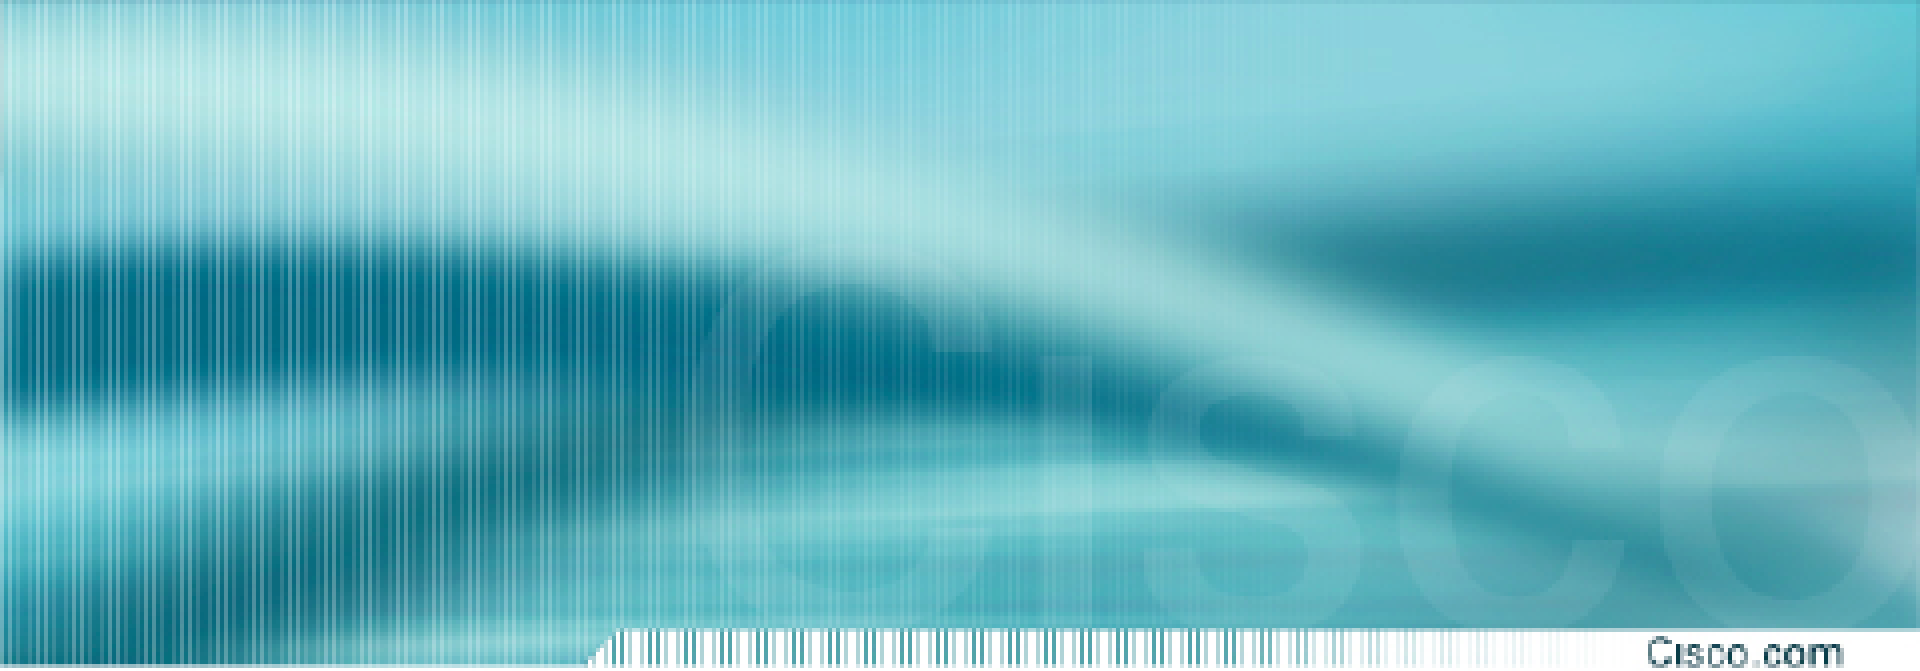

IPv6 Addressing

Objectives

Cisco.com

- **Upon completion of this module, you will be able to perform the following tasks:**
  - **Describe and use IPv6 addressing**
  - **Describe various IPv6 address types**
  - **Describe Path MTU discovery process**
  - **Describe Neighbor Discovery process**
  - **Configure IPv6 on Cisco routers**

IPv6 Address Representation

Cisco.com

- 16-bit fields in case insensitive, colon delimited, hexadecimal representation

  2031:0000:130F:0000:0000:09C0:876A:130B
- Leading zeros in a field are optional:

  2031:0:130F:0:0:9C0:876A:130B
- Successive fields of 0 represented as ::, but only once in an address:
  - 2031:0:130F::9C0:876A:130B
  - 2031::130F::9C0:876A:130B 
  - 0:0:0:0:0:0:0:1 => ::1
  - 0:0:0:0:0:0:0:0 => ::
- IPv4-compatible address representation
  - 0:0:0:0:0:0:192.168.30.1 = ::192.168.30.1 = ::C0A8:1E01

IPv6 Unicast Address Types

Cisco.com

- Unicast addresses are used in a one-to-one context.
- IPv6 Unicast addresses are:

  Link-Local Addresses

  Site-Local Addresses

  Aggregatable Global Unicast Addresses

  Unspecified, loopback and IPv4 compatible

Link-Local Address Format

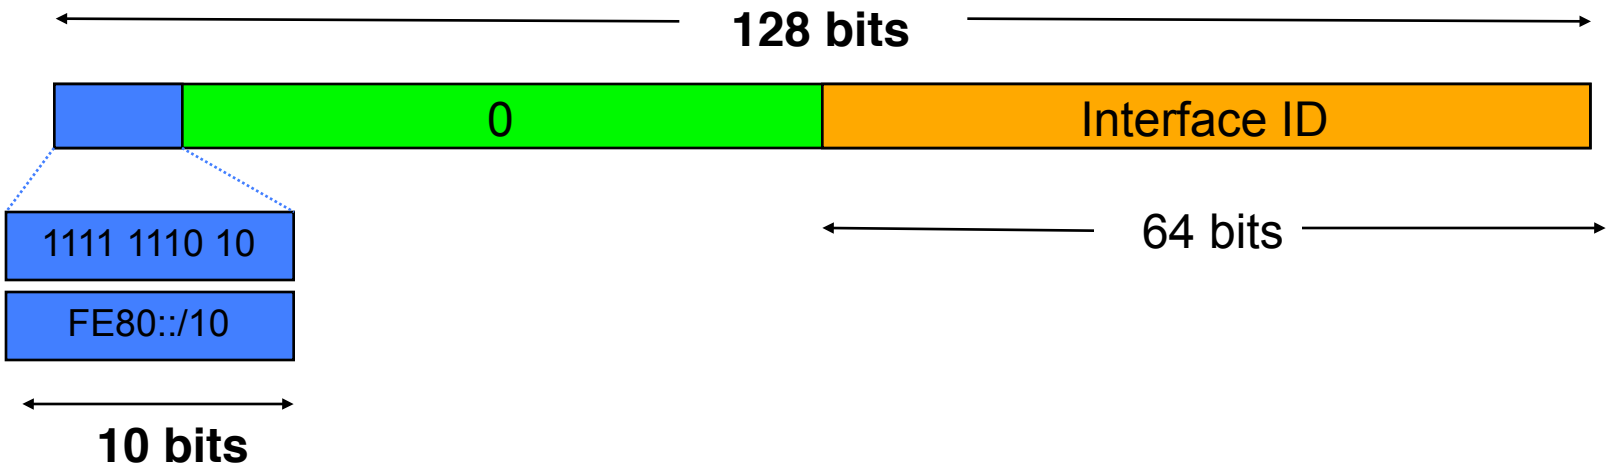

- **Link-local addresses:**

  Have a limited scope of the link

  Are automatically configured with the interface ID

Site-Local Address Format

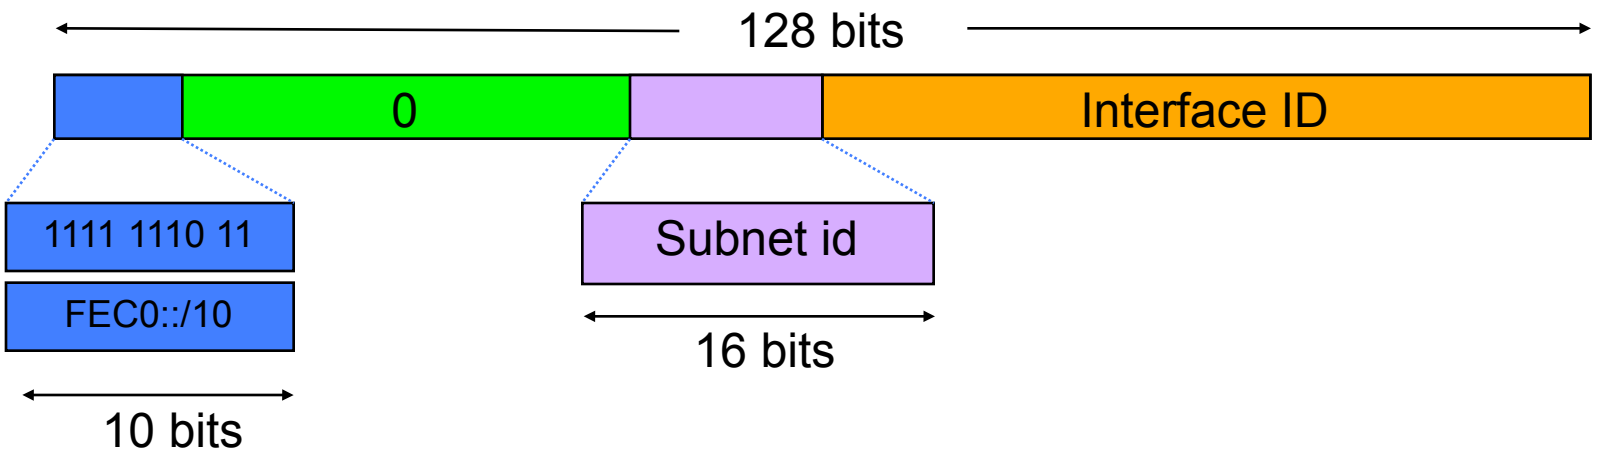

- **Site-local addresses:**

  Have a limited scope of the site

  Contain the inside topology of the site with the subnet ID

Aggregatable Global Unicast Addresses

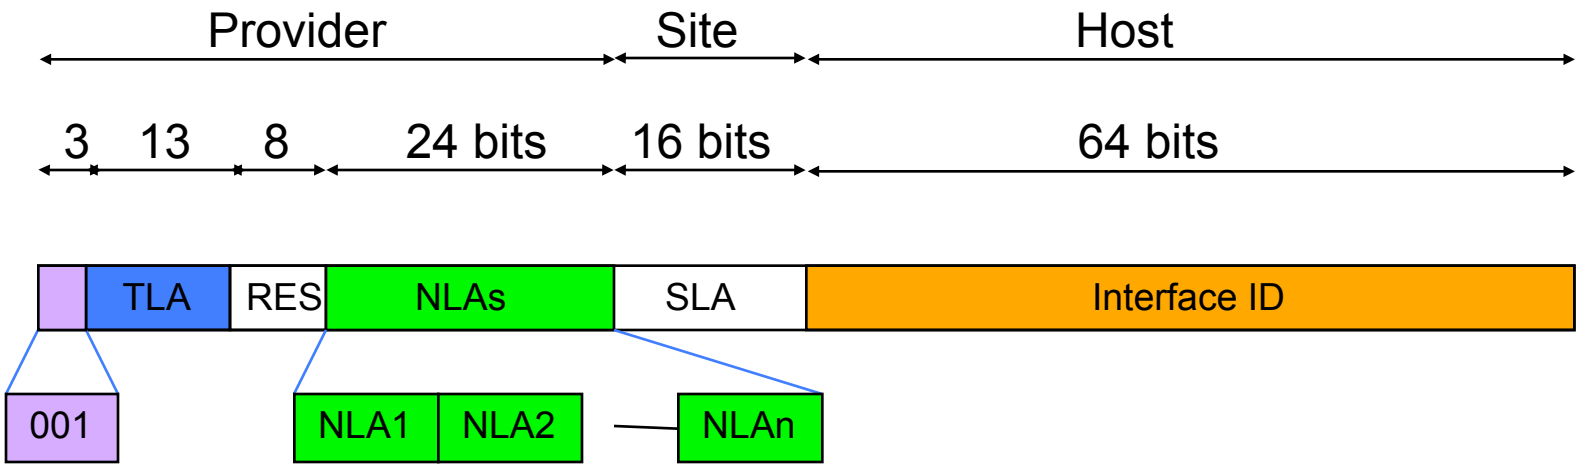

- **Aggregatable Global Unicast addresses are:**
  - **Addresses for generic use of IPv6**
  - **Structured as a hierarchy to keep the aggregation**

Address Allocation

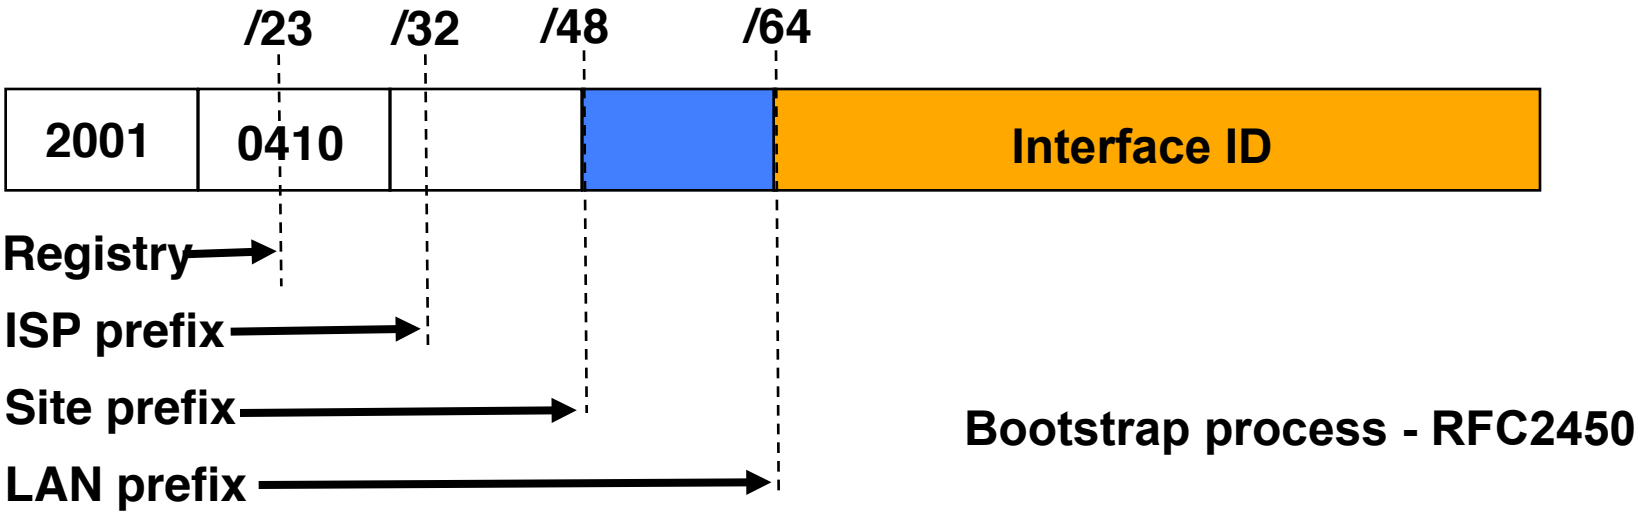

- The allocation process is under review by the Registries:
  - IANA allocates 2001::/16 to registries
  - Each registry gets a /23 prefix from IANA
  - Formely, all ISP were getting a /35
  - With the new proposal, Registry allocates a /36 (immediate allocation) or /32 (initial allocation) prefix to an IPv6 ISP
  - Policy is that an ISP allocates a /48 prefix to each end customer
  - ftp://ftp.cs.duke.edu/pub/narten/ietf/global-ipv6-assign-2002-04-25.txt

Hierarchical Addressing & Aggregation

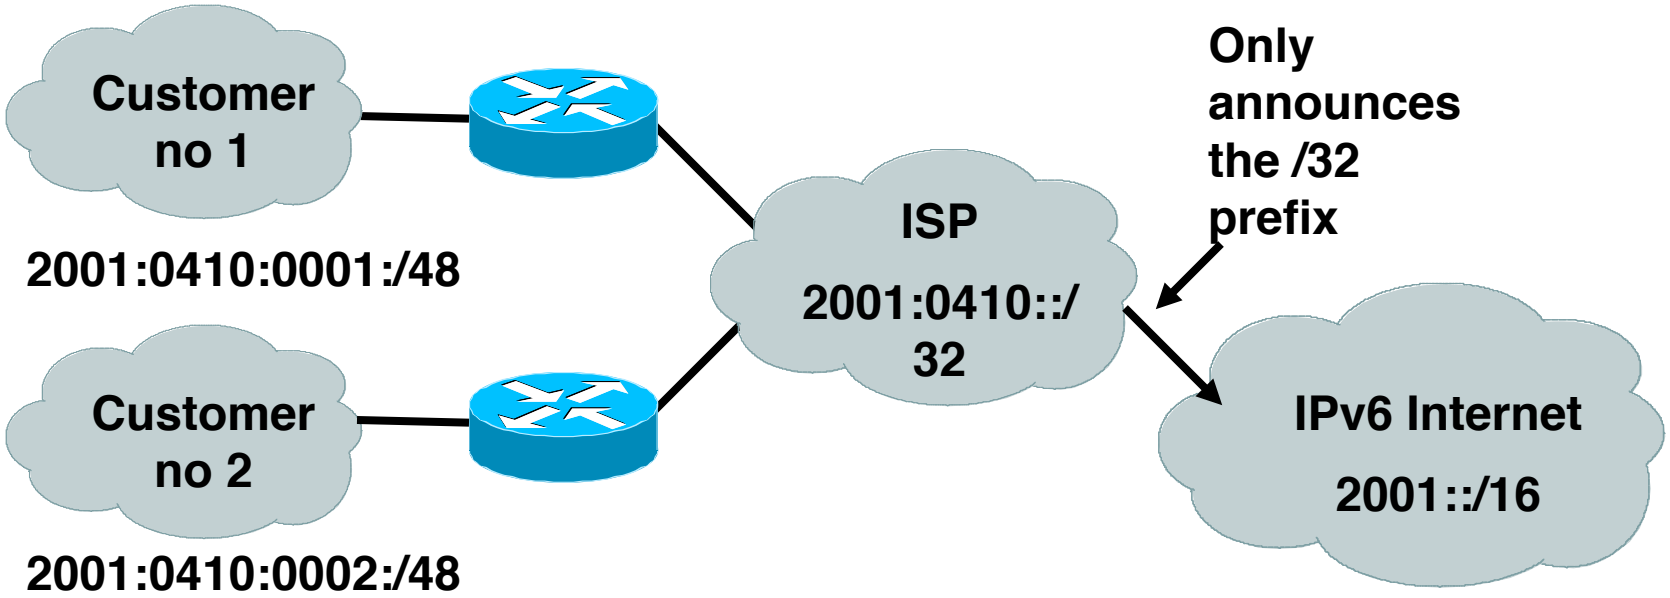

Larger address space enables:

- Aggregation of prefixes announced in the global routing table.
- Efficient and scalable routing.

But current Multi-Homing schemes break the model

IPv4 Compatible Addresses

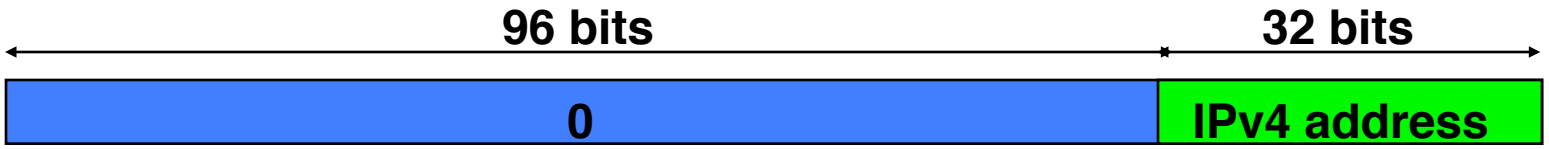

0:0:0:0:0:0:192.168.30.1

= ::192.168.30.1

= ::C0A8:1E01

- **IPv4 compatible address:**

  Is a way to insert the IPv4 address in an IPv6 address

  Enables easy automatic tunnelling

Address Representation

Cisco.com

- **Format:**
  - **IPv4-compatible:**
    - **0:0:0:0:0:0:192.168.30.1**
    - **= ::192.168.30.1**
    - **= ::C0A8:1E01**
  - **In a URL, it is enclosed in brackets**
    - **http://[2001:1:4F3A::206:AE14]:8080/index.html**
    - **Cumbersome for users**
    - **Mostly for diagnostic purposes**
    - **Use fully qualified domain names (FQDN)**

Anycast Address Format

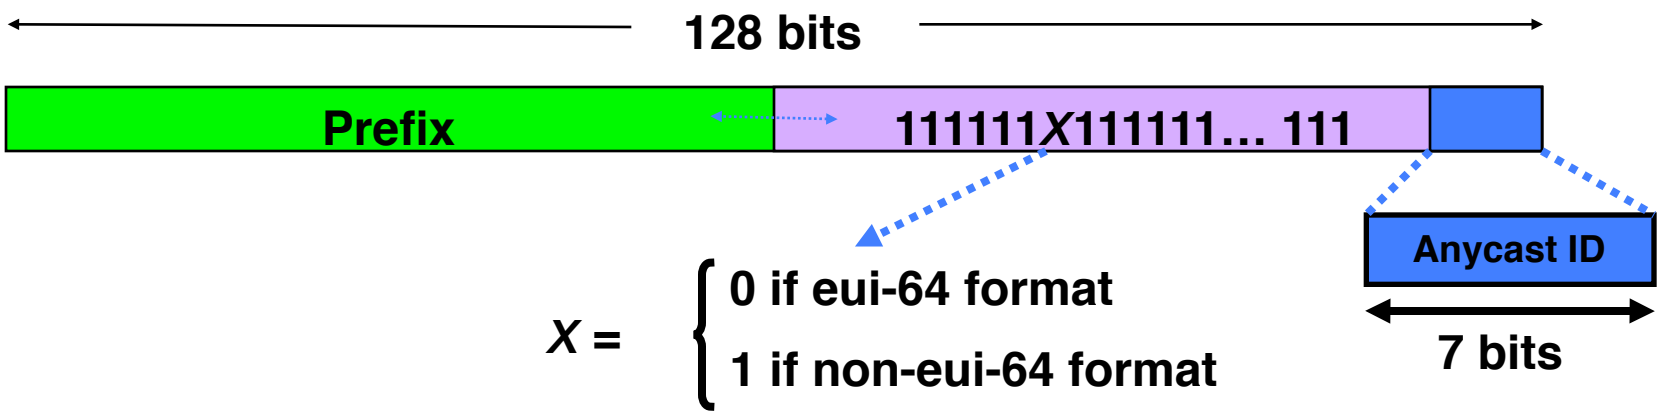

- **Anycast:**
  - **Is one-to-nearest type of address**
  - **Has a current limited use**

Unspecified and Loopback Addresses

- **Unspecified address:**

  0:0:0:0:0:0:0:0

  Used as a placeholder when no address available (initial DHCP request, DAD)

- **Loopback address:**

  0:0:0:0:0:0:0:1

  Same as 127.0.0.1 in IPv4

  Identifies self

IPv6 Addressing

Cisco.com

- **Prefix Format (PF) Allocation**

  PF = 0000 0000 : Reserved

  PF = 001 : Aggregatable Global Unicast Address

  PF = 1111 1110 10 : Link Local Use Addresses (FE80::/10)

  PF = 1111 1110 11 : Site Local Use Addresses (FEC)::/10)

  PF = 1111 1111 : Multicast Addresses (FF00::/8)

  Other values are currently Unassigned (approx. 7/8th of total)

- **All Prefix Formats have to support EUI-64 bits Interface ID setting**

  But Multicast

Device Address Requirements

Required Node Addresses

Cisco.com

- Any IPv6 node should have the following addresses enabled:

  Link-local address for each interface

  Loopback address

  Assigned unicast addresses

  All-nodes multicast address

  Solicited-node multicast address for each of its assigned unicast and anycast address

  Multicast address of all other groups to which the host belongs

  Site-local address if used

Required Router Addresses

Cisco.com

- **Any IPv6 router should have the following addresses enabled:**
  - **All the required node addresses**
  - **Subnet-router anycast addresses for the interfaces configured to act as forwarding interfaces**
  - **Other anycast configured addresses**
  - **All-routers multicast addresses**
  - **Specific multicast addresses for routing protocols**

Multicast Addressing

Multicast

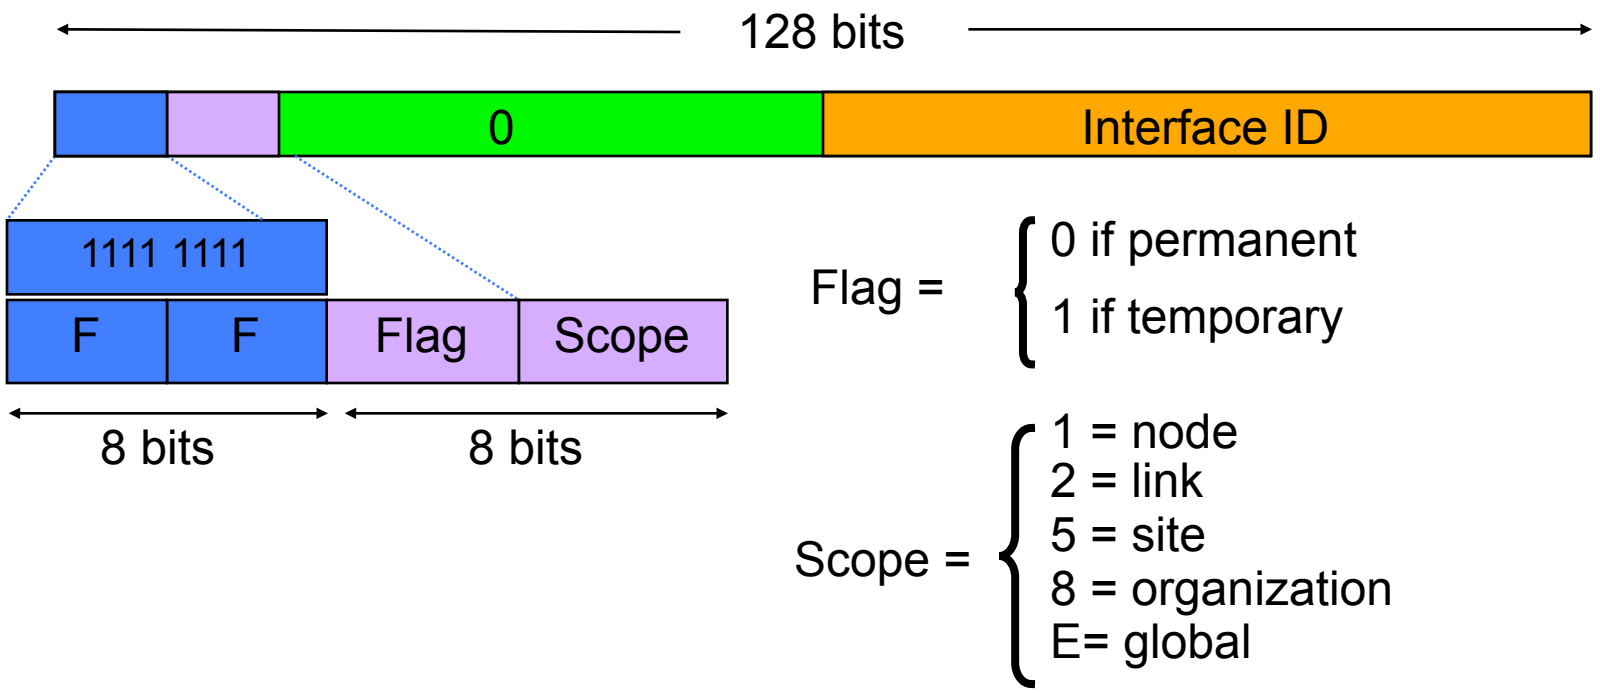

- **Multicast is used in the context of one-to-many. A multicast scope is new in IPv6.**

Multicast Assigned Addresses

FF0X:: is reserved (X=0..F)

Inside this range, the following are assigned:

Address	Meaning	Scope
FF01::1	All Nodes	Node-local
FF02::1	All Nodes	Link-local
FF01::2	All Routers	Node-local
FF02::2	All Routers	Link-local
FF05::2	All Routers	Site-local
FF02::1:FFXX:XXXX	Solicited-Node	Link-local

Multicast mapping over Ethernet

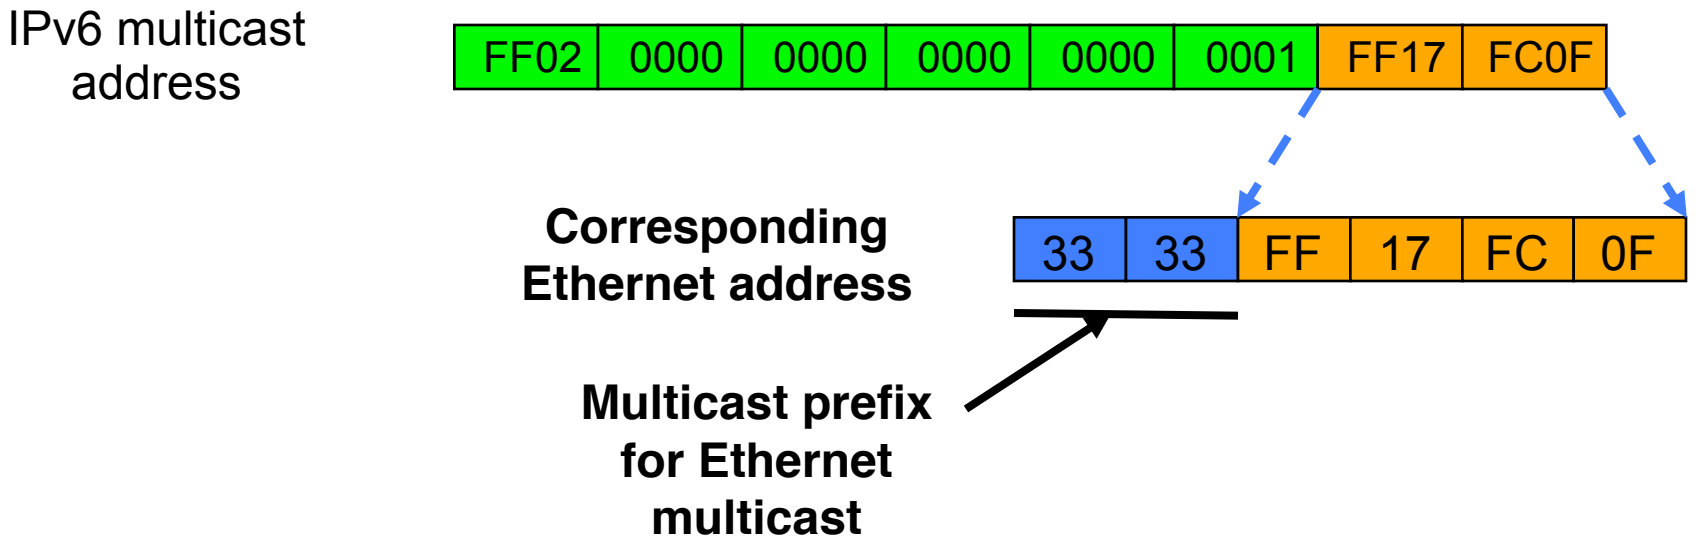

- **Mapping of IPv6 multicast address to Ethernet address is:**

  33:33:<last 32 bits of the IPv6 multicast address>

IPv6 Addressing Examples

IPv6 Addressing Examples

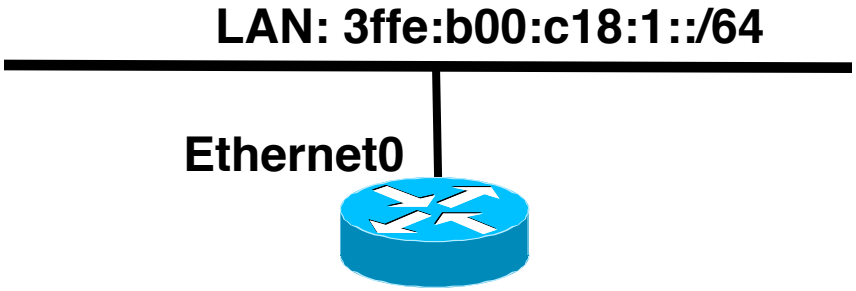


```
interface Ethernet0
 ipv6 address 2001:410:213:1::/64 eui-64
```

MAC address: 0060.3e47.1530

```
router# show ipv6 interface Ethernet0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::260:3EFF:FE47:1530
Global unicast address(es):
    2001:410:213:1:260:3EFF:FE47:1530, subnet is 2001:410:213:1::/64
  Joined group address(es):
    FF02::1:FF47:1530
    FF02::1
    FF02::2
  MTU is 1500 bytes
```

Eui-64

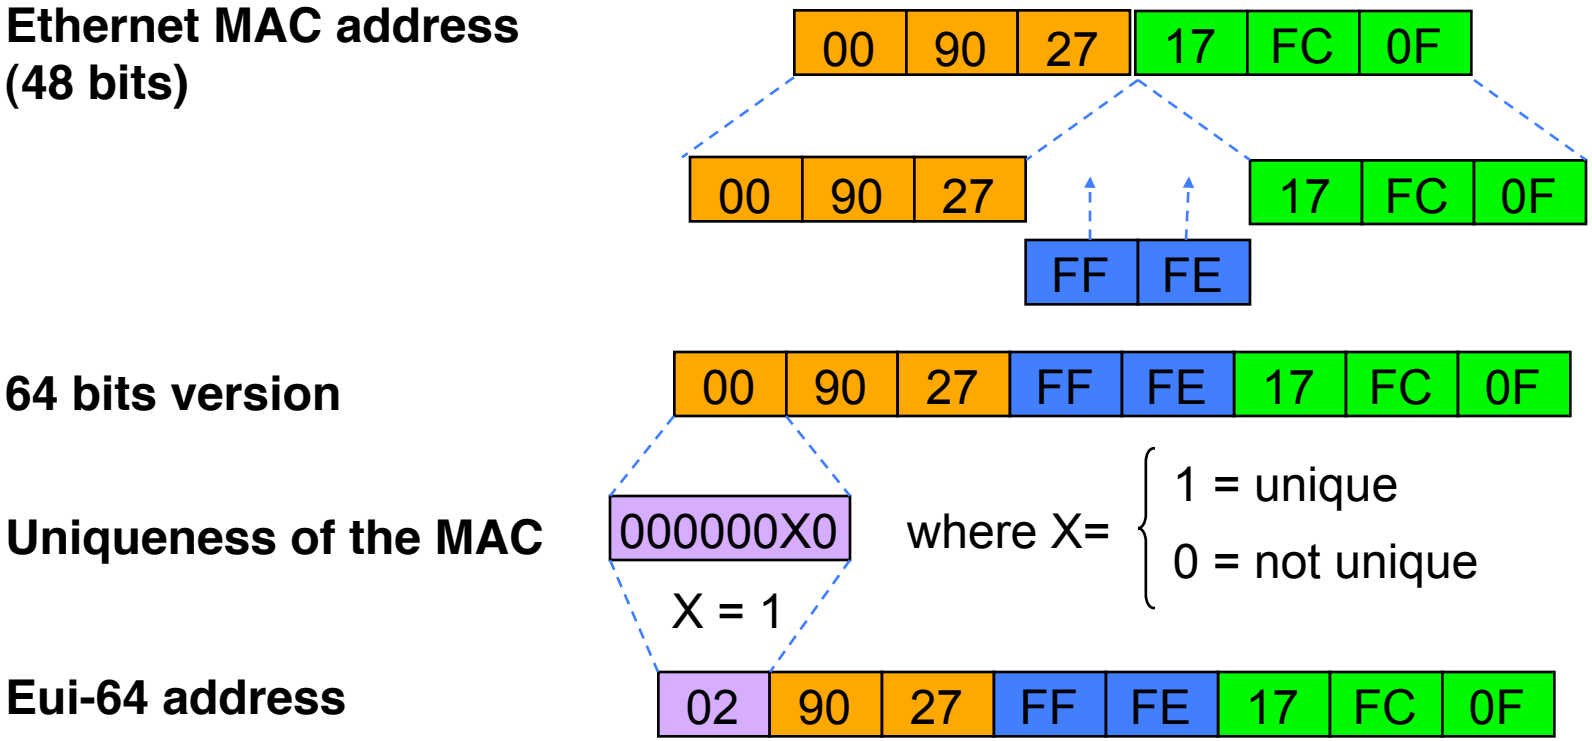

- **Eui-64 address is formed by inserting "FFFE" and ORing a bit identifying the uniqueness of the MAC address.**

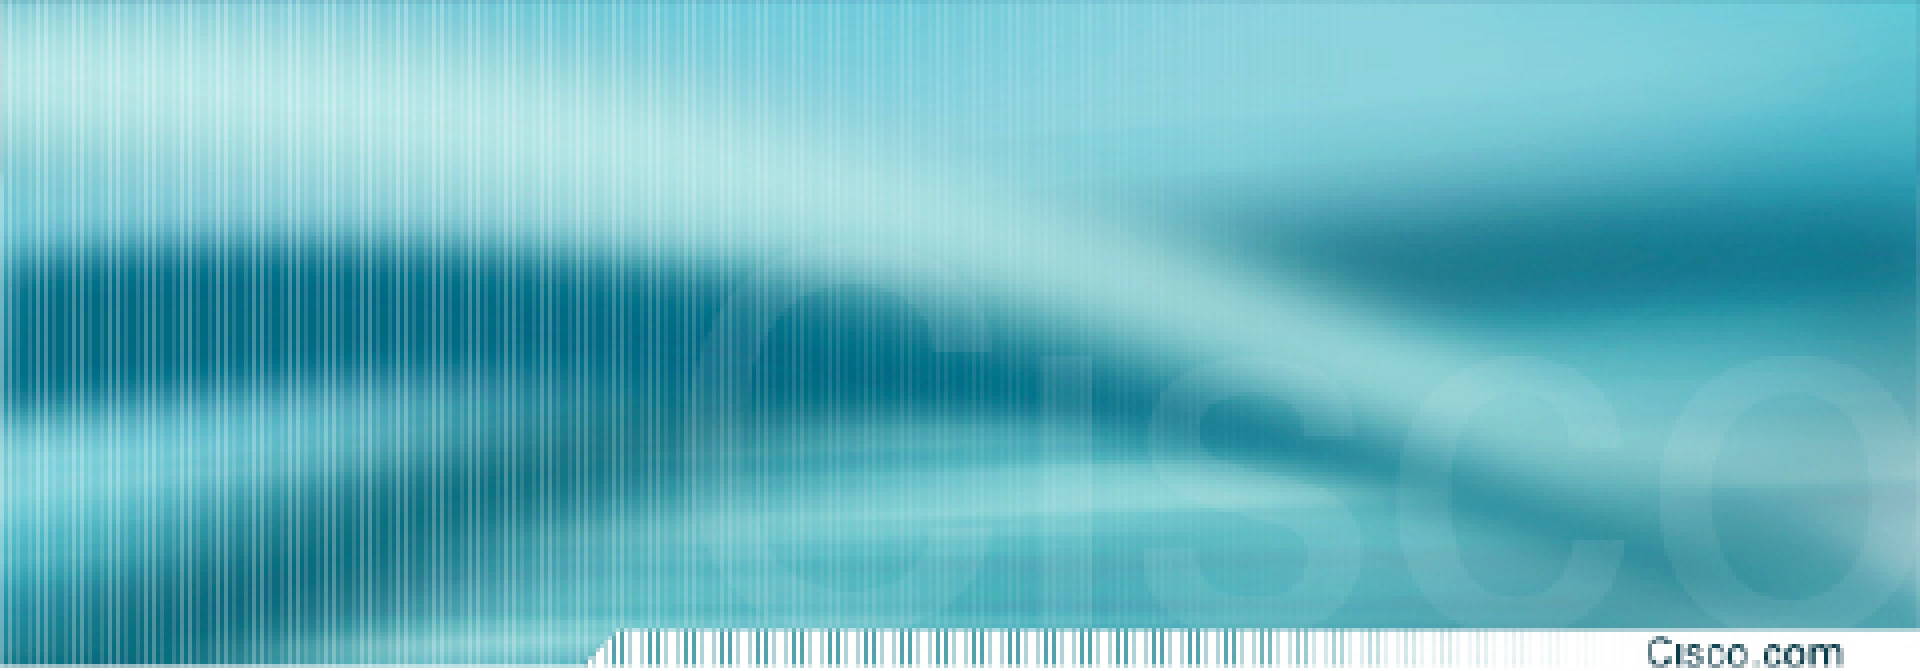

Data Link

IPv6 over Data Link layers

Cisco.com

- **IPv6 is defined for most data link layers:**

Ethernet

FDDI

Token Ring

Arcnet

PPP

Non-Broadcast Multiple Access (NBMA)

ATM

Frame Relay

IPv6 over Ethernet

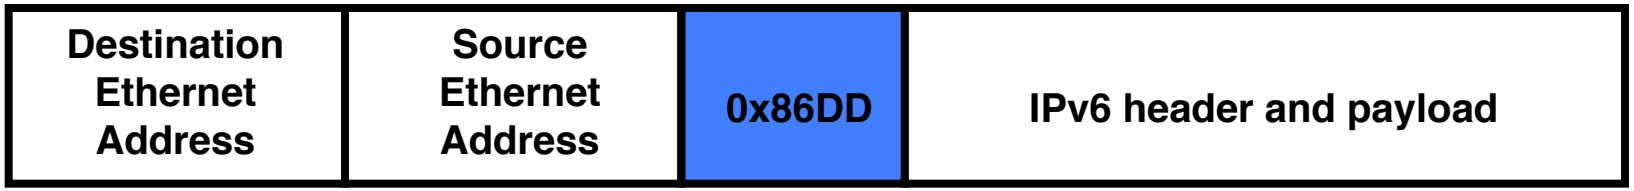

IPv6 has a specific Ethernet protocol ID

Different from the IPv4 one

MTU Considerations

MTU Issues

- minimum link MTU for IPv6 is 1280 octets (versus 68 octets for IPv4)

  => on links with MTU < 1280, link-specific fragmentation and reassembly must be used

- implementations are expected to perform path MTU discovery to send packets bigger than 1280
- minimal implementation can omit PMTU discovery as long as all packets kept ≤ 1280 octets
- a Hop-by-Hop Option supports transmission of "jumbograms" with up to 2^{32} octets of payload

Neighbour Discovery

Neighbour Discovery (RFC 2461)

- **Protocol built on top of ICMPv6 (RFC 2463)**
- **combination of IPv4 protocols (ARP, ICMP, IGMP,…)**
- **Fully dynamic, interactive between Hosts & Routers**
- **defines 5 ICMPv6 packet types**

  Router Solicitation / Router Advertisements

  Neighbor Solicitation / Neighbor Advertisements

  Redirect

Neighbour Discovery (RFC 2461)

- **defined mechanisms between nodes attached on the same link**
  - **Router discovery**
  - **Prefix discovery**
  - **Parameters discovery, ie: link MTU, hop limit,…**
  - **Address autoconfiguration**
  - **Address Resolution (same function as ARP)**
  - **Next-hop determination**
  - **Neighbor Unreachability Detection (useful for default routers)**
  - **Duplicate Address Detection**
  - **Redirect**

IPv6 Auto-Configuration

- **Stateless (RFC2462)**

  Host autonomously configures its own Link-Local address

  Router solicitation are sent by booting nodes to request RAs for configuring the interfaces.

- **Stateful**

  DHCPv6 (under definition at IETF)

- **Renumbering**

  Hosts renumbering is done by modifying the RA to announce the old prefix with a short lifetime and the new prefix.

  Router renumbering protocol (RFC 2894), to allow domain-interior routers to learn of prefix introduction / withdrawal

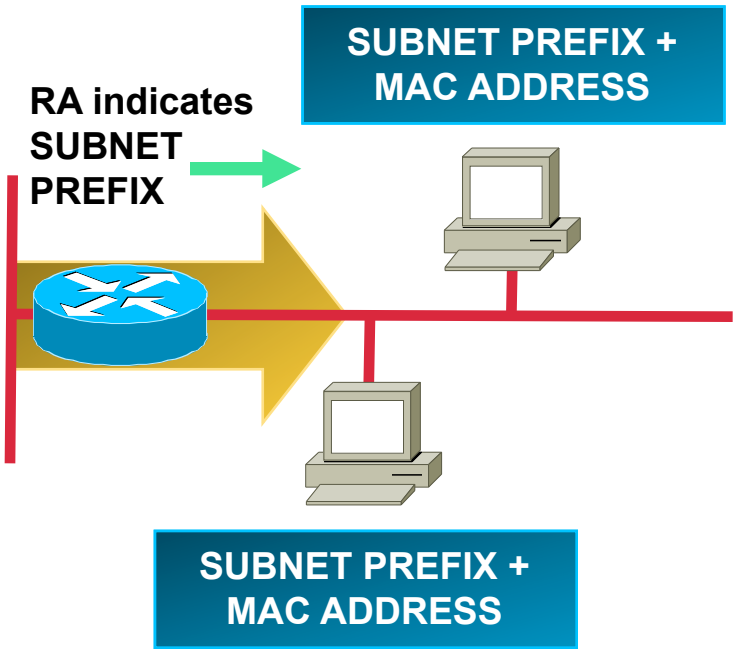

At boot time, an IPv6 host build a Link-Local address, then its global IPv6 address(es) from RA

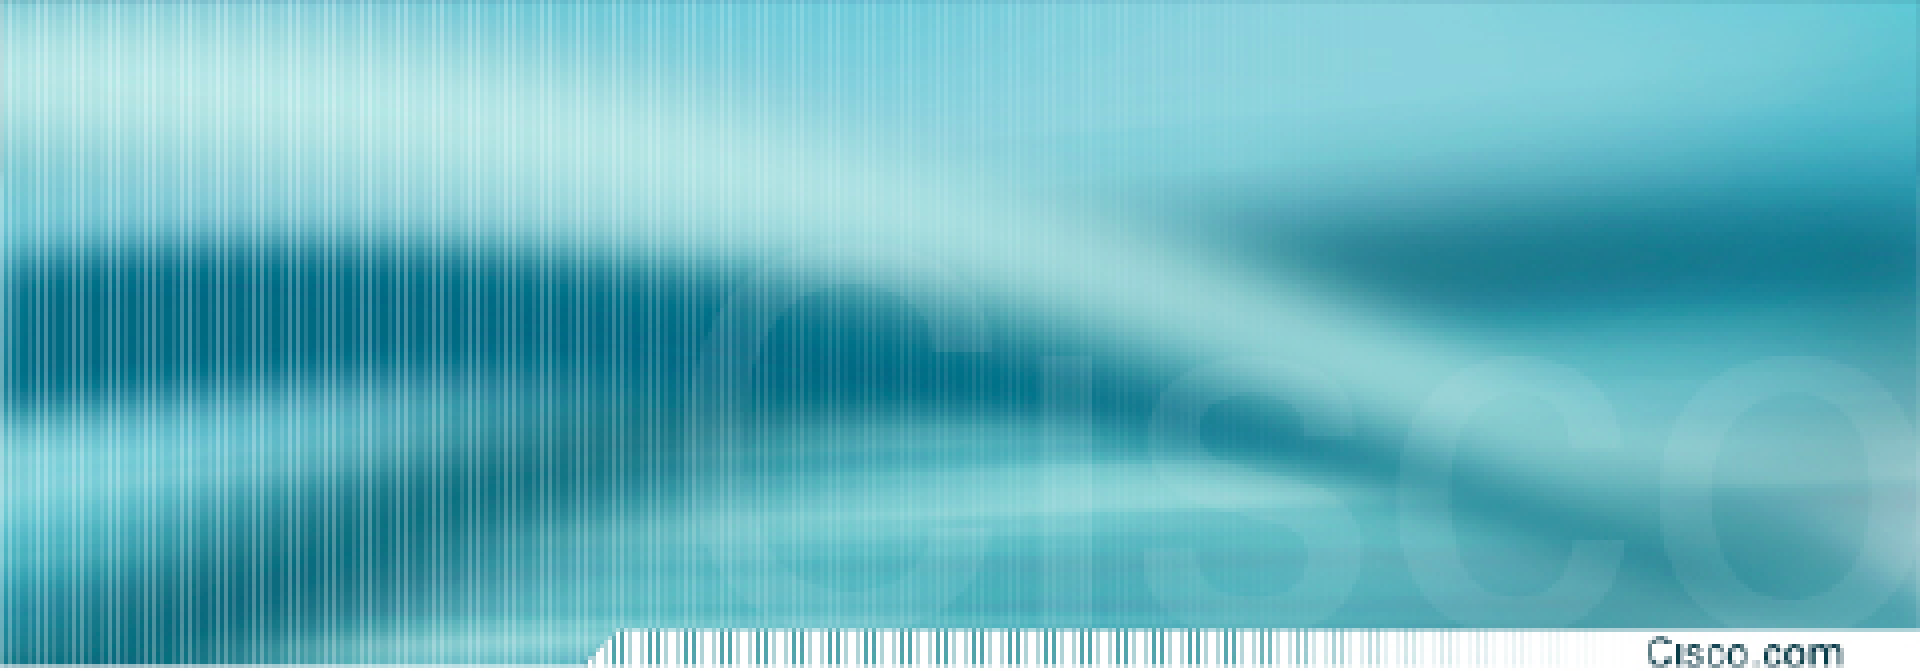

Routing and IPv6

Routing in IPv6

- As in IPv4, IPv6 has 2 families of routing protocols: IGP and EGP

  IGP are RIPng (RFC 2080), Cisco EIGRP for IPv6, OSPFv3 and Integrated IS-ISv6 EGP is MP-BGP4 (RFC 2858 and RFC 2545)
- IPv6 still uses the longest-prefix match routing algorithm.
- i/IS-ISv6 (draft-ietf-isis-ipv6-02)

  Integrated IGP for IPv4 & IPv6
- OSPFv3 (RFC 2740)

  « Ships in the Night » routing, has to run OSPFv2 for IPv4
- IPv6 tunnels & Routing considerations, eg. 6to4 tunnels

Configuring IPv6 Routing

Cisco.com

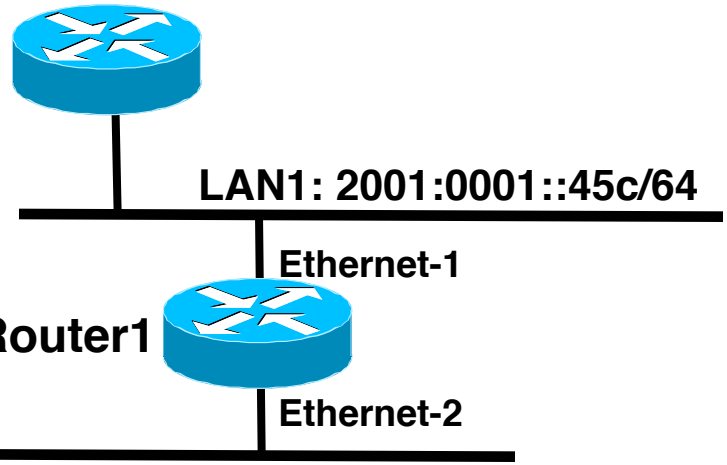

OSPFv3

Integrated IS-ISv6

```
Router1#
 interface loopback 0
     ip address 192.222.222.1 255.255.255.0
 interface ethernet-1
     ipv6 address 2001:0001::45c/64
     ipv6 ospf 1 area 1 enable

  interface ethernet-2
     ipv6 address 2001:0002::45a/64
     ipv6 ospf 1 area 1 enable

  ipv6 router ospf 1
          redistribute static
```

```
Router1#
 interface ethernet-1
     ipv6 address 2001:0001::45c/64
     ipv6 router isis

  interface ethernet-2
     ipv6 address 2001:0002::45a/64
     ipv6 router isis

  router isis
      address-family ipv6
      redistribute static
      exit-address-family
      net 42.0001.0000.0000.072c.00
```

Initial Network Setup

Why this Network Design

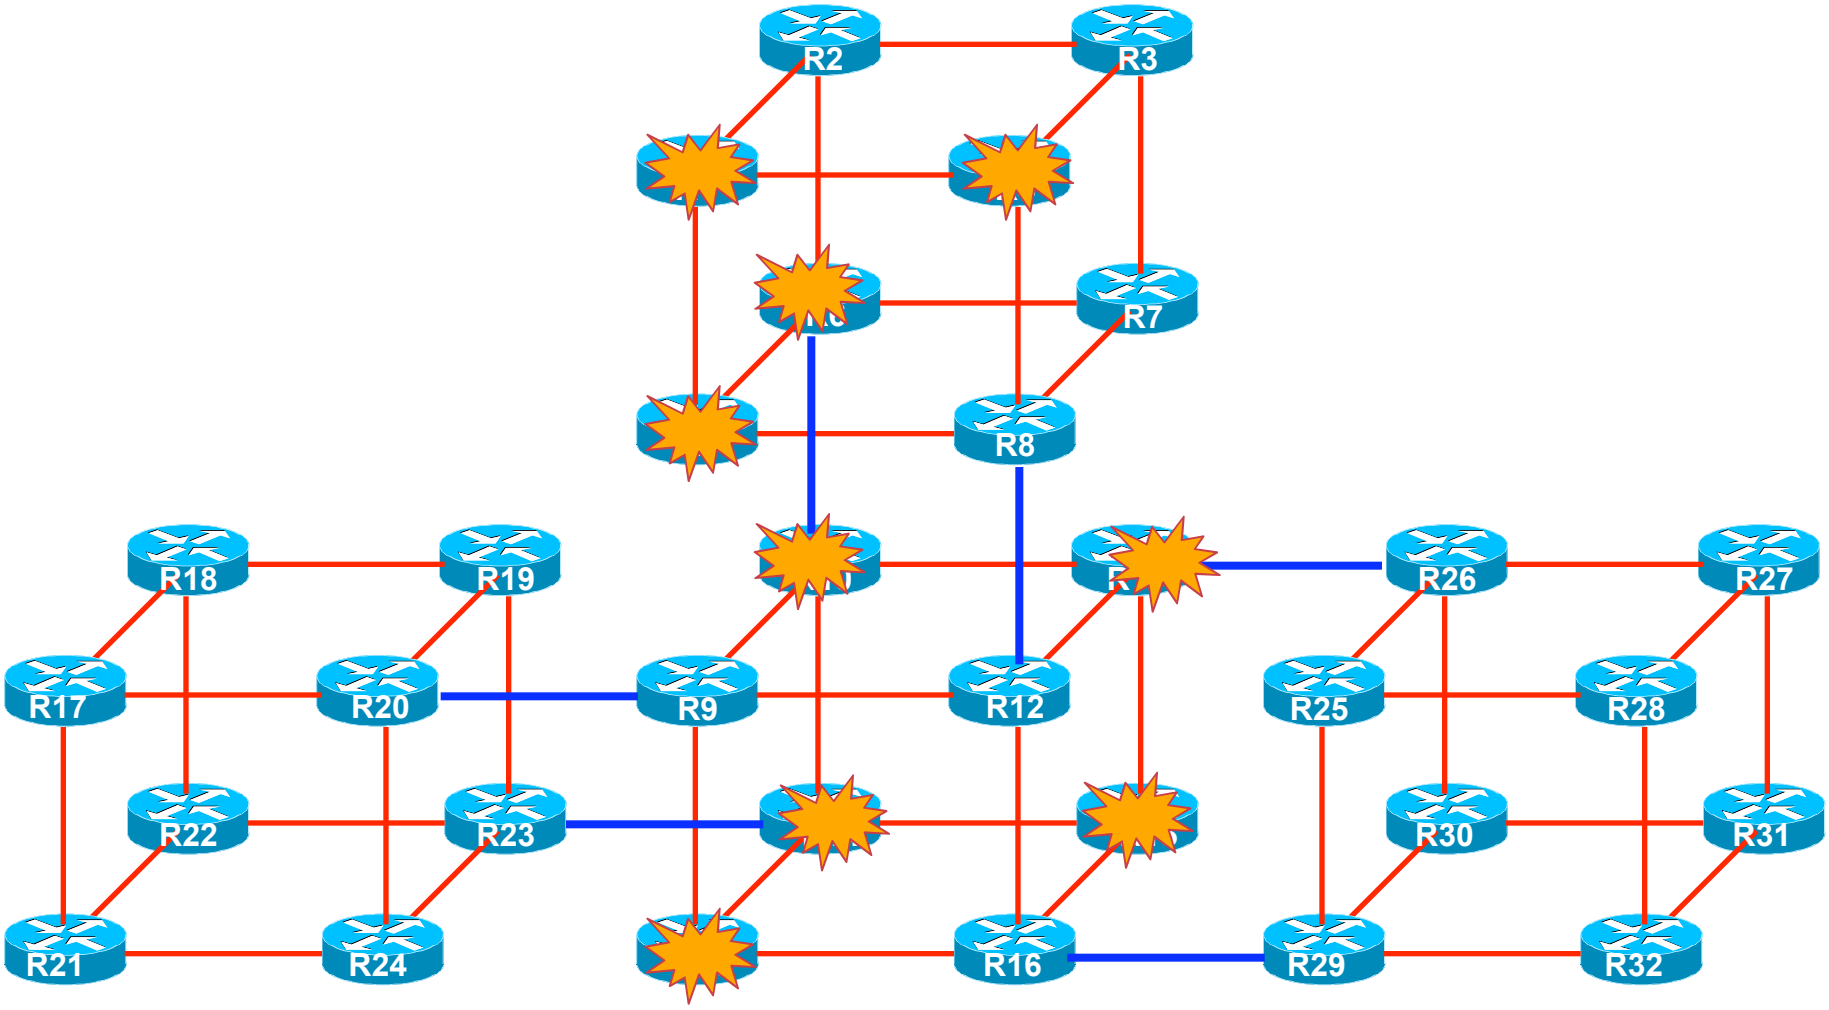

Network (Interface Map) Topology

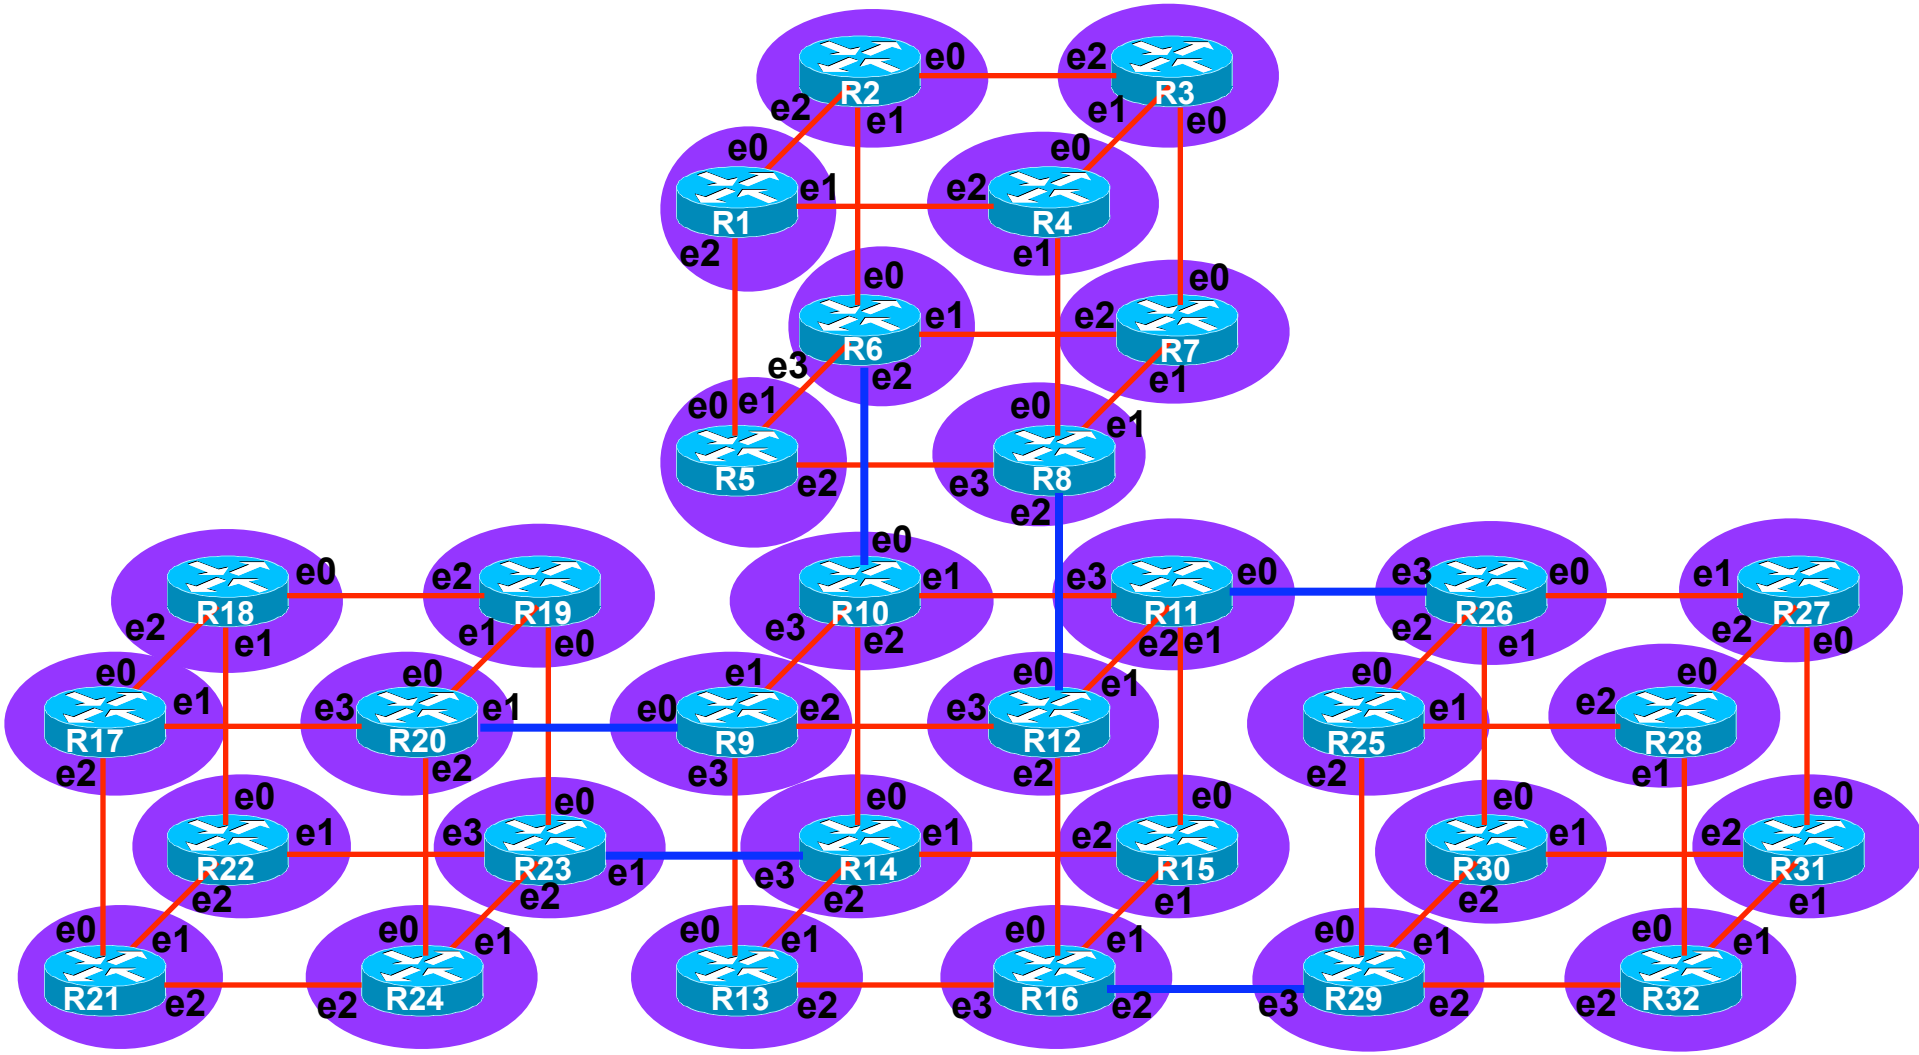

IPv4 OSPF Area Layout

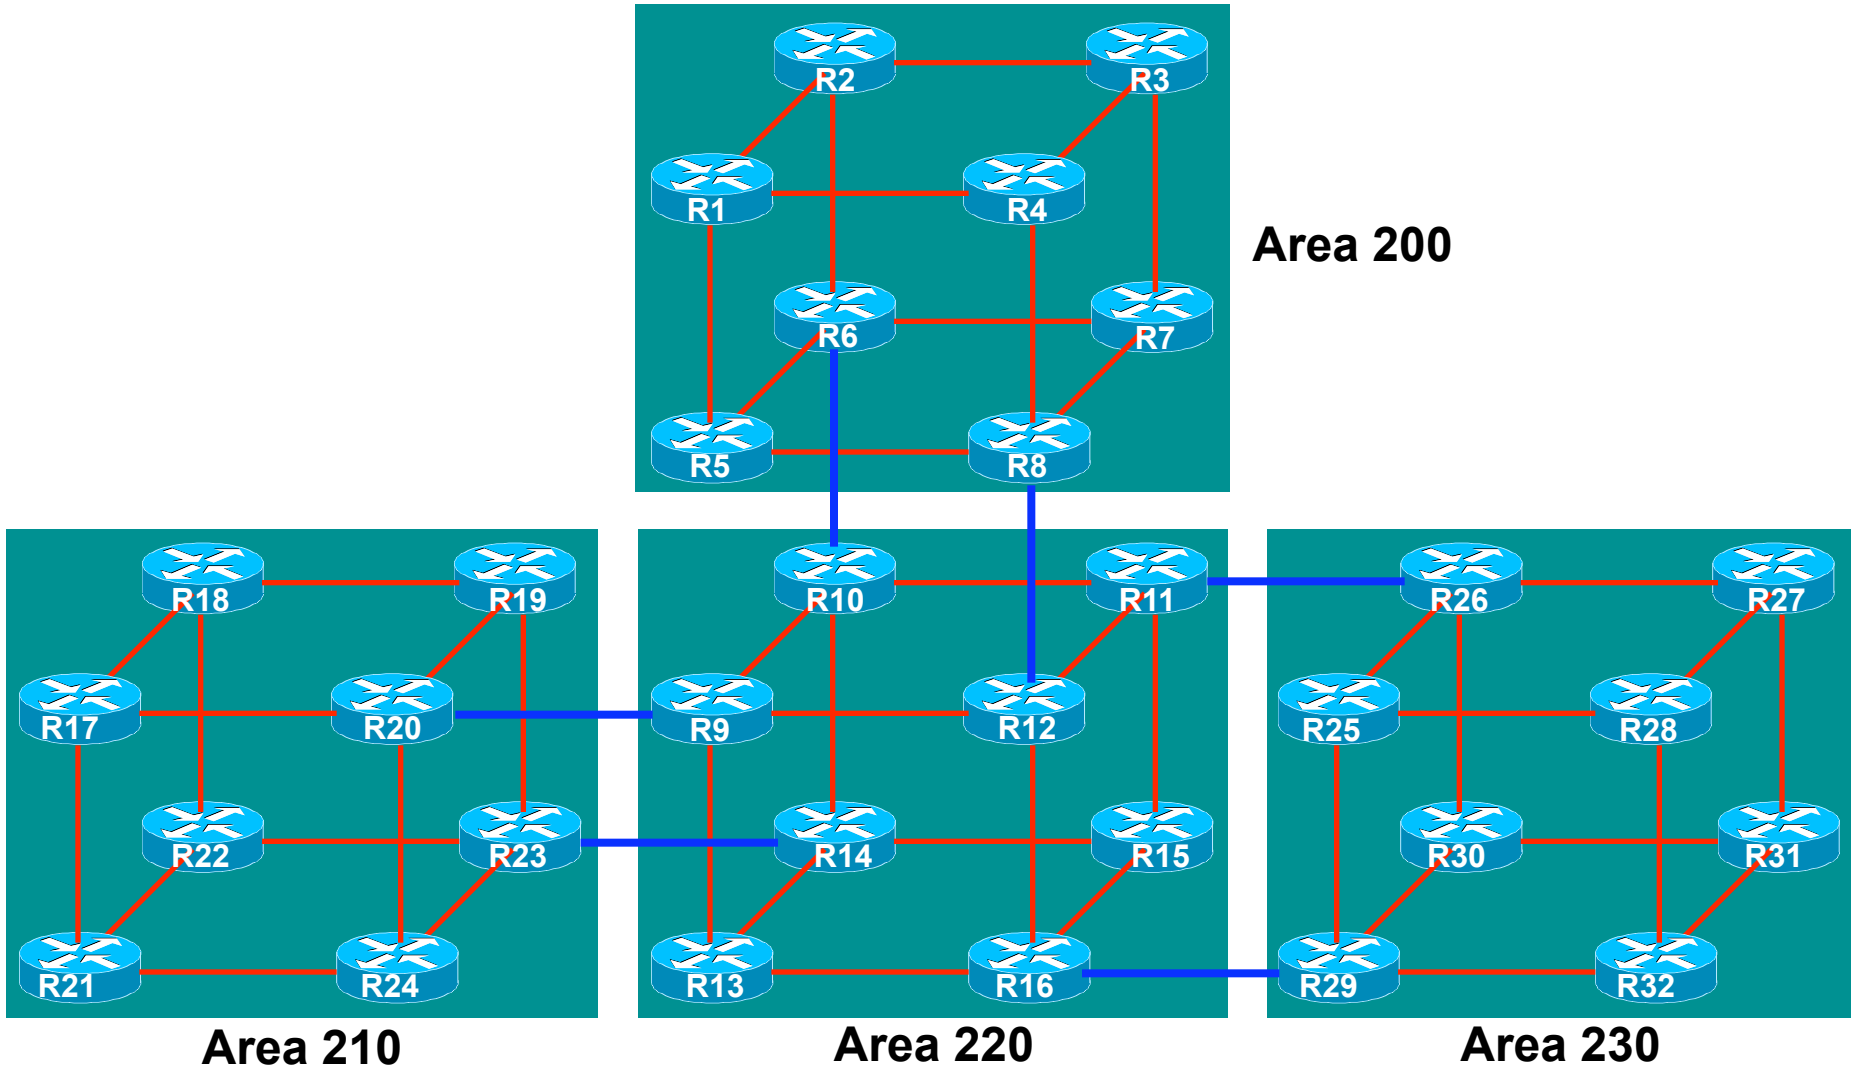

IPv4 Addressing Plan

Area Nets	Routers	Links
200.200.0.0/16	/22	/30
210.210.0.0/16	/22	/30
220.220.0.0/16	/22	/30
230.230.0.0/16	/22	/30

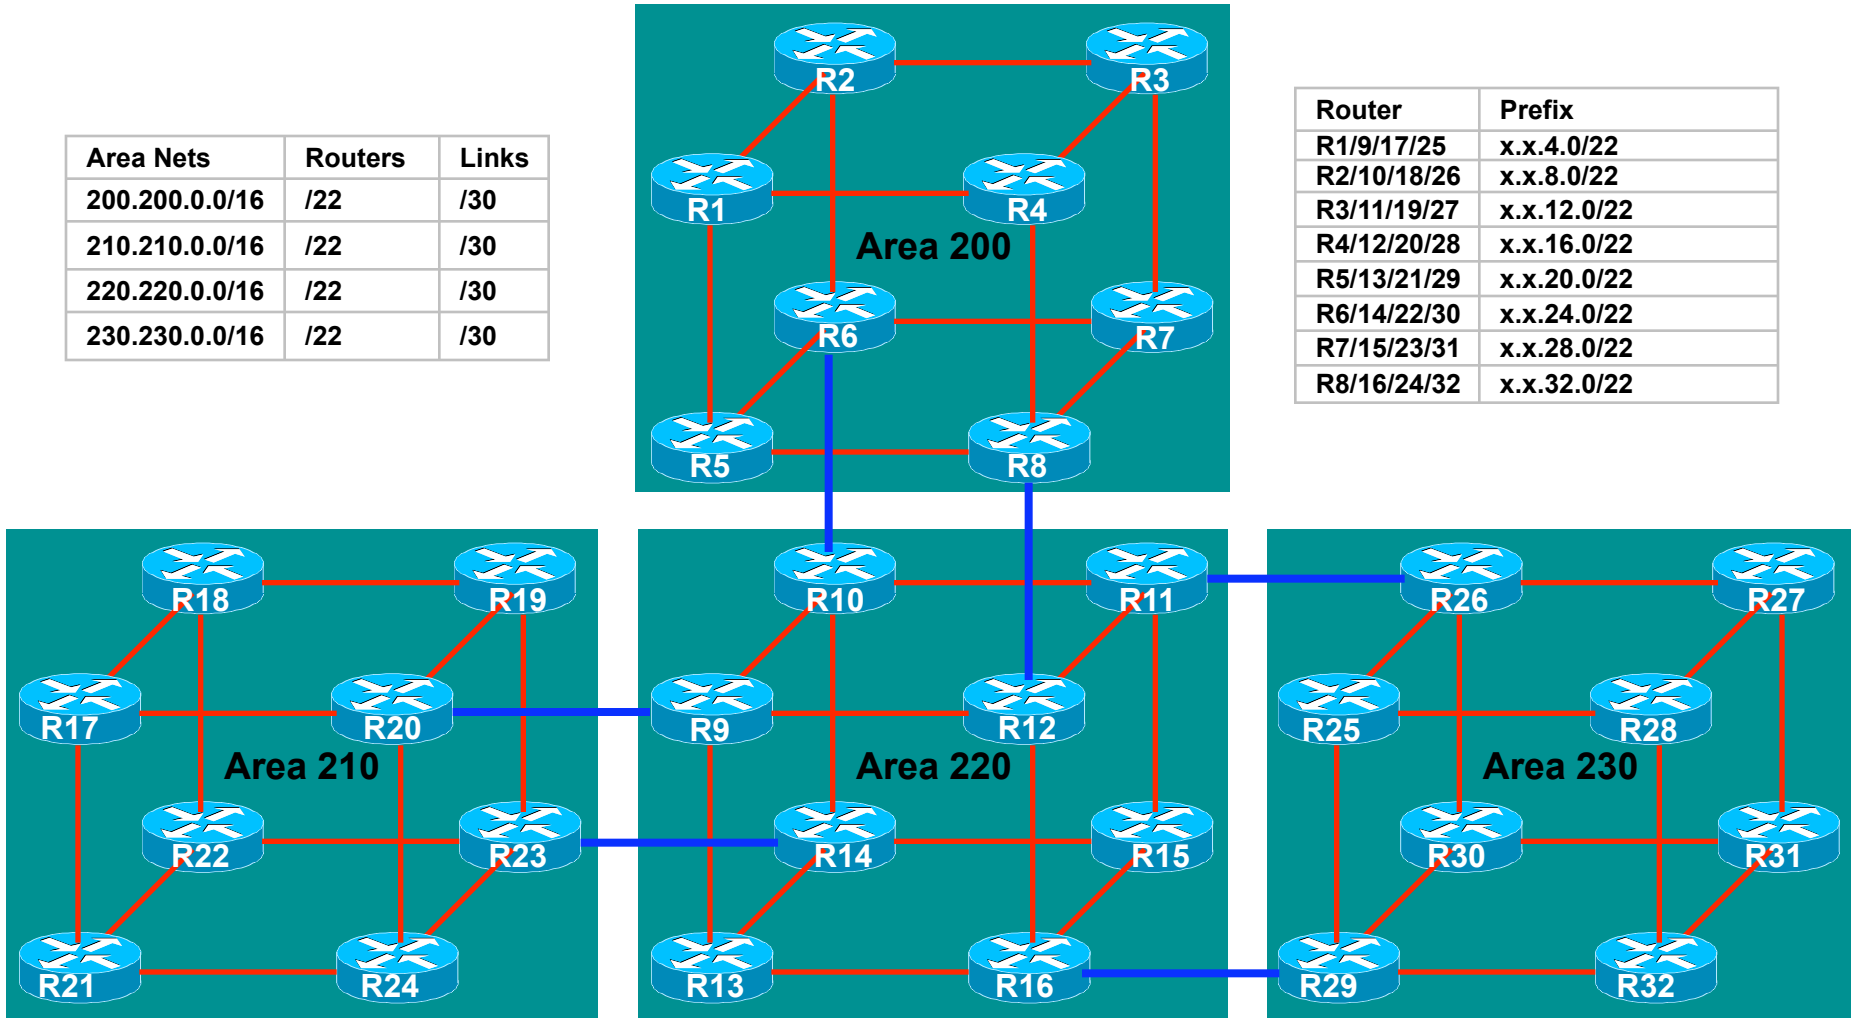

Router	Prefix
R1/9/17/25	x.x.4.0/22
R2/10/18/26	x.x.8.0/22
R3/11/19/27	x.x.12.0/22
R4/12/20/28	x.x.16.0/22
R5/13/21/29	x.x.20.0/22
R6/14/22/30	x.x.24.0/22
R7/15/23/31	x.x.28.0/22
R8/16/24/32	x.x.32.0/22